



THE
POWER
TO KNOW.

SAS[®] 9.2 OLAP Server User's Guide



The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2009. *SAS® 9.2 OLAP Server: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® 9.2 OLAP Server: User's Guide

Copyright © 2009 SAS Institute Inc., Cary, NC, USA

ISBN 978-1-59994-338-1

All rights reserved. Produced in the United States of America.

For a hard-copy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a Web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

U.S. Government Restricted Rights Notice. Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions set forth in FAR 52.227-19 Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st electronic book, February 2009

1st printing, February 2009

2nd electronic book, September 2009

2nd printing, September 2009

3rd electronic book, May 2010

3rd printing, May 2010

SAS® Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at support.sas.com/publishing or call 1-800-727-3228.

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

Contents

<i>What's New</i>	<i>ix</i>
Overview	ix
OLAP Cube Building Features	ix
OLAP Cube Maintenance Features	x
Cube Aggregation Tuning Features	xi
Incremental Cube Update Features	xi
Security For Cubes	xi
SAS Tree View Features	xii
New OLAP Procedure Options	xii
Chapter 1 △ OLAP Introduction and Overview	1
What Is OLAP?	1
What Is a Cube?	3
Understanding the Cube Structure	4
Why You Should Use Cubes	5
Chapter 2 △ Defining the SAS OLAP Environment	7
SAS OLAP Environment — Overview	7
SAS Metadata Server	8
SAS Application Server	9
SAS Workspace Server	9
SAS OLAP Server	9
SAS Management Console	10
SAS OLAP Server Monitor	10
SAS OLAP Cube Studio	10
OLAP Schema	10
SAS Libraries	11
Data Tables	11
Authorization Permissions	11
Chapter 3 △ Organization and Management of Your Data	13
Cube Data Organization	13
OLAP Schemas	14
SAS Libraries and Tables	16
Managing Folders In SAS OLAP Cube Studio	19
SAS OLAP Cube Jobs	20
Chapter 4 △ Planning for SAS OLAP Cubes	23
Overview	23
Data Tables Used to Define SAS OLAP Cubes	24
Aggregation Design	26
Aggregation Storage	28

SAS OLAP Cube Size Specifications	29
Naming Guidelines and Rules for the SAS OLAP Server	30
SAS Formats Available for Measures	32
Statistics Available for Measures	34
Chapter 5 \triangle Building Cubes and Administering Cubes	35
SAS OLAP Cube Studio and the OLAP Procedure	36
Connecting and Reconnecting to a Metadata Server	36
Cube Designer Wizard	38
Defining Data Sources For a Cube	38
Defining Drill-Through Tables	39
Defining Dimensions and Levels	40
Specifying the TIME Dimension	40
Specifying GIS Map Information for a Dimension	41
Defining Cube Hierarchies	42
Defining Measures For A Cube	47
Defining Member Properties	49
Defining Aggregations While Building a Cube	50
Saving the Cube Metadata or Creating the Physical Cube	51
Saving the OLAP Procedure (Long Form versus Short Form)	52
Viewing Cubes in SAS OLAP Cube Studio	53
Chapter 6 \triangle Modifying and Maintaining Cubes	55
Editing a Cube	56
Renaming A Cube Object	56
Deleting Cubes and Cube Objects	57
Refreshing Cube Metadata	57
Tuning Cube Aggregations	58
Specifying Tuning and Performance Options in Cube Aggregations	61
Specifying Calculated Members and Measures	63
Multiple Language Support and Dimension Table Translations	65
Adding SAS System Options to a Cube	66
Synchronizing Column Changes	67
Accessing OLAP Cubes from SAS: SQL Pass-Through Facility for OLAP	67
Impact Analysis and Reverse Impact Analysis	71
Disabling and Enabling Cubes	72
Cube Security	73
Chapter 7 \triangle Updating SAS OLAP Cubes	79
Overview	79
Updating a Cube In-Place	80
Incremental Updates of Cubes and Cube Generations	80
Coalescing Cube Aggregations	82
Updating a Cube In a Production Environment	82
Archiving and Deleting Cube Generations	84
Updating the Captions and Descriptions for a Cube	85

Adding New Members to an Incrementally Updated Cube	85
Reorganization of Cube Levels	85
Updating Member Properties	86
Specifying Drill-Through Tables	87
NWAY Considerations	87
Updating Multiple Language Support Cubes	87
Format Search Path and SAS Source Code Considerations	87
Exporting Cubes that Have Been Updated	88
Input Data Tables for Cube Updates	88
Schema and Repository Considerations	89
Physical Storage and Metadata Considerations	90
Connecting to a Workspace Server	90
Proc OLAP Options	90
Proc OLAPOPERATE Options and SAS OLAP Monitor	90
Updating Cubes in SAS OLAP Cube Studio	90
Chapter 8 \triangle Cube Building and Modifying Examples	93
Defining A Connection Profile	94
Building a Cube from a Detail Table	96
PROC OLAP Example for a Detail Table	110
Building a Cube from a Star Schema	114
Building a Cube from a Summary Table	133
Tuning Aggregations For a Cube	151
Adding Data to a Cube with Cube Update	160
Adding Calculated Members To A Cube	170
Setting Member Authorizations On A Dimension	180
Setting Identity Driven Security	187
Viewing a Cube in SAS OLAP Cube Studio	192
Creating a Time Dimension in SAS OLAP Cube Studio	197
Synchronizing Column Changes	199
Specifying an ESRI GIS Map For a Cube Dimension	200
Creating Multiple Hierarchies For a Cube	205
Set IGNORE_MISSING_DIMKEYS for a Star Schema	208
Implementing Drill-through to Detail Data in a SAS OLAP Cube	210
Exporting a Cube From SAS OLAP Cube Studio	216
Importing a Cube Into SAS OLAP Cube Studio	219
Chapter 9 \triangle Using SAS OLAP Cubes	225
Using a Cube with ADO MD	225
Using a Cube with OLE DB for OLAP	225
Using a Cube with Additional SAS Products	226
Using a Cube with Third-Party Clients	229
Chapter 10 \triangle Importing and Exporting SAS OLAP Cubes	231
Importing and Exporting SAS OLAP Cubes	231
ExportCubes and ImportCubes Batch Tools	232

Export SAS Package and Import SAS Package	235
Validating Data After It Is Moved	240
Cube Promotion and Migration Resources	240

Appendix 1 \triangle The OLAP Procedure 243

The OLAP Procedure	244
Syntax: OLAP Procedure	244
PROC OLAP Statement	245
METASVR Statement	256
DIMENSION Statement	257
LEVEL Statement	263
PROPERTY Statement	266
HIERARCHY Statement	267
MEASURE Statement	269
AGGREGATION Statement	274
DROP_AGGREGATION Statement	276
DEFINE Statement	276
UNDEFINE Statement	278
USER_DEFINED_TRANSLATIONS Statement	279
REORGANIZE_LEVEL Statement	281
Loading Cubes	281
Maintaining Cubes	286
Specialized Options for PROC OLAP	288

Appendix 2 \triangle The OLAPOPERATE Procedure 291

The OLAPOPERATE Procedure	291
Syntax: OLAPOPERATE Procedure	292
Connection Options	292
Disconnecting from an OLAP Server	293
Listing Active Sessions on an OLAP Server	293
Listing Session Queries on an OLAP Server	294
Listing Session Rowsets on an OLAP Server	294
Closing a Session	295
Cancelling a Query Result Set	296
Disabling a Cube	296
Enabling a Cube	296
Enabling and Disabling Cubes for a Cube Update	296
Refreshing a Cube	296
Stopping an OLAP Server	297

Appendix 3 \triangle Resources and Reference Topics for SAS OLAP Cubes 299

OLAP Server Resources	299
OLAP Schema Resources	300
Cube Resources	300
Table Resources	300
Dimension Resources	301

Level Resources	302
Hierarchy Resources	302
Measure Resources	303
Member and Member Properties Resources	303
Calculated Members and Measures Resources	303
Aggregation Resources	304
Incremental Cube Update Resources	304
Drill-Through to Detail Resources	305
GIS Mapping Resources	305
Appendix 4 \triangle SAS OLAP Cube Studio Accessibility Features	307
SAS OLAP Cube Studio Accessibility Features	307
Appendix 5 \triangle Recommended Reading	309
Recommended Reading	309
Glossary	311
Index	319

What's New

Overview

The SAS OLAP Server contains various new functions, tools, and features for SAS 9.2.

- OLAP cube building features
- OLAP cube maintenance features
- cube aggregation tuning features
- incremental cube update features
- security for cubes
- SAS tree view features
- OLAP procedure options

These functions are accessed from SAS OLAP Cube Studio. In addition, some of these functions have equivalent OLAP procedure statements or options.

OLAP Cube Building Features

- The data table selection process in the Cube Designer wizard has been simplified and reorganized into an easier and more intuitive order. Tables are now grouped based on the libraries that the tables are registered to. The table selection pages also contain functions that enable you to access the properties for a table and view the data for a table.
- You can now create multiple hierarchies for a dimension with the Define a Hierarchy dialog box in the Dimension Designer. You can also designate one of the hierarchies as a default hierarchy.
- You can now build a time dimension with user-supplied time hierarchies that auto-populate the levels and formats for the Time dimension. The Add-Supplied dialog box enables you to select from suggested time hierarchies that are based on a single date-formatted column.

The COLUMN= option is used on the LEVEL statement when the Add-Supplied function is used.

- If the metadata server connection is lost during a SAS OLAP Cube Studio session, a new metadata server connection is automatically attempted during the same SAS OLAP Cube Studio session. This enables you to continue your work in that same session without losing any data changes you are making.

OLAP Cube Maintenance Features

- The Calculated Members Wizard is now fully integrated into SAS OLAP Cube Studio. The addition of the advanced expression builder enables you to specify user-defined formats, set the solve order, and build a custom MDX formula for a calculated member.
- The View Cube function enables you to verify the cube build process and visualize the data structure of a cube.
- SAS OLAP cube metadata can now be exported and imported as part of a SAS package or in batch mode. Cubes built with SAS OLAP Cube Studio 4.2 can be copied or moved between SAS 9.2 systems with the Export SAS Package and Import SAS Package functions that are part of SAS OLAP Cube Studio 4.2. Cubes that were built with a version of SAS that is earlier than 9.2 can also be moved to the SAS 9.2 environment by using the ExportCubes and ImportCubes batch tools.

Note: The ExportCubes batch tool is supported in a hot fix that is applied to SAS 9.1.3 SP4. Cube metadata that is exported with the ExportCubes tool can then be imported in the second maintenance release after SAS 9.2 with the ImportCubes batch tool. △

- The Synchronize Column Changes function enables you to synchronize a cube when the input table for an existing cube has encountered a column name change. This function finds the name differences between the cube and its input table. It then changes the internal cube files to match the input table column names.
- The Refresh Cube function enables the SAS OLAP Server to access calculated members that were created or updated since the cube was last accessed by the SAS OLAP Server. This function notifies running SAS OLAP Server instances to capture calculated member changes made to the cube since it was last loaded by those instances.
- You can now quickly change the OLAP schema for a cube from within SAS OLAP Cube Studio.
- Cubes can now be disabled to make necessary changes and updates and then enabled again. Disabling a cube makes it transparent to the users connected to any active OLAP Server instances. A cube is disabled in order to make edits to the cube or to perform administrative tasks on the cube. Enabling a cube brings it back into production.
- You can now choose to export the code generated by PROC OLAP in either a long form or short form. The Export Code function enables you to store in a text file the PROC OLAP code that is used to edit or build the cube. The long form contains all PROC code that is used to create a cube. The short form text file contains only the PROC OLAP statement, the cube name, and the METASVR statement and obtains the remainder of the cube structure from an existing registration on the metadata server.
- You can now perform impact analysis for a cube. This function enables you to view the relationship of a cube to the objects that it is associated with. Impact analysis shows the potential impact of changes that you might make to a cube.
- A cube or cube-related object can now have documents and notes attached to provide additional information about the cube or object.

- Cube jobs are now automatically generated when a cube is built. A cube job is a collection of SAS tasks that create output. When you deploy a cube job, the code needed to create the cube is generated.
- You can now change the name of a cube or other object (such as a job, document, or note) by selecting that object in the tree view and then selecting the Rename function.

Cube Aggregation Tuning Features

- The newly designed Aggregation Tuning wizard enables you to build, edit, and customize aggregations for a cube. You can generate new aggregations for a cube based on the ARM log or cardinality, or you can use the manual aggregation builder.

Incremental Cube Update Features

- You can now add data and members to a cube without having to rebuild the cube. To support this feature, new PROC OLAP statements and options have been added. Within SAS OLAP Cube Studio, the Incremental Update wizard has been added.
- The Coalesce Aggregations function is part of the Incremental Update function. It enables you to combine all of the individual rack tables that result from incremental updates into a single aggregation table. The COALESCE_AGGREGATIONS option has been added to support this function. In SAS OLAP Cube Studio, the Coalesce Aggregations wizard has been added.

Security For Cubes

- You can now define member-level security in SAS OLAP Cube Studio. The Add Authorization dialog box enables you to add MDX expressions that limit access to a cube dimension. This ensures that only designated portions of the data are visible to a user or group of users.

Starting in the third maintenance release for SAS 9.2, you can use the Add Authorization dialog box to grant or deny Read permissions for a member only or for a member and its descendents.

The SECURITY_SUBSET= option is now available in the OLAP procedure. The option determines whether security totals will be used. The subset that is defined is then applied when calculating the cell values.

- Identity-driven security enables you to substitute identity values in a permission condition. It enables you to insert a placeholder into the permission condition that, at query time, gets resolved to a string that represents the user identity.
- Starting in the third maintenance release for SAS 9.2, you can specify batch security in SAS OLAP Cube Studio with the Manage Permission Tables function. The Manage Permission Tables function enables you to create a special SAS data set known as a permission table that contains cube access controls for submitting in bulk. A permission table is a table of access control information that can later be applied to a cube with batch SAS code. The Manage Permission Tables dialog box enables you to create and modify permission tables as well as import access controls (permissions) from a cube or an OLAP schema. You can also execute the code interactively or export the code to a file for use in a stored process or deployed job flow.

SAS Tree View Features

- The tree view is now divided into two tabs, the **Folders** tab and **Inventory** tab. The **Inventory** tab lists the cubes and their components, such as schemas, tables, and libraries. The **Folders** tab displays a folder view of the components of the cube. This view shows which folders a cube, source tables, and libraries are located in.
- You can now access functions that are specific to an object in the tree view. The functions that are listed in the context menu are specific to that object and the current status of that object. You can access the context menu by right-clicking an object.
- In SAS OLAP Cube Studio, you can now copy or move objects between folders. You can also find the folder location of an object in the **Inventory** tab with the Find In Folders function.

New OLAP Procedure Options

The following OLAP procedure options are new for SAS OLAP 9.2:

PROC OLAP Statement Options

- ADD_DATA
- ASYNCINDEXLIMIT=
- COALESCE_AGGREGATIONS
- ESRI_MAP_SERVER=
- ESRI_REPLACE
- MAX_RETRIES=
- MAX_RETRY_WAIT=
- MIN_RETRY_WAIT=
- NONUPDATEABLE
- OUTCUBE=
- OUTSCHEMA=
- RENAME
- SECURITY_SUBSET= YES|NO
- SYNCHRONIZE_COLUMNS (replacing SYNCHRONIZE_LEVELS)
- UPDATE_DISPLAY_NAMES
- UPDATE_IN_PLACE

Dimension Statement Options

- MAP_SERVICE=
- UPDATE_DIMENSION=

LEVEL Statement Options

- COLUMN=


- ESRI_MAP_LAYER=
- FORMAT=

MEASURE Statement Options

- INCLUDE_CALCULATED_MEMBERS
- NOINCLUDE_CALCULATED_MEMBERS

PROPERTY Statement Options

- ESRI_MAP_FIELD=



CHAPTER

1

OLAP Introduction and Overview

<i>What Is OLAP?</i>	1
<i>Data Storage and Access</i>	1
<i>Benefits of OLAP</i>	2
<i>OLAP Variations</i>	3
<i>What Is a Cube?</i>	3
<i>Understanding the Cube Structure</i>	4
<i>Dimensions, Levels and Hierarchies</i>	4
<i>Members</i>	4
<i>Measures</i>	4
<i>Calculated Measures</i>	5
<i>Aggregations</i>	5
<i>Why You Should Use Cubes</i>	5
<i>Cube Usage and Storage Space Reduction</i>	5
<i>Multi-Threading Capabilities</i>	5
<i>Easy Setup and Maintenance</i>	6
<i>Data Management: Choosing Your Own Tool</i>	6

What Is OLAP?

Online Analytical Processing (OLAP) is a technology that is used to create decision support software. OLAP enables application users to quickly analyze information that has been summarized into multidimensional views and hierarchies. By summarizing predicted queries into multidimensional views prior to run time, OLAP tools provide the benefit of increased performance over traditional database access tools. Most of the resource-intensive calculation that is required to summarize the data is done before a query is submitted.

Data Storage and Access

Decision makers are asked to make timely and accurate decisions that are based on the past performance and behavior of an organization as well as on future trends and directives. To make effective business decisions, business analysts must have access to the data that their company generates and responds to. This access must include timely queries, summaries, and reviews of numerous levels and combinations of large, recurrent amounts of data. The information that business analysts review determines the quality of their decisions.

Organizations usually have databases and data stores that maintain repeated and frequent business transaction data. These data storage systems provide simple yet detailed storage and retrieval of specific data events. However, these systems are not

well suited for analytical summaries and queries that are typically generated by decision makers. For decision makers to reveal hidden trends, inconsistencies, and risks in a business, they must be able to maintain a certain degree of momentum when querying the data. An answer to one question usually leads to additional questions and review of the data. Simple data stores do not successfully support this type of querying.

A second type of storage, the data warehouse, is better suited for this. Data is maintained and organized so that complicated queries and summaries can be run. OLAP further organizes and summarizes specific categories and subsets of data from the data warehouse. This results in a robust and detailed level of data storage with efficient and fast query returns. SAS OLAP cubes can be built from either partially or completely denormalized data warehouse tables. Stored, precalculated summarizations called *aggregations* can be added to the cube to improve cube access performance. Aggregations can either be pre-built relational tables, or you can let the cube create its own optimized aggregates.

Benefits of OLAP

The ability to have coherent, relevant, and timely information is the reason OLAP has gained in popularity. OLAP systems can help reveal evasive inconsistencies and trends in data that might not have been seen before. OLAP users can intuitively search data that has been consolidated and summarized within the OLAP structure. In addition, OLAP tools allow for tasks such as sales forecasting, asset analysis, resource planning, budgeting, and risk assessment. OLAP systems also provide the following benefits:

- fast access, calculations, and summaries of an organization's data
- support for multiple user access and multiple queries
- the ability to handle multiple hierarchies and levels of data
- the ability to presummarize and consolidate data for faster query and reporting functions
- the ability to expand the number of dimensions and levels of data as a business grows

To fully understand the benefits of OLAP and the details of its effective implementation, it helps to examine the technology from two perspectives— that of the users, and that of the information technology (IT) administrators who are responsible for OLAP implementation. The users, typically business analysts and executives, expect the data to be organized according to categories that reflect the way they think about the enterprise. For IT administrators, OLAP can present a long list of technical issues, including these concerns:

- storage requirements and associated costs
- client and server capabilities
- maintenance activities such as update and backup
- performance considerations such as the amount of time that is required to build a multidimensional model
- the ability of the OLAP solution to integrate with current or planned data warehouse strategies and architectures
- security requirements for cube data

The SAS OLAP Server and SAS OLAP Cube Studio provide resources and functionality to address these concerns. When building SAS OLAP cubes, you can perform functions and specify settings that affect the following:

- cube aggregation storage and query performance
- cube dimension security and identity-driven security

- updates of cube data
- maintenance of cubes (such as adding calculated members or changing an OLAP schema.)

Because SAS OLAP is a component of the SAS Intelligence Platform, it works in conjunction with other SAS applications to provide an overall solution to the access and maintenance of a company's data.

OLAP Variations

OLAP technology can be further defined by the methods for storing and accessing data and by the performance of queries against that data. SAS OLAP supports three different variations of OLAP technology:

MOLAP (multidimensional OLAP)

is a type of OLAP that stores summaries of detail data (aggregates) in multidimensional database structures. MOLAP cubes are most suited for slicing and dicing of data and are used when performance and query speed is critical. For further information on MOLAP, see the topic “MOLAP Aggregation Storage” on page 28.

ROLAP (relational OLAP)

is a type of OLAP in which multidimensional data is stored in a relational database such as a SAS table or an ORACLE table. ROLAP is more scalable than other OLAP types and handles extensive amounts of data well. Although performance can be somewhat slow, ROLAP is limited only by the size of the relational database it is identified with. For further information on ROLAP, see the topic “ROLAP Aggregation Storage” on page 28.

HOLAP (hybrid OLAP)

is a type of OLAP in which relational OLAP (ROLAP) and multidimensional OLAP (MOLAP) are combined. In HOLAP, the source data is usually stored using a ROLAP strategy, and aggregations are stored using a MOLAP strategy. It combines the best features of both ROLAP and MOLAP. This combination usually results in the smallest amount of storage space. In HOLAP, aggregates can be precalculated and can be linked into a hybrid storage model.

What Is a Cube?

One of the advantages of OLAP is how data and its relationships are stored and accessed. OLAP systems house data in structures that are readily available for detailed queries and analytics. Cubes are central to the OLAP storage process.

A *cube* is a set of data that is organized and structured in a hierarchical, multidimensional arrangement. The cube is usually derived from a subset of a data warehouse. Unlike relational databases that use two-dimensional data structures (often in the form of columns and rows in a spreadsheet), OLAP cubes are logical, multidimensional models that can have numerous dimensions and levels of data. Also, an organization typically has different cubes for different types of data.

One of the challenges of OLAP cube data storage and retrieval is the growth of data and how that growth affects the number of dimensions and levels in a cube hierarchy. As the number of dimensions increases over time, so does the number of data cells on an exponential scale. To maintain the efficiency and speed of the OLAP queries, the cube data is often presummarized into various consolidations and subtotals (aggregations).

Note: The SAS OLAP Server term *cube* is synonymous with the terms *hyper-cube* and *multi-cube*. \triangle

Understanding the Cube Structure

A SAS OLAP cube stores data in a method that enables fast retrieval of summarized data. Data summarization in this context means condensing large numbers of detail records into meaningful numbers such as counts, sums, averages, or other statistical measures. The structure of a cube is hierarchical in nature and is derived from the associations between the different columns and rows of data in a data source. SAS OLAP cubes are comprised of dimensions, levels, hierarchies, members, and member properties. This structure enables you to easily select data subsets and navigate the cube structure when querying the cube.

Dimensions, Levels and Hierarchies

SAS OLAP cubes organize data in a hierarchical arrangement, according to dimensions and measures. *Dimensions* group the data along natural categories and consist of one or more levels. Each level represents a different grouping within the same dimension. For example, a time dimension can include levels such as years, months, and days. Or an organization dimension of a bank's customer service centers can include levels such as branches, states, and regions.

Levels are organized into one or more *hierarchies*, typically from a coarse-grained level (for example, Year) down to the most detailed one (for example, Day). The individual category values (for example, 2002 or 21Jan2002) are called *members*.

A dimension can also have *multiple hierarchies* to provide different sequences of groupings. For example, a "Time" dimension can have a "Fiscal Year" hierarchy and a "Calendar Year" hierarchy.

Members

Each combination of values within a dimension is called a member. Some examples of members are shown here.

[Time].[2003]

[Time]. [2004].[January]

[Time]. [2004].[February].[12th]

For each dimension, there is also the special member called the *ALL member*, which represents the total for all members (for example, [Time].[All Time]). Not all categorical data attributes need to become a member of a hierarchy level. Some grouping information is needed only as additional information for a member or for applying subsets to data. These attributes can be loaded into *member properties*. Member properties can be associated with any level in a hierarchy.

Measures

Measures are the cube data values that are summarized and analyzed. A measure is the combination of a numeric input column with a roll-up rule or statistic. Measures are loaded from the data source that you summarize from. One input column can load one or more measures. For example, you can create the measures "Sum of Amount" and "Maximum of Amount" from the input column "Amount."

Calculated Measures

Not all measures are directly derived from input columns. You can create *calculated measures*, which are formulas that are based on the values of other measures. Calculated measures are designed and stored with the cube.

Aggregations

An aggregation is a summary of detail data that is stored with or referred to by a cube. They are the basis for fast response to data queries in OLAP applications. An aggregation is possible at each intersection of a level of one or more dimensions. Any combination of dimension levels can become a stored aggregation, as long as it is appropriate within the defined hierarchies. One of the major factors that influences query response time is which aggregations you create and use to query your cube. The aggregations that are being stored with the cube affect cube build time, the absolute cube file size, SAS OLAP Server CPU usage, and query response times. As a result, determining and building your cube aggregations is a crucial component of good cube design.

Why You Should Use Cubes

SAS cubes are designed to offer efficient data storage, fast data access, easy data maintenance, and flexibility in data management. The following sections explore cubes and multidimensional storage.

Cube Usage and Storage Space Reduction

While cubes are the format of choice to guarantee fast query response times against your data warehouse, SAS OLAP cubes are also often a very space-efficient choice for data storage. In many cases, a basic cube without additional aggregations can be smaller than the input data because the process of creating the cube consolidates records. SAS OLAP cubes use the hierarchy information for efficient aggregations storage. SAS OLAP cubes also deal efficiently with data sparsity by using virtual placeholders for empty cells. This removes the need for any physical representation of empty cells. A good rule of thumb is, the larger your input data, the greater the storage gain by loading data into a cube.

Multi-Threading Capabilities


Loading data into cubes and executing queries against the cube take advantage of the multi-threading capabilities of your server machine. Aggregations are created in parallel at cube build time. The creation of individual aggregations takes advantage of the Parallel Group-By capabilities of the SAS data engine. At query execution, the multi-threading capabilities of your server machine are fully used to concurrently serve queries by multiple users. Both query evaluation and data access are executed in parallel. To further increase query performance and reduce disk access, you can allocate additional memory on your server to be used for an in-memory aggregation cache.

Easy Setup and Maintenance

A cube is the physical representation of your logical dimensional model. The tools that are provided to update and maintain the cube reflect the multidimensional model, which makes both setup and maintenance of your cube as intuitive as possible. SAS Management Console, a Web-based administrator interface, enables you to set up and manage OLAP servers. SAS OLAP Cube Studio provides the workspace and cube designer tools that you need to create and maintain cubes. You can also use the SAS OLAP procedure to create and maintain cubes in a batch environment.

Data Management: Choosing Your Own Tool

If you create your own aggregations by using data management tools such as SQL, PROC SUMMARY, or the tools of your preferred relational database management system (RDBMS), then you can link those aggregations to your cubes without replicating the data within the cube. Any queries against those aggregations are executed by the appropriate SQL engine, and take advantage of any capabilities that engine might have. This allows you the flexibility to use the data management tools of your choice. It also allows you to distribute your data for your cube aggregations across multiple database systems, servers, and platforms. If you choose to let the cube builder create the aggregations, then you can control where to store the data and index files for each aggregation.



CHAPTER

2

Defining the SAS OLAP Environment

<i>SAS OLAP Environment — Overview</i>	7
<i>SAS Servers</i>	7
<i>SAS Clients</i>	8
<i>Data Organization and Management Components</i>	8
<i>User and User Group Authorization Permissions</i>	8
<i>SAS Metadata Server</i>	8
<i>SAS Application Server</i>	9
<i>SAS Workspace Server</i>	9
<i>SAS OLAP Server</i>	9
<i>SAS Management Console</i>	10
<i>SAS OLAP Server Monitor</i>	10
<i>SAS OLAP Cube Studio</i>	10
<i>OLAP Schema</i>	10
<i>SAS Libraries</i>	11
<i>Data Tables</i>	11
<i>Authorization Permissions</i>	11

SAS OLAP Environment — Overview

Before you can begin developing SAS OLAP cubes, you must define and install the servers, tools, and user interfaces that enable you to create and manage your SAS OLAP cubes and related data. When you install your SAS OLAP environment, you will need to plan for and include the following items:

SAS Servers

SAS OLAP uses a combination of SAS servers to store cube metadata, to store the physical cube structure, and to query cubes after they are created. The term *server* refers to a program or programs that wait for and fulfill requests from client programs for data or services. Several types of SAS servers are available to handle different workload types and processing intensities.

- a “SAS Metadata Server” on page 8
- a “SAS Application Server” on page 9
- a “SAS Workspace Server” on page 9
- a “SAS OLAP Server” on page 9

SAS Clients

The SAS OLAP environment uses the following Java desktop applications:

- “SAS Management Console” on page 10
- “SAS OLAP Server Monitor” on page 10
- “SAS OLAP Cube Studio” on page 10

Data Organization and Management Components

The following data organization and management components are necessary for and used by the SAS OLAP environment.

- “OLAP Schema” on page 10
- “SAS Libraries” on page 11
- “Data Tables” on page 11

User and User Group Authorization Permissions

- “Authorization Permissions” on page 11.

SAS Metadata Server

The SAS Metadata Server is a multi-user server that enables users to read metadata from and write metadata to one or more SAS metadata folders. This server is a centralized resource for storing, managing, and delivering metadata for all SAS applications across the enterprise. It enables all users to access consistent and accurate data. Some types of metadata objects that can be stored in the SAS metadata include data libraries, tables, cubes, users and user groups, user authorization permissions, and server definitions. The default metadata server that is installed during your system installation and configuration is the SASMeta server.

You can define and administer SAS Metadata Servers in SAS Management Console. The Server Manager and Metadata Manager plug-ins contain various administrative functions. The Server Manager provides functions to manage SAS server definitions stored in the SAS Metadata. It enables you to define and edit information regarding server locations and connections. The Metadata Manager plug-in enables you to maintain and define the various SAS metadata for your SAS environment.

Note: For more information on the SAS Metadata Server, see the topic "SAS Metadata Server" in the *SAS Intelligence Platform: Overview*. \triangle

The SAS Metadata Server uses the Integrated Object Model (IOM) that is provided by SAS Integration Technologies. IOM provides distributed object interfaces to Base SAS software features. It enables you to use industry-standard languages, programming tools, and communication protocols to develop client programs that access these services on IOM servers. The SAS Open Metadata Interface (OMI) is an object-oriented application programming interface (API) that interacts with the SAS Metadata Server. SAS OLAP Cube Studio is an example of an application that is compliant with the SAS Open Metadata Interface.

With SAS OLAP Server, all relevant structural information is contained within the cube and most of it is also replicated within the SAS Open Metadata Architecture. This is done so you can do the following:

- disassociate the cube definition process from cube creation, thus enabling you to create a cube by using its stored definition
- define and enforce security at the SAS Open Metadata Architecture level
- manage and control the data source in the centralized SAS Metadata Repository

You can find documentation about the SAS Open Metadata Architecture in the *SAS Intelligence Platform: System Administration Guide* and the *SAS Intelligence Platform: Data Administration Guide*.

SAS Application Server

During installation and setup, a metadata object is defined that represents the SAS server tier in your environment. This is called the SAS Application Server. A SAS Application Server is not an actual server that can execute SAS code submitted by clients. Rather, it is a logical container for a set of application server components, which do execute code. It is identified as SASApp in the SAS Management Console interface. For a SAS OLAP installation, a SAS Workspace Server and a SAS OLAP Server are contained in the application server.

Note: For further information on the SAS Application Server, see “Understanding the SAS Application Server” in the *SAS Intelligence Platform: Application Server Administration Guide*. △

SAS Workspace Server

The SAS Workspace Server provides access to SAS software features such as the SAS language, SAS libraries, and the server file system. The SAS Workspace Servers interact with SAS by creating a server process for each client connection. Each server process (workspace) represents a Foundation SAS session. This enables client programs to access SAS libraries, perform tasks by using the SAS language, and retrieve the results.

A program called the object spawner runs on a workspace servers host machine. The spawner listens for incoming client requests and launches server instances as needed. For further information on SAS Workspace Servers, see "Understanding Workspace Server and Stored Process Servers" in the *SAS Intelligence Platform: Application Server Administration Guide*.

SAS OLAP Server

The SAS OLAP Server is a scalable server that provides multi-user access to the data that is stored in SAS OLAP cubes. The purpose of the SAS OLAP Server is to respond to queries from cube viewers. SAS OLAP queries are performed by using the Multidimensional Expressions (MDX) query language in client applications that are connected to the OLAP server by using OLE DB for OLAP (an extension of OLE DB that is used by COM-based clients), or through a similarly designed Java interface. A SAS OLAP Server runs in the background on specially configured host computers. It processes data by using a multi-threaded kernel that enables you to take advantage of your server’s parallel processing abilities.

Note: For further information on the SAS OLAP Server, see "Administering SAS OLAP Servers" in the *SAS Intelligence Platform: Application Server Administration Guide*. Δ

SAS Management Console

SAS Management Console provides a single point of access and control to perform the administrative tasks required to create and maintain an integrated SAS environment across multiple platforms. SAS Management Console utilizes a plug-in architecture that is specific to your SAS installation and environment. With SAS Management Console you can perform the following tasks:

- manage the metadata for your SAS data and environment
- maintain library and server definitions
- create, manage and maintain authorization information for your users and groups
- administer resource access control templates

Note: For further information on SAS Management Console see the Help in SAS Management Console. Δ

SAS OLAP Server Monitor

The SAS OLAP Server Monitor plug-in for SAS Management Console enables you to manage tasks that are specific to the SAS OLAP Server. You can configure SAS OLAP Servers, administer the connection to a SAS OLAP Server, monitor sessions that are running on a SAS OLAP Server, and monitor queries that are being processed by a SAS OLAP Server. You can also enable and disable cubes and refresh cube metadata. For further information on SAS OLAP Server Monitor, see the Help for SAS OLAP Server Monitor in SAS Management Console.

SAS OLAP Cube Studio

SAS OLAP Cube Studio provides cube designers with an easy-to-use graphical user interface for creating SAS OLAP cubes. You can build SAS OLAP cubes, edit cubes, incrementally update cubes, tune aggregations, and make various other modifications to existing cubes. The *SAS OLAP Server: User's Guide* covers the different functions and features available for creating SAS OLAP cubes. See Chapter 5, "Building Cubes and Administering Cubes," on page 35 for more information.

OLAP Schema

A SAS OLAP schema is an organization container for SAS OLAP cubes. A SAS OLAP Schema specifies the location of a set of cubes. It communicates with a SAS OLAP Server about which cubes can be accessed by the server and then queried. When you install your SAS OLAP Server, a default SAS OLAP schema is created named *SASMain - OLAP Schema*. If needed, you can have multiple schemas. However, you can only have one schema assigned to an OLAP server at a given time. For more information on

SAS OLAP schemas, see “OLAP Schemas” on page 14. Also, see “Managing OLAP Cube Data” in the *SAS Intelligence Platform: Data Administration Guide*.

SAS Libraries

A SAS library is a collection of one or more files that are referenced and stored as a unit. For SAS OLAP, libraries serve as organization containers for the data tables that you use to create SAS OLAP cubes. You can have numerous libraries to organize your data tables with. To define your libraries, you can use the New Library wizard in SAS OLAP Cube Studio or in SAS Management Console. For more information on defining and managing libraries, see “SAS Libraries and Tables” on page 16 and “Assigning Libraries” in the *SAS Intelligence Platform: Data Administration Guide*.

Data Tables

Data tables are frequently used to define data stores. They can be used to define data stores, summary data, a join, or a table that holds information that does not conform to any other data storage type. A detail, or base, table is any data table defined in the SAS metadata that contains the measures and levels for a cube. The detail table consists of unsummarized data that must include one column for each level and one numeric analysis column for each set of measures that will be generated. There are different types of data tables that can be used to build a SAS OLAP cube. You can use the following tables to construct SAS OLAP cubes:


- detail tables (base tables)
- star schemas (a fact table and associated dimension tables)
- summary tables (N-way data set)
- drill-through tables
- aggregation tables

For detailed information on the different data tables, see “Data Tables Used to Define SAS OLAP Cubes” on page 24. Also, see the *SAS Intelligence Platform: Data Administration Guide* for further information on managing data table sources.

Authorization Permissions

When creating and editing SAS OLAP cubes, you must establish and apply security permissions to ensure the correct access to cube data among the different users on a system. SAS authorization permissions can be set at general levels or fine-grain levels and are conveyed across user identities (user, group) and objects (folder, table, library). Authorization permission settings are then inherited from parent identities or objects.

SAS OLAP cubes and the different components of a cube (dimensions, hierarchies, levels, measures, calculated measures, and cube jobs) can have authorization permissions set using the authorization functions that are available in SAS OLAP Cube Studio and in the Authorization Manager plug-in in SAS Management Console. For further information on establishing security and authorization permissions for your SAS OLAP environment, see “Cube Security” on page 73. In addition, see the *SAS Intelligence Platform: Security Administration Guide* for detailed information on security administration for your SAS environment.



CHAPTER

3

Organization and Management of Your Data

<i>Cube Data Organization</i>	13
<i>OLAP Schemas</i>	14
<i>Defining an OLAP Schema</i>	14
<i>Editing an OLAP Schema</i>	14
<i>Changing the OLAP Schema For a Cube</i>	15
<i>Deleting OLAP Schemas</i>	15
<i>Viewing the OLAP Schema Properties</i>	15
<i>Assigning the OLAP Schema with the OLAP Procedure</i>	16
<i>SAS Libraries and Tables</i>	16
<i>Defining Libraries and Tables</i>	16
<i>Creating a New Library Definition for Source Data Tables</i>	16
<i>Defining Tables Used to Build Cubes</i>	17
<i>Renaming Tables</i>	18
<i>Renaming a Table From the Tree View</i>	18
<i>Renaming a Table from the Properties Dialog Box</i>	18
<i>Deleting Tables</i>	18
<i>Viewing Table Data in SAS OLAP Cube Studio</i>	19
<i>Managing Folders In SAS OLAP Cube Studio</i>	19
<i>SAS Metadata Folders</i>	19
<i>Finding an Object in a SAS Folder</i>	19
<i>Copying an Object to a Different SAS Folder</i>	19
<i>Moving Objects Between SAS Folders</i>	19
<i>SAS OLAP Cube Jobs</i>	20
<i>SAS OLAP Cube Jobs</i>	20
<i>Cube Job Deployment and Redeployment</i>	20
<i>Job and Deployed Job Properties</i>	21
<i>Job Properties</i>	21
<i>Deployed Job Properties</i>	21

Cube Data Organization

When working with SAS OLAP Cube data you must manage the storage and organization of your cube input tables and the cube data, including the cube metadata and the corresponding physical cube files. During the course of developing and executing SAS OLAP Cubes, you can do the following:

- define the cube source data and libraries
- establish the OLAP schema assignment for your cubes
- create folders in which to store your cubes
- manage any cube jobs that you deploy

OLAP Schemas

A SAS OLAP schema represents a group of cubes in the SAS metadata and specifies which group of cubes an OLAP server can access. An initial OLAP schema is assigned to a SAS OLAP Server when that server is defined in the metadata. Although you can have multiple schemas, a server can only access the cubes in one schema at a time. As a result, you do not need to create more OLAP schemas than there are OLAP servers on your system.

Once an OLAP Schema is available, a cube can be assigned to that OLAP schema during the cube building process. Whereas there are no absolute restrictions for the number of cubes per OLAP Schema, assigning a large number of cubes to a schema should be avoided. This is because the cubes compete for the OLAP server's cube cache and data cache at the time they are accessed and handled by the server.

You can create and maintain OLAP schemas in SAS OLAP Cube Studio or SAS Management Console. Some of the tasks you can perform include the following tasks:

Defining an OLAP Schema

In addition to SAS Workspace Server and SAS OLAP Server definitions, you must also have an OLAP schema defined in the active metadata server. The standard SAS deployment process creates an OLAP schema named SASApp - OLAP Schema and assigns it to the SAS OLAP Server that is also configured during deployment. To define new schemas, you can use the New OLAP Schema Wizard, which is available from SAS OLAP Cube Studio and SAS Management Console. In SAS Management Console, you can launch the OLAP Schema Wizard when you view the OLAP Schema tab on the Properties dialog box for an existing server. These steps explain how to use SAS OLAP Cube Studio to launch the OLAP Schema Wizard and define a new schema:

- 1 Connect to the SAS Metadata Server.
- 2 Select **File ► New ► OLAP Schema**.
- 3 On the General page of the New OLAP Schema Wizard, enter the schema name and description. Click **Next**.
- 4 On the Server Assignment page, specify the OLAP servers that can access the schema. This step is optional. Click **Next**.

Note: If you choose not to specify the server using the OLAP Schema Wizard, you can add that information later by modifying the schema's properties. △

- 5 On the Finish page, click **Finish**.

Note: When defining a new SAS OLAP Server, if you accept the default definition settings, then an OLAP schema is automatically created and assigned to the server. To change that assignment to a different OLAP schema, you must edit the server definition. In SAS Management Console, select the Server Manager plug-in to make these changes. △

Editing an OLAP Schema

The OLAP schema specifies which group of cubes an OLAP server can access. To edit an OLAP schema in SAS OLAP Cube Studio, complete these steps:

- 1 In the Tree View, select the OLAP Schema you want to edit. Then select **Properties** from the **Edit** menu or from the schema's context menu.
- 2 In the Properties dialog box, make changes on the **General** and **Server Assignment** tabs.

You can view cubes assigned to the schema on the **Cubes** tab. To change the OLAP schema assignment for a cube, you must select the properties for a cube and select the **Location** tab. From here you can enter the name of a different OLAP schema.

- 3 When you are finished, click **OK** on the Properties dialog box.

Changing the OLAP Schema For a Cube

When working with SAS OLAP Cubes, it might be necessary to assign a cube to a different schema. You can change the OLAP schema assignment for a cube by selecting **Change OLAP Schema** from the **Actions** menu or the cube's context menu. The Change OLAP Schema dialog box opens. The selected cube is currently assigned to the OLAP schema that is listed in the dialog box.

Select an OLAP schema from the drop-down list of OLAP schemas. If needed, you can also create a new OLAP schema by selecting the **New** button. This opens the New SAS OLAP Schema wizard. After you have selected an OLAP schema, select the **OK** button. The cube is now assigned to the selected schema.

Note: When reassigning a SAS OLAP cube to a different OLAP Schema, be careful to note which OLAP server the schema is assigned to. △

Deleting OLAP Schemas

To delete an OLAP schema, complete the following steps:

- 1 In the Tree View, select the Schema you want to delete.
- 2 Select **Delete** from the **Edit** menu or from the schema's context menu.
- 3 Click **Yes** on the Confirm Delete dialog box.

Note: You cannot delete a schema that currently has OLAP cubes assigned to it. In addition, you must have the necessary authorizations.

Viewing the OLAP Schema Properties

In SAS OLAP Cube Studio, you can view the properties for a SAS OLAP schema by selecting the schema from the Tree View and then selecting the **Properties** function. The Properties function can be accessed from the **Edit** menu and from the schema's context menu. The OLAP Schema Properties dialog box contains tabs that provide general information as well as information that is specific to SAS OLAP schemas. The following tabs are specific to OLAP schemas:

Server Assignment	lists the OLAP servers defined in the current metadata. On the Server Assignment tab you can select an OLAP server from the Available list and move it to the Selected list.
Cubes	lists cubes that are assigned to the current schema. The Cubes tab also displays whether the cubes physically exist or exist as metadata only.

Assigning the OLAP Schema with the OLAP Procedure

To assign an OLAP schema, you must use the OLAP_SCHEMA= option with the METASVR statement. (See “METASVR Statement” on page 256.) Here is an example of the METASVR statement and the OLAP_SCHEMA= option:

```
METASVR
  HOST    = "hw4195.ec.sas.com"
  PORT    = 8561
  OLAP_SCHEMA = "SASApp - OLAP Schema";
```

Note: For more information on SAS OLAP schemas see “Managing OLAP Cube Data” in the *SAS Intelligence Platform: Data Administration Guide*. \triangle

SAS Libraries and Tables

Defining Libraries and Tables

When you set up the SAS OLAP Cube Studio environment, you must define the data tables that are used to build cubes and the libraries that they are assigned to. The following tasks must be completed:

- A metadata server must be started.
- A workspace server must be started and registered in the metadata.
- The metadata for the data tables and libraries must be registered in the metadata and stored on the workspace server.

Note: You must have ReadMetadata and WriteMetadata permissions to perform these tasks. \triangle

The tables and libraries can be defined in SAS applications such as SAS Data Integration Studio or SAS Management Console before you set up your SAS OLAP Cube Studio environment. If however, this is not done beforehand, you can define them in SAS OLAP Cube Studio. You need to have your tables and libraries defined if you plan to do either of these tasks:

- create the physical cube in addition to registering its metadata
- manually add, modify, or drop specific aggregations for a cube

You can use the Source Designer wizard to define your data tables. Use the New Library wizard to define your libraries.

Creating a New Library Definition for Source Data Tables

You can create a new library definition for your source data tables after you start a SAS Metadata Server and define a SAS Workspace Server in a SAS metadata folder. You can create new library definitions using the New Library Wizard, which is available from SAS OLAP Cube Studio and SAS Management Console.

If you have more than one SAS Workspace Server defined, you should assign the library to all of the workspace servers that might be used to create cubes. Otherwise, SAS OLAP Cube Studio attempts to download the data to the server where the cube is being built. This requires SAS/CONNECT and may not be the most efficient way to build a cube.

These steps explain how to use SAS OLAP Cube Studio to launch the New Library Wizard and define a new library:

- 1 Connect to the SAS Metadata Server.
- 2 Select **File** ► **New** ► **Library**. This opens the New Library Wizard.
- 3 Select a Resource Template and then click **Next** to continue.

Note: If you connected to the SAS Metadata Server with an unrestricted metadata profile, you are prompted to enter your user ID and password. Δ
- 4 In the New Library Wizard, enter the library name. You also can enter an optional description. Click **Next**.
- 5 Enter values for the libref, the Engine (BASE is the default), the Content Server, and the Path Specification fields. Follow these guidelines:
 - The libref is a short name (or alias) for the full physical name of a SAS library (for example, sasuser).
 - The path specifies the physical location of the tables contained in the library that you are defining. Select an existing path from the box or click **New** to enter a new path.
 - If your data is accessed through a WebDAV content server, select the **Enable webDAV Support** check box. The content server specifies the HTTP server that is used to access the data.
 - Click the **Advanced Options** button to set host-independent options such as file encoding, as well as host-specific options. Click **OK** to close the dialog box and return to the New Library Wizard.
 - Click **Next**.
- 6 Select the SAS Workspace Servers on which the new library is to reside. Click **Next**.
- 7 Click **Finish** to complete the new library definition.

Note: For further information on SAS Libraries see “Assigning Libraries” in the *SAS Intelligence Platform: Data Administration Guide*. Δ

Defining Tables Used to Build Cubes

You define tables using the Source Designer wizard, which is available from SAS OLAP Cube Studio and SAS Data Integration Studio. Following is a list of the tables that can be used to define a cube:

- detail tables (unsummarized data)
- fact tables and dimension tables (for cubes based on star schemas)
- aggregation tables (fully summarized external tables)
- drill-through tables (views maintained by the user that represent all of the data used to define a cube)

These steps explain how to use SAS OLAP Cube Studio to launch the Source Designer wizard and define new tables:

- 1 Connect to the SAS Metadata Server.
- 2 Select **File** ► **Register** ► **Table**. This opens the Source Designer wizard. The Source Designer wizard is also available from the Cube Designer wizard.
- 3 Select the SAS data source and click **Next**.
- 4 From the **Server** drop-down list, select the name of a SAS Application Server context that contains a logical SAS Workspace Server. Click **Next**.

- 5 Enter your user name and password for the selected SAS Workspace Server. Click **OK** to return to the Source Designer wizard. Click **Next**.
- 6 Select the name of the SAS library that points to the tables that you are importing from the current SAS Workspace Server. Click **Next** to see a list of SAS data sets in the selected library. (You also can click **New** to create a new library.)
- 7 Select the data sets that you want to load into the metadata, and click **Next**.
- 8 Click **Finish**.

Note: If you previously selected a server context and tested the SAS Workspace Server connection, you might not be prompted for the name of the server context. △

Renaming Tables

In SAS OLAP Cube Studio, you can change the name of a data table that you have registered in the SAS metadata folder. The table name that you see in the Tree View is the display name for the table, similar to a label or short description. You must have WriteMetadata permission on the table for which you want to view properties.

Renaming a Table From the Tree View

To rename a table from the Tree View, complete the following steps.

- 1 In the Tree View, select the **Inventory** tab. Select a data table from the **Table** node.
- 2 Select the **Rename** function from the **Edit** menu or from the table's context menu. The table name in the Tree View is highlighted.
- 3 Enter a new name and select **ENTER**.

Renaming a Table from the Properties Dialog Box

To rename a table from the Properties dialog box, complete the following steps.

- 1 In the Tree View, select the **Inventory** tab. Select a data table from the **Table** node.
- 2 Select the **Properties** function from the **Edit** menu or from the table's context menu.
- 3 Select the **General** tab.
- 4 Enter a new name in the **Name** field.
- 5 Click **OK** to save the change and close the dialog box, or click **Apply** to save the change and remain in the dialog box.

For information on naming tables, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30.

Deleting Tables

You can delete tables that you have registered in the SAS metadata. In order to delete a table you must have been granted WriteMetadata permission on that table. Complete the following steps to delete a table.

- 1 In the Tree View, select the **Inventory** tab.
- 2 Select **Tables** and open the list of tables.
- 3 Select the table that you want to delete.
- 4 Select **Delete** from the **Edit** menu or from the table's context menu.

- 5 Click **yes** in the Confirm Delete message box. The table is deleted.

Viewing Table Data in SAS OLAP Cube Studio

In SAS OLAP Cube Studio, you can view the underlying data of the tables that are available to build SAS OLAP cubes with. Tables are listed under the **Tables** node on the **Inventory** tab of the Tree View. You can view the underlying data for a table by selecting **View Data** from the **Actions** menu or from the table's context menu.

Note: For further information on managing tables see “Managing Table Metadata” in the *SAS Intelligence Platform: Data Administration Guide*. △

Managing Folders In SAS OLAP Cube Studio

SAS Metadata Folders

In order to define cubes, you must be connected to a running metadata server and have a SAS folder to store your cubes to. A SAS folder is a location in which data, metadata, or programs are stored, organized, and maintained. SAS folders are accessible to users either directly or through a network. You can create and edit SAS folders in SAS OLAP Cube Studio, SAS Management Console, or SAS Data Integration Studio.

To create a new folder in SAS OLAP Cube Studio, select **File ► New ► Folder**. The New Folder dialog box opens. Enter the name and location of the folder and select **OK**.

Finding an Object in a SAS Folder

When you are working in the **Inventory** tab of the Tree View, you can quickly locate which folder an object is stored in with the Find In Folders function. Select an object in the **Inventory** tab of the Tree View. Select **Find In Folders** from the **Edit** menu or from the object's context menu. The Tree View then displays the folder location of the object.

Copying an Object to a Different SAS Folder

In SAS OLAP Cube Studio, you can copy a cube object from one folder to another. Select an object in the Tree View. Select **Copy to Folder** from the **Edit** menu or from the object's context menu. The Select a Location dialog box opens. Select a folder from the folders list and click **OK**. The object is copied to the new folder.

Note: You must have WriteMemberMetadata (WMM) on the folder that you are copying objects to. For further information on permissions for different objects, see the *SAS Intelligence Platform: Security Administration Guide*. △

Moving Objects Between SAS Folders

In SAS OLAP Cube Studio, you can move an object from one folder to another. Select an object in the Tree View. Select **Move to Folder** from the **Edit** menu or from the object's context menu. The Select a Location dialog box opens. Select a folder from the folders list and click on **OK**. The object is moved to the new folder.

Note: You must have WMM on the folder that you are moving objects to. For further information on permissions for different objects, see the *SAS Intelligence Platform: Security Administration Guide*. △

Note: For more information on working with SAS folders, see “Working with SAS Folders” in the *SAS Intelligence Platform: System Administration Guide*. △

SAS OLAP Cube Jobs

SAS OLAP Cube Jobs

When you create a cube definition in SAS OLAP Cube Studio, a cube job is automatically created for that cube. A job is a collection of SAS tasks that create output. The cube job is the metadata that links a cube to its load data table. This can be a detail table, a fact table, or a fully summarized table. The cube job also contains a classifier map in the metadata which links the columns in the load table to the levels, measures, and member properties for the cube.

Note: The cube job does not link the star schema dimension tables or the aggregation tables to the cube. △

After a cube definition and the cube job is created, you can perform the following tasks:

- deploy and redeploy a cube job for scheduling
- export a SAS package
- view the properties for a cube job

When you deploy a cube job, the code needed to create the cube is generated. It is then stored in a file in the specified deployment directory. In addition, a deployed job object is created. This new object is linked to the original cube job and it can be added to a SAS flow for scheduling. You can also now redeploy the job as needed.

Cube Job Deployment and Redeployment

SAS OLAP cube jobs are deployed to create cubes and are listed in the Job (cube) folder under the **Inventory** tab of the Tree View. To deploy a cube job, select the job in the Tree View and select **Actions ► Deploy**. The Deploy a job for scheduling dialog box opens. From here you can enter information about the job. If the job is deployed successfully, you receive a confirmation message.

The deployed job is then listed in the Deployed job folder in the Tree View. Each time you deploy a job, a separate deployed job object is created and listed in the Deployed job folder. For example: If you deploy a SAS cube job based on the cube "Sales", the deployed job is labeled as "Sales". If you deploy that same cube job again, the next deployed job is labeled "Sales000". Deploy the cube job a third time and it is labeled "Sales001".

Note: Deployed jobs for a cube job can also be accessed from that job's context menu. Right-click on the cube job and select **Scheduling**. The deployed jobs are listed in the **Scheduling** drop-down list. A deployed cube job can also be redeployed after the initial deployment. Select a deployed cube job from the Deployed job folder in the Tree View. Then, select **Actions ► Redeploy**. △

Job and Deployed Job Properties


Job Properties

In SAS OLAP Cube Studio, you can view the properties for an OLAP cube job or a deployed job. Select a job from the Tree View and select **Properties** from the **Edit** menu or from the job's context menu. The Properties dialog box then opens. A cube job has the same standard dialog box tabs as other cube objects. In addition, it also contains the following tabs:

Process	The Process tab contains options to specify how the job's code is generated (Code Generation) and what Execution Host is used.
Pre and Post Process	The Pre and Post Process tab contains options to insert user-written code at the beginning or end of the current job.

Deployed Job Properties

A deployed job has the same standard dialog box tabs as other cube objects. It also contains the **Scheduling Details** tab for the deployment. This tab includes the command that is used to execute the current job in batch mode. It also includes the settings for the SAS Application Server that is used to deploy the job.



CHAPTER

4

Planning for SAS OLAP Cubes

<i>Overview</i>	23
<i>Data Tables Used to Define SAS OLAP Cubes</i>	24
<i>Detail Tables</i>	24
<i>Fact Tables and Dimension Tables (Star Schema)</i>	24
<i>Summary Tables</i>	25
<i>Aggregation Tables</i>	25
<i>Drill-Through Tables</i>	26
<i>Aggregation Design</i>	26
<i>Aggregation Size</i>	26
<i>User Query Patterns and ARM Logging</i>	26
<i>Administering ARM Logs</i>	27
<i>Aggregation Performance Settings</i>	27
<i>Aggregation Storage</i>	28
<i>MOLAP Aggregation Storage</i>	28
<i>ROLAP Aggregation Storage</i>	28
<i>Choosing MOLAP or ROLAP Aggregation Storage</i>	29
<i>SAS OLAP Cube Size Specifications</i>	29
<i>Defining the Number of Dimensions for A Cube</i>	29
<i>Defining the Number of Hierarchies for A Cube</i>	30
<i>Defining the Number of Levels, Measures, Members, and Properties</i>	30
<i>Naming Guidelines and Rules for the SAS OLAP Server</i>	30
<i>General Naming Guidelines</i>	30
<i>SAS OLAP Cubes</i>	31
<i>Dimensions, Levels, Hierarchies, and Measures</i>	31
<i>OLAP Cube Aggregations</i>	32
<i>OLAP Schemas</i>	32
<i>Calculated Measures and Members</i>	32
<i>SAS Formats Available for Measures</i>	32
<i>Statistics Available for Measures</i>	34

Overview

The goal of an OLAP system is to have data that is organized, available, and presented as relevant information to decision makers. To successfully build and query SAS OLAP cubes, the data in your input files must be internally consistent. In addition, the columns of input data sets need to reflect the dimensional levels of the cubes. When building an OLAP application, it is beneficial to assume that building the actual cube is, the smaller part of the project. The larger task includes finding and collecting the data and making that data consistent.

When planning for a SAS OLAP cube you should consider the following items:

- The data tables you use to build the cube
- Aggregation design
- Aggregation storage considerations; MOLAP or ROLAP
- Size specifications for SAS OLAP cubes and cube objects
- Naming guidelines and rules for SAS OLAP cubes and cube objects
- SAS formats that are available for cube measures
- Statistics that are available for cube measures

Data Tables Used to Define SAS OLAP Cubes

A cube is always loaded from data in relational tables. This data can be stored either in SAS tables or in external RDBMSs and can be accessed through a wide selection of SAS data engines, including the Base SAS engine, SPD Server, SPD Engine, and the SAS/ACCESS engines to external RDBMSs. SAS software enables you to be independent of the physical storage format of the data. When building a SAS OLAP cube, you must consider the format that your source data is stored in and how that source data is to be used in the SAS OLAP cube. Cubes can be loaded from data that is contained in any of the following types of tables:

- detail tables
- fact tables and dimension tables (star schema)
- summary tables
- aggregation tables
- drill-through tables

Detail Tables

A detail, or base, table is any table defined in the SAS metadata that contains the measures and levels for a cube. A detail table consists of unsummarized data that must include one column for each level and one numeric analysis column for each set of measures that will be generated. A detail table includes all the columns that are needed to load the hierarchy levels, the level properties, and measures. Specifically, a detail table must contain the following elements:

- one column per dimension level (char or num)
- one column for each property (char or num)
- numeric analysis variables

Fact Tables and Dimension Tables (Star Schema)

A star schema consists of a single fact table and one or more dimension tables. The fact table must contain one numeric analysis column for each set of measures that will be generated. You cannot have more than one property value for each distinct value of the corresponding level. To successfully load a cube, all foreign keys in the fact table need to have a corresponding primary key in a dimension table. A star schema configuration is organized with the following requirements and considerations:

- The input columns for data measures must be stored in the fact table.
- The input columns for dimension levels and properties must be stored in one table for each dimension.

- The data records of the fact table and the dimension tables must be linked via primary and foreign keys.
- If the dimension levels are defined in a dimension table, all the level columns for that dimension must be contained in the same dimension table.
- Both the dimension keys and fact keys are single columns, not combinations of columns.
- A dimension can be in the fact table. In this case, all the level columns are in the fact table and no fact or dimension key is required.
- The dimension key can also be a level in the dimension.

The fact table requires the following conditions:

- one key column with a foreign key for each dimension (character or numeric)
- numeric analysis variables

A dimension table requires the following conditions:

- one key column with a primary key (character or numeric)
- one column for each dimension level (character or numeric)
- one column for each property (character or numeric)

Summary Tables

The summary table is also known as the N-way data set. It is a table that is already summarized and requires the following elements:

- one column per dimension level (character or numeric)
- one column for each property (character or numeric)
- one column per stored measure, summarized by the appropriate statistic for the measure

Aggregation Tables

In addition to a summary table for the N-way aggregation, summary tables can also be provided for other aggregations. Aggregation tables contain presummarized data for any combination of dimension levels. This feature enables cubes to access summarized numbers quickly. Aggregated data can be created and automatically stored in MOLAP (Multidimensional OLAP) tables as the cube is built or manually stored in ROLAP (Relational OLAP) tables and linked into the cube.

All aggregation tables must contain a column for each measure in the cube where the statistic for the measure is one of the following: N, NMISS, SUM, MAX, MIN, or USS. An aggregation table can be used in two ways:

- as an NWAY data source for the cube. In this case, the table must contain a column for every level in the cube and a column for every stored measure.
- as a subaggregation for the cube. In this case, the table must include a column for each level of the aggregation and a column for every stored measure.

When planning your aggregation tables, note that the column names and attributes (type, format) must be consistent across all input tables for a cube. Aggregation tables require the following structure and organization.

- one column for each dimension level in the aggregation
- one column per measure

Drill-Through Tables

Many SAS OLAP applications give users the ability to select a cell or a range of cells and then view the input data that the cell data was summarized from. Drill-through tables can be used by client applications to provide a view from processed data into the underlying data source. Drill-through tables are views that represent all of the input data used to define a cube. Assigning a drill-through table to a cube is optional.

When you are evaluating the data source for a drill-through table you should consider the following requirements:

- The drill-through table must have the same columns with the same attributes (name, type, format) as the tables that the cube was loaded from.
- In many cases, the detail table is also the drill-through table.
- For a star schema, a drill-through table is created with a view that fully joins the star schema. For a summary table, that summary table can be used as a drill-through table.

Note: When you select a data table for drill-through, you may need to consider user access and security restrictions for that table. For further information, see “Security for Drill-through Tables” on page 76. \triangle

Aggregation Design

Efficient drilling or traversing of the cube data is a key factor in flexible and swift decision making and analysis. In order to maintain speed and consistency in reporting, data is usually pre-calculated or aggregated. An important factor in query performance is good aggregation design, which includes decisions about total storage space, available build time, storage location, and storage format.

Aggregation Size

When planning your data storage and design, it is helpful to approximate the size of aggregations. A basis for estimating aggregation size is the number of distinct values in a dimension level, otherwise known as cardinality. The other factor that determines aggregations size is density. Density is a measure of how many members of each dimension in an aggregation occur in combination with the members of the other dimensions (For example, there might not be sales of a specific product on a specific date). The total cube size, as well as the resources that are available for the cube build process, determine the build time that is needed. It is also important to note that build time should not exceed the cube update interval.

Aggregation size and available hardware influence your choices for aggregation partitioning. You can separate aggregations into multiple files. A reduced file size might accelerate OLAP server access time, particularly if multiple processors are available for multi-threaded processing. You can use either pre-aggregated summary tables, the cube’s own efficient aggregation storage, or a combination of both. Using indexes on either storage type might increase query performance, while also increasing storage space and build time.

User Query Patterns and ARM Logging

In choosing the best aggregations to summarize and store with the cube, the most important factor to take into account is user query behavior. It is recommended that

you start with an initial aggregation design that is based on a minimal set of aggregations or on your best assumptions about usage patterns.

After the cube is deployed to users, you can use the ARM (Application Response Management) logging capabilities of the SAS OLAP Server to collect data about the usage pattern and the performance of individual queries. You can analyze the collected data to find out which cube aggregations are used most, which aggregations you can safely eliminate without harming query performance, and which aggregations are often requested but don't exist in the cube. The ARM data provides all the information that you need in order to get the optimal query improvement for your build time and cube storage space.

Administering ARM Logs

When using ARM logs to generate aggregations for a cube, you must consider the administration and maintenance of the ARM logs. If you are using multiple load balanced servers, you must plan for the naming of the ARM logs that these servers write to. Each load balanced server that you use writes to a designated ARM log file. If the ARM log file name is the same for the different load balanced servers, then the servers will write to the same ARM log simultaneously. This can produce an invalid ARM log. It is recommended that you provide a unique ARM log name for each load balanced server to write ARM log data to. You can then select each ARM log file to analyze in the Aggregation Tuning dialog box.

Note: For further information about ARM logging, see the topic “ARM Logging” in the chapter “Administering SAS OLAP Servers” in the *SAS Intelligence Platform: Application Server Administration Guide* and the *SAS Application Response Measurement (ARM) Reference*. △

Aggregation Performance Settings

When planning and creating aggregations for a cube, you can customize various performance settings for a single aggregation or for all aggregations in the cube. Below are aggregation specific settings that you can change:

- ASYNCINDEXLIMIT=
- COMPACT_NWAY
- COMPRESS | NOCOMPRESS
- CONCURRENT=
- DATAPATH=
- INDEXPATH=
- INDEXSORTSIZE=
- MAXTHREADS=
- INDEX | NOINDEX
- PARTSIZE=
- SEGSIZE=

Note: ASYNCINDEXLIMIT=, CONCURRENT=, INDEXSORTSIZE=, and MAXTHREADS= are available only in the PROC OLAP statement. △

You can change these settings using either the OLAP procedure or in SAS OLAP Cube Studio. In SAS OLAP Cube Studio, you can apply these settings from these functions:

- Cube Designer – Aggregations dialog box
- Aggregation Tuning dialog box

Aggregation Storage

When determining how to efficiently deliver data in a multidimensional cube, how data is stored is an important factor. Summarizing data with aggregations improves query response. Any combination of dimension levels can become a stored aggregation. Which aggregations are being stored has a direct effect on the SAS OLAP Server CPU usage, file I/O, and query response times. The aggregations that are being stored also affect cube build time and the absolute cube file size. Therefore, it is a trade-off between a single instance of resource use at cube build time and multiple instances of resource use at cube query time.

The aggregated data values for SAS OLAP cubes can be stored either with the cube in the cube's internal format or external to the cube in relational summary tables.

MOLAP	MOLAP aggregation storage is the cube-internal storage for aggregations. MOLAP aggregation tables are created as part of the cube creation step.
ROLAP	ROLAP aggregation storage is the cube-external summary tables. ROLAP summary tables need to be pre-calculated from the input data (using tools such as SQL or PROC MEANS/PROC SUMMARY) and made known to the cube at cube creation time.

MOLAP Aggregation Storage

SAS MOLAP aggregation storage maintains the cube data in the same table format as the format that is used by the SAS Scalable Performance Data (SPD) Engine. MOLAP aggregation storage takes advantage of key contraction and allows data access by using the cube's internal data representation directly.

MOLAP aggregation storage of SAS OLAP cubes has the same threading and scalability features as the files used by the SAS SPD Engine. The data and the index section of the files are stored in different physical files. This enables parallel access to the data and index sections. The data and index files themselves are stored in partitions, enabling parallel data retrieval within the same file. The partitions of the data and the index section can be distributed over multiple disc controllers, thus adding a further boost to the threaded and partitioned architecture by reducing contention and possible bottlenecks in the physical I/O.

ROLAP Aggregation Storage

ROLAP tables used in SAS OLAP cubes can be SAS data sets, SAS data views, or any tables or views accessible through a SAS engine. This extends the choice of available storage options for SAS OLAP cubes to include SPDE, SPDS, and any RDBMS product for which a SAS/ACCESS software product is available.

ROLAP aggregation tables must conform to the structure of the input data. The columns that feed the dimension levels must have the same column names and attributes that were used in the input data when loading the cube. In addition, all aggregations must be stored in fully de-normalized form. Here are some guidelines to make aggregation columns for measures available:

- Each ROLAP aggregation table must include all columns for the cube's measures with stored statistics.
- SAS OLAP cube aggregations store the following statistics: SUM, N, NMISS, USS, MIN, and MAX. Other available statistics are derived from the stored statistics by

internal calculations. For example, in order to include a measure for the AVG statistic in your cube, you need to make columns available in your ROLAP aggregation tables that were generated by using SUM and N (count).

ROLAP data is stored in either a flat file or with a star schema. With ROLAP, each instance of slicing and dicing of data is part of an SQL query (or multiple SQL queries) and is comparable to a WHERE clause in the SQL statement.

ROLAP data requests can also run against the data that was used to create and load the cube, whether from a detail table or a star schema. The cube's input data can be used in place of the aggregation with the combination of the lowest level of all dimensions (often called NWAY or NWAY aggregation, which is a name borrowed from PROC MEANS/PROC SUMMARY where it denotes the combination of all CLASS variables).

Choosing MOLAP or ROLAP Aggregation Storage

MOLAP aggregation storage is optimized for SAS OLAP Server internal processing and has a minimal data-size footprint. It uses threaded, parallel data access and is tunable to any hardware environment. MOLAP aggregation storage is convenient because it doesn't require additional data management steps.

ROLAP aggregation storage enables you to use existing ROLAP schemas and reuse legacy SAS OLAP Server, SAS 8 HOLAP structures. ROLAP aggregation storage enables users to use database systems and data servers of their choice to store and serve cube aggregation data. Existing processes can be used to create and access aggregation data to off-load and distribute data access, I/O, and rollup to server systems of the user's choice.

A hybrid approach is possible. For example, users with existing ROLAP structures can build a "light" SAS OLAP cube with no additional stored aggregations and add MOLAP aggregations to further tune the cube performance.

SAS OLAP Cube Size Specifications

When you are creating SAS OLAP cubes, there are some size specifications for the various components of the cube.

Defining the Number of Dimensions for A Cube

The maximum number of dimensions that can be defined in a cube is determined by combining the number of dimensions with the number of multiple hierarchies. The maximum value of that sum is 128. Mathematically, the sum is expressed as follows:

$$\text{MaxDims} = \text{NumDims} + \text{NumMultipleHeirarchies} = 128$$

All hierarchies other than the first hierarchy in each dimension apply to the total. The following are some examples of cubes that meet the maximum number of dimensions:

- 128 dimensions, each dimension has 1 hierarchy
- 127 dimensions, 1 dimension has 2 hierarchies
- 126 dimensions, 1 dimension has 3 hierarchies
- 126 dimensions, 2 dimensions have 2 hierarchies

Defining the Number of Hierarchies for A Cube

When specifying hierarchies, you must define at least one hierarchy for each dimension. The maximum number of hierarchies that can be defined in a cube is determined by combining the number of multiple hierarchies with the number of dimensions. The maximum value of that sum is 128. Mathematically, the sum is expressed as follows:

$$\text{MaxHeirs} = \text{NumMultHiers} + \text{NumDimensions} = 128$$

Specifically, each DIMENSION statement must identify at least one unique HIERARCHY statement. The maximum number of hierarchies that can be defined in a dimension is determined only by the maximum number of hierarchies that can be defined in a cube. All hierarchies other than the first hierarchy in each dimension apply to the total. Levels in the same dimension can be shared between hierarchies.

Defining the Number of Levels, Measures, Members, and Properties

The following are maximum values for the specified cube element:

Levels

The maximum number of levels for a cube is 256. There can be up to 19 levels per hierarchy.

Measures

The maximum number of measures per cube is 1024.

Members

The maximum number of members is 2^{32} members per hierarchy.

Properties

The maximum number of properties per level is 256.

Naming Guidelines and Rules for the SAS OLAP Server

General Naming Guidelines

When you are creating SAS OLAP cubes, there are naming guidelines and restrictions that you should adhere to. Below are general guidelines for naming SAS OLAP cubes and related objects:

Name Uniqueness	Name uniqueness is case insensitive. Therefore, a name can contain mixed-case letters. SAS stores and writes the variable name in the same case that is used in the first reference to the variable. However, when SAS processes a variable name, SAS internally converts it to uppercase. You cannot, therefore, use the same variable name with a different combination of uppercase and lowercase letters to represent different variables. For example, cat, Cat, and CAT all represent the same variable. Names must be unique within a folder by type but not by subtype.
Macro Variables or Functions	When building or editing a SAS OLAP cube, you should be aware that the fields within the Cube Designer wizard do not support macro variables or functions. The fields are purely character strings and are not resolved before being written to the cube's metadata.

Name Literals If the name has embedded blanks or characters other than letters of the Latin alphabet, numbers, or underscores, then PROC OLAP formats the name as a name literal. This means that it is enclosed within quotation marks followed by the letter n. (Name literals enable you to use special characters or blanks that are not otherwise allowed in SAS names. Here are some examples:

- CUBE='Financials@HQ'n
- DIMENSION 'Product@Work Dimension'n hierarchies=(Product@Work Hierarchy'n)
- HIERARCHY 'Product@Work Hierarchy'n levels=(prodtype product)

VALIDVARNAME=Option If you need to use special characters in a name, you must use the VALIDVARNAME= option. Ideally, the SAS Workspace Server should have the VALIDVARNAME= system option set to ANY. If the SAS Workspace Server is not running with VALIDVARNAME=ANY, then you can set the option on the Cube Designer wizard. In this case, select the **Advanced** button on the Cube Designer-General page. On the **Submit SAS Code** tab, enter the VALIDVARNAME= option. Modify your names to meet the naming requirements that the SAS Workspace Server is running with.

Note: For more information about the VALIDVARNAME= option, see “VALIDVARNAME= System Option” in the *SAS Language Reference: Dictionary*. Δ

SAS OLAP Cubes

In addition to the general requirements, SAS OLAP Server names for cubes follow these general rules:

- A cube name can be up to 32 characters in length.
- A cube name must be a SAS name.
- A cube name cannot contain square brackets. This is an MDX restriction.
- A cube name cannot have non-printable characters.
- A cube name cannot have leading white spaces because cubes are public objects that are contained in metadata folders.
- A cube name cannot contain forward slashes or backward slashes.
- If you need to use special characters in a name you must use the VALIDVARNAME = ANY option. This applies to special characters except forward slashes, backward slashes, and square brackets.

Dimensions, Levels, Hierarchies, and Measures

SAS OLAP Server names for dimensions, levels, hierarchies, and measures, follow these general rules:

- A name can be up to 32 characters in length.
- A name must be a SAS name.
- If you need to use special characters in a name you must use the VALIDVARNAME = ANY option.

Note: Cube dimensions cannot have the same name as any level name within the dimension. However, if there is only one hierarchy for the dimension it is automatically assigned the same name as the dimension. Δ

OLAP Cube Aggregations

Aggregation names can be up to 32 characters in length. They can also include slashes, periods, and square brackets.

OLAP Schemas

SAS OLAP Server names for OLAP schemas follow these general rules:

- A schema name can be up to 60 characters in length.
- A schema name cannot have forward slashes or backward slashes.
- A schema name cannot have non-printable characters.
- A schema name cannot have leading white spaces because schemas are public objects that are contained in metadata folders.

Calculated Measures and Members

SAS OLAP Server names for calculated measures and members follow this general rule:

- A calculated measure or member name can be up to 60 characters in length. However, it is recommended that you maintain a 32 character limit for calculated measures.

SAS Formats Available for Measures

For each measure that you define for a cube, you can select the SAS format that is used to display the value of the measure. The following formats are available from the **Format** drop-down list on the Measure Details page of the Cube Designer wizard in SAS OLAP Cube Studio:

Table 4.1 SAS Formats Available for Measures

SAS Format	Description
15.0	Designates the width of the resulting string will be a maximum of 15 characters after the value has been formatted.
BEST	Enables SAS to choose the best numeric notation.
BEST15	Enables SAS to choose the best numeric notation. The width of the resulting string will be a maximum of 15 characters.
COMMA	Writes numeric values with commas and decimal points.
COMMA15	Writes numeric values with commas and decimal points. The width of the resulting string will be a maximum of 15 characters.
COMMAX	Writes numeric values with periods and commas.
COMMAX15	Writes numeric values with periods and commas. The width of the resulting string will be a maximum of 15 characters.

SAS Format	Description
DOLLAR	Writes numeric values with dollar signs, commas, and decimal points.
DOLLAR15	Writes numeric values with dollar signs, commas, and decimal points. The width of the resulting string will be a maximum of 15 characters.
DOLLARX	Writes numeric values with dollar signs, periods, and commas.
DOLLARX15	Writes numeric values with dollar signs, periods, and commas. The width of the resulting string will be a maximum of 15 characters.
EURO	Writes numeric values with a leading euro symbol, a comma that separates every three digits, and a period that separates the decimal fraction.
EURO15	Writes numeric values with a leading euro symbol, a comma that separates every three digits, and a period that separates the decimal fraction. The width of the resulting string will be a maximum of 15 characters.
EUROX	Writes numeric values with a leading euro symbol, a period that separates every three digits, and a comma that separates the decimal fraction.
EUROX15	Writes numeric values with a leading euro symbol, a period that separates every three digits, and a comma that separates the decimal fraction. The width of the resulting string will be a maximum of 15 characters.
NEGPAREN	Writes negative numeric values in parentheses.
NEGPAREN15	Writes negative numeric values in parentheses. The width of the resulting string will be a maximum of 15 characters.
PERCENT	Writes numeric values as percentages.
PERCENT15	Writes numeric values as percentages. The width of the resulting string will be a maximum of 15 characters. W.D writes standard numeric data one digit per byte.
WORDF	Writes numeric values as words with fractions that are shown numerically.
WORDF10	Writes numeric values as words with fractions that are shown numerically. The width of the resulting string will be a maximum of 10 characters.
WORDS	Writes numeric values as words.
WORDS10	Writes numeric values as words. The width of the resulting string will be a maximum of 10 characters.
YEN	Writes numeric values with yen signs, commas, and decimal points.
YEN15	Writes numeric values with yen signs, commas, and decimal points. The width of the resulting string will be a maximum of 15 characters.

Note: For more information about SAS formats, see "Formats" in the SAS Language Reference: Dictionary. Δ

Statistics Available for Measures

When you define a cube, you select the statistics that are used to calculate the cube's measures. Here is the list of available statistics.

Base Stored Statistics:


- Count
- Sum
- Maximum
- Minimum
- Count of Missing Values
- Uncorrected Sum of Squares

Derived Statistics:

- Average
- Range
- Corrected Sum of Squares
- Variance
- Standard Deviation
- Standard Error of Mean
- Coefficient of Variance
- T Value
- Probability of Greater Absolute Value
- Lower Confidence Limit
- Upper Confidence Limit

New cubes that are based on a data source that contains existing summarized data must include measure statements for the stored statistics. Measure statements are required for each derived statistic that you want to create for the new cube. For example, if you want to calculate AVG, you must create measures for N and SUM, as well as AVG. The following table indicates which stored statistics are required for each derived statistic.

Derived Statistics	Required Stored Statistics
AVG	N, SUM
RANGE	MIN, MAX
CSS	N, SUM, USS
VAR, STD, STDERR, CV, T, PRT, LCLM, UCLM	N, SUM, USS



CHAPTER

5

Building Cubes and Administering Cubes

<i>SAS OLAP Cube Studio and the OLAP Procedure</i>	36
<i>SAS OLAP Cube Studio</i>	36
<i>SAS OLAP Procedure</i>	36
<i>Connecting and Reconnecting to a Metadata Server</i>	36
<i>Creating a Connection Profile</i>	37
<i>Connecting with an Existing Profile</i>	38
<i>Reconnecting to a Metadata Server</i>	38
<i>Cube Designer Wizard</i>	38
<i>Defining Data Sources For a Cube</i>	38
<i>Defining Drill-Through Tables</i>	39
<i>Defining Dimensions and Levels</i>	40
<i>Defining Dimensions</i>	40
<i>Defining Levels for a Dimension</i>	40
<i>Specifying the TIME Dimension</i>	40
<i>Designating a TIME Dimension</i>	40
<i>Add Levels</i>	41
<i>Add Supplied Time Hierarchies</i>	41
<i>Specifying GIS Map Information for a Dimension</i>	41
<i>Specifying a GEO Type Dimension in PROC OLAP</i>	42
<i>Defining Cube Hierarchies</i>	42
<i>Defining a Default Hierarchy</i>	42
<i>Defining Multiple Hierarchies for a Dimension</i>	42
<i>HIERARCHY Statement</i>	43
<i>Cube Designer</i>	43
<i>Defining Ragged and Unbalanced Hierarchies for a Dimension</i>	44
<i>Defining Ragged and Unbalanced Hierarchies in SAS OLAP Cube Studio</i>	44
<i>Defining Ragged and Unbalanced Hierarchies with PROC OLAP</i>	45
<i>Ragged Hierarchies and Unique Member Names</i>	46
<i>Defining Measures For A Cube</i>	47
<i>Selecting Measures in SAS OLAP Cube Studio</i>	47
<i>Unique Member Count Measures</i>	47
<i>NUNIQUE Statistic</i>	47
<i>Defining Stored and Derived Measures for a Fully Summarized Cube</i>	48
<i>Stored Measures</i>	48
<i>Derived Measures</i>	48
<i>Statistics Available for Measures</i>	48
<i>Base Stored Statistics</i>	48
<i>Derived Statistics</i>	49
<i>Defining Member Properties</i>	49
<i>Property Statement</i>	50
<i>Cube Designer</i>	50

<i>Defining Aggregations While Building a Cube</i>	50
<i>Adding Aggregations to a Cube</i>	50
<i>Defining Stored Aggregations For a Fully Summarized Cube</i>	50
<i>Saving the Cube Metadata or Creating the Physical Cube</i>	51
<i>New Cubes</i>	51
<i>Cubes Only Defined in the Metadata</i>	51
<i>Cubes That Physically Exist</i>	51
<i>Saving the OLAP Procedure (Long Form versus Short Form)</i>	52
<i>Export Code</i>	53
<i>Viewing Cubes in SAS OLAP Cube Studio</i>	53
<i>SAS OLAP Server Connection</i>	54
<i>Cube Permissions</i>	54
<i>Disabled Cubes</i>	54

SAS OLAP Cube Studio and the OLAP Procedure

SAS OLAP Cube Studio

SAS OLAP Cube Studio contains several tools and functions that enable you to build, modify, and maintain OLAP cubes. Specifically, the Cube Designer wizard guides you through the cube building process. You can specify the data used to build the cube, the structure of the cube, and the measures and aggregations that will be used to process queries for the cube. After you have built your cube in the Cube Designer wizard, you can view a simple display of the cube and verify your cube data and structure.

SAS OLAP Cube Studio also provides tools and functions for modifying and maintaining cubes. After a cube is built you can make various changes and updates to the cube. Here are some functions that you can perform on a cube after it is created:

- tune the aggregations for the cube
- add and edit calculated members and measures for the cube
- incrementally update the cube
- export and then import a SAS package between SAS environments

For more information on functions and tasks that you can perform on a cube after it is created, see Chapter 6, “Modifying and Maintaining Cubes,” on page 55.

SAS OLAP Procedure

In addition to using SAS OLAP Cube Studio, you can define and build SAS OLAP cubes with the OLAP procedure. The OLAP procedure contains various statements and options to help you define a cube in batch mode if needed or preferred. See Appendix 1, “The OLAP Procedure,” on page 243 for detailed information on the OLAP procedure.

Connecting and Reconnecting to a Metadata Server

The SAS Metadata Server enables users to write metadata objects to and read metadata objects from SAS folders. It maintains information about users, groups, data libraries, servers, and various user-created products such as cubes and information

maps. The SAS Metadata Server contains a metadata identity for every user of the SAS Intelligence Platform. This includes each user's login information, including a user ID and an encrypted password. When a user logs on to a SAS application, the application verifies the user's identity by checking it against the metadata identity.

To create SAS OLAP cubes, you must create a connection profile for the metadata server that you access while using SAS OLAP Cube Studio. A connection profile is a user's view of a metadata server and is used to connect the application to a metadata server. The Connection Profile wizard enables you to specify the connection profile. When you define a connection profile, you specify the server and the port to which the application connects. You can also specify a user ID and password for all connections to the server or allow users to specify a user ID and password each time they log on. When an application opens a connection profile, it has access to the metadata on the specified server based on the security credentials of the user ID used to connect.

A connection profile is a client-side definition of where a SAS Metadata Server is located. The connection profile definition includes a host name and a port number, and can contain a user's login information and instructions for connecting to the metadata server automatically.

Creating a Connection Profile

In order to define cubes, you must be connected to a running metadata server. Connection information is stored in a connection profile, which is a client-side definition of where the SAS Metadata Server is located. The profile definition includes a host name and port number. In addition, the connection profile can contain a user's login information and instructions for connecting to the metadata server automatically.

You can define new connection profiles using the Connection Profile wizard. This wizard is available from within various SAS products, including SAS OLAP Cube Studio, SAS Management Console, and SAS Data Integration Studio. To create a new connection profile, follow these steps:

- 1 In SAS OLAP Cube Studio, select **File ► Connection Profile** to display the Connection Profile dialog box. The Connection Profile dialog box also appears automatically when you launch SAS OLAP Cube Studio, unless you have set a default profile.
- 2 Select the **Create a new connection profile** radio button to launch the Connection Profile wizard.
- 3 On the Connection Profile wizard, enter connection information to create the new profile, including the following items:
 - profile name
 - machine name
 - port number
 - user name
 - password
 - authentication domain
 - whether to save user ID and password in this profile
 - whether to use Integrated Windows authentication (single sign-on)
- 4 Click **Finish** to return to the Connection Profile dialog box.
- 5 If you want to always use the selected profile when SAS OLAP Cube Studio is launched, click the **Set this connection profile as the default** check box.
- 6 Click **OK** to open the selected profile.

Connecting with an Existing Profile

Existing connection profiles are listed on the Connection Profile dialog box, which appears automatically when you launch SAS OLAP Cube Studio (unless you have set a default profile). If the dialog box does not appear automatically, select **File ► Connection Profile**. Complete these steps in the Connection Profile dialog box:

- 1 Select the **Open an existing connection profile** radio button.
- 2 Select the name of the profile from the drop-down selection list.
- 3 If you want to always use the selected profile when SAS OLAP Cube Studio is launched, select the **Set this connection profile as the default** check-box.
- 4 Click **OK** to open the selected profile.

Reconnecting to a Metadata Server

If the connection to the metadata server is broken, a dialog box appears and asks if you want to attempt reconnection. Click **Try Now**. SAS OLAP Cube Studio then attempts to reconnect to the metadata server. If the reconnection is successful, you can continue your work. The user credentials from the previous session are used.

Cube Designer Wizard

Within SAS OLAP Cube Studio, the Cube Designer wizard guides you through the process of building a cube. You can create cube definitions as well as build complete cubes. The following are some tasks that you can perform:

- define the data sources used to load a cube
- specify any drill-through tables used by the cube
- define the cube dimensions, levels, and hierarchies
- select measures and measure details for the cube
- specify unique member count measures (NUNIQUE)
- specify member properties
- configure aggregations
- select language translation tables for a star schema's dimension tables
- define user-written formats to create the cube
- define how to handle missing values for ragged hierarchies

Defining Data Sources For a Cube

When you create a cube you must specify the data source tables for the cube. Those data source tables must be registered in the SAS metadata. Within SAS OLAP Cube Studio, the Register Table wizard enables you to identify a data source and register it in the SAS metadata. Any data source that you plan to use for your cube must be registered in the SAS metadata so that it is accessible in SAS OLAP Cube Studio.

The type of data table(s) you are using determines the type of cube that you can build. There are three types of SAS OLAP cubes that you can build:

- 1 detail table
- 2 star schema

3 fully summarized table

On the Cube Designer –General page, you first specify the cube name, the storage location for the cube, and the type of input data you are using. The type of cube you create (detail, star schema, fully summarized) determines the information that you must specify in the Cube Designer wizard. A star schema cube requires different information than a detail cube or a fully summarized cube. As a result, the wizard presents a slightly different arrangement of steps for each type of cube that you build.

After you complete entering information for the Cube Designer – General page, you then select the specific input data table for the cube. All cube types require an input data table. The Cube Designer – Input page enables you to select a datasource table for the cube. You can also view the underlying table data and the table properties from this page. Star schema and fully summarized cubes have additional tables that must be specified as follows:

- For star schema cubes you must also define the dimension tables for the cube. Dimension tables are defined after input tables on the Cube Designer – Dimension Tables page.
- For fully summarized cubes, you also define aggregation tables on the Cube Designer – Aggregation Tables page. You do this when you define aggregations for the cube.

Defining Drill-Through Tables

Drill-through tables enable you to display, at query time, the unsummarized detail data that a table cell or selected cells were summarized from. Using the drill-through capability, you can view the specific observations from the underlying data that make up an aggregate value. Many OLAP applications give you the ability to select a cell or a range of cells and then view the input data that the cell data were summarized from. Drill-through capability enables companies to access data that is not stored on an OLAP server and make it accessible to end users of an OLAP application. When the OLAP server receives a request for this additional data at query time, it automatically submits a query and retrieves the data from a data warehouse or from an OLTP (online transaction processing) system. In order to provide this capability, a drill-through table must be specified for a cube.

A drill-through table can be a view, data set, or other data file that contains data that is used to define a cube. For you to successfully use a drill-through table, it must have the same columns with the same attributes as the table that the cube was loaded from. Detail data tables are most commonly used as drill-through tables. However, if the cube was loaded from a star schema, a view that fully joins the star schema can be used as a drill-through table.

In the Cube Designer wizard, you can select a drill-through table on the Cube Designer – Drill-Through page. Select a table from the Available tables list and move it to the Selected table list.

You can also use the OLAP procedure, `DRILLTHROUGH_TABLE | DT_TABLE | DT_TBL=` option to define a drill-through table for a cube. To understand how drill-through is implemented in a SAS OLAP cube, see the cube building example “Implementing Drill-through to Detail Data in a SAS OLAP Cube” on page 210.

Note: When you select a data table for drill-through, you may need to consider user access and security restrictions for that table. For further information see “Security for Drill-through Tables” on page 76. Δ

Defining Dimensions and Levels

Within the Cube Designer wizard, you can specify dimensions and levels for a cube on the Dimension Designer wizard. This wizard enables you to select the type of dimension you are creating and then select the levels for that dimension.

Defining Dimensions

The Dimension Designer wizard enables you to add new dimensions to a cube and edit existing dimensions. On the Dimension Designer – General page, you can specify the following information for a dimension:

- the type of dimension you are creating (Standard, TIME, or GEO)
- the sort order for level values in that dimension
- whether to allow new members for the dimension during an incremental update of the cube
- star schema mapping details, such as the dimension table, key, and fact key
- missing values for ragged or unbalanced data

Defining Levels for a Dimension

After you specify the general information for a dimension, you can select the levels for that dimension. On the Dimension Designer – Level page, select the **Add** button. The Add Levels dialog box opens. From here you can select the input columns for the levels that you are adding to the dimension. You can have a maximum of 256 levels per cube.

If you have defined a dimension as a TIME type dimension, the **Add** button is changed to a drop-down list of options. A second option for adding levels to the dimension is then available. You can select either of the following options:

- Add levels**
- Add supplied time hierarchies**

Note: See “Specifying the TIME Dimension” on page 40 for more details. \triangle

After you have selected the levels for a dimension, you can select the attributes for each level on the Dimension Designer – Level page. You can specify the following information for a level:

- format
- caption
- sort order
- missing member string
- description

You can also add more levels to the dimension, duplicate a level, or delete a level.

Specifying the TIME Dimension

Designating a TIME Dimension

The TIME dimension is a unique type of dimension for SAS OLAP cubes. When building a cube, and specifically, when defining a dimension, you must select the type of

dimension that you are creating. You can create Standard, TIME, or GEO type dimensions. For the TIME dimension, you can also provide time hierarchies with the input table for the cube. Supplied time hierarchies can help you build the dimension and auto-populate the levels and some level properties for the cube.

On the Dimension Designer – General page select the TIME type dimension. On the Dimension Designer – Level page, the **Add** button becomes a drop-down list of options. The **Add levels** and **Add supplied time hierarchies** options are now available for selection.

Note: The Add (levels) function is always available for selection regardless of the type of dimension you are creating. When you specify a TIME type dimension, the **Add** button is converted to a drop-down list, where you can select which of the Add (level) specific functions you can use. Δ

Add Levels

The Add Levels dialog box enables you to select input columns for the levels that you are adding to the dimension. When you select the **Add levels** option, the Add Levels dialog box opens. You can select from a list of available input columns and add them to the list of input columns for new levels.

Add Supplied Time Hierarchies

In the Dimension Designer, the **Add supplied time hierarchies** option becomes available when a dimension is defined as a TIME type dimension and supplied time hierarchies are available. When you select the **Add supplied time hierarchies** option, the Add Supplied dialog box opens. For time-specific dimensions, you can select from a list of supplied time hierarchies. These supplied time hierarchies can help you build the dimension and auto-populate the levels and some level properties for the cube. In this dialog box, you select from a list of supplied time hierarchies for the dimension. The Add Supplied dialog box is available only if the TIME type is selected on the Dimension Designer - General page.

Note: For more information on applying time hierarchies, see the cube building example “Creating a Time Dimension in SAS OLAP Cube Studio” on page 197. Δ

Specifying GIS Map Information for a Dimension

The SAS OLAP Cube Studio GIS Map function enables you to store ESRI Geographic Information System (GIS) spatial map information in the SAS metadata. This GIS information can then be read by the SAS OLAP Server and returned during a cube query.

When building or editing a cube, the GIS Map function enables you to identify a geography-based dimension and then assign ESRI spatial map information to that dimension. To define GIS information for a SAS OLAP cube, you identify a dimension as a geographic-type dimension (GEO) in the Dimension Designer - General window. You can have only one GEO dimension per cube. After a dimension has been marked as a GEO-type dimension, the **GIS Map** button becomes active on the Cube Designer - Dimensions page.

The GIS Map dialog box enables you to assign ESRI spatial map information to levels of the GEO-type dimension. You can add map information to an existing cube, modify map information for a cube, or delete the map information from a cube. When

ESRI map information is added to a cube, the property objects for the mapped cube levels are listed in the Cube Designer’s Member Property window. From here you can modify the format, caption, and description for the member property. These member properties are named SAS_SPATIAL_ID by default.

You can set up the ESRI ArcGIS server information in the metadata by using the Map Service Manager plug-in in SAS Management Console. Users can access the functionality through the SAS Web OLAP Viewer.

Note: If a dimension is changed from type GEO to Standard or TIME, all the map information is removed. Δ

For further information on the GIS Map window, see the SAS OLAP Cube Studio Help topic “GIS Maps”. For information on the ESRI Map Component, see the “Configuring the ESRI Map Component” chapter in the *SAS Intelligence Platform: Web Application Administration Guide*

Specifying a GEO Type Dimension in PROC OLAP

If you create a cube with the OLAP procedure, you can define a cube dimension as a GEO type. The TYPE= option, on the DIMENSION statement, enables you to identify a dimension as a GEO type dimension. This can only be used when you initially create a cube with the OLAP procedure. In order to add geographic information to an existing cube, you must use SAS OLAP Cube Studio. For more details, see the TYPE= option.

Defining Cube Hierarchies

Defining a Default Hierarchy

When you define cube dimensions, levels, and hierarchies in SAS OLAP Cube Studio, a default hierarchy for a dimension is automatically created if a hierarchy is not explicitly defined. This default hierarchy includes all levels that were specified for the current dimension and the order they were listed in for the dimension. In addition, if you define multiple hierarchies and do not select a default, then the default is automatically assigned to the first hierarchy that is created for the dimension. On the Dimension Designer – Hierarchy window, you can click the **Default** button to set a selected hierarchy as the default for the dimension.

Defining Multiple Hierarchies for a Dimension

SAS OLAP cubes are organized into dimensions and levels of data. The levels are then arranged into hierarchies. After an initial hierarchy has been created, you can define additional hierarchies for a single dimension of a cube. This enables you to have multiple possible drill paths for the same data. When you create more than one hierarchy for a dimension, the levels have some restrictions:

- A level in a dimension might be used in more than one hierarchy within that dimension. However, levels cannot be used in hierarchies that are not defined within the dimension that the level is defined in.
- Each level must be used in at least one hierarchy.
- Levels from the same dimension that are picked for an aggregation must be in the drill order for at least one hierarchy in that dimension.

- You cannot share levels between dimensions.

You can arrange the levels in a hierarchy in any order. The one exception to this is the TIME dimension. Levels in hierarchies in the TIME dimension must follow a prescribed order that is determined by the numerical value that is assigned to the type. This order is from the smallest value (Years, 16) to the greatest value (Seconds, 3,096). You can only have one TIME dimension for a cube. The dimension hierarchies also have some restrictions:

- The first hierarchy that is defined for the dimension is designated as the default. When there are multiple hierarchies, you can designate the default hierarchy for the dimension.
- Hierarchy names must be unique across the cube. If there is a single hierarchy for a dimension, then its name must be the name of the dimension. Also, dimension and hierarchy names cannot be the same as a level name within that dimension.
- For any cube that is loaded with a star schema, in which a dimension table represents multiple hierarchies for that dimension, the dimension key that is used to join the dimension table to the fact table will be used for all hierarchies of that dimension.

HIERARCHY Statement

The HIERARCHY statement is used with the PROC OLAP statement when you define a cube:

```
hierarchy campaigns
  levels=(campaign_type campaign sub_campaign);
```

Cube Designer

You can establish multiple hierarchies by using the Cube Designer – Dimension page, which is located in the SAS OLAP Cube Studio Cube Designer. To add a hierarchy to an existing dimension, select a dimension, and then click **Modify**. This opens the Dimension Designer – General page. It is populated with the values for the selected dimension. Select **Next** until you reach the Dimension Designer - Hierarchy page. Select **Add** to create an additional hierarchy.

Note: You can modify existing hierarchies by selecting a hierarchy and clicking **Modify**. You can also assign a default hierarchy by selecting a hierarchy and clicking **Default**. The first hierarchy is automatically the default hierarchy. Δ

Note: An exception to defining multiple hierarchies for a dimension is the TIME dimension. Levels in hierarchies in the TIME dimension must follow a prescribed order that is determined by the numeric value that is assigned to the type. This order is from the smallest value (Year, 16) to the greatest value (Seconds, 3,096). Δ

On the Dimension Designer – Define a Hierarchy page, you can define a new hierarchy and select the different levels and their order for the hierarchy.

Defining Ragged and Unbalanced Hierarchies for a Dimension

Dimension levels are arranged in one or more hierarchies. Hierarchies, by process of ordering, have a branching arrangement, and the different member levels have parent and child relationships. For instance, at company X the sales staff are located in different regions and cities in different countries. A balanced hierarchy might look like this:

- global sales president (top of hierarchy)
- sales presidents (per country)
- regional sales managers
- city sales managers

Because of differences in the cube data, hierarchies are often not balanced and possibly have missing members. For example, some sales regions might not have sales managers assigned to a specific city. Or some countries might not have sales regions, just cities. These real-world scenarios would create hierarchies that have missing member data and possibly ragged hierarchies. This affects the drillpath of the cube data.

You can also drill to missing members within a path and continue to drill down to members that are present.

Note: Existing SAS 9.2 OLAP cubes that have been updated with one or more new members for a hierarchy can possibly contain ragged hierarchies and must be rebuilt in the second maintenance release after SAS 9.2. Δ

Defining Ragged and Unbalanced Hierarchies in SAS OLAP Cube Studio

The Cube Designer in SAS OLAP Cube Studio enables you to specify the missing members for a hierarchy and the type of data that is missing. Here are the Cube Designer pages that enable you to specify missing member information:

Ragged Hierarchies	Located on the Advanced Cube Options dialog box, on the Cube Designer – General page, this tab enables you to specify character and numeric missing member information. By default, no missing member information is indicated with the value None .
Dimension Designer – Level Properties	<ul style="list-style-type: none"> <input type="checkbox"/> Ragged – Ignore Missing Members specifies whether to ignore or use global or hierarchy-specific ragged hierarchy settings. To ignore settings, set this property to True. To use the settings, set this property to False. By default, this is set to False. <input type="checkbox"/> Ragged – Designate Missing Members specifies that the Cube Designer use the specified string to identify missing values and override any global or hierarchy-specific ragged hierarchy settings. You can use up to 256 characters. The value of the True/False setting in Ragged – Ignore Missing Members controls whether or not you override any global or hierarchy-specific ragged hierarchy settings.
Dimension Designer – Define a Hierarchy	<p>You can select one of these options from the Ragged Hierarchies tab:</p> <ul style="list-style-type: none"> <input type="checkbox"/> Ignore. From this drop-down list, select True to ignore the global missing member settings that you entered at the Advanced Cube Options dialog box. Select False to use the global settings for the current hierarchy.

Defining Measures For A Cube

Selecting Measures in SAS OLAP Cube Studio

When you build a cube in the Cube Designer wizard, you must specify the measures that are used to help query the cube. After you have defined the data input source and the structure of the cube (dimensions, levels, hierarchies) you can define measures on the Select Measures page of the Cube Designer wizard. You can view the cube's currently defined measures, add new measures, and remove existing measures.

To define a measure, you select from the **Available** list of columns and statistics. This list contains the numeric columns in the input data source and the statistics available for each column. Move the selected items to the **Selected** list. You can have a maximum of 1024 measures per cube.

Note: If you are including aggregated data from tables other than the input data source, considerations for stored statistics should be made. This applies when you are creating a cube from a detail table or a star schema. You must include measures for the stored statistics that are required for each derived statistic that you create in the cube. For example, if you want to calculate AVG, you must create measures for N and SUM, as well as AVG. △

Unique Member Count Measures

The Unique Member Count Measures dialog box enables you to store NUNIQUE (unique member count) statistics as measures with a cube. You can define unique member count measures on the Unique Member Count Measures dialog box. This is accessed from the Select Measures page of the Cube Designer wizard. Select the **Define** button to open the Unique Member Count Measures dialog box.

You can add or delete a unique member count measure on this dialog box. To create the unique member count measure, select the level and hierarchy for the new measure. Selecting the level in a tree structure (dimension, hierarchy, level) automatically assigns the dimension and hierarchy to that measure. A name is automatically generated for that unique member count measure consisting of the level name, "NUNIQUE", and the parent hierarchy name. It is important to note that only one combination of level and hierarchy can be defined for a measure. After a combination has been used to create a unique member count measure, the combination cannot be used again. In addition, if the level or hierarchy is deselected, then any associated defined unique member count measures are deleted from the cube. For a cube to build, at least one non-unique member count (non-NUNIQUE) measure must be defined for the cube.

NUNIQUE Statistic

You can define a distinct count statistic using the MEASURE statement and the NUNIQUE statistic. The LEVEL and HIERARCHY options for the MEASURE statement are used with the NUNIQUE statistic and are ignored for non-NUNIQUE statistics if specified.

Note: The LEVEL name is optional. If it is omitted, then the level name is assumed to be the name specified for the NUNIQUE measure. The HIERARCHY name is only required if the level is in multiple hierarchies. For further information about using the NUNIQUE statistic see the "MEASURE Statement" on page 269 . △

Defining Stored and Derived Measures for a Fully Summarized Cube

When you are building a cube with fully summarized data, you can select predefined (stored) measures that are actually stored with the summarized data source. You can also then create derived measures from those stored measures. In the Cube Designer wizard, pages for stored and derived measures are available when you have selected a fully summarized data source to build a cube from.

Stored Measures

On the Stored Measures page of the Cube Designer wizard, you can select the NWAY columns that contain stored measures data. Select from the **Available** list and move the needed measures to the **Selected** list. On the Assign Stored Measures page of the Cube Designer wizard, you can then assign a statistic to each stored measure and each measure to an analysis group.

Derived Measures

After you have selected stored measures from the summarized data source, you can define any needed derived measures for the cube. You can store only derived statistics that can be calculated from the available stored statistics. For example, to use the derived statistic AVG, you must have stored statistics for N and SUM with the same assigned group name.

On the Select Derived Measures page of the Cube Designer wizard, select the **Add** button to create a new derived statistic or the **Modify** button to change an existing derived measure. To define a derived measure, you specify the following options:

Analysis Group	displays a list of analysis groups assigned to the cube's stored measures. Select a group from this drop-down list to see available statistics in the Derived Statistics drop-down list. Only those groups that contain stored measures that can be used to derive a statistic are included in the drop-down list.
Derived Statistics	displays a list of statistics that you can derive from the statistics that you are storing with the cube for the selected analysis group. Select the derived statistic that you want to use. Only those statistics that can be derived from the stored statistics are included in the drop-down list.
Measure Name	displays a name for the derived measure. The name must be a valid SAS name (up to 32 characters). See "Naming Guidelines and Rules for the SAS OLAP Server" on page 30 for more information.

Note: A derived statistic can be used only once per cube per analysis group. \triangle

Statistics Available for Measures

When you define a cube, you select the statistics that are used to calculate the cube's measures. Here is the list of available statistics.

Base Stored Statistics

- Count
- Sum
- Maximum

- Minimum
- Count of Missing Values
- Uncorrected Sum of Squares

Derived Statistics

- Average
- Range
- Correct Sum of Squares
- Variance
- Standard Deviation
- Standard Error of Mean
- Coefficient of Variance
- T Value
- Probability of Greater Absolute Value
- Lower Confidence Limit
- Upper Confidence Limit

New cubes that are based on a data source that contains existing summarized data must include measure statements for the stored statistics, which are required for each derived statistic that you want to create for the new cube. For example, if you want to calculate AVG, you must create measures for N and SUM, as well as AVG. The following table indicates which stored statistics are required for each derived statistic.

Table 5.1 Stored Statistics That Are Used to Create Derived Statistic

Derived Statistics	Required Stored Statistics
AVG	N, SUM
RANGE	MIN, MAX
CSS	N, SUM, USS
VAR, STD, STDERR, CV, T, PRT, LCLM, UCLM	N, SUM, USS

Defining Member Properties

When you create a SAS OLAP cube, the information that is relevant to the cube is defined with the cube hierarchy, measures, and aggregations (summaries) that will be stored with the cube. Additional information that is part of the cube member data can be included in the cube definition as a member property. Member properties are attributes of dimension members that provide an additional gradation of information to users of the cube data. Member property information is usually not as significant as the levels and members within a dimension, and therefore does not qualify as a level or member. However, it can have additional analytical value that can be useful at query time.

A member property is assigned to a level within a hierarchy, and a level can have multiple properties that are assigned to it. For hierarchy placement, a member property is assigned (by default) to all hierarchies that the select level is included in. However, you can remove one or more (but not all) of the hierarchies that the member property is assigned to. When you create a member property, you must specify the

property name, column, and level. Member property names can be shared across a cube but must be unique for a specific level within a specific hierarchy. You can also specify a caption, description, and format. The format that you specify here is used instead of the format in the data set.

Property Statement

The PROPERTY statement is used with the PROC OLAP statement when you define a cube:

```
PROPERTY zipcode-region
    column=post_code
    hierarchy=geographic
    level=region;
```

Cube Designer

You can also establish member properties with the Member Property dialog box that is part of the Cube Designer wizard in SAS OLAP Cube Studio. This is accessed after measures are defined, when you create a cube or edit a cube. On the Member Properties page, select **Add** to create a member property. In the Define a Member Property dialog box, enter the member property name, level, column, and caption.

Defining Aggregations While Building a Cube

Adding Aggregations to a Cube

The Cube Designer wizard enables you to add aggregations to a cube when the cube is created. You can define aggregations that are built with the cube on the Aggregations page of the Cube Designer wizard. From here you can add and modify aggregations, specify performance options, and choose whether to create the default NWAY aggregation. Selecting the **Add** button opens the Add Aggregation page of the Cube Designer wizard. It enables you to select from available levels to create an aggregation.

Defining Stored Aggregations For a Fully Summarized Cube

You can also specify stored aggregations for a cube. When you are building a cube in the Cube Designer wizard, you can choose to build from a detail table, a star schema, or a fully summarized table. If you build a cube from a fully summarized table, you can *automatically* specify aggregations that are stored in aggregation tables. If you are building a cube from a detail table or a star schema, you can also select to include aggregations that are stored in aggregation tables. The General page of the Cube Designer wizard provides the check-box option **Cube will use aggregated data from other tables**. This option enables you to create your cube using aggregated data from tables other than the detail table or star schema that you selected as the input.

The Aggregation Tables page of the Cube Designer wizard enables you to select aggregation tables for the cube. From these aggregation tables, you can then select stored aggregations to include with the cube. On the Stored Aggregations page of the

Cube Designer wizard, you can add and modify a stored aggregation. After you have defined the stored aggregations, you can affix further aggregations that are built with the cube.

Saving the Cube Metadata or Creating the Physical Cube

When you are building a cube in the Cube Designer wizard you can choose to save only the cube's definition to the active metadata or to save the cube's definition and build the cube. Your options depend on whether the cube:

- is a new cube
- is defined only in the metadata
- physically exists

New Cubes

If you are creating a new cube, the Finish page of the Cube Designer wizard contains these options:

Save the metadata and create the cube

saves the cube's definition to the current metadata and creates the physical cube.

Save the metadata but do not create the cube

registers only the cube definition to the current metadata.

Cubes Only Defined in the Metadata

If you are editing a cube that exists only in the metadata, the Finish page contains these items:

Save the metadata and create the cube

saves definition changes to the current metadata and creates the physical cube.

Save the metadata but do not create the cube

registers the modified cube definition in the current metadata.

Cubes That Physically Exist

If you are editing a cube that physically exists, the Finish page contains these items:

Delete the cube, save the metadata updates, and re-create the cube

saves definition changes to the current metadata and creates the physical cube.

Any existing physical cube is deleted before the new cube is built. If the cube deletion fails, then the cube metadata is not updated.

Delete the cube, save the metadata updates, but do not re-create the cube

registers only the modified cube definition in the current metadata. Any existing physical cube is deleted, but the modified metadata is not used to build a new cube. If the deletion fails, then the cube metadata is not updated. You can use this option to delete a physical cube without changing the cube's definition; simply navigate to the Finish window without making any changes to the cube's definition.

The **Export Code** button opens the Export Code dialog box. Click **Finish** to process the cube and close the Cube Designer wizard. After processing, a standard SAS log window displays the short form of PROC OLAP that was submitted to create the cube.

Note: To perform tasks on the physical cube (such as deleting, rebuilding, and tuning), you must have the appropriate file access permissions at the operating system level. If you do not have access, contact your system administrator for more information. Δ

Saving the OLAP Procedure (Long Form versus Short Form)

The OLAP procedure is one of the SAS tools that you can use to create, update, and delete cubes. It provides statements that enable you to add dimensions, levels, hierarchies, measures and aggregations to a cube as it is being created. These statements, when combined with the PROC OLAP statement and METASVR statement, are often referred to as the long form of PROC OLAP. The long form of PROC OLAP is used when there is no metadata definition for a cube in the SAS metadata. When submitted, the long form of the procedure first creates the metadata definition of the cube in the SAS metadata and then creates the physical files for the cube. If a metadata definition for the cube specified on the PROC OLAP statement already exists, an error message is printed in the log.

In addition to basic cube creation tasks, you can use a shorter form of PROC OLAP to perform maintenance tasks on an existing cube metadata definition. The short form of PROC OLAP must include the PROC OLAP statement and the METASVR statement. Other statements, such as the DEFINE, AGGREGATION, DROP_AGGREGATION and UNDEFINE statements can also be included, depending on the task you want to perform. These include the following:

- delete the physical files for the cube while leaving the metadata definition intact (PROC OLAP with the DELETE_PHYSICAL option and METASVR statements)
- create the physical files for the cube based on the information in the existing metadata definition (PROC OLAP and METASVR statements)
- define new calculated members for a cube or remove existing calculated members (PROC OLAP, METASVR and DEFINE or UNDEFINE statements)
- define new aggregations or delete existing aggregations (PROC OLAP, METASVR, AGGREGATION or DROP_AGGREGATION statements)

You can use SAS OLAP Cube Studio to modify the metadata definition of a cube and then use the short form of PROC OLAP to rebuild the cube, including the changes. For example, if you want to add a dimension to an existing cube, edit the cube structure in SAS OLAP Cube Studio to add the dimension. Then, from the Finish panel of the Cube Designer, submit the short form of PROC OLAP to delete the physical files for the cube, save the changes to the metadata for the cube, and then submit the short form of PROC OLAP to re-create the physical files. An advantage to using the short form in this fashion is that links to access control entries, information maps, and reports will not be broken as the cube is updated.

To accomplish the same task using the long form of PROC OLAP, you would still use the short form with the DELETE option to delete the physical cube files and to delete the metadata definition. The appropriate DIMENSION statements would be added to the long form and the code submitted to create the new metadata and the new physical files. But now the access controls for the cube would need to be re-created and information maps and reports would need to be adjusted to account for the newly created cube.

Export Code

The Export Code function is available in SAS OLAP Cube Studio. The Export Code function enables you to store in a text file the PROC OLAP code that is used to edit or build the cube. There are two options to exporting the PROC OLAP code: the long form or the short form. The Export Code dialog box appears when you perform one of these steps:

Click **Export Code** on the Finish page in the Cube Designer wizard.

Select **Export Code** from the **Actions** menu.

On the Export Code dialog box, you can select one of the following options:

- 1 **Long form** – enables you to save the full version of the OLAP procedure code to a text file.
- 2 **Short form** – enables you to save the short version of the OLAP procedure code to a text file.

The Export Code function is also available from the Aggregation Tuning and Incremental Update functions. When accessed from these functions, the text file that is generated contains the short form of PROC OLAP.

For both the long form and short form options, enter a fully qualified path, including file extension, to the location for the text file. The text file that is created stores the PROC OLAP code that is used to create the currently selected cube. Use the **Browse** button to navigate to a location on the file system. You can also open the drop-down box for either the long or short form options. SAS OLAP Cube Studio maintains a record of the five most recently used files for you to select from. The most recently used file is at the top of the list.

The PROC OLAP code that is generated can contain MDX expressions. If the MDX expression contains mixed quotation marks (both single and double quotes), do not place quotation marks around the MDX string. You must resolve this issue before submitting the PROC OLAP code. When working with quotation marks follow these guidelines:

- If the MDX string contains only single quotation marks, then place double quotation marks around the MDX string.
- If the MDX string contains only double quotation marks, then place single quotation marks around the MDX string.

Note: The PROC OLAP code that is generated by the Export Code function does not include the user ID and password due to security reasons. You must manually add the user ID and password to the exported file before submitting. △

Viewing Cubes in SAS OLAP Cube Studio

After you have built a cube in SAS OLAP Cube Studio, you can examine the contents of the cube with the View Cube function. The View Cube function is available from the **Actions** menu and from the cube context menu. The View Cube function is available only for use with physically built cubes.

SAS OLAP Server Connection

When you select the View Cube function, a connection to the SAS OLAP Server must be made. You can identify the SAS OLAP Server for a cube by selecting the cube's OLAP schema in the tree view. Then select the **Server Assignment** tab on the Properties dialog box for that schema. If there is only one SAS OLAP Server assigned to the OLAP schema, then that server is used. If a SAS OLAP Server is not selected for the schema, a dialog box is displayed that enables you to select a SAS OLAP server. If a connection to the SAS OLAP Server cannot be made, then the View Cube dialog box cannot be launched.

Cube Permissions

In order to use the View Cube function, a cube must have the appropriate Read and ReadMetadata permissions. Read and ReadMetadata permissions can be granted for the cube with the Authorization function that is available in the Properties dialog box. If you do not have permissions to read the cube, the View Cube function fails when you attempt to load it. The following authorization error message is displayed:

```
View Cube results are not available. Possible reasons are
1) an error occurred in reading the data or
2) the permissions may not allow viewing
```

You can select **OK** to close the message dialog box or **Details**. The Details panel displays the following message:

```
The set is empty
```

In addition, if the Read permissions for a component of the cube (dimension, hierarchy, etc.) have been denied, those components are not available for selection in the cube dimension tree.


Disabled Cubes

If a cube has been disabled, you can see the same restriction as if the cube has been denied Read permissions. The View Cube function fails when you attempt to load it. The following authorization error message is displayed:

```
View Cube results are not available. Possible reasons are
1) an error occurred in reading the data or
2) the permissions may not allow viewing
```

You can select **OK** to close the message dialog box or **Details**. The Details panel then displays the following message:

```
The set is empty
```



CHAPTER

6

Modifying and Maintaining Cubes

<i>Editing a Cube</i>	56
<i>Renaming A Cube Object</i>	56
<i>Deleting Cubes and Cube Objects</i>	57
<i>Deleting a Cube in SAS OLAP Cube Studio</i>	57
<i>Deleting Cube Objects in SAS OLAP Cube Studio</i>	57
<i>Deleting the Physical Cube in SAS OLAP Cube Studio</i>	57
<i>Using the DELETE and DELETE_PHYSICAL OPTIONS</i>	57
<i>Refreshing Cube Metadata</i>	57
<i>MDX DDL REFRESH Statement</i>	58
<i>Tuning Cube Aggregations</i>	58
<i>SAS OLAP Cube Studio — Tune Aggregations Function</i>	58
<i>Aggregation Tuning dialog box</i>	59
<i>ARM Log</i>	60
<i>Cardinality</i>	60
<i>Manual</i>	60
<i>Using PROC OLAP to Tune Aggregations</i>	61
<i>Monitoring OLAP Server Performance</i>	61
<i>Specifying Tuning and Performance Options in Cube Aggregations</i>	61
<i>Setting Options on the Aggregation Tuning Dialog Box</i>	61
<i>Default Tab</i>	62
<i>Aggregation Tab</i>	62
<i>Setting Options with PROC OLAP</i>	62
<i>Specifying Calculated Members and Measures</i>	63
<i>Using the Calculated Members Wizard</i>	63
<i>Selecting Calculation Types</i>	64
<i>Simple Calculations</i>	64
<i>Time Calculations</i>	64
<i>Custom Calculations</i>	65
<i>Multiple Language Support and Dimension Table Translations</i>	65
<i>SAS OLAP Cube Studio and Dimension Table Translations</i>	66
<i>PROC OLAP and the USER_DEFINED_TRANSLATIONS Statement</i>	66
<i>SAS Servers and Character Encoding</i>	66
<i>Adding SAS System Options to a Cube</i>	66
<i>Synchronizing Column Changes</i>	67
<i>Accessing OLAP Cubes from SAS: SQL Pass-Through Facility for OLAP</i>	67
<i>Conversion Issues</i>	67
<i>VALIDVARNAME</i>	68
<i>Data Types</i>	68
<i>Specifying Member Names</i>	68
<i>SQL Pass-Through Example</i>	69
<i>PROC SQL Syntax</i>	69

<i>Connecting Directly to a SAS OLAP Server</i>	70
<i>Connecting to an OLAP Server Using the OLEDB Protocol</i>	70
<i>Specifying User Authentication Information</i>	71
<i>Impact Analysis and Reverse Impact Analysis</i>	71
<i>Disabling and Enabling Cubes</i>	72
<i>Disabling Cubes</i>	72
<i>Enabling Cubes</i>	72
<i>Cube Security</i>	73
<i>Setting Permission Conditions on Cube Dimensions</i>	73
<i>SECURITY_SUBSET PROC OLAP Option</i>	74
<i>Identity-Driven Security</i>	74
<i>Build Formula Dialog Box</i>	75
<i>Integrated Windows Authentication and Single Sign-On</i>	75
<i>Single Sign-On (SSO) to SAS IOM Bridge servers</i>	75
<i>SSO and SAS OLAP Cube Studio</i>	75
<i>Integrated Windows Authentication</i>	75
<i>Trusted Peer Connection</i>	76
<i>Exporting Code</i>	76
<i>Security for Drill-through Tables</i>	76
<i>Applying Batch Security with Permission Tables</i>	77

Editing a Cube

After a cube is created, it might be necessary to make changes to the cube. Whether you need to change a cube that exists only as metadata or has physical files built, you can modify the cube with the Edit Cube Structure function. This function launches the Cube Designer wizard, where you can make various changes to the cube including the following:

- changing the source data for the cube
- adding or removing structural data to the cube (dimensions, levels, hierarchies)
- changing properties for cube components (level properties, measure attributes, etc.)
- adding or removing cube measures, member properties, or aggregations

The Cube Designer wizard gives you the opportunity to edit and build the cube based on the existing cube metadata. When you are finished making changes to the cube in the Cube Designer wizard, you can choose to update the metadata or update the metadata and build the cube.

To access the Edit Cube Structure function, select the cube you need to edit in the SAS OLAP Cube Studio tree view. Then select **Edit Cube Structure** from the **Actions** menu or from the cube's context menu. This opens the Cube Designer wizard. The selected cube is opened and can now be edited.

Renaming A Cube Object

It might be necessary at times to rename a cube or other object (such as a job, document, or note) in SAS OLAP Cube Studio. You can change the name of an object by selecting that object in the tree view and then selecting **Rename** from the **Edit** menu or the object's context menu. Enter a new name for the object and click outside the object name. The object is now renamed. For further information on SAS naming conventions see the topic "Naming Guidelines and Rules for the SAS OLAP Server" on page 30.

Note: You must have WriteMetadata permissions to rename a cube. For more information on permissions, see the *SAS Intelligence Platform: Security Administration Guide*. Δ

Deleting Cubes and Cube Objects

Deleting a Cube in SAS OLAP Cube Studio

While working in SAS OLAP Cube Studio, you can delete a cube. The Delete function deletes the cube metadata definition and the cube's physical files. If you are deleting a SAS OLAP cube, you should note the difference between the Delete function and the Delete Physical Cube function. The Delete function deletes the entire cube, whereas the Delete Physical function deletes the physical cube files but maintains the cube metadata.

The Delete function should be used on a cube only if you are re-creating the cube from a completely different data set or table. The use of the Delete function will remove all information about a cube, including security information and information maps.

Deleting Cube Objects in SAS OLAP Cube Studio

The tree view in SAS OLAP Cube Studio lists the different objects that are used to create and manage cubes, including tables, libraries, schemas, and jobs. As with cubes, you can delete an object by selecting **Delete** from either the **Edit** menu or the object's context menu. After selecting **Delete**, select **Yes** to delete the object.

Deleting the Physical Cube in SAS OLAP Cube Studio

When building and maintaining SAS OLAP cubes, it might be necessary at times to delete the physical files for a cube but maintain the cube metadata. You can delete just the physical files for a cube by using the Delete Physical Cube function in SAS OLAP Cube Studio. The cube metadata is maintained and the cube is then listed in the tree view as a cube metadata definition.

Using the DELETE and DELETE_PHYSICAL OPTIONS

If you are manually editing and deleting cubes, you can use the DELETE or DELETE_PHYSICAL options with the PROC OLAP statement.

Refreshing Cube Metadata

When a cube is accessed by a SAS OLAP Server, the cube data is read by the OLAP Server and then reported when queries against the cube are made. It is possible that changes and updates to a cube can occur after the OLAP Server accesses the cube. To make certain the data that is being queried is the most recent, the Refresh Cube function can be used. The Refresh Cube function enables the SAS OLAP Server to access cube data that was created or updated since the cube was last accessed by the SAS OLAP Server. You might also want to refresh a cube that has had changes or

functions performed since you opened the cube in SAS OLAP Cube Studio. For example, if a cube has had calculated members or named sets defined by other users, the Refresh Cube function would make the changes visible with the Cube Viewer function.

To refresh a cube, select the cube in the tree view. Select **Refresh Cube** from the **Actions** menu or from the cube's context menu. The Refresh Cubes function then updates the OLAP server metadata for the cube. If the refresh process is successful, you receive a message in the lower left corner of the SAS OLAP Cube Studio window stating: "Cube Refresh of "*cubename*" is successful".

You must be connected to a server and have administrative permissions in order to select the Refresh Cubes function.

MDX DDL REFRESH Statement

The REFRESH statement can be sent manually. You can send the REFRESH statement for each additional server that the schema is associated with.

```
refresh cube (cubename | "_ALL_" )
```

Cubename specifies a single cube to refresh for the current server connection. `_ALL_` specifies that all cubes are refreshed for the current server connection. Here are some examples.

This example uses the REFRESH statement to refresh the metadata associated with a cube named OrionStar.

```
refresh cube [OrionStar]
```

This example uses the REFRESH statement to refresh the metadata for all cubes managed by the connected server.

```
refresh cube _ALL_
```

You can use the OLAP MDX SQL Pass-Through facility to send the DDL REFRESH statement to a server. Here is an example.

```
proc sql noerrorstop;
  connect to olap (&olapcon);
  execute
  (
    refresh cube [OrionStar]
  ) by olap;
```

```
proc sql noerrorstop;
  connect to olap (&olapcon);
  execute
  (
    refresh cube _ALL_
  ) by olap;
```

Tuning Cube Aggregations

SAS OLAP Cube Studio — Tune Aggregations Function

After you have built a cube in SAS OLAP Cube Studio, you can add new aggregations to the cube. You can also fine-tune the existing aggregations that were

built with the cube and delete any unnecessary aggregations. The Aggregation Tuning function enables you to generate new aggregations based on cube data cardinality or ARM log data. You can also manually build an aggregation, selecting the exact hierarchies and levels that you want to generate aggregations from. To modify a cube's aggregations, select the cube from the tree view and select **Aggregation Tuning** from the **Actions** menu or from the cube's context menu.

Aggregation Tuning dialog box

The Aggregation Tuning dialog box provides functionality to either automatically generate aggregation recommendations or to manually define cube aggregations. Aggregations that are defined with the Aggregation Tuning function are added to a list of any existing aggregations. This list of aggregations can then be reviewed and edited. You can modify performance options for aggregations and, if needed, you can remove any unnecessary aggregations. Finally, you can build the aggregations for the cube. You can also export the resulting PROC code to a text file.

The Aggregations table displays the aggregations for the cube. This includes both existing aggregations and newly created aggregations that can be built. You can select individual aggregations to edit, or you can select multiple, sequential rows of aggregations. The information displayed on the Aggregations table includes the following columns:

Status	This column indicates the current status of the individual aggregations.
Count	This is a statistic that is generated from the ARM log. It shows the number of queries that match the aggregation definition.
Total Time	This is a statistic that is generated from the ARM log. It shows the total wall time (for internally accessing the cube data) for queries that match the aggregation definition. The time value is displayed in hr : min : sec : millisecond format.
Average Time	This is a statistic that is generated from the ARM log. It shows the average wall time (for internally accessing the cube data) for queries that match the aggregation definition. The time value is displayed in hr : min : sec : millisecond format.

Note: If an aggregation already exists, it can have its performance values updated from the Arm log. Δ

The Aggregation Tuning Dialog also contains the following options:

Export Code

This option opens the Export Code dialog box. You can preview and save to a text file the PROC OLAP code that is used to update the cube aggregations. The aggregations that are new, modified, or dropped in the Aggregations table are written to the text file. This button is inactive until you add a new aggregation to the Aggregations table or drop an existing aggregation from the table.

Update Aggregations

This option enables you to update the aggregations that are listed as new or modified in the Aggregations table. When **Update Aggregations** is selected, the cube is updated to include the new aggregations and changes to any existing aggregations. This button is inactive until you add a new aggregation to the Aggregations table, modify the name of an existing aggregation, or drop an existing aggregation from the table. The Aggregations Tuning dialog box closes after each build.

In the Aggregations Tuning dialog box, you can select one of the following methods to define aggregations for a cube:

ARM Log

ARM analysis is used to monitor and diagnose the performance of various SAS applications. It enables you to measure and record application performance and query response times in a designated log file. When using an ARM log to perform ARM analysis of cube aggregations, you can select either of the following options on the ARM :

Create aggregation recommendations and update existing performance values based on the ARM log.

You can add aggregations from the ARM log and update the statistics for the aggregations that have entries in the ARM log. This is the default option.

Update performance values based on the ARM log.

You can update the statistics for aggregations that have entries in the ARM log. This option is used when you want to examine the information in the ARM log and verify that the existing aggregations are being used. This can help determine which aggregations to drop.

The **Analyze** button opens the Analysis Recommendations dialog box. From here you can select one or more recommended aggregations. The aggregations you select are added to the bottom of the Aggregations table as highlighted rows. The **Analyze** button is inactive until text is entered in the **Enter an ARM log file** field.

For ARM log analysis, the generated aggregations list displays all aggregations that can be determined from the cube query data that is provided in the ARM log. If you select all of the generated aggregations and add them to the Aggregations table, you cannot generate further aggregations from that same ARM log. If you reselect the **Analyze** button on the **ARM Log** tab, you receive a message stating: "The ARM analysis did not recommend any additional aggregations to add to the cube".

Note: For further information about ARM logging, see "SAS OLAP Server Monitoring and Logging" in the *SAS Intelligence Platform: System Administration Guide* and the *SAS Interface to Application Response Measurement (ARM): Reference*. Δ

Cardinality

The **Cardinality** tab enables you to add aggregations that are recommended based on the relative cardinality (number of members) of the cube levels. This method of adding aggregations is used when a cube is first created and before an ARM log can be generated. For each aggregation build, the cardinality algorithm generates up to 100 aggregation recommendations that are based on the cardinality ratio between the cube levels. The aggregations with the highest cardinality are recommended. All dimensions are included in the analysis, and any dimension of type TIME has up to the first two of its levels included on each aggregation.

Manual

The **Manual** tab enables you to select the exact hierarchies and levels that you want to generate aggregations from. On the **Manual** tab, the hierarchies for the cube are listed individually as columns. Levels for the hierarchies are numerically listed in drop-down lists on the columns. The default selection value of each column is **None**. When you select an individual level from a hierarchy, you are selecting that level and its parent levels.

Using PROC OLAP to Tune Aggregations

To modify an aggregation through PROC OLAP, use the DROP_AGGREGATION statement to delete the aggregation, and then use the AGGREGATION statement to define the new aggregation.

- DROP_AGGREGATION *level-name1* < *level-name2* ...*level-nameN* > /
NAME='aggregation-name' ;
- AGGREGATION *level-name1* / < NAME='aggregation-name' > ;

For more information about the DROP_AGGREGATION and AGGREGATION statements, see “The OLAP Procedure” on page 244.

Monitoring OLAP Server Performance

SAS OLAP Server performance is monitored and logged with the Application Response Measurement (ARM) interface. The ARM interface provides built-in logging capabilities and generates log records that indicate query content and query start and completion times. From this data, information regarding aggregation usage, individual query response times, or throughput can be determined. Traditionally, ARM enables system administrators to monitor application executions, run times, performance, and completion. SAS OLAP Server uses ARM to monitor the following:

- application behavior
- user behavior and usage
- server loads
- cube optimization (query response time)
- cube metrics—counts of connections and queries

For more information, see “Monitoring Performance Using Application Response Measurement (ARM)” in *SAS Language Reference: Concepts*.

Specifying Tuning and Performance Options in Cube Aggregations

When you build cubes, you can set various options that improve and optimize cube creation and query performance. These options can be set for all aggregations in a cube or for a specific aggregation. Additionally, these options can be set by using the PROC OLAP options or in SAS OLAP Cube Studio. These options are stored with the cube metadata in the SAS metadata.

Setting Options on the Aggregation Tuning Dialog Box

In the Cube Designer – Aggregation Tuning dialog box, the **Options** button is provided for access to tuning options. Select the **Options** button to open the Performance Options dialog box. There are two tabs for setting tuning options, the **Default** tab and the **Aggregation** tab.

Default Tab

The default performance options are applied to all aggregations for the cube. These performance options include the following:

- amount of memory (in megabytes) that is available for aggregation creation
- maximum number of threads that are used to create an aggregation index
- number of aggregations to create in parallel
- partition size (in megabytes) of aggregation table partitions
- number of observations (in kilobytes) to include in the index component file segment
- location of index component files
- location of partitions in which to place aggregation table data
- aggregation tables that are stored in compressed format.

For specific information about these functions, see the Performance Options - Default tab Help page in SAS OLAP Cube Studio Help.

Aggregation Tab

The aggregation-specific performance options are applied to an individual aggregation for the cube and override the global option settings for that aggregation. You can define and modify performance options for an aggregation or delete options for an aggregation. The aggregation-specific performance options include the following:

- partition size (in megabytes) of aggregation table partitions
- number of observations (in kilobytes) to include in the index component file
- location of index component files
- location of partitions in which to place aggregation table data
- aggregation tables stored in compressed format
- aggregations created with indexes

For specific information about these functions, see the Performance Options - Define tab Help page in SAS OLAP Cube Studio Help.

Setting Options with PROC OLAP

You can set options for all aggregations in a cube or for a specific aggregation. To set options for all aggregations, set the options in the PROC OLAP statement. To set options for a single aggregation, set the options in the PROC OLAP AGGREGATION statement. The options include the following:

ASYNINDEXLIMIT= *n*

specifies a limit on the number of indices that will be created in parallel during the cube build process.

COMPRESS | NOCOMPRESS

specifies whether or not to store the aggregation tables in a compressed format on disk.

CONCURRENT=*n*

specifies the maximum number of aggregations to create in parallel.

DATAPATH=(*pathname1' ...'pathnameN'*)

specifies the location of one or more partitions in which to place aggregation table data.

INDEXPATH=(*'pathname1' ...'pathnameN'*)

specifies the locations of the index component files that correspond to each aggregation table partition as specified by the DATAPATH= option.

INDEXSORTSIZE=*n*

specifies the amount of memory in megabytes that is available when aggregations are created. The default is the system's available memory.

MAXTHREADS=*n*

specifies the maximum number of threads that are used to asynchronously create the aggregation indexes.

INDEX | NOINDEX

specifies whether or not to create the aggregations with indexes.

PARTSIZE=*partition-size*

specifies the partition size in megabytes of the aggregation table partitions and their corresponding index components.

SEGSIZE=*rows-per-segment*

specifies the number of observations (table rows) in kilobytes to include in the index component file segment.

WORKPATH=*pathname*

specifies one or more locations for temporary work files.

Note: ASYNCINDEXLIMIT=, CONCURRENT=, INDEXSORTSIZE=, and MAXTHREADS= are available only in the PROC OLAP statement. △

For more information about these options, see “PROC OLAP Statement” on page 245 and “AGGREGATION Statement” on page 274.

Specifying Calculated Members and Measures

After you have built a cube, you might find it necessary to add members or measures to the cube. You can add calculated members or measures to a cube in the following ways:

- manually by submitting PROC OLAP code with the DEFINE statement. See “DEFINE Statement” on page 276.
- using the Calculated Members Wizard

The Calculated Members function enables you to add new members for the Measures dimension. A calculated member is a definition (for a dimension member) that you create and store with the cube. The calculated member value is generated later during query time. You can also define a calculated member as a measure.

Using the Calculated Members Wizard

The Calculated Members function can be accessed from the SAS OLAP Cube Studio **Actions** menu or from the context menu for a cube. Click **Calculated Members** to open the Calculations dialog box.

In the Calculations dialog box, you can add new calculated members or modify existing members for the selected cube. The list box displays all calculated members that are defined in the metadata for the selected cube. Select **Add** to launch the New Member Wizard and define a new calculated member.

To modify an existing member, select the member and then click **Edit**. The Edit a Calculated Member dialog box opens.

Selecting Calculation Types

When you click **Add** in the Calculations dialog box, the New Member wizard is launched. In this wizard, you select the type of calculation that you want to create. You can select one of the following types:

- Simple calculations** to create a simple calculation formula. See “Simple Calculations” on page 64.
- Time analysis calculations** to create a time-based calculation formula. See “Time Calculations” on page 64.
- Custom calculations** to enter a custom calculation formula. See “Custom Calculations” on page 65.

Simple Calculations

On this page, you select the variables to create a simple calculation. You can select one of the following calculation types:

- Sum
- Difference
- Ratio
- Percent Increase
- Percent Decrease
- Distinct Count

The **Formula** drop-down lists for the members are populated with all measures, including calculated measures. The **Formula** panel changes depending on the calculation-type radio button you select. When a formula is selected, the **Formula** panel is populated with drop-down lists appropriate for the selected calculation. A member must be selected in each list. The member selections for the different formulas are preserved across formulas within a type.

If you select **Distinct Count**, a selection tree is displayed. Each dimension has a node that expands into its hierarchies and then into the levels for each hierarchy. One variable in the hierarchy tree must be selected.

For further information on the Simple Calculations page, see the SAS OLAP Cube Studio Help.

Time Calculations

On this page, you select the variables to create a time calculation. You can select one of the following calculation types:

- Opening Balance
- Closing Balance
- Rolling Total
- Average Over Time
- Compare Parallel Periods
- Compare Consecutive Periods

The **Formula** panel changes depending on the time-calculation radio button that you select. The **Existing measure** drop-down list is populated with all measures, including calculated measures. The **Time period** drop-down list is populated with members from

the Time dimension. A member must be selected in each list. The member selections for the different formulas are preserved across formulas within a type.

Note: This option is not available if there are no TIME-type dimensions for the cube. Δ

For further information on the Time Calculations page, see the SAS OLAP Cube Studio Help.

Custom Calculations

On this page, you enter the variables to create a custom calculation. You can enter the following:

Parent dimension

specifies a parent dimension for the new calculated member that you are creating. The default dimension is Measures. If you select a dimension other than Measures, this field is initially populated with the unique name for the default hierarchy's ALL member.

Parent member

specifies a parent member for the new calculated member. Use the **Browse** button to display a tree of valid members for the dimension. If you select the **Browse** button and the cube physically exists, you are prompted to log on to an OLAP server, if you haven't already.

Note: If the cube exists, you must connect to the OLAP server so that the valid members can be retrieved from the server. If the cube does not exist or a connection cannot be made to the OLAP server, the valid members displayed in the tree will be the ALL members for each of the hierarchies for the dimension. Δ

Name

specifies a name for the custom calculation.

Formula

specifies an MDX formula for the new calculated member.

Click **Verify** to validate the MDX formula entered in the **Formula** text box. This function is only available for cubes that physically exist. When you select **Verify**, the Log on to an OLAP server dialog box opens. Select a server and enter a user ID and password. A connection is made to the selected OLAP server. Click **Clear** to clear the contents of the **Formula** text box.

The **Generated MDX** text area allows you to view the MDX code that currently exists for the selected cube and any new MDX code that is entered in the **Formula** text box.

For further information on the Custom Calculations page, see the SAS OLAP Cube Studio Help.

Multiple Language Support and Dimension Table Translations

OLAP cube data is often generated in one language but needed in other languages. For example, a company's OLAP cube data might be stored in English, but users who speak Spanish and Turkish need access to it. So, the member values as well as the captions that are assigned to dimensions, hierarchies, levels, and so on need to be translated. Multiple language support is available only for cubes that are loaded from star schemas. It is used to read your alternate locale data sets and create locale-specific metadata for use at query time. Query results are returned in the language of the requested locale.

You can specify language support when building a cube either in the Cube Designer wizard or with PROC OLAP code. There are 56 possible language locales, and English is the default language.

SAS OLAP Cube Studio and Dimension Table Translations

In the Cube Designer – General page, select the **Advanced** button. If you selected **Star Schema** as the input type in the Cube Designer – General window, you will see the **Dimension Table Translations** tab. From the **Available** list, select the needed languages for the translation tables and move it to the **Selected** list. The first language in the **Selected** list is the default language.

PROC OLAP and the USER_DEFINED_TRANSLATIONS Statement

The USER_DEFINED_TRANSLATIONS statement is used in conjunction with the DIMENSION statement options DIMTABLEMEMMPREF= and DIMTABLELIBREF=. For more information, see the “DIMENSION Statement” on page 257.

SAS Servers and Character Encoding

If your server metadata contains characters other than those typically found in English, then you must be careful to start your server with an ENCODING= or LOCALE= system option that accommodates those characters. For example, a SAS server started with the default US English locale cannot read metadata that contains Japanese characters. SAS will fail to start and log a message indicating a transcoding failure.

In general, different SAS jobs or servers can run different encodings (such as ASCII/EBCDIC or various Asian DBCS encodings) as long as the encoding that is used by the particular job or server can represent all the characters of the data being processed. In the context of server start up, this requires that you review the characters used in the metadata describing your server (as indicated by the SERVER= object server parameter) to ensure that SAS runs under an encoding that supports those characters.

Adding SAS System Options to a Cube

When you build an OLAP cube, it is often necessary to include additional SAS code that must run prior to the creation of the cube. This includes the creation of user-written formats, PROC statements, and format search paths for the formats that are used on input tables. The Advanced Cube Options dialog box is accessed from the Cube Designer - General dialog box. It provides two entry tabs, **Submit SAS Code** and **Format Search Path**. Both tabs provide entry fields for SAS code. You can enter any text in the fields. There is no validation of the text that is entered. However, error messages are sent to the SAS log.

The text is saved to the cube metadata in the SAS Metadata Repository and is used every time the cube is created. You can edit or remove the text after it is initially entered. Highlight the text and make any needed changes, or use the DELETE key to remove the text. Select **OK** to save your changes.

Submit SAS Code

You can use this tab to enter a PROC statement or any SAS code that you want to submit before the cube is created. SAS code is submitted before any format search path.

Format Search Path

You can use this tab to enter names of catalogs or libraries for the format search path. The catalogs and libraries must be separated by a blank and will be searched in the order in which they are listed. You use the SAS system option FMTSEARCH= here.

Note: For more information about SAS formats, see “Formats” in *SAS Language Reference: Dictionary*. △

Note: When you build a cube with SAS OLAP Cube Studio, the format search path is saved with the cube metadata in the SAS Metadata Repository and used every time the cube is recreated. However, if you submit PROC OLAP code through a SAS session, outside of SAS OLAP Cube Studio, the format search path is ignored. PROC OLAP does not read the information from the SAS Metadata Repository or write the information to the SAS Metadata Repository. △

Synchronizing Column Changes

The Synchronize Column Changes function in SAS OLAP Cube Studio enables you to synchronize a cube when a table for an existing cube has encountered a column name change. This function finds the name differences between the cube and its input table and updates the internal cube files with the name change. The Synchronize Column Changes function is available for a cube if the cube physically exists and you have WriteMetadata permissions for that cube.

You can access the Synchronize Column Changes function in SAS OLAP Cube Studio from the **Actions** menu or from the context menu for a cube. If you are not connected to a workspace server when you select a cube, you are prompted to select a workspace server.

Note: You cannot update an unregistered cube. △

Accessing OLAP Cubes from SAS: SQL Pass-Through Facility for OLAP

The SQL Pass-Through facility for OLAP enables a SAS user to connect to an OLAP server and execute queries against OLAP data within the PROC SQL environment. PROC SQL establishes a connection to an OLAP server by using the PROC SQL CONNECT statement.

After a connection is made to the OLAP server, multiple queries can be submitted by using the OLAP multidimensional expressions (MDX) query language. These queries are run against existing OLAP cubes. While using the SQL Pass-Through facility, you can use the MDX DDL statement to create and delete named sets and calculated members, scoped either globally, or for the current server connection. A PROC SQL query is then closed after all observations (rows) of data are returned. To disconnect from the server, you must submit the PROC SQL DISCONNECT statement.

Conversion Issues

OLAP cube data is multidimensional and flexible in regard to data name lengths and restrictions. However, when PROC SQL sends a query to the OLAP server, data is

returned in a flattened, tabular format that contains rows (observations) and columns (variables).

The SAS OLAP Server has unique naming conventions that specify valid column names, lengths, and types. Column names that are returned from the SAS OLAP Server can contain characters (periods, spaces, brackets) and can be unrestrained in length. Additionally, OLAP data types can be variable-length strings, floating-point numbers, or integers. This differs from SAS data set naming conventions, and some conversion is necessary.

VALIDVARNAME

The SQL Pass-Through facility supports the existing SAS option `VALIDVARNAME`. You can specify the `VALIDVARNAME` option to control variable names. The current default setting for `VALIDVARNAME` is `V7`, and variable names can be a maximum of 32 characters in length. Each variable must start with a letter or the underscore character and can contain letters, underscores, and numbers. Uppercase and lowercase letters are also allowed.

When converting column names to SAS variable names, the SQL Pass-Through facility for OLAP will do the following:

- truncate the column name to the maximum size that is allowed
- replace any invalid characters with an underscore
- use a numeric suffix to differentiate between duplicate variable names that are generated during the data conversion

Note: For additional information about naming restrictions for the SAS OLAP Server, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30. \triangle

Note: For further information about the `VALIDVARNAME=` system option, see “`VALIDVARNAME=System Option`” and “Names in the SAS Language” in the *SAS Language Reference: Dictionary*. \triangle

Data Types

OLAP query results that contain member names or strings are converted to a fixed length `CHAR` type. All OLAP numeric types are converted to standard SAS numeric types (8-byte floating point). Missing values are handled by standard SAS conventions.

Specifying Member Names

When specifying member names, single quotation marks and double quotation marks are uniquely processed by the SQL Pass-Through facility. The following exceptions occur when using quotation marks in member names.

A member name contains a single quotation mark.	If a member name in the OLAP query contains a single quotation mark (for example, Adam’s), the entire query must be enclosed in double quotation marks when passed to the SQL Pass-Through facility. This ensures that the single quotation mark is maintained in the query text. The double quotation marks are then removed before the OLAP query is processed on the OLAP server.
---	--

Note: Macro variables are not expanded within strings that contain single quotation marks. This includes MDX statements enclosed in single quotation marks. \triangle

A member name contains double quotation marks.	If a member name in the OLAP query contains double quotation marks, the entire query must be enclosed in single quotation marks when passed to the SQL Pass-Through facility. This ensures that
--	---

the double quotation marks are maintained in the query text. The single quotation marks are then removed before the OLAP query is processed on the OLAP server.

Here is an example of the SQL Pass-Through facility with quotation marks. The member name *Adam's* requires the double quotation marks around the SELECT statement.

```
proc sql;
connect to olap (host=localhost port=5451 user=user pass=pass);
  select * from connection to olap (
    "select [A].[All A].[M].[Adam's] on Rows,
      crossjoin([Measures].defaultMember, [B].defaultMember)
    on Columns
      from [QuoteTest]"
  );
disconnect from olap;
quit;
```

SQL Pass-Through Example

In the following example, PROC SQL connects to the SQL Pass-Through facility for OLAP to create a new data set named temp, which contains all the variables that are returned from the multidimensional expression (MDX) defined in the SELECT query. The OLAP server returns query results in a tabular format known as a flattened rowset. The table rows become the observations of the output data set, and the table columns become the variables. After all the rows are returned, PROC SQL closes the query. The server connection is terminated when the program encounters a DISCONNECT statement or when the PROC SQL step ends.

Note: Because the OLAP server does not impose the same restrictions on column names, types, and lengths that SAS imposes on data sets, some conversion might be required. Δ

```
100      proc sql;
101          connect to olap (host=localhost service=olap1);
102          create table temp as select * from connection to olap
103              (
104                  select { dealers.dealer.members } on 0,
105                      { [cars].[Car].members,
106                        [cars].[Color].members } on 1
107                  from mddbcars
108              );
109          disconnect from olap;
110          quit;
```

PROC SQL Syntax

CONNECT TO OLAP|OLEDB

establishes a connection to an OLAP server. This statement is required. Specify OLAP to connect to a SAS OLAP server. Specify OLEDB to connect to all other

OLEDB for OLAP-compliant servers. When connecting to an OLAP server, implicit connection is not supported. Therefore the CONNECT TO OLAP statement is required. However, when connecting to an OLEDB for OLAP compliant server, implicit connection is supported. For further information about the CONNECT statement, see “Pass-Through Facility for Relational Databases” in the *SAS/ACCESS for Relational Databases: Reference*.

DISCONNECT FROM OLAP | OLEDB
ends the connection to the OLAP server.

EXECUTE (MDX DDL statement) BY OLAP | OLEDB
specifies an MDX DDL statement for creating or deleting calculated members or named sets. This option is executed by the SAS OLAP Server. For further information, see “Basic MDX DDL Syntax” in the *SAS OLAP Server: MDX Guide*.

SELECT . . . FROM CONNECTION TO OLAP | OLEDB (MDX statement);
specifies the MDX query that is sent to the connected SAS OLAP server.

Connecting Directly to a SAS OLAP Server

When connecting directly to a SAS OLAP server, you specify a SAS OLAP server in the CONNECT statement. For example:

```
CONNECT TO olap (host=localhost service=olap1);
```

You can also submit MDX commands to query SAS OLAP cubes and create various temporary OLAP components that can be used during a PROC SQL session including named sets and drill-through paths.

Here are the connection options for a SAS OLAP Server that is started with bridge protocol:

HOST=*machine-name*
specifies either the DNS name or the IP address of the machine that is hosting the OLAP server.

PORT=*port-number* | SERVICE=*service-name*
either the *port-number* or *service-name* is required. The *port-number* specifies the numeric value of the port on which the OLAP server resides. The *service-name* is used to look up the port number of the machine that is hosting the OLAP server.

USER=*userid*
a string that specifies the user’s identification for the specified SAS OLAP server. If included, this option is enclosed with parentheses with the required arguments and any options.

PASS=*password*
a string that specifies the password for the user who is identified with the USER= option. If included, this option is enclosed with parentheses with the required arguments and any options.

Connecting to an OLAP Server Using the OLEDB Protocol

When connecting to OLE DB for OLAP data the SAS/ACCESS interface to OLE DB is used. With this approach you specify OLEDB in the CONNECT statement. For example:

```
CONNECT TO oledb (provider=msolap prompt=yes);
```

With this approach you can create a PROC SQL view of the data or specify MDX statements to access the OLE DB for OLAP data directly. In this approach the MDX

statements can be used for read-only access of OLE DB for OLAP data. For further information about the CONNECT statement, see “Syntax for the Pass-Through Facility for Relational Databases” in the *SAS/ACCESS for Relational Databases: Reference*. Also see “Accessing OLE DB for OLAP Data” in the *SAS/ACCESS 9.1 Supplement for OLE DB (SAS/ACCESS for Relational Databases)*.

Note: For detailed information about PROC SQL syntax, see “Overview of the Pass-Through Facility” in the *SAS/ACCESS for Relational Databases: Reference* and “Syntax for Remote SQL Pass-Through (RSPT) Facility” in the *SAS/SHARE User’s Guide*. Δ

Specifying User Authentication Information

When you use the SQL Pass-Through facility to connect to an OLAP server, you can choose how the user authentication information is identified. You can choose to store the user information in the PROC SQL statement or enter the information at the time of connection to an OLAP server.

- Enter the user information with the connect statement.** – The user name and password are included and stored in the PROC SQL syntax. For example:

```
connect to olap (host=localhost port=5451 user="user" pass="pass");
```

- Enter the user information at the time of connection to an OLAP server.** – Authentication with an OLAP server occurs at the time of connection and is accomplished with *Integrated Windows authentication*. You do not have to store a user name or password in the connection string. For example:

```
connect to olap (host=localhost port=5451);
```

This option works in two instances:

- 1 If both the OLAP server and the SAS client are started with the -SSPI option, then authentication is attempted using Integrated Windows authentication. This requires that both the OLAP server and the user’s computer be Windows PCs that have access to a shared domain server. This usually occurs when both machines are on the same network.
- 2 If the -SSPI option is not specified on either the client or the OLAP server, or if the Integrated Windows authentication fails, an attempt is then made to contact a metadata server to create a one-time password. This is possible when the OLAP server and the client use the same metadata server. The access information for the metadata server can either pre-exist in a SAS profile or otherwise be queried using a dialog box at runtime.

Note: For more information on Integrated Windows authentication, see the *SAS Intelligence Platform: Security Administration Guide*. Δ

Impact Analysis and Reverse Impact Analysis

The Impact Analysis function enables you to view the relationship of a cube to the objects that it is associated with. Impact analysis shows the potential impact of changes that you might make to a cube. In SAS OLAP Cube Studio, you can perform impact analysis only on a cube. If a cube has an information map, it is also displayed. The Reverse Impact Analysis function enables you to see how changes to the cube’s associated objects could affect the cube. This view shows all tables that the cube is

composed of, including the input, detail, and aggregation tables. Impact Analysis and Reverse Impact Analysis display only the objects for which you have ReadMetadata permission.

To perform impact analysis on a cube, select a cube in the tree view and select **Impact Analysis** from the **Tools** menu or from the cube's context menu. The Impact Analysis dialog box opens. The dialog box contains the **Impact Analysis** tab and the **Reverse Impact Analysis** tab. Each tab contains two options that enable you to toggle between a tree view and a diagram view of the cube and its objects. The diagram view shows a process flow diagram of the cube.

Each node of the cube analysis also contains a context menu that is available on both the **Impact Analysis** and **Reverse Impact Analysis** tabs. This context menu enables you to see the object's properties, analyze columns, and open (view the data in) a table. You can also print the analysis view at any time.

Note: You can use SAS Data Integration Studio to perform impact analysis on other objects, such as tables and jobs. For further information on the use of the Impact Analysis function in SAS Data Integration Studio, see "Using Impact Analysis" in the *SAS Data Integration Studio: User's Guide*. \triangle

Disabling and Enabling Cubes

Disabling Cubes

In SAS OLAP Cube Studio, you can disable and enable cubes to perform various cube changes and adjustments. You can disable a SAS OLAP cube if you need to perform administrative tasks or make edits to the cube. You might disable a cube to perform the following tasks:

- tuning cube aggregations
- adding calculated members to a cube
- changing the security settings for a cube
- renaming a cube
- editing and rebuilding a cube
- changing the OLAP schema for a cube

Disabling a cube removes it from production. As a result, queries against the cube are affected. Queries that are currently running against the cube are completed. However, new queries against the cube are not accepted by the OLAP server. To disable a cube in SAS OLAP Cube Studio, select a cube from the tree view. Then select **Disable** from either the **Actions** menu or the cube's context menu. The cube is then disabled in the tree view.

Note: You must have Administer permissions on the OLAP server to disable a cube. For more information on permissions, see the *SAS Intelligence Platform: Security Administration Guide*. \triangle

Enabling Cubes

You can enable a SAS OLAP cube that you have previously disabled. A cube is disabled in order make edits to the cube or to perform administrative tasks on the cube. Enabling a cube brings it into production, so that new queries against the cube are accepted by the OLAP Server. In order to enable a cube in SAS OLAP Cube Studio,

select a disabled cube from the tree view. Then select **Enable** from either the **Actions** menu or the cube's context menu. The cube is then enabled and listed as an enabled cube in the tree view.

After the cube is enabled, you can resume queries to the cube that require SAS OLAP Server sessions. In SAS OLAP Cube Studio, you can also use the View Cube function.

Note: You must have Administer permissions on the OLAP server to enable a cube. For more information on permissions, see the *SAS Intelligence Platform: Security Administration Guide*. Δ

Cube Security

Setting Permission Conditions on Cube Dimensions

SAS security enables you to set authorization permissions on a cube and the various components of a cube, such as hierarchies and levels. You can also apply member-specific filters to cube data by using MDX expression filters known as *permission conditions*. Permission conditions limit access to members of a cube dimension so that only designated portions of the data are visible to a user or group of users. With SAS OLAP data, permission conditions impose only explicit grants of the Read permission and can be specified only on dimension objects.

Permission conditions rely on the SAS OLAP Server for enforcement. At query time, the server performs the filtering to determine which dimension members should be returned to each requesting user. The members that are returned by the MDX expression must all belong to the dimension on which the permission condition is defined. Each filter consists of an MDX expression which subsets the data in a dimension as appropriate for a particular user or group. The filters are stored in the cube metadata.

You can set permission conditions from within SAS OLAP Cube Studio or from within SAS Management Console using the **Authorization** tab. You can define new permission conditions and modify existing permission conditions for a user or group identity. To create a permission condition, follow these steps:

- 1 From within SAS OLAP Cube Studio, select a cube in the tree view and drill to a dimension. In the dimension's Properties dialog box, select the **Authorization** tab.
- 2 Select (or add) the user or group whose Read access you want to limit. In the permissions list, add an explicit grant of the Read permission for that user or group.
- 3 If the selected user or group does not already have a permission condition defined, the **Add Authorization** button is now enabled. Click **Add Authorization** to open the Add Authorization dialog box. You can now create an MDX expression filter by specifying members and permissions for the expression filter. From this dialog box, you can also launch the Build Formula dialog box and create an MDX expression filter that limits the current dimension as appropriate for the selected user or group.
- 4 If the selected user or group does already have a permission condition defined, the **Edit Authorization** button is now enabled. Click **Edit Authorization** to open the Build formula dialog box and create an MDX expression filter.

See the example "Setting Member Authorizations On A Dimension" on page 180 for more details.

Note: For more information on OLAP member-level security, see the SAS OLAP Member Authorization Help in SAS Management Console. Also, see "OLAP

Member-Level Permissions” in the *SAS Intelligence Platform: Security Administration Guide*. \triangle

SECURITY_SUBSET PROC OLAP Option

After you have set the necessary permission conditions for dimension members, you must indicate whether the OLAP Server includes these permission conditions when processing cube queries. In order for a cube to control the roll-up values for designated members, the SECURITY_SUBSET PROC OLAP option must be set to YES when the cube is built. The SECURITY_SUBSET= option is used in conjunction with permission conditions to specify whether the roll-up values include only the members in the permission condition or all members in the cube. If you set SECURITY_SUBSET= YES, cell values are recalculated at query time based on the security subset defined by the active permission conditions for the given user. If you set SECURITY_SUBSET=NO, the OLAP Server does not recalculate the cell values. The default value (NO) includes all members within a total.

Note: For more information, see the SECURITY_SUBSET PROC OLAP option and the topic "SAS OLAP Security Totals and Permission Conditions" in the *SAS OLAP Server: MDX Guide*. \triangle

Identity-Driven Security

It is sometimes necessary to substitute identity values in a permission condition to further refine member-level security. Identity-specific values are dynamically derived according to the user ID with which a client is authenticated. Those values are then used to filter the target data. The identity-specific values are derived from identity-driven properties that are stored in the metadata for each user and group.

The identity-driven properties are as follows:

SAS.ExternalIdentity

This property translates to optional, site-specific values such as Employee ID. Those values are not automatically stored in the metadata repository and need to be loaded and maintained.

SAS.IdentityGroupName

This property resolves to the name of the requesting group identity (for example, Portal Admins Group).

SAS.PersonName

This property resolves to the name of the requesting user identity (for example, SAS Demo User).

SAS.IdentityName

This property returns the name of either the requesting group identity or the requesting user identity, depending on whether the user ID is a group login or a personal login.

SAS.Userid

This property translates to the authenticated user ID, normalized to one of the uppercase formats USERID or USERID@DOMAIN (for example, SASDEMO@LXXXXX).

SAS.IdentityGroups

This property resolves to the names of the groups of which a user is a member.

Note: For more information on OLAP member level security, see the topic "OLAP Member-Level Permissions" in the *SAS Intelligence Platform: Security Administration Guide*. △

Build Formula Dialog Box

You can apply identity values to a permission condition from the Build Formula dialog box. To access this dialog box, select a cube dimension from within SAS Management Console or SAS OLAP Cube Studio. The dimension Properties dialog box contains an **Authorization** tab. From here, you can assign Read access to the appropriate user or group. You can then select the **Edit Authorization** button. This opens the Build Formula dialog box. The identity values are then located on the **Data Sources** tab. On the **Data Sources** tab, you can select an identity property and insert it into the MDX expression that you are building.

See "Setting Identity Driven Security" on page 187 for more details.

Note: For more information on identity-driven properties for OLAP member authorization see the SAS OLAP Member Authorization Help in SAS Management Console. Also, see "Identity-Driven Properties" in the *SAS Intelligence Platform: Security Administration Guide*. △

Integrated Windows Authentication and Single Sign-On

Single Sign-On (SSO) to SAS IOM Bridge servers

Single sign-on (SSO) is an authentication model that enables users to access a variety of computing resources without being repeatedly prompted for their user IDs and passwords. This functionality enables a client application to connect to a SAS IOM Bridge Server as the currently logged-on user. For example, single sign-on can enable a user to access SAS servers that run on different platforms without interactively providing the user's ID and password for each platform. Single sign-on can also enable someone who is using one application to launch other applications based on the authentication that was performed when the user was initially logged on to the system.

SSO and SAS OLAP Cube Studio

When you have logged on to a metadata server in SAS OLAP Cube Studio, you can perform a variety of functions, including building a cube and saving the metadata for that cube, editing cube metadata, and viewing cube properties. Some functions within SAS OLAP Cube Studio also require a workspace server or an OLAP server in addition to the metadata server. These functions include Build Physical Cube, Incremental Cube Update, Calculated Members, and Cube View. Single sign-on enables you to seamlessly access these servers without having to provide your login credentials. Two technologies are available in SAS 9.2 to accomplish this:

- 1 *Integrated Windows authentication* is available when the client and server are both running on Windows.
- 2 *a trusted peer connection* is available otherwise.

Note: Before using single sign-on, you must have a user defined in the metadata repository for any user ID that you want to connect to the metadata server with. △

Integrated Windows Authentication

A client and server engaged in Integrated Windows authentication use Microsoft's Security Support Provider Interface (SSPI) to choose the best security package for their

configuration. When you first open SAS OLAP Cube Studio you are prompted to log on to a metadata server with a connection profile that you have previously created. This connection profile contains information about the metadata server, including the machine ID, port ID, user ID, and password. There is also a check box labeled **Use Integrated Windows authentication (single sign-on)**. If this check box is selected, then you will use Integrated Windows authentication (IWA) to connect to the servers (metadata server, workspace server, and OLAP server). You will not need to provide a user ID or password when you open SAS OLAP Cube Studio. The user ID and password of your current Windows session will be used. By default this option is not selected when you create a connection profile in SAS OLAP Cube Studio.

Note: For more information on Integrated Windows authentication see "Integrated Windows Authentication" in the *SAS Intelligence Platform: Security Administration Guide*. Δ

Trusted Peer Connection

The other authentication technology is a trusted peer connection. This method depends heavily on the integrity of the client environment. The IOM Bridge protocol client obtains the identity of the user of the client application from the system environment. When a trusted peer connection is being used, you are prompted during an initial login for your user ID and password when connecting to a metadata server. This is the behavior of previous versions of SAS OLAP Cube Studio. However, when a connection to a workspace server (submitting procedure code) or a connection to an OLAP server (cube viewer) is made in SAS 9.2 OLAP Cube Studio, you are no longer prompted to log on. The connection is created using the metadata server login credentials.

Note: For more information on trusted peer connections, see "Trusted Peer" in the *SAS Intelligence Platform: Security Administration Guide*. Δ

Exporting Code

In SAS OLAP Cube Studio, you can export PROC OLAP code that you have created while using various functions and wizards, including Cube Designer, Aggregations Tuning, and Incremental Update. The Export Code function enables you to save PROC OLAP code to a designated file. When using the Export Code function, some of the metadata connection profile information is saved to the file. The host and port information for the metadata server is saved to the exported file. In addition, the information is displayed on the Finish page of the different wizards. However, the user ID and password are not saved to the exported file or displayed on the Finish page. This is because single sign-on authentication removes the need to display or save the user ID and password.

Security for Drill-through Tables

Different users of cube data may have different security and access restrictions that must be adhered to and applied when querying the underlying data for a cube. When selecting a data table for drill-through, you may need to define user restrictions for certain data in the drill-through table. The SAS Metadata LIBNAME engine is used to assign the drill-through table library on a per-session basis. This allows client credentials to be used when determining which columns the user can see. Columns which have ReadMetadata permissions denied on the drill-through table will not be visible to the user. If access is denied for a column on the drill-through table, that level must also be denied Read access in the cube. Conversely, if a level is denied Read access

in the cube, that column in the drill-through table must have ReadMetadata access denied. For more information on the SAS Metadata LIBNAME engine see the topic “Pre-assigning Libraries” in the *SAS Intelligence Platform: Data Administration Guide*.

Because of this enforcement, the following changes could affect your site:

- User-defined formats will not work if the FORMATS catalog is accessed by the same libref as the data. You must specify a different libref for the FORMATS catalog or move the user-defined formats into a different location and assign a new library.
- If SAS Trusted User does not have ReadMetadata permissions on the library definition, then the drill-down functionality fails. You can grant ReadMetadata permission to the SAS Trusted User by using SAS Management Console.
- You cannot drill down in an OLAP cube if there are discrepancies between the physical table and the metadata that is defined for the table. You can update the metadata for the table by using the Update Metadata function in SAS Management Console or SAS Data Integration Studio. If the table is part of a job in SAS Data Integration Studio, then check the code for the job to verify that there are no LENGTH statements that would cause a mismatch to occur.


Note: For further information on drill-through tables, see “Defining Drill-Through Tables” on page 39. Δ

Applying Batch Security with Permission Tables

When applying permissions to a cube, you might need to address permissions for different combinations of users, groups, SAS OLAP Servers, schemas, and cubes, as well as different elements of the cube, including the dimensions, hierarchies, levels, and measures. For example, you might need to grant ReadMetadata and Read access to the group that contains specific cube users. Or you might need to restrict Read access for different components of a cube (dimension, hierarchy, level, or measure) using MDX conditions for each cube component, per user, consumer, or group. These various combined permission settings can be easily created and managed with batch security that is applied through permission tables.

Starting in the third maintenance release for SAS 9.2, you can specify batch security in SAS OLAP Cube Studio and SAS Data Integration Studio with the Manage Permission Tables function. The Manage Permission Tables function enables you to create a special SAS data set known as a permission table that contains cube access controls for submitting in bulk. A permission table is a table of access control information that can later be applied to a cube with batch SAS code. The Manage Permission Tables dialog box enables you to create and modify permission tables as well as import access controls (permissions) from a cube or an OLAP schema. You can also execute the code interactively or export the code to a file for use in a stored process or deployed job flow.

When a cube is created, security for that cube is determined by the permission settings that are found in the cube metadata. In SAS OLAP Cube Studio, permission tables will appear in the metadata tree as a table. You must have WriteMetadata access to create and modify permission tables. To access the Manage Permission Tables function in SAS OLAP Cube Studio, select **Tools ► Manage Permission Tables**. The Manage Permission Tables dialog box opens.



CHAPTER

7

Updating SAS OLAP Cubes

<i>Overview</i>	79
<i>Updating a Cube In-Place</i>	80
<i>Incremental Updates of Cubes and Cube Generations</i>	80
<i>OUTCUBE Option</i>	81
<i>OUTSCHEMA Option</i>	81
<i>OUTCUBE and OUTSCHEMA Options</i>	81
<i>Coalescing Cube Aggregations</i>	82
<i>Updating a Cube In a Production Environment</i>	82
<i>Disable the Production Cube</i>	82
<i>Rename the Cubes</i>	83
<i>Enable the New Production Cube</i>	83
<i>Example 1</i>	83
<i>Example 2</i>	84
<i>Archiving and Deleting Cube Generations</i>	84
<i>Updating the Captions and Descriptions for a Cube</i>	85
<i>Adding New Members to an Incrementally Updated Cube</i>	85
<i>Reorganization of Cube Levels</i>	85
<i>Why Reorganize?</i>	86
<i>Scenario Example</i>	86
<i>OLAP Procedure</i>	86
<i>Multiple Language Support Cubes</i>	86
<i>Updating Member Properties</i>	86
<i>Specifying Drill-Through Tables</i>	87
<i>NWAY Considerations</i>	87
<i>Updating Multiple Language Support Cubes</i>	87
<i>Format Search Path and SAS Source Code Considerations</i>	87
<i>Exporting Cubes that Have Been Updated</i>	88
<i>Input Data Tables for Cube Updates</i>	88
<i>Schema and Repository Considerations</i>	89
<i>Physical Storage and Metadata Considerations</i>	90
<i>Connecting to a Workspace Server</i>	90
<i>Proc OLAP Options</i>	90
<i>Proc OLAPOPERATE Options and SAS OLAP Monitor</i>	90
<i>Updating Cubes in SAS OLAP Cube Studio</i>	90

Overview

After a SAS OLAP cube is created, it might be necessary to add further data to the cube without completely re-creating the cube. If a cube is in production, you must consider what the best process to update your cube is. You should also consider an

approach that will have as little impact as possible on users. It is important to limit the amount of time that a cube is unavailable to users while the cube is being updated with new data. It is also important that the updated cube be checked for correctness before the cube is made available for queries. To address these concerns, you can choose either to update a SAS OLAP cube in-place or to create an incremental copy of the cube.

You can update SAS OLAP cubes that physically exist and that are generated from either detail tables or star schemas. You cannot update a fully summarized cube. Furthermore you must have ReadMetadata and WriteMetadata permissions for all elements of the cube

New data to add to a cube can include new members and new data for the cube or additional data for existing members. You can update a cube using SAS OLAP Cube Studio or manually using “The OLAP Procedure” on page 244 and “The OLAPOPERATE Procedure” on page 291.

Updating a Cube In-Place

Depending on the business needs for a cube, it might be more suitable to update a cube in-place. A separate version of the cube is not created. The process of updating in-place enables you to add data to a cube while that cube is still online. The cube maintains its name and OLAP schema assignment. This process might be more appropriate for some users and their reporting needs, as it is less complicated than creating incremental generations of a cube. However, there is no opportunity to test and verify the update in a separate process.

You can update a cube in-place by using the UPDATE_INPLACE PROC OLAP option. This option is used with the ADD_DATA option to signal that an update will occur. Here is an example of the UPDATE_INPLACE option:

```
proc olap data=mylib.newdata cube=cubeA add_data
    update_inplace;
metasvr host="myhost" port=8561 protocol=BRIDGE repository=my repository
    olap_schema=prodSchema;
run;
```

The process of updating in-place does not create a new metadata registration when the update is complete. The existing metadata registration points to the new cube generation and the old generation is deleted. In addition, you do not have to disable the cube prior to beginning the in-place update. OLAP server sessions referencing the old cube will continue to do so until the last session closes the old cube. After that, new sessions will begin to reference the updated cube. However, it might be desirable to have a specific “cross-over” to the new cube at some point after the update completes.

See “Update a Cube In-Place” on page 160 for a discussion of updating a cube in-place with the SAS OLAP Cube Studio Incremental Update Wizard.

Incremental Updates of Cubes and Cube Generations

When you create a SAS OLAP cube, data from a table is summarized and sub-aggregates are created for the cube. The Incremental Update function enables you to create an incremental update or copy of the original cube. This is known as a generation of the cube. A generation is a complete, independent version of the original cube, with a separate metadata registration and separate sets of member and property trees. The new generation shares the original data (previous aggregation data) rather than maintaining duplicate copies.

The new cube generation enables the original cube to remain online for queries during the entire update process. This is because the original cube is not modified. The original cube continues in production as before while the new cube is available for administrative review and updates (such as security updates or new global calculated members). This provides the administrator with an opportunity to examine and verify the correctness of the updates before bringing the cube online for production without affecting the original cube. A new cube generation also provides the additional benefit of keeping historical data (in the form of the previous generations) available for as long as necessary.

You can update a cube incrementally by using the OUTCUBE or the OUTSCHEMA options. These options are used with the ADD_DATA option to signal that an update will occur and a new generation of the cube will be created. You can use either option separately or both together.

OUTCUBE Option

Here is an example of the OUTCUBE option. In this example, the cube PRODSHEMA is updated and a cube generation (TSTCUBE) is created in the same schema as the original cube.

```
proc olap fact=olapsio.factcars cube=PRODCUBE
      OUTCUBE=TSTCUBE
      ADD_DATA;
  metasvr host=&host port=&port protocol=&protocol userid=&userid pw=&pw
      repository=&repos olap_schema=PRODSHEMA;
run;
```

OUTSCHEMA Option

You can also create a generation of a cube by specifying a different schema than the original cube. In this example, the new cube retains the name of the original cube PRODCUBE (OUTCUBE was not specified). However, the new cube now resides in the schema TESTSCHEMA, as specified in the OUTSCHEMA option.

```
proc olap fact=olapsio.factcars cube=PRODCUBE
      OUTSCHEMA=TESTSCHEMA
      ADD_DATA;
  metasvr host=&host port=&port protocol=&protocol userid=&userid pw=&pw
      repository=&repos olap_schema=PRODSHEMA;
run;
```

OUTCUBE and OUTSCHEMA Options

You can use both OUTCUBE and OUTSCHEMA options to create a new cube in a different schema. In this example, a new cube called TSTCUBE is created and it resides in the schema TESTSCHEMA.

```
proc olap fact=olapsio.factcars cube=PRODCUBE
      OUTSCHEMA=TESTSCHEMA
      OUTCUBE=TSTCUBE
      ADD_DATA;
  metasvr host=&host port=&port protocol=&protocol userid=&userid pw=&pw
      repository=&repos olap_schema=PRODSHEMA;
run;
```

The possible functions that you can perform on a SAS OLAP cube depend on whether the cube is the original cube or a generation of a cube. These functions can also depend on which generation of the original cube it is. Any generation of a cube can have calculated members added to or deleted from it. However, only the most current cube generation can be edited, rebuilt, tuned, or updated. You can add data to cubes that are created in SAS 9.2. A cube built in SAS 9.1.3 must be rebuilt in SAS 9.2 in order to use the Incremental Update function. Furthermore, you must have ReadMetadata and WriteMetadata permissions for all elements of the cube. See “Generating a New Cube” on page 164 for a discussion of creating a cube generation with the SAS OLAP Cube Studio Incremental Update Wizard.

Note: The SAS OLAP Cube Studio right-click menu options for Edit, Create, Manual Tuning, Advanced Tuning, and Delete Physical Cube are not active if the cube is not the current generation. Δ

Coalescing Cube Aggregations

The Incremental Update function enables you to add new members to a cube and update data. Every time a cube update is done, the incremental aggregation data is written to a new SPDE table (rack). The more cube updates that have been completed, the more tables or racks there are that must be read during a query, and the slower the access time can be.

Periodically, it can be beneficial for performance reasons to combine all of the individual tables into a single aggregation table. Aggregations are coalesced (combined) in a separate job from adding data to a cube. Only aggregations with more than one rack are coalesced. The COALESCE_AGGREGATIONS option can be executed only for the most current generation of the cube. In addition, only MOLAP aggregations (those aggregations created by PROC OLAP) can be coalesced. The coalesce task is processed in-place. A new generation of the cube is not created.

After the cube generation is coalesced, it is no longer dependent on the aggregation data from previous generations. If those generations have already been deleted, then the old aggregation tables are also deleted.

You can coalesce all the aggregations by using the option on the PROC OLAP statement. In addition, the Coalesce wizard is available in SAS OLAP Cube Studio.

See “Coalesce Incremental Data for a Cube” on page 167 for a discussion of coalescing cube aggregations with the SAS OLAP Cube Studio Incremental Update wizard.

Updating a Cube In a Production Environment

When a new generation of a cube is created, it can be designated as the production version of the original cube. The new generation must be given a different name from the original cube or reside in a different OLAP schema. In order to update a cube in production, you must disable the existing production cube, rename the cubes, and then enable the new production cube.

Disable the Production Cube

To disable a cube, you can use either the SAS OLAP Monitor plug-in to SAS Management Console or the PROC OLAPOPERATE statement (see “Disabling a Cube” on page 296). When a cube is disabled, it is taken offline and is not available for new

queries. All existing query results are closed, but the sessions are left open and can be closed by the SAS OLAP Server Monitor. If a query is being processed, however, the disable action will fail.

Rename the Cubes

Rename the original cube to another name (for example, Sales_OLD). Then rename the cube generation to the original cube name (Sales). To rename a cube, you can use either SAS OLAP Cube Studio or the PROC OLAP RENAME option. The RENAME option enables you to reassign cube generations. It is used with the OUTCUBE or the OUTSCHEMA options. Renaming a cube updates the metadata for the cube but does not change the file structure or physical location of the cube. To rename a cube, you must have ReadMetadata and WriteMetadata permissions on all parts of the cube and the cube must be disabled.

Enable the New Production Cube

To enable a cube, you can use either the SAS OLAP Monitor plug-in to SAS Management Console or the PROC OLAP OPERATE statement (see “Enabling a Cube” on page 296). Enabling the cube enables queries to resume on the server for the new production cube. Existing reports will continue to work as they reference the cube by name.

Example 1

Here is an example where an existing production cube is replaced with an updated version in the same schema. The old cube is then deleted.

Perform an incremental update on production cube A.

```
proc olap cube=A outcube=A_New add_data data=lib1.data1;
metasvr .... ;
run;
```

Disable the production cube.

```
proc olapoperate <server-connection-options>;
disable cube A;
run;
```

Rename the production cube to a temporary name.

```
proc olap rename cube=A outcube=A_Old;
metasvr ...;
run;
```

Rename the new cube to the original production cube name.

```
proc olap rename cube=A_New outcube=A;
metasvr ...;
run;
```

Enable the new production cube (make available for queries).

```
proc olapoperate <server-connection-options>;
enable cube A;
run;
```

Delete the old cube.

```
proc olap cube=A_Old delete;
metasvr ...;
run;
```

Example 2

In this second example, an existing production cube is replaced with an updated version from another schema and the old cube is archived.

Perform an incremental update on production cube A. This creates a cube A located in OLAP Schema TestSchema.

```
proc olap cube=A outschema=TestSchema add_data data=lib1.data1;
metasvr .... olap_schema=ProductionSchema ;
run;
```

Disable the production cube.

```
proc olapoperate <server-connection-options>;
disable cube A;
run;
```

Move the old cube to an archive schema.

```
proc olap rename cube=A outcube=A_September2005 outschema=ArchiveSchema;
metasvr ... olap_schema=ProductionSchema;
run;
```

Move the new cube to the original production schema.

```
proc olap rename cube=A outschema=ProductionSchema;
metasvr ...olap_schema=TestSchema;
run;
```

Enable the new production cube (make available for queries).

```
proc olapoperate <server-connection-options>;
enable cube A;
run;
```

Archiving and Deleting Cube Generations

If multiple generations of a cube have been created, special considerations are needed when archiving or deleting these generations. It is possible to archive an old cube generation into another schema or under another name. The cube will then be available for access to historical data. If you are deleting cube generations, it is important to note that a cube that is an earlier generation must be deleted with PROC OLAP. Deleting only the metadata will not work for generations of cubes.

Because cube generations generally look through to aggregation data in previous generations, cube data cannot be physically deleted as long as another generation is dependent upon it. Only those physical files that are not needed by the new cube (for example, member and property trees) are actually deleted. The aggregation files remain available to subsequent generations. In addition, the metadata registration in the metadata repository is deleted.

Updating the Captions and Descriptions for a Cube

You can update the captions and descriptions for a cube without adding data to the cube. This is considered a nonstructural update and can be performed on a cube with the UPDATE_DISPLAY_NAMES option for PROC OLAP. This option enables you to identify captions and descriptions on a cube that will change. Here is an example:

```
proc olap cube=cubename description="new description" UPDATE_DISPLAY_NAMES
  dt=newdtname;
  metasvr options;
  dimension time description="New TimeDescription" caption="New TimeCaptions";
```

Note: The only description that can change when a cube is renamed is the cube description (the drill-through table name might also be changed during a rename). Other descriptions and captions must be altered in a separate step after the cube is renamed. △

Adding New Members to an Incrementally Updated Cube

By default, the Incremental Update function expects new members and new data to be added to the cube. However, when adding new data to a star schema, you can specify that only data will be added with the NONUPDATEABLE option. When this option is selected, the Incremental Update process does not check the dimension tables for new members. The NONUPDATEABLE option can be specified on the PROC statement or on an individual DIMENSION statement. If the NONUPDATEABLE option is specified on the PROC statement, none of the dimension tables are read. Alternately, you can specify NONUPDATEABLE on the DIMENSION statement if you are adding new members to a specific dimension.

Note: You should set the IGNORE_MISSING_DIMKEYS=VERBOSE option when using the NONUPDATEABLE option. △

IN SAS OLAP Cube Studio, you can check the options **Cube Designer - Dimensions** and **Dimension Designer- General**.

Note: Existing SAS 9.2 OLAP cubes that have been updated with one or more new members for a hierarchy can possibly contain ragged hierarchies and must be rebuilt in the second maintenance release after SAS 9.2. △

Reorganization of Cube Levels

When you are administering SAS OLAP Cubes, it might be necessary to perform multiple updates of a cube for various business reasons. After a cube has been updated several times, it might be necessary to reorganize the levels for the cube. The deciding factor for this would most likely be a failure of the cube to update successfully, accompanied with error messages that are similar to the following messages:

```
ERROR: The new member name "64 " belonging to the "type" level cannot be
      added to cube because the limit for inserting new captions between
      existing members " 63" and "1001" has been reached. Consider using
      the proc olap reorganization options to create additional space.
ERROR: A problem was encountered when attempting to update the cube's
      hierarchies with new members.
```

ERROR: Cube "cubename " cannot be updated.

If you cannot update your cube and receive this message, you must reorganize your cube level before you can resume updating your cube.

Why Reorganize?

Reorganization of a cube must be done after a cube update fails when accompanied by an error message that suggests using "proc olap reorganization options" to create additional space for that level. A level may run out of space for new level members after multiple cube update events. This occurs when new level members have been added to the same sorted location for each event. If too many new level members are inserted into the same sorted location, then eventually cube update fails. Unfortunately, this limit can be reached rather quickly if multiple cube update events take place and, for each event, new level members are inserted into the same location between existing members due to their relative sort order. Reorganization changes the internal structure of the cube so that cube update can once again proceed successfully.

Scenario Example

Consider that you have a sales data cube that contains a level for the customer name, and you have numerous data records for that level where the customer name starts with the letter S. Each time a cube update takes place, a new customer name is added whose name starts with the letter S. There is a finite amount of reserved space for new last names when a cube is first made. If all new names added sort in the same location of the alphabet, rather than being evenly distributed throughout the alphabet, then the odds of running out of reserved space increase with each new cube update event.

Eventually, there is no longer any reserved space left to fit in yet another last name that starts with S. However, there might be ample room to add a name that starts with the letter Q, X, or I. Nonetheless, the customer name level will still have to be reorganized before new names starting with the letter S can be added to the cube.

OLAP Procedure

The following PROC OLAP items are available to reorganize SAS OLAP cubes:

- REORGANIZE_LEVELS | REORG_LEVELS option for PROC OLAP
- "REORGANIZE_LEVEL Statement" on page 281

In SAS 9.2, reorganization can only be performed with the PROC OLAP commands.

Multiple Language Support Cubes

You cannot reorganize multiple language support (MLS) cubes. You can, however, update an MLS cube.

Updating Member Properties

When you add data to a cube, the values for member properties are normally not updated. However, you can specify the DIMENSION statement option UPDATE_DIMENSION and the MEMBERS_AND_PROPERTIES value to update the member properties for a dimension. When updating member properties, the new values

must be contained in either the ADD_DATA table for a detail load cube or the update table for the dimension for a star schema load.

For a star schema cube, it is possible to change the member property values without adding data to the cube. The table of new data is not required. You can select to update member properties on all the dimensions or on just certain dimensions. In both instances, the original dimension table for a dimension is scanned for new values unless a different table is provided. The Cube Designer – Dimensions page and the Dimension Designer - General page have the option **Allow new members during incremental update**. See the SAS OLAP Cube Studio Help for further details.

Specifying Drill-Through Tables

Because new data and possibly new members are being added to the cube, a new drill-through table might be required. You can specify a drill-through table that will be associated with the cube after it is updated with the DRILLTHROUGH_TABLE PROC OLAP option or you can specify a new drill-through table in SAS OLAP Cube Studio. Follow these steps to enter a new drill-through table in SAS OLAP Cube Studio.

- 1 Select **Incremental Update**.
- 2 Select either **Update In-Place** or **Generate New Cube**.
- 3 On the Update General page, enter the new drill-through table in the **New drill-through table** field.

NWAY Considerations

When you use the Incremental Update function on a cube, the cube should have a MOLAP NWAY. If the cube is loaded from a fully summarized table, the Incremental Update function is not allowed. The NO_NWAY option is also of importance. When using PROC OLAP to add data the cube, you must have the NO_NWAY option set to NO. If the NO_NWAY option is set to YES, it is assumed that the cube does not have a MOLAP NWAY and adding data is not allowed.

Updating Multiple Language Support Cubes

For multiple language support star schema cubes, the original dimension tables are always used (see “Multiple Language Support and Dimension Table Translations” on page 65). When you update a non-MLS star schema, if a table is not specified for the dimension table, the original dimension table is scanned for new members.

If you are updating a MLS cube, you cannot modify the languages that are selected for the cube. To add new members, the original dimension tables (including the different language tables) must be updated with the new members. For Multiple Language Support cubes, the original tables are always re-read.

Format Search Path and SAS Source Code Considerations

When updating cubes that have added syntax for a user written formatted search path or SAS code, you must consider how that information will be updated as well. Here are some possible scenarios.

Table 7.1 Formatted Search Path Scenarios

Update Process	Update Results
You perform an update in-place and do not make any changes to the format search path or SAS source code.	The metadata for the cube does not need to be updated. When the update runs, the current format search path and SAS source code must be submitted with the code.
You perform an update in-place and change the format search path or the SAS source code.	The new values for the format search path and SAS source code will be updated in the metadata for the cube in addition to being submitted with the code or saved with the exported code.
You create a new cube in the update and do not make any changes to the format search path or the SAS source code.	The new generation of the cube will be created using the same format search path and SAS source code as the original cube, and its metadata should also include the same format search path and SAS source as the original cube. The submitted code or exported code will contain the current format search path and SAS source code. PROC OLAP will copy the current format search path and SAS source code to the new cube as part of the update.
You create a new cube in the update and enter a new format search path or new SAS source code.	The next generation of the cube will be created with the new format search path and SAS source code, and its metadata will include the new format search path and SAS source code. The submitted or exported code will contain the new format search path and new SAS source code. However, PROC OLAP will copy what is currently in the metadata for the original cube.

Exporting Cubes that Have Been Updated

After you update a cube, it might be necessary to export that cube for use in another SAS OLAP environment. The following rules apply to exporting cubes that have been updated:

- If a cube exists and has been updated but not coalesced, it cannot be exported.
- If a cube exists and has been updated, and it has been coalesced, it can be exported. This is possible even if it is not the most current generation of the cube.
- The export template should not include the NextGeneration or PrevGeneration associations.

Input Data Tables for Cube Updates

When you add data to a cube, you must select a table that contains the new data for the cube. This might be new data for existing members or new members with new data. The table must be defined in the metadata repository and it must contain all the columns for the measures. Depending whether you are incrementally updating the cube or rebuilding the cube, you must manage the input data accordingly. If you are updating the cube incrementally, then your table does not need to contain all the

cumulative data. However, if you are rebuilding the cube, then your input tables must contain all the data from the original build and all the incremental updates.

If you are rebuilding a cube, it is very important that the data used as input to the rebuilt cube contain all the incremental updates that have been made since the original cube was created. This can be done one of two ways:

- 1 The input base or fact table definition in the SAS metadata should be a view that combines the original input table plus all subsequent update tables into a single logical table. This is the recommended approach.
- 2 The input base or fact table itself is physically updated to add data from all update tables.

It is not possible for PROC OLAP or SAS OLAP Cube Studio to verify the correctness of the input data. This is the responsibility of the administrator or cube builder. Depending on what data source the cube is loaded from, some requirements apply:

Table 7.2 Input Table Considerations

Data Source	Requirement
Detail table	All the level columns must be in the new data table.
Star schema	<p>The foreign key columns must be in the new data table. However, the level columns must be in the dimension table along with the unique key columns.</p> <p>For star schema cubes, you do not have to select a source table. If a source table is not selected, then only the dimension tables will be processed and only if the UPDATE_DIMENSION option is set to MEMBERS or MEMBERS_AND_PROPERTIES for the dimension. No new data is added to the cube. In this case, the ADD_DATA option should not be added to the PROC OLAP statement.</p> <p>You can, however, still specify new dimension tables and specify a value for the UPDATE_DIMENSION option. This is only valid for dimensions in a star schema that originate from a dimension table. Dimensions in a star schema that come from the fact table can only be updated if a source table is selected and ADD_DATA is specified on the PROC OLAP statement.</p>
Fully summarized table	The cube cannot be updated.

When you add data to a cube, the cube and its corresponding tables must be kept in synch. When you add new members to the cube using the Incremental Update process, the original tables are re-read by PROC OLAP unless a different table is provided.

Note: If the original dimension table for a star schema is updated with the new members, you do not need to select a table for a dimension. Δ

Schema and Repository Considerations

When you add data to a cube, a new name or a new OLAP schema must be provided for the updated cube. By default, the OLAP schema of the original cube is selected. In addition, the new generation of the cube should be created in the same repository as the original cube. If you are defining the repository with the PROC OLAP METASVR statement, use the repository for the cube that is being updated. Do not use the default repository.

Physical Storage and Metadata Considerations

The physical files that compose SAS OLAP cubes are contained within two subfolders:

- the cube path folder
- the generation path folder

The cube path folder name matches the cube name, under the folder defined by the PATH= option. The aggregation racks (tables) are located in the cube folder. The remaining internal cube files are placed in a generation folder. Each generation will have a new generation folder, with subsequent names (gen0001, gen0002, gen0003, and so on). These folders will be created under the existing cube folder.

Connecting to a Workspace Server

Because adding data to a SAS OLAP cube involves submitting SAS code to accomplish the actual update, a connection to the workspace server must be established. If a connection cannot be made, you can continue using the Incremental Update function and save the generated PROC OLAP code to submit later in a SAS session.

Proc OLAP Options

Depending on the type of cube update you are performing (in-place or incremental), the following PROC OLAP statements and options can be used:

- ADD_DATA
- NONUPDATEABLE
- OUTCUBE
- OUTSCHEMA
- RENAME
- UPDATE_DIMENSION
- UPDATE_DISPLAY_NAMES
- UPDATE_INPLACE

Proc OLAPOPERATE Options and SAS OLAP Monitor

In addition to PROC OLAP options, it might be necessary to use PROC OLAPOPERATE options when promoting an incremental update of a cube online. You can use the PROC OLAPOPERATE options DISABLE CUBE and then ENABLE CUBE, or you can perform these functions with the SAS OLAP Monitor plug-in for SAS Management Console.

- “Disabling a Cube” on page 296
- “Enabling a Cube” on page 296

Updating Cubes in SAS OLAP Cube Studio

The Incremental Update function in SAS OLAP Cube Studio enables you to add new members and new data to a cube as well as update captions for dimensions, hierarchies,

levels, member properties, and measures. The Incremental Update function enables you to specify cube updates and submit the generated procedure code without rebuilding the cube. Although you can add data to a cube, the cube structure cannot be modified. Adding or removing dimensions, levels, measures, hierarchies and properties is not part of the Incremental Update function.


When you initially create a cube, you can establish whether the cube will allow new members to be added to the cube at a later point. You can also allow new members to be added to a specific dimension on a cube. The Cube Designer – Dimensions page and the Dimension Designer – General page on the Cube Designer wizard have the option **Allow new members during incremental update**. See the SAS OLAP Cube Studio Help for further details.

The Incremental Update wizard is available from the **Actions** menu and the context menu in SAS OLAP Cube Studio. It is active on the menu if the selected cube meets the requirements for adding data. If you select the Incremental Updates function, a drop-down menu opens. Depending on the type of update you need to perform, select one of the following options.

- Update In-Place**
- Generate New Cube**
- Coalesce Incremental Data**

After you have updated your cube, you can check to see a list of the generations for the cube. Select the **Generations** tab on the Properties dialog box for the cube.

See the SAS OLAP Cube Studio Help for further details on the Incremental Update function.



CHAPTER

8

Cube Building and Modifying Examples

<i>Defining A Connection Profile</i>	94
<i>Building a Cube from a Detail Table</i>	96
<i>Enter General Cube Information</i>	96
<i>Select A Detail Table</i>	96
<i>Drill-Through Table</i>	97
<i>Table Options</i>	97
<i>Define Dimensions, Levels, and Hierarchies</i>	97
<i>Creating a Time Dimension</i>	101
<i>Define Measures</i>	103
<i>Define Member Properties</i>	104
<i>Define Aggregations</i>	105
<i>Build the Cube</i>	107
<i>Save a Cube's PROC OLAP Code</i>	108
<i>PROC OLAP Statements and Options For a Detail Table</i>	108
<i>PROC OLAP Example for a Detail Table</i>	110
<i>Building a Cube from a Star Schema</i>	114
<i>Enter General Cube Information</i>	114
<i>Select Fact and Dimension Tables</i>	116
<i>Drill-Through Table</i>	117
<i>Table Options</i>	118
<i>Define Dimensions, Levels and Hierarchies</i>	118
<i>Define Measures</i>	122
<i>Define Member Properties</i>	123
<i>Define Aggregations</i>	124
<i>Build the Cube</i>	125
<i>PROC OLAP CODE for the Star Schema Example</i>	126
<i>PROC OLAP Statements and Options For a Star Schema</i>	131
<i>Building a Cube from a Summary Table</i>	133
<i>Enter General Cube Information</i>	133
<i>Select an Input Table</i>	134
<i>Drill-Through Table</i>	134
<i>Table Options</i>	134
<i>Define Dimensions, Levels and Hierarchies</i>	135
<i>Define Measures</i>	139
<i>Stored Measures</i>	139
<i>Derived Measures</i>	140
<i>Edit Measure Details</i>	141
<i>Define Member Properties</i>	142
<i>Define Aggregation Tables</i>	142
<i>Define Stored Aggregations</i>	143
<i>Define Aggregations</i>	145

<i>Build the Cube</i>	145
<i>PROC OLAP CODE for the Summary Table Example</i>	146
<i>PROC OLAP Statements and Options For a Summary Table</i>	149
<i>Tuning Aggregations For a Cube</i>	151
<i>Cardinality Tuning</i>	151
<i>Manual Tuning</i>	156
<i>Arm Log Tuning</i>	158
<i>Adding Data to a Cube with Cube Update</i>	160
<i>Update a Cube In-Place</i>	160
<i>Generating a New Cube</i>	164
<i>Coalesce Incremental Data for a Cube</i>	167
<i>Adding Calculated Members To A Cube</i>	170
<i>Creating a Simple Calculation</i>	170
<i>Creating a Time Analysis Calculation</i>	175
<i>Creating a Custom Calculation</i>	177
<i>Setting Member Authorizations On A Dimension</i>	180
<i>Setting Identity Driven Security</i>	187
<i>Viewing a Cube in SAS OLAP Cube Studio</i>	192
<i>Adding Level Data to a Cube View</i>	192
<i>Replacing Level Data On A Cube View</i>	195
<i>Creating a Time Dimension in SAS OLAP Cube Studio</i>	197
<i>Synchronizing Column Changes</i>	199
<i>Specifying an ESRI GIS Map For a Cube Dimension</i>	200
<i>Creating Multiple Hierarchies For a Cube</i>	205
<i>Set IGNORE_MISSING_DIMKEYS for a Star Schema</i>	208
<i>Implementing Drill-through to Detail Data in a SAS OLAP Cube</i>	210
<i>Exporting a Cube From SAS OLAP Cube Studio</i>	216
<i>Importing a Cube Into SAS OLAP Cube Studio</i>	219

Defining A Connection Profile

Each time you log onto SAS OLAP Cube Studio, a connection is made to a metadata server. Before you can build a cube in SAS OLAP Cube Studio, you must specify a connection profile that contains the information for the metadata server that you want to build your cube with. In SAS OLAP Cube Studio, select **File ► Connection Profile**. On the Open a Connection Profile dialog box, you can choose to either create a new connection profile or edit an existing one. If you choose to create a new connection profile, the Connection Profile wizard will open. The following display shows the information that is entered for a connection profile.

Enter the machine information for the metadata server that you will connect to and retrieve a data source from.

- In the New Connection Profile dialog box, enter the machine ID, port, your user ID, and password.

Note: These fields are the equivalent of the following METASVR statement options:

- HOST=
- PORT=
- USERID=
- PW=.

△

- Here is an example of the PROC OLAP code:

```
metasvr host=localhost
        port=9999
        protocol=bridge
        userid=userid
        pw=pw
        repository=Foundation
        olap_schema="OLAP Schema"
```

After you have entered the basic connection information, you can select to do one of the following:

- enter the authentication domain for the SAS Metadata Server
- save the user ID and password with the connection profile *or* use Integrated Windows authentication (single sign-on). The Integrated Windows authentication option specifies that access is based on the user's previous authentication to the Windows desktop.

Note: For more information on authentication and single sign-on, see "Understanding Authentication" in the *SAS Intelligence Platform: Security Administration Guide*. △

Building a Cube from a Detail Table

A detail, or base, table is a table whose data pertains to a single area of interest. It is any table defined in a SAS Metadata Repository that contains the measures and levels for a cube. You can build an OLAP cube from a detail table by using the Cube Designer wizard in SAS OLAP Cube Studio. In this example, you use data from a recent product marketing campaign. You establish measures and summaries of various aspects of the data, such as product statistics, geographic location of potential customers, and revenue summaries.

Enter General Cube Information


After you have established a connection profile, you can begin to create a cube. Select **File ► New ► Cube**. On the Cube Designer – General page, enter the basic cube information. For this example input type, you select **Detail Table**. The following display shows fields that you enter information for.

- Enter information in the following fields:
 - Name**
 - Description**
 - OLAP schema**
 - Location** (SAS folder)
 - Physical cube path** (path in the file system to store the cube)
 - Work path** (path for temporary work files)
 - Input Type** (detail table)

Select A Detail Table

At the Cube Designer – Input page, select a detail table for your cube. If one does not exist for your data, select **Define Table**, and then define the source that you will

import your metadata from. The following display shows the input table Detail that is selected for the example cube.



Drill-Through Table

At the Cube Designer - Drill-Through dialog box, you can select or define an optional drill-through table. Drill-through tables can be used by client applications to provide a view from processed data into the underlying data source.

If a drill-through table does not exist for your data, select **Define Table**, and then define the source that you will import your metadata from.

Table Options

The **Table Options** button is available in both the Cube Designer - Input and the Cube Designer - Drill-Through dialog boxes. It opens the Table Options dialog box. It enables you to specify data set options that are used to open the data set. For example, you could enter a WHERE clause or subsetting information that is then applied to the selected table when it is opened. The options are stored as part of the cube and then reapplied when the data is accessed at run time. You can also specify data set options in the Dimension Designer – General dialog box (for use with star schemas) and the Stored Aggregates dialog box (for use with summarized tables). For more information, see “Data Set Options” in *SAS Language Reference: Concepts*.


Define Dimensions, Levels, and Hierarchies

Now that your basic metadata server and cube information has been entered, you can define the different dimensions and their respective levels and hierarchies. For this example, the following dimensions are created:


- Products
- Dates
- Geography

- Customers
- Orders

At the Cube Designer – Dimensions page, select **Add**.




This opens the Dimension Designer – General page, as seen in the following display.




Enter the information in the following fields:

- **Name**
- **Caption**
- **Description**
- **Type** (Standard, GEO, or TIME)
- **Sort Order**.


Select **Next**. This opens the Dimension Designer - Level page. Next, select **Add** to open the Add Levels page, as seen in the following display.




Select the levels you want to add to the dimension. Select **OK** to return to the Dimension Designer – Level page, where the selected levels are listed. You can now define properties such as format, time type, and sort order for the levels that you have selected. See the following display.




Next, define hierarchies for the levels on the Dimension Designer - Hierarchy page. You can select **Add** to open the Define a Hierarchy page and individually select the levels for the hierarchy.



Or you can select **Finish** on the Dimension Designer - Hierarchy page to accept the order of the levels that are defined on the previous Dimension Designer - Level page. If you select this option, the hierarchy is assigned the same name as the dimension. See the following display.




Repeat this process for each dimension. After you create each dimension, it is listed in the **Dimensions** panel of the Cube Designer - Dimensions page. See the following display.




Creating a Time Dimension


When you create the Dates dimension, you must specify the **TIME** dimension type on the Dimension Designer – General page. See the following display.




Specifying the **TIME** dimension type enables **Add supplied time hierarchies** on the Dimension Designer – Level page. The **Add** button is converted to a drop-down list of options. The **Add levels** and **Add supplied time hierarchies** options are now available for selection. See the following displays.




This display shows the **Add** button when selected.




Select **Add supplied time hierarchies**. This opens the Add Supplied dialog box. Select from the list of supplied time hierarchies to create the time levels. This also creates the hierarchies for the dimension. See the following display.



You can then define properties such as time type and sort order for the levels that you have selected. See the following display.




The hierarchy or hierarchies that are selected on the Add Supplied – dialog box are listed in the **Hierarchies** panel on the Dimension Designer – Hierarchy page. If there is only one hierarchy, as with this example, the hierarchy name is changed to match the dimension name. See the following display.




Define Measures

You can now define the measures for the cube. In this example, you define measures for the CostPrice Per Unit. Define the measures for the cube at the Cube Designer – Select Measures page, as shown in the following display.




Modify any measure attributes such as measure captions and formats at the Cube Designer – Measure Details page, as shown in the following display.




Define Member Properties

You can now define the member properties for any needed cube members. A *member property* is an attribute of a dimension member. A member property is also an optional cube feature that is created in a dimension to provide users with additional information about members. For this example, you can define the customer gender as a member property. Define member properties in the Cube Designer - Member Property dialog box, as seen in the following display.




At the Define a Member Property page, enter the member property name, level, column, format, and caption.




Define Aggregations

You can now define the aggregations for the cube. Aggregations are summaries of detailed data that are stored with a cube or referred to by a cube. They can help

contribute to faster query response. Define the aggregations for the cube from the Cube Designer - Aggregations page, as shown in the following display.




Select **add** to specify a user-defined aggregation. This opens the Add Aggregation dialog box, as shown in the following display. In this dialog box you can select levels to add to the aggregation that you are defining.



Select **OK** to return to the Cube Designer - Aggregations page, where the new aggregation is listed. Select **Next** to go to the Cube Designer - Finish page.

Build the Cube


You can now build the cube. On the Cube Designer - Finish page, review the settings for the cube, and then select one of the cube creation options, as shown in the following display.



You can choose to do one of the following:

- 1 save the metadata and create the cube
- 2 save the metadata but do not create the cube

You can also select whether to save the generated PROC OLAP code. Select **Export Code**. This opens the Export Code dialog box, as shown in the following display.



You can select to save either the long or short form of the code, or both. Enter the file location(s) where you want to save the resulting code. Select **OK** when finished to return to the Cube Designer - Finish page.

On the Cube Designer - Finish page, select **Finish** to complete the wizard. If the cube builds successfully, you will receive progress messages about the cube build, as seen in the following display.



Save a Cube's PROC OLAP Code

In SAS OLAP Cube Studio, you can elect to save the PROC OLAP code that is generated when a cube is built. The code is saved to a text file that you specify. The information saved in the file includes the following items:

- the SAS LIBNAME statement
- any FMTSEARCH statements
- any additional SAS code
- the PROC OLAP statement
- the METASVR statement
- all other PROC OLAP statements

You can access the Save PROC OLAP Code dialog box by using one of the following methods:

- 1 On the tree view in SAS OLAP Cube Studio, right-click on a cube and select **Export Code**.
- 2 In the Cube Designer – Finish page in the Cube Designer wizard, right-click **Export Code**.

The Export Code dialog box opens. You can select to save either the long or short form of the code, or both. Enter the file location or locations where you want to save the resulting code. Select **OK** when finished.

PROC OLAP Statements and Options For a Detail Table

The PROC OLAP code that is generated when a detail cube is built is listed below. A detail cube is unique in that it uses the DATA= option to specify the data source for the cube. The statements each have options that are either required or optional, depending on the cube structure.

Table 8.1 Statements and Options Used to Load Cubes from a Detail Table

Statements	Options	Required or Optional?
PROC OLAP	DATA=	Required
	CUBE=	Required
	PATH=	Required
	DESC=	Optional

Statements	Options	Required or Optional?
METASVR	NO_NWAY	Optional
	WORKPATH=	Optional
	DT_TABLE	Optional
	OLAP_SCHEMA=	Required
	REPOSITORY=	Optional
	HOST=	Optional
	PORT=	Optional
	PROTOCOL=	Optional
	USERID=	Optional
DIMENSION	PW=	Optional
	HIERARCHIES=	Required
	DESC=	Optional
	CAPTION=	Optional
	TYPE=TIME	Required only for TIME dimensions
LEVEL	SORT_ORDER=	Optional
		The LEVEL statement is optional unless you want to specify time periods for each level in a TIME dimension. If you specify a time period for one level, then you must specify a time period for all levels. To specify a time period, you use the TYPE= option.
HIERARCHY	LEVELS=	Required
	DESC=	Optional
	CAPTION=	Optional
MEASURE	STAT=	Required
	COLUMN ANALYSIS=	Required
	AGGR_COLUMN=	Required if you use the AGGREGATION statement with the /TABLE= option
	DESC=	Optional
	CAPTION=	Optional
	UNITS=	Optional
	FORMAT=	Optional
	DEFAULT=	Optional
PROPERTY	prop-name	Required

Statements	Options	Required or Optional?
	LEVEL=	Required
AGGREGATION		The AGGREGATION statement is optional unless you are creating additional aggregations. In that case, you must specify the names of the contiguous levels to be used to create the aggregation. Use the /TABLE= option for cubes that contain aggregated data from tables other than the input data source.

PROC OLAP Example for a Detail Table

In addition to building a cube in SAS OLAP Cube Studio, you can build an OLAP cube with PROC OLAP code and execute the code in a SAS session. Running PROC OLAP registers your cube and its sources in a metadata repository. It also creates the files that make up the cube. These are the *possible* input types for an OLAP cube that is built from a detail table:

- a data table (specified in the PROC OLAP statement DATA= option)
- dimension tables (specified in the DIMENSION statement DIMITBL= option)
- presummarized tables (specified in the AGGREGATION statement TABLE= option)

In this example, you use data from a recent product marketing campaign. For this cube, you establish measures and summaries of various aspects of the data, such as product statistics, geographic location of potential customers, and revenue summaries.

- 1 **Define the metadata profile and general information.** You use the PROC OLAP and METASVR statements here. The detail table is specified in the PROC OLAP statement DATA= option. The CUBE= option is used to specify the *metadata* folder location of a cube and the name of the cube. The folder must already exist in the metadata. The PATH= option specifies the physical or logical path to the location of a new cube. Within the specified path, the cube is stored in a directory that uses the name of the cube in uppercase letters.

The METASRV statement is used to establish the metadata connection. It identifies the metadata repository in which existing cube metadata information exists or in which metadata about a new cube should be stored. The statement is also used to provide a user's identification and password for the identified repository. Also, the DRILLTHROUGH_TABLE= option is used here to indicate the drill-through table. Drill-through tables are optional and can be used by client applications to provide a view from processed data into the underlying data source.

```
proc olap data=olapsio.detail
  cube=/cubefolders/mycubefolder/Campaign1'
  path="c:\cubes"
;
metasvr host=localhost
  port=9999
  protocol=bridge
```

```

userid=userid
pw=pw
repository=Foundation
olap_schema="OLAP Schema"
;

```

- 2 Define dimensions, levels, and hierarchies.** Now that your basic metadata server and cube information has been entered, you can define the different dimensions and their respective levels and hierarchies. You use the DIMENSION, HIERARCHY, and LEVEL statements here.

Note: For time-specific levels in a dimension, the LEVEL statement is required. Also, there can be only one time-specific dimension. △

```

dimension products
  hierarchies=(products)
  caption="Products"
  sort_order=ascunformatted
;
hierarchy products
  levels=(product_category product_group product_id
product_line product_name)
;
dimension date
  hierarchies=(date)
  caption="Date"
  type=time
;
hierarchy date
  levels=(year quarter month)
;
level year
  type=year
;
level quarter
  type=quarters
;
level month
  type=months
;
dimension geography
  hierarchies=(geography)
  caption="Geography"
;
hierarchy geography
  levels=(Country Region City)
;
dimension customer
  hierarchies=(customer)
  caption="Customer"
;
hierarchy customer
  levels=(customer_age customer_group customer_type)
;
dimension orders

```

```

    hierarchies=(orders)
    caption="Orders"
    ;
  hierarchy orders
    levels=(order_date total_retail_price costprice_per_unit)
    ;

```

- 3 Define measures.** You can now define the measures for the cube. A *measure* is an input column and a roll-up rule (statistic). Only certain measures are physically stored. Other measures are derived from the stored measures at run time. In this example, you want measures for the product CostPrice_per_unit average and a range.

You use the MEASURE statement here.

```

measure costprice_per_unitrange
  column=costprice_per_unit
  stat=range
  format=dollar10.2
  ;
measure costprice_per_unitavg
  column=costprice_per_unit
  stat=avg
  format=10.2
  ;

```

- 4 Define member properties.** You can now define the member properties for any needed cube members. A *member property* is an attribute of a dimension member. A member property is also an optional cube feature that is created in a dimension to provide users with additional information about members. For this example, you can define the customer genders as a member property.

You use the PROPERTY statement here.

```

property gender
  column=customer_gender
  hierarchy=customer
  level=customer_id
  ;

```

- 5 Define aggregations.** You can now define the aggregations for the cube. Aggregations are summaries of detailed data that is stored with a cube or referred by a cube. Their existence can reduce cube query time. If all aggregations are to be generated at the time of cube creation (MOLAP cube), then you can select specific aggregations that must be created in addition to the NWAY, which is the only aggregation that PROC OLAP makes by default.

You use the AGGREGATION statement here.

```

aggregation product_group product_line
  product_category country year
  name='ProductYear'
  ;

```

- 6 Build the cube.** You can now build the cube. Execute the PROC OLAP statement within the SAS System or in batch-mode. Here is the complete PROC OLAP code.

```

proc olap data=olapsio.detail
  cube=/cubefolders/mycubefolder/Campaign1'
  path="c:\cubes"

```

```

;
metasvr host=localhost
port=9999
protocol=bridge
userid=userid
pw=pw
repository=Foundation
olap_schema="OLAP Schema"
;
dimension products
hierarchies=(products)
caption="Products"
sort_order=ascunformatted
;
hierarchy products
levels=(product_category product_group product_id
product_line product_name)
;
dimension date
hierarchies=(date)
caption="Date"
type=time
;
hierarchy date
levels=(year quarter month)
;
level year
type=year
;
level quarter
type=quarters
;
level month
type=months
;
dimension geography
hierarchies=(geography)
caption="Geography"
;
hierarchy geography
levels=(Country Region City)
;
dimension customer
hierarchies=(customer)
caption="Customer"
;
hierarchy customer
levels=(customer_age customer_group customer_type)
;
dimension orders
hierarchies=(orders)
caption="Orders"
;
hierarchy orders

```

```

        levels=(order_date total_retail_price costprice_per_unit)
        ;
measure costprice_per_unitrange
        column=costprice_per_unit
        stat=range
        format=dollar10.2
        ;
measure costprice_per_unitavg
        column=costprice_per_unit
        stat=avg
        format=10.2
        ;
property gender
        column=customer_gender
        hierarchy=customer
        level=customer_id
        ;
aggregation product_group product_line
        product_category country year
        name='ProductYear'
        ;
run;

```

Note: Any libraries must be specified before you run the PROC OLAP code. \triangle

Note: When a SAS OLAP cube is created, a directory for that cube is also created. This directory is assigned the same name as the cube, but in uppercase letters. For example, when you save a cube in **C:\olapcubes** and name the cube Campaigns, the cube is saved in the directory **C:\olapcubes\CAMPAIGNS**. \triangle

Building a Cube from a Star Schema

In this example, you build a cube from a star schema. A *star schema* is a data source that contains tables in a database in which a single fact table is connected to multiple dimension tables. With a cube based on a star schema, you identify the fact table, the dimension tables, and the keys that map the tables together. In this example, you use data from a recent product marketing campaign to establish measures and summaries of various aspects of the data, such as product statistics, geographic location of potential customers, and revenue summaries.

Enter General Cube Information

After you have established a connection profile, you can begin to create a cube. Select **File ► New ► Cube**. On the Cube Designer – General page, enter the basic cube information. For the input type in this example, you select **Star Schema**. The following display shows fields that you enter information for.

Cube Designer - General
Provide information about the cube that you want to create, and specify where the cube and its metadata will be stored.


Name: OrderCube
 Description: Star Schema Cube
 OLAP schema: SASApp - Olap Schema
 Location: /Shared Data
 Physical cube path: C:\
 Work path (optional): C:\tmp

Input Type: Star schema
 Cube will use aggregated data from other tables
 Include secured member values in presummarized computations

Buttons: New..., Select..., Browse..., Advanced, < Back, Next >, Cancel, Help


- Enter information in the following fields:
 - **Name**
 - **Description**
 - **OLAP schema**
 - **Location** (SAS folder)
 - **Physical path** (path in the file system to store the cube)
 - **Work path** (path for temporary work files)
 - **Input Type** (star schema)

After you have selected the star schema input type, you can select **Advanced** and define star schema options and dimension table translation languages for the cube. The following display shows the Advanced Cube Options dialog box. For star schema cubes you can specify options for missing keys and whether to compact the NWAY aggregation at build time.




Select Fact and Dimension Tables

On the Cube Designer – Input page, select a fact table for your cube. If one does not exist for your data, select **Define Table**, and then define the source that you will import your metadata from. The following display shows the fact table ORDFACT that is selected for the example cube.



After a fact table has been chosen, you select the dimension tables for the cube on the Cube Designer – Dimension Tables page. See the following display.



Drill-Through Table

In the Cube Designer - Drill-Through dialog box, you can select or define an optional drill-through table. Drill-through tables can be used by client applications to provide a view from processed data into the underlying data source.

If a drill-through table does not exist for your data, select **Define Table**, and then define the source that you will import your metadata from.

Table Options


The **Table Options** button is available in both the Cube Designer - Input and the Cube Designer - Drill-Through dialog boxes. It opens the **Table Options** dialog box. It enables you to specify data set options that are used to open the data set. For example, you could enter a WHERE clause or subsetting information that is then applied to the selected table when it is opened. The options are stored as part of the cube and then reapplied when the data is accessed at run time. You can also specify data set options in the Dimension Designer – General dialog box (for use with star schemas) and the Stored Aggregates dialog box (for use with summarized tables). For more information, see “Data Set Options” in *SAS Language Reference: Concepts*.

Define Dimensions, Levels and Hierarchies

Now that your basic metadata server and cube information has been entered, you can define the different dimensions and their respective levels and hierarchies. For this example, the following dimensions are created:

- Product
- Time
- Geography
- Customer
- Organization

On the Cube Designer - Dimensions page, select **Add**.



This opens the Dimension Designer—General page, as seen in the following display.

Dimension Designer

Dimension Designer - General
Specify the information for the dimension:

Name: Customer

Caption: Customer Dimension

Description: Dimension Table for Order cube

Type: STANDARD Allow new members during dimensional update

Sort order: Ascending Uniformly

Star Schema Table

Use the fact table

Table: CUSTDIM

Key: Customer_Id

Fact key: Customer_Id

Table options:

< Back Next > Cancel Help


Enter information in the following fields:

- Name**
- Caption**
- Description**
- Type** (Standard, GEO, or TIME)
- Sort order**


Because you are building a cube from star schema data, you must identify the table for the dimension, and the keys that link the dimension and fact tables together. For each dimension that you define, specify the following:

- Table** (for the dimension you are creating)
- Key** (for the dimension you are creating)
- Fact key**
- any **Table options** for the selected dimension table


Select **Next**. This opens the Dimension Designer - Level page. Next, select **Add** to open the Add Levels page, as seen in the following display.




Select the levels you want to add to the dimension. Select **OK** to return to the Dimension Designer - Level page, where the selected levels are listed. You can now define properties such as format, time type, and sort order for the levels that you have selected. See the following display.




Next, define hierarchies for the levels on the Dimension Designer - Hierarchy page. Select **Add** to on the Dimension Designer – Define a Hierarchy page and individually select the levels for the hierarchy.



You can also select **Finish** on the Dimension Designer - Hierarchy page to accept the order of the levels that are defined on the previous Dimension Designer - Level page. If you select this option, the hierarchy is assigned the same name as the dimension. See the following display.




Repeat this process for each dimension. After you create each dimension, it is listed in the **Dimensions** panel of the Cube Designer - Dimensions page. See the following display.




Note: If you are defining a TIME dimension you can select **Supplied** on the Dimension Designer - Level page. See “Creating a Time Dimension in SAS OLAP Cube Studio” on page 197 for further information. △

Define Measures

You can now define the measures for the cube. In this example, you define measures for the CostPrice Per Unit. Define the measures for the cube on the Cube Designer - Select Measures page, as shown in the following display.




Modify any measure attributes such as measure captions and formats on the Cube Designer - Measure Details page, as shown in the following display.



Define Member Properties

You can now define the member properties for any needed cube members. A *member property* is an attribute of a dimension member. A member property is also an optional cube feature that is created in a dimension to provide users with additional information about members. For this example, you can define the customer gender as a member property. Define member properties in the Cube Designer - Member Property dialog box, as seen in the following display.



On the Define a Member Property page, enter the member property name, level, column, format, and caption.

Define a Member Property

Name:

Level:

Column:

Format:

Caption:

Description:


Available:

Selected:


Define Aggregations

You can now define the aggregations for the cube. Aggregations are summaries of detailed data that are stored with a cube or referred to by a cube. They can help contribute to faster query response. You define the aggregations for the cube from the Cube Designer - Aggregations page.

Select **Add** to specify a user-defined aggregation. This opens the Add Aggregation dialog box, as shown in the following display. In this dialog box you can select levels to add to the aggregation that you are defining.




Select **OK** to return to the Cube Designer - Aggregations page, where the new aggregation is listed as shown in the following display.



Select **Next** to go to the Cube Designer - Finish page.

Build the Cube


You can now build the cube. On the Cube Designer - Finish page, review the settings for the cube, and then select one of the cube creation options, as shown in the following display.



You can choose to do one of the following:

- 1 save the metadata and create the cube
- 2 save the metadata but do not create the cube

You can also select whether to save the generated PROC OLAP code. Select **Export Code**. This opens the Export Code dialog box as shown in the following display.



PROC OLAP CODE for the Star Schema Example

```
LIBNAME olapsio BASE "\\olap\tmp\libolap" ;

PROC OLAP
  CUBE           = "/Shared Data/OLAPSchemas/OrderCube"
  PATH           = 'c:\v9cubes'
  DESCRIPTION    = 'starschemacube'
  FACT           = olapsio.ORDER_FACT
;

METASVR
  HOST           = "J12345.na.abc.com"
  PORT           = 8561
  OLAP_SCHEMA   = "SASApp - OLAP Schema";
```

```

DIMENSION Customers
  CAPTION          = 'Customers'
  SORT_ORDER      = ASCENDING
  DIMTBL          = olapsio.CUSTDIM
  DIMKEY          = Customer_Id
  FACTKEY         = Customer_ID
  HIERARCHIES     = (
    Customers
  ) /* HIERARCHIES */;

HIERARCHY Customers
  ALL_MEMBER = 'All Customers'
  CAPTION    = 'Customers'
  LEVELS     = (
    Customer_Type Customer_Group
    Customer_Gender Customer_Age
  ) /* LEVELS */
  DEFAULT;

LEVEL Customer_Type
  CAPTION      = 'Customer Type'
  SORT_ORDER  = ASCENDING;

LEVEL Customer_Group
  CAPTION      = 'Customer Group'
  SORT_ORDER  = ASCENDING;

LEVEL Customer_Gender
  FORMAT       = $GENDER.
  CAPTION      = 'Customer Gender'
  SORT_ORDER  = ASCENDING;

LEVEL Customer_Age
  CAPTION      = 'Customer Age'
  SORT_ORDER  = ASCENDING;

DIMENSION Geography
  CAPTION          = 'Geography'
  TYPE            = GEO
  SORT_ORDER      = ASCENDING
  DIMTBL          = olapsio.GEOGDIM
  DIMKEY          = Street_Id
  FACTKEY         = Street_ID
  HIERARCHIES     = (
    Geography
  ) /* HIERARCHIES */;

HIERARCHY Geography
  ALL_MEMBER = 'All Geography'
  CAPTION    = 'Geography'
  LEVELS     = (
    Country Region State
  ) /* LEVELS */

```

```

        DEFAULT;

LEVEL Country
    FORMAT          = $COUNTRY.
    CAPTION         = 'Country'
    SORT_ORDER     = ASCENDING;

LEVEL Region
    CAPTION         = 'Region Name'
    SORT_ORDER     = ASCENDING;

LEVEL State
    CAPTION         = 'State'
    SORT_ORDER     = ASCENDING;

DIMENSION Organization
    CAPTION         = 'Organization'
    SORT_ORDER     = ASCENDING
    DIMTBL         = olapsio.ORGDIM
    DIMKEY         = Employee_Id
    FACTKEY        = Employee_ID
    HIERARCHIES    = (
        Organization
    ) /* HIERARCHIES */;

HIERARCHY Organization
    ALL_MEMBER     = 'All Organization'
    CAPTION        = 'Organization'
    LEVELS         = (
        Company Group Department
    ) /* LEVELS */
    DEFAULT;

LEVEL Company
    CAPTION         = 'Company'
    SORT_ORDER     = ASCENDING;

LEVEL Group
    CAPTION         = 'Group'
    SORT_ORDER     = ASCENDING;

LEVEL Department
    CAPTION         = 'Department'
    SORT_ORDER     = ASCENDING;

DIMENSION Product
    CAPTION         = 'Product'
    SORT_ORDER     = ASCENDING
    DIMTBL         = olapsio.PRODIM
    DIMKEY         = Product_ID
    FACTKEY        = Product_ID
    HIERARCHIES    = (
        Product
    ) /* HIERARCHIES */;

```

```

HIERARCHY Product
  ALL_MEMBER = 'All Product'
  CAPTION    = 'Product'
  LEVELS     = (
    Product_Category Product_Group Product_Line
  ) /* LEVELS */
  DEFAULT;

LEVEL Product_Category
  CAPTION      = 'Product Category'
  SORT_ORDER   = ASCENDING;

LEVEL Product_Group
  CAPTION      = 'Product Group'
  SORT_ORDER   = ASCENDING;

LEVEL Product_Line
  CAPTION      = 'Product Line'
  SORT_ORDER   = ASCENDING;

DIMENSION Time
  CAPTION      = 'Time'
  TYPE         = TIME
  SORT_ORDER   = ASCFORMATTED
  DIMTBL       = olapsio.TIMEDIM
  DIMKEY       = Date
  FACTKEY      = Order_Date
  HIERARCHIES  = (
    Time
  ) /* HIERARCHIES */;

HIERARCHY Time
  ALL_MEMBER = 'null'
  LEVELS     = (
    Year Quarter Month
  ) /* LEVELS */
  DEFAULT;

LEVEL Year
  COLUMN      = Date
  FORMAT      = YEAR4.
  TYPE        = YEAR
  CAPTION     = 'Year'
  SORT_ORDER  = ASCFORMATTED;

LEVEL Quarter
  COLUMN      = Date
  FORMAT      = QTR1.
  TYPE        = QUARTERS
  CAPTION     = 'Quarter'
  SORT_ORDER  = ASCFORMATTED;

LEVEL Month

```

```

COLUMN          = Date
FORMAT          = MONNAME9.
TYPE            = MONTHS
CAPTION         = 'Month'
SORT_ORDER      = ASCFORMATTED;

```

```
PROPERTY Ages
```

```

LEVEL          = Customer_Age
COLUMN         = Customer_Birth_Date
CAPTION        = 'Customer Birth Date'
HIERARCHY      = (
  Customers
) /* HIERARCHIES */;

```

```
MEASURE Total_Retail_PriceMAX
```

```

STAT          = MAX
COLUMN        = Total_Retail_Price
CAPTION       = 'Maximum Total_Retail_Price'
FORMAT        = DOLLAR13.2
DEFAULT;

```

```
MEASURE Total_Retail_PriceAVG
```

```

STAT          = AVG
COLUMN        = Total_Retail_Price
CAPTION       = 'Average Total_Retail_Price'
FORMAT        = DOLLAR13.2;

```

```
MEASURE CostPrice_Per_UnitMAX
```

```

STAT          = MAX
COLUMN        = CostPrice_Per_Unit
CAPTION       = 'Maximum CostPrice_Per_Unit'
FORMAT        = DOLLAR13.2;

```

```
MEASURE CostPrice_Per_UnitAVG
```

```

STAT          = AVG
COLUMN        = CostPrice_Per_Unit
CAPTION       = 'Average CostPrice_Per_Unit'
FORMAT        = DOLLAR13.2;

```

```
AGGREGATION /* Default */
```

```

/* levels */
Company Country Customer_Age Customer_Gender Customer_Group
Customer_Type Department Group Month Product_Category
Product_Group Product_Line Quarter
Region State Year
/ /* options */
NAME          = 'Default';

```

```
AGGREGATION /* ProductCustomer */
```

```

/* levels */
Customer_Type Customer_Group Customer_Gender
Customer_Age Product_Category
/ /* options */
NAME          = 'ProductCustomer';

```

```

FORMAT Customer_Birth_Date DATE9.;

RUN;

```

PROC OLAP Statements and Options For a Star Schema

The following table lists the PROC OLAP statements and options that you use to load a cube from a star schema. A star schema is a set of input tables that are defined in a repository. The set of tables includes a single fact table and one or more dimension tables. The fact table must contain at least one numeric analysis column for each set of measures that is generated. To specify the data source for a star schema, you must use the FACT=, DIMTBL=, DIMKEY=, and FACTKEY= options.

Table 8.2 Statements and Options Used to Load Cubes from a Star Schema

Statements	Options	Required or Optional?
PROC OLAP	FACT=	Required
	CUBE=	Required
	PATH=	Required
	DESC=	Optional
	NO_NWAY	Optional
	WORKPATH=	Optional
	DT_TABLE	Optional
METASVR	OLAP_SCHEMA=	Required
	REPOSITORY=	Optional
	HOST=	Optional
	PORT=	Optional
	PROTOCOL=	Optional
	USERID=	Optional
DIMENSION	PW=	Optional
	HIERARCHIES=	Required
	DIMTBL=	Required for cubes that support one locale. If the cube will contain multiple national languages, replace this option with DIMTABLELIBREF= and DIMTABLEMEMMPREF=.
	FACTKEY=	Required
	DIMKEY=	Required
	DIMTABLELIBREF	Required if you build a cube that will contain multiple national languages. Replaces DIMTBL=.

Statements	Options	Required or Optional?
	DIMTABLEMEMMPREF	Required if you build a cube that will contain multiple national languages. Replaces DIMTBL=.
	DESC=	Optional
	CAPTION=	Optional
	TYPE=TIME	Required only for TIME dimensions
	SORT_ORDER=	Optional
LEVEL		The LEVEL statement is optional unless you want to specify time periods for each level in a TIME dimension. If you specify a time period for one level, then you must specify a time period for all levels. To specify a time period, you use the TYPE= option.
HIERARCHY	LEVELS=	Required
	DESC=	Optional
	CAPTION=	Optional
MEASURE	STAT=	Required
	COLUMN ANALYSIS=	Required
	AGGR_COLUMN=	Required if you use the AGGREGATION statement with the TABLE= option
	DESC=	Optional
	CAPTION=	Optional
	UNITS=	Optional
	FORMAT=	Optional
	DEFAULT=	Optional
PROPERTY	prop-name	Required


Statements	Options	Required or Optional?
AGGREGATION	LEVEL=	Required
		The AGGREGATION statement is optional unless you are creating additional aggregations. In that case, you must specify the names of the contiguous levels to be used to create the aggregation. Use the TABLE= option for cubes that contain aggregated data from tables other than the input data source.

Building a Cube from a Summary Table

In this example, you build a cube with fully summarized data. A *summary table* is a data source that contains a crossing of all dimensions for a cube. In this example, furniture company sales data for the PRDNWYPR table is used. The table contains columns for location, sales, date, and product values. It also contains stored measures for actual and predicted sales figures.

Enter General Cube Information

After you have established a connection profile, you can begin to create a cube. Select **File ► New ► Cube**. On the Cube Designer – General page, enter the basic cube information. For this example input type, you select **Fully Summarized Table**. The following display shows fields that you enter information for.




- Enter information in the following fields:
 - **Name**
 - **Description**
 - **OLAP schema**
 - **Location** (SAS folder)
 - **Physical cube path** (path in the file system to store the cube)
 - **Work path** (path for temporary work files)
 - **Input Type** (fully summarized table)

Note: When the **Fully summarized table** input type is selected, the option **Cube will use aggregated data from other tables** is automatically selected. △

Select an Input Table

On the Cube Designer – Input page, select an input table for your cube. If one does not exist for your data, select **Define Table**, and then define the source that you will import your metadata from. The following display shows the summarized table PRDNWPR that is selected for the example cube.



Drill-Through Table

In the Cube Designer - Drill-Through dialog box, you can select or define an optional drill-through table. Drill-through tables can be used by client applications to provide a view from processed data into the underlying data source.

If a drill-through table does not exist for your data, select **Define Table**, and then define the source that you will import your metadata from.

Table Options

The **Table Options** button is available in both the Cube Designer - Input and the Cube Designer - Drill-Through dialog boxes. It opens the Table Options dialog box. It


enables you to specify data set options that are used to open the data set. For example, you could enter a WHERE clause or subsetting information that is then applied to the selected table when it is opened. The options are stored as part of the cube and then reapplied when the data is accessed at run time. You can also specify data set options in the Dimension Designer – General dialog box (for use with star schemas) and the Stored Aggregates dialog box (for use with summarized tables). For more information, see “Data Set Options” in *SAS Language Reference: Concepts*.

Define Dimensions, Levels and Hierarchies

Now that your basic metadata server and cube information has been entered, you can define the different dimensions and their respective levels and hierarchies. For this example, the following dimensions are created:

- Products
- Dates
- Geography

On the Cube Designer - Dimensions page, select **Add**.



This opens the Dimension Designer – General page, as seen in the following display.

Dimension Designer - General
Specify the information for the dimension.

Name:

Caption:

Description:

Type: Allow new members during incremental updates

Sort order:

Star Schema Table

Use the fact table

Table:

Key:

Factkey:

Table column:

Enter information in the following fields:

- Name**
- Caption**
- Description**
- Type** (Standard, GEO, or TIME)
- Sort order**.

Select **Next**. This opens the Dimension Designer - Level page. Next, select **Add** to open the Add Levels dialog box, as seen in the following display.

Add Levels

Available input columns:


- COUNTRY
- DIVISION
- MONTH
- PRODTYPE
- PRODUCT
- QUARTER
- REGION
- YEAR
- _FREQ_
- _TYPE_
- actn
- actsum
- co_hq
- co_pop
- predn

Input columns for new levels:


- COUNTRY
- DIVISION

Select the levels you want to add to the dimension. Select **OK** to return to the Dimension Designer - Level page, where the selected levels are listed. You can now


define properties such as format, time type, and sort order for the levels that you have selected. See the following display.



Next, define hierarchies for the levels on the Dimension Designer - Hierarchy page. You can select **Add** to open the Define a Hierarchy page and individually select the levels for the hierarchy.




Or you can select **Finish** to accept the order of the levels that are defined on the previous Dimension Designer - Level page. If you select this option, the hierarchy is assigned the same name as the dimension. See the following display.



Repeat this process for each dimension. For this example, the following dimensions are created:

- Geography*: (levels: Country, Division)
- Products*: (levels: Product, Product Type)
- Dates*: (Year, Quarter, Month)

After you create each dimension it is listed in the **Dimensions** panel of the Cube Designer - Dimensions page. See the following display:



Define Measures


For this example, stored measures and derived measures are created. Stored measures are base measures that are loaded from the fully summarized table. When you are creating a cube from a fully summarized table, the table must have a corresponding single column for each stored measure that you want to include in the cube. The base statistics are SUM, N, MIN, MAX, NMISS, and USS.

Derived measures are measures that are built from the stored measures that you have selected for the cube. Derived measures are assigned to an analysis group when they are created. An analysis group is used to identify the numeric column in the original unsummarized data source that was used as the analysis variable for the stored measure. It can also be a name that identifies a logical association between several stored measures.


Note: For further information on measure statistics, see “Statistics Available for Measures” on page 48. △

Stored Measures

You can now select the stored (base) measures for the cube in the Cube Designer - Select Stored Measures page. From the list of available measures, select the stored measures that you want to include in the cube. For this example **actn**, **actsum**, **predn**, and **predsum** are included.




On the Cube Designer - Assign Stored Measures page, you can specify the **Statistic** and **Analysis Group** options for the stored measures. For the **Statistic**, select the appropriate statistic from the drop-down menu. For the **Analysis Group**, enter a name that identifies a logical association between the stored measures. The **Analysis Group** name can identify the numeric column in the original unsummarized data source that was used as the analysis variable for the stored measure. If the table contained two measures from the same analysis column, both of the base measures should have the same analysis group specified.




Derived Measures


On the Cube Designer – Select Derived Measures page, specify the measures that are derived from the stored measures. Each derived measure is based on a set of required stored measures. If the stored measures for an analysis group do not include all those required for a specific derived measure, then that measure cannot be included in the cube. On the **Define a derived measure** panel, select the **Analysis group**, **Derived statistic**, and **Measure name** for the derived measure that you are creating.



When you are finished, select **Apply**. The derived measure is listed in the **Derived Measures** list.




For this example, two derived measures are created: ActualDerived and Predicted Derived.



Edit Measure Details

On the Cube Designer – Edit Measure Details page, you can select a default measure and modify measure details for the stored and derived measures.




Define Member Properties

You can now define the member properties for any needed cube members. A *member property* is an attribute of a dimension member. A member property is also an optional cube feature that is created in a dimension to provide users with additional information about members. Define member properties in the Cube Designer - Member Property dialog box.

Define Aggregation Tables


On the Cube Designer – Aggregation Tables page, you associate aggregation tables with the summarized data source that you specified as the input data source for the cube. When you open the Cube Designer – Aggregation Tables page, the input data table that you selected to build your cube with is listed in the **Selected tables** list. You can then select a table to use as the aggregation table from the **Available tables** list and move it to the **Selected tables** list. For this example, the table PGEOTIME is used as the aggregation table.

Note: If the cube is loaded from a fully summarized data source, then the measure names within the selected aggregation tables must match the measure names in the input data source. If the cube is loaded from a detail table or a star schema, then all of the selected aggregation tables must use the same measure names. For all cubes, the levels must be the same as those in the input data source. Δ




Define Stored Aggregations


You can now define stored aggregations for the cube. Stored aggregations are aggregations that are stored in the aggregation tables. On the Cube Designer - Stored Aggregations page, select **Add** to create a stored aggregation.



In the **Specify an aggregation** dialog box, enter the aggregation name, aggregation table, and any necessary table options for the aggregation. Then select the levels that are used in the aggregation. For this example, the stored aggregations named Aggregation1 and Aggregation2 are created.



The aggregations are listed on the Cube Designer – Stored Aggregations page.




Define Aggregations

If needed, you can now define aggregations that are stored with the cube. Define the aggregations for the cube from the Cube Designer - Aggregations page.

Build the Cube

After you have finished entering information for the summarized cube, you can build the cube. On the Finish page, you can review the details of the cube that you just defined. You also choose whether to only save the cube's definition to the active metadata or to save the cube's definition and build the cube. Select the option **Save the metadata and create the cube**. You can also choose to export the PROC OLAP code to a text file. When you are done, select **Finish**. The cube is processed and the Cube Designer wizard will close.



PROC OLAP CODE for the Summary Table Example

```

OPTIONS VALIDVARNAME=ANY;
LIBNAME olapsio BASE '\\\olap\tmp\libolap" ;

PROC OLAP
  CUBE              = "/Shared Data/OLAPSchema/SumCube"
  PATH              = 'C:\v9cubes'
  DESCRIPTION       = 'Fully Summarized Cube'
  NONUPDATEABLE
  MAXTHREADS       = 5000
;

METASVR
  HOST              = "J12345.na.sas.com"
  PORT              = 8561
  OLAP_SCHEMA       = "SASApp - OLAP Schema";

DIMENSION Geography
  CAPTION           = 'Geography'
  SORT_ORDER        = ASCENDING
  HIERARCHIES       = (
    Geography
  ) /* HIERARCHIES */;

HIERARCHY Geography
  ALL_MEMBER        = 'All Geography'
  CAPTION           = 'Geography'
  LEVELS            = (
    COUNTRY DIVISION
  ) /* LEVELS */

```

```

    DEFAULT;

LEVEL COUNTRY
    FORMAT          = $CHAR10.
    CAPTION         = 'Country'
    SORT_ORDER     = ASCENDING;

LEVEL DIVISION
    FORMAT          = $CHAR10.
    CAPTION         = 'Division'
    SORT_ORDER     = ASCENDING;

DIMENSION Products
    CAPTION         = 'Products'
    SORT_ORDER     = ASCENDING
    HIERARCHIES    = (
        Products
    ) /* HIERARCHIES */;

HIERARCHY Products
    ALL_MEMBER     = 'All Products'
    CAPTION        = 'Products'
    LEVELS         = (
        PRODUCT PRODTYPE
    ) /* LEVELS */
    DEFAULT;

LEVEL PRODUCT
    FORMAT          = $CHAR10.
    CAPTION         = 'Product'
    SORT_ORDER     = ASCENDING;

LEVEL PRODTYPE
    FORMAT          = $CHAR10.
    CAPTION         = 'Product type'
    SORT_ORDER     = ASCENDING;

DIMENSION Dates
    CAPTION         = 'Dates'
    TYPE           = TIME
    SORT_ORDER     = ASCENDING
    HIERARCHIES    = (
        Dates
    ) /* HIERARCHIES */;

HIERARCHY Dates
    ALL_MEMBER     = 'All Dates'
    CAPTION        = 'Dates'
    LEVELS         = (
        YEAR QUARTER MONTH
    ) /* LEVELS */
    DEFAULT;

LEVEL YEAR

```

```

        FORMAT          = 4.
        TYPE            = YEAR
        CAPTION         = 'Year'
        SORT_ORDER     = ASCENDING;

LEVEL QUARTER
        FORMAT          = 8.
        TYPE            = QUARTERS
        CAPTION         = 'Quarter'
        SORT_ORDER     = ASCENDING;

LEVEL MONTH
        FORMAT          = MONNAME3.
        TYPE            = MONTHS
        CAPTION         = 'Month'
        SORT_ORDER     = ASCENDING;

MEASURE predsumSUM
        STAT            = SUM
        ANALYSIS        = PredictGroup
        AGGR_COLUMN    = predsum
        CAPTION         = 'Sum of predsum'
        FORMAT          = Best12.
        DEFAULT;

MEASURE prednN
        STAT            = N
        ANALYSIS        = PredictGroup
        AGGR_COLUMN    = predn
        CAPTION         = 'Number of values for predn'
        FORMAT          = 12.;

MEASURE actsumSUM
        STAT            = SUM
        ANALYSIS        = ActualGroup
        AGGR_COLUMN    = actsum
        CAPTION         = 'Sum of actsum'
        FORMAT          = Best12.;

MEASURE actnN
        STAT            = N
        ANALYSIS        = ActualGroup
        AGGR_COLUMN    = actn
        CAPTION         = 'Number of values for actn'
        FORMAT          = 12.;

MEASURE PredictedDerived
        STAT            = AVG
        ANALYSIS        = PredictGroup
        CAPTION         = 'Average PredictGroup'
        FORMAT          = Best12.;

MEASURE ActualDerived
        STAT            = AVG

```

```

ANALYSIS      = ActualGroup
CAPTION       = 'Average ActualGroup'
FORMAT        = Best12.;

AGGREGATION /* Default */
/* levels */
COUNTRY DIVISION MONTH PRODTYPE
PRODUCT QUARTER YEAR
/ /* options */
TABLE        = olapsio.PRDNWYPR
NAME         = 'Default';

AGGREGATION /* Aggregation1 */
/* levels */
COUNTRY YEAR
/ /* options */
TABLE        = olapsio.PGEOTIME
NAME         = 'Aggregation1';

AGGREGATION /* Aggregation2 */
/* levels */
COUNTRY DIVISION PRODTYPE
PRODUCT QUARTER YEAR
/ /* options */
TABLE        = olapsio.PRDNWYPR
NAME         = 'Aggregation2';

RUN;
```

PROC OLAP Statements and Options For a Summary Table

The following table lists the PROC OLAP statements and options that you use to load cubes from a fully summarized data source (a crossing of all dimensions also known as an NWAY). Unlike a detail table or star schema, a fully summarized cube does not use either the DATA= or FACT= option to specify the data that is used to build the cube. Instead the TABLE= option is used on the AGGREGATION statement.

Table 8.3 Statements and Options Used to Load Cubes from Fully Summarized Data

Statements	Options	Required or Optional?
PROC OLAP	CUBE=	Required
	PATH=	Required
	DESC=	Optional
	NO_NWAY	Optional
	WORKPATH=	Optional
	DT_TABLE	Optional
METASVR	OLAP_SCHEMA=	Required
	REPOSITORY=	Optional
	HOST=	Optional

Statements	Options	Required or Optional?
	PORT=	Optional
	PROTOCOL=	Optional
	USERID=	Optional
	PW=	Optional
DIMENSION	HIERARCHIES=	Required
	DESC=	Optional
	CAPTION=	Optional
	TYPE=TIME	Required only for TIME dimensions
	SORT_ORDER=	Optional
LEVEL		The LEVEL statement is optional unless you want to specify time periods for each level in a TIME dimension. If you specify a time period for one level, then you must specify a time period for all levels. To specify a time period, you use the TYPE= option.
HIERARCHY	LEVELS=	Required
	DESC=	Optional
	CAPTION=	Optional
MEASURE	STAT=	Required
	COLUMN ANALYSIS=	Required
	AGGR_COLUMN=	Required
	DESC=	Optional
	CAPTION=	Optional
	UNITS=	Optional
	FORMAT=	Optional
	DEFAULT=	Optional
PROPERTY	prop-name	Required
	LEVEL=	Required

Statements	Options	Required or Optional?
AGGREGATION	Names of the contiguous levels to be used to create the aggregation	Required (additional AGGREGATION statements without the TABLE= option can be used to create aggregations other than the automatically defined NWAY).
	TABLE=	Required


Tuning Aggregations For a Cube

After you have built a cube, you can modify the aggregations for the cube with the Aggregation Tuning function in SAS OLAP Cube Studio. The Aggregation Tuning function enables you to either automatically generate aggregation recommendations or to manually define cube aggregations. In the SAS OLAP Cube Studio tree view, select a cube and select **Aggregation Tuning** from the **Actions** menu. In the Aggregations Tuning dialog box, you can select from three different methods of aggregation tuning:

- 1 “Cardinality Tuning” on page 151
- 2 “Manual Tuning” on page 156
- 3 “Arm Log Tuning” on page 158


Cardinality Tuning

The **Cardinality** tab enables you to add aggregations that are recommended based on the relative cardinality (number of members) of the cube levels. This method of adding aggregations is used when a cube is first created and before an ARM log can be generated. On the **Cardinality** tab, select **Analyze**.




The Analysis Recommendations dialog box opens. A list of recommended aggregations is displayed. This list can contain up to 100 aggregation recommendations

for the cube. Select the aggregation recommendations that you want to add to the Aggregations table. Click **Select All** to choose all of the aggregation recommendations. Select **OK** to execute the analysis. The selected aggregation recommendations are added to the list of aggregations in the **Aggregations** table.




In the Aggregation Tuning dialog box, the **Aggregations** table displays the newly recommended aggregations. In the **Status** column, the new aggregation symbol is listed for each new aggregation. The **Method** column lists the recommended aggregations as **Cardinality** type aggregations. You can scroll through the list of aggregations to analyze which aggregations to keep, edit properties for aggregations, or remove aggregations.




You can also view the levels that compose the aggregation recommendations. Select the **Preview levels** checkbox on the Aggregation Tuning dialog box. The **Preview levels** panel opens, displaying the levels for the currently selected aggregation recommendation.


Note: The **Preview levels** panel is read-only. Δ




In the Aggregation Tuning dialog box, you can select to edit or drop aggregations that loaded with the cube. When you select an existing aggregation, the **Drop** button becomes active.



If you select **Drop**, the drop icon appears in the **Status** column for that aggregation. In addition, the **Clear Drop** button becomes active. This button enables you to clear the drop status for an aggregation. Those aggregations that are marked with the drop icon are dropped from the cube the next time you select the **Build Aggregations** button.




You can also review those aggregations that were recommended by cardinality analysis, and if needed delete them. When you select a newly recommended aggregation in the **Aggregations** table, the **Delete** button becomes active.




If you select **Delete**, a message dialog box opens. Select **Yes** to delete the selected aggregations. Select **No** to cancel the Delete function.




You can also modify the performance options for the cube aggregations. In the Aggregation Tuning dialog box, select **Options**. This opens the Performance Options dialog box. On the **Default** tab, you can set default performance option settings for all generated aggregations for the cube.



On the **Aggregation** tab you can select an individual aggregation and change performance option settings for that specific aggregation. Select **OK** when you are finished changing the performance options. The Performance Options dialog box closes.



If needed, you can export the SAS code that is used to build the aggregations that are listed in the **Aggregations** table. Select **Export Code** on the Aggregation Tuning dialog box. The Export Code dialog box opens. Enter the path of the text file that you are exporting the code to. Click **OK** to create the file. The code is stored in a text file that you can further review, edit, and use to recreate the cube later.




When you have finished updating and modifying the aggregations listed in the **Aggregations** table, you can select **Build Aggregations**. The Aggregation Tuning function builds the recommended aggregations, and drops any existing aggregations that are marked with the drop icon. A build success confirmation message opens and the Aggregation Tuning dialog box closes. Click **OK** on the message box.




Manual Tuning

The **Manual** tab enables you to select the exact hierarchies and levels that you want use to generate aggregation recommendations from. The manual tuning option is used when you have dimension levels that are frequently used with other dimension levels. It is also used when a level has unique member counts (NUNIQUE) set.


Select the **Manual** tab on the Aggregation Tuning dialog box. The hierarchies for the cube are listed individually as columns. Levels for the hierarchies are numerically listed in drop-down lists on the columns. When you select an individual level from a hierarchy, you are selecting that level and its parent levels.




From the drop-down lists, select the levels that you want to use to create the aggregation recommendation. Both the **Add to List** and the **Reset All** buttons are enabled whenever one or more of the drop-down lists are changed to a selection other than the default selection of **(None)**.



For this example, select the levels for **Campaign_Type**, **Campaign_Start_Year**, **Product_Group**, and **Division**. Select **Add to List**.




The aggregation recommendation that you created is listed in the **Aggregations** table. On the row for that aggregation, the levels that were manually selected for the aggregation are displayed and the new aggregation icon is listed in the **Status** column. When you are finished adding aggregation recommendations, select **Build Aggregation** to build the aggregations for the cube.




Arm Log Tuning

The **ARM Log** tab enables you to add aggregations that are based on the query analysis records that are stored in an Application Response Measuring (ARM) log. ARM analysis is recommended for optimal cube tuning. When the ARM API is used, a log file can be created that can be gleaned for query performance details. The ARM log is used to analyze query patterns and determine which aggregations to generate that will most likely have a positive impact on query performance. In order to create an ARM log, ARM logging must be turned on and queries must be performed against the cube. This function is only as effective as the amount of ARM data that is provided and it also depends on whether the ARM log data truly reflects the future cube query patterns. For more information on creating the ARM log see “Using ARM to Monitor SAS OLAP Server Performance” in the *SAS 9.2 Intelligence Platform: System Administration Guide*.


Select the **ARM Log** tab and select the option **Create aggregation recommendations and update existing performance values based on the ARM log**. In the **Enter an ARM log** text field, enter the file path for the ARM log that you will use to generate aggregation recommendations. Use the **Browse** button to search for an ARM log if needed. After you have selected an ARM log file, select **Analyze**.



After you have selected **Analyze**, the Analysis Recommendations dialog box opens. A list of recommended aggregations is displayed. Select the aggregation recommendations that you want to add to the **Aggregations** table. Click **Select All** to choose all of the aggregation recommendations. Click **OK** to execute the analysis. The selected aggregation recommendations are added to the list of aggregations in the **Aggregations** table.



In the Aggregation Tuning dialog box, the aggregation recommendations are listed. The **Method** column displays the ARM log status for these aggregation recommendations. Select **Build Aggregations** to build the aggregations for the cube.




Alternately, you can also choose to update only the statistics for the existing aggregations that have entries in the ARM log. The **Update performance values based on the ARM log** option is used when you want to examine the information in the ARM log and verify that the existing aggregations are being used. This can help determine which aggregations to drop. Select **Update** to execute the update.

Note: The label of the **Analyze** button changes to **Update** when you select this radio button. By default, the button is labeled **Analyze**. △



The existing aggregations that are listed in the **Aggregations** table have their **Count**, **Total Time**, and **Average Time** column values updated. The **Method** column also lists **Arm Log** as the update method.



Adding Data to a Cube with Cube Update

After you have built a cube with either a detail table or a star schema, you can update the cube with the Incremental Update function in SAS OLAP Cube Studio. You can choose to update the original cube or generate an updated generation (copy) of the cube. You can then choose to coalesce (consolidate) the aggregation racks (SPDE tables) that have been created during the cube update into a single SPDE table.

Note: For more information on updating SAS OLAP cubes, see Chapter 7, “Updating SAS OLAP Cubes,” on page 79 . Δ

Update a Cube In-Place

The Update In-Place function enables you to update the original cube. In SAS OLAP Cube Studio select a cube from the tree view that you want to update. You can then right-click on the cube and select **Incremental Update** \rightarrow **Update In-Place**. You can also select the function from the **Actions** menu.

On the Update General page of the Incremental Update wizard, select the source table that you will use to update the cube with. If you are updating a star schema cube, the check box **Add data during update** is available. If selected, it activates the **New source table** box. Select a new source table. You can also click on **Select** and choose a table that is not in the drop-down list. If a drill-through table is included in the cube update, click **Select** beside the **New drill-through table** text field and choose a

drill-through table. If the cube was originally built without a drill-through table, the option **(none)** is listed in the **Current drill-through table** text field.


If you are updating a star schema cube, you can also specify how missing dimension keys are handled if encountered. Select one of the three radio button options under **When searching through new members in a table and a missing key is encountered**.

Note: Because you are updating the original cube with the Update In-Place function, the fields for **Cube Name**, **OLAP Schema**, and **Location** are inactive. Δ

The screenshot shows a dialog box titled "Incremental Update - Update In-Place - CUBEDL2". The main heading is "Update General" with the instruction "Select or enter values to use in this incremental cube update." The dialog contains several input fields and buttons:


- Cube name:** CUBEDL2
- OLAP schema:** SASApp - OLAP Schema (with a "New..." button)
- Location:** /Users/Donna/My Folder (with a "Select..." button)
- Current source table:** olapsio.ORDFACT
- New source table:** olapsio.ORDFACT (with a "Select..." button)
- Add data during update
- Current drill-through table:** (none)
- New drill-through table:** (with a "Select..." button)
- When searching for new members in a dimension table and a missing dimension key is encountered:**
 - Print the error and stop processing
 - Print a summary of missing keys and continue processing
 - Print details about missing keys and continue processing
- Advanced** button
- Navigation buttons at the bottom: < Back, Next >, Cancel, Help

You can also select **Advanced** and specify aggregation performance options, any needed SAS code, and a search path for user-written formats to include with the cube update. Select **Next** when you are finished.





On the Dimension Selection page, you can click a dimension in the **Select a dimension** list. You can then select one of the radio buttons under **Select an update process method**. When the cube was originally built, you could also specify whether a dimension could be updated at a later time or not. Any dimensions that cannot be updated are listed in the **Blocked from new member updates** list and are not available for selection.

If you are updating a star schema cube, you can select a dimension table that can be used to update an individual dimension. Each dimension has a default table assigned from the original cube build. Using the left and right arrow keys, you can move tables between the **Available tables** and the **Selected table** lists. You can also view statistics about the table you are selecting with the **Properties**, **View Data**, and **Table Options** buttons. If the table you want to use is not included in the **Available Tables** list, you can register the table with the **Register Table** button. Select **Next** when you are finished.



On the **Summary of Selections** page, review the generated PROC OLAP code that is used to update the cube. You can save the PROC OLAP code to a file by selecting **Export Code**. Select the **Finish** to run the update of the cube.





Generating a New Cube

The Generate New Cube function enables you to create an updated copy of the original cube without modifying the original cube. It is useful for cubes that are in a production environment, as it enables the system administrator to check the cube generation for accuracy before making it available for user queries. In SAS OLAP Cube Studio, select a cube from the tree view that you want to update. You can then right-click the cube and select **Incremental Update** → **Generate New Cube**. You can also select the function from the **Actions** menu. This opens the Update General page of the Incremental Update wizard.

Because you are creating a generation of the original cube, you must enter a cube name that is different from the original cube in the **Cube Name** text box. You can also change the OLAP schema and the folder location for the new cube.

You can then select the source table that you will use to update the cube with. If you are updating a star schema cube, the check box **Add data during update** is available. If selected, it activates the **New source table** box. Select a new source table. You can also click **Select** and choose a table that is not in the drop-down list.

If a drill-through table is included in the cube update, click **Select** beside the **New drill-through table** text field and choose a drill-through table. If the cube was originally built without a drill-through table, the option (**none**) is listed in the **Current drill-through table** text field.

If you are updating a star schema cube, you can also specify how missing dimension keys are handled if encountered. Select one of the three radio button options under **When searching through new members in a table and a missing key is encountered**.

Incremental Update - Generate New Cube - CUBEDL2

Update General
Select or enter values to use in this incremental cube update.

Cube name: CUBEDL21129315

OLAP schema: SASApp - OLAP Schema New...

Location: JUsers/Donna/My Folder Select...

Current source table: olapsto.ORDFACT

New source table: olapsto.ORDFACT Select...

Add data during update

Current drill-through table: (none)

New drill-through table: Select...

When searching for new members in a dimension table and a missing dimension key is encountered:

- Print the error and stop processing
- Print a summary of missing keys and continue processing
- Print details about missing keys and continue processing

Advanced

< Back Next > Cancel Help

You can also select **Advanced** and specify aggregation performance options, any needed SAS code, and a search path for user-written formats to include with the cube update. Select **Next** when you are finished.

Advanced

Global Options | Submit SAS Code | Format Search Path

Default performance option settings for all generated aggregations:

Memory available during creation (megabytes):

Maximum number of threads:

Number to create in parallel:

Physical path to aggregation tables:

Maximum partition size (megabytes):

Store in compressed format:

Create with indexes:


Physical path to indexes:

Segment size (columns):


OK Cancel Help

On the Dimension Selection page, you can click a dimension in the **Select a dimension** list. You can then select one of the radio button options under **Select an update process method**. When the cube was originally built, you could also specify whether a dimension could be updated at a later time or not. Any dimensions that cannot be updated are listed in the **Blocked from new member updates** list and are not available for selection.

If you are updating a star schema cube, you can select a dimension table that can be used to update an individual dimension. Each dimension has a default table assigned from the original cube build. Using the left and right arrow keys, you can move tables between the **Available tables** and the **Selected table** lists. You can also view statistics about the table you are selecting with the **Properties**, **View Data**, and **Table Options** buttons. If the table you want to use is not included in the **Available Tables** list, you can register the table with the **Register Table** button. Select **Next** when you are finished.




On the Summary of Selections page, review the generated PROC OLAP code that is used to update the cube. You can save the PROC OLAP code to a file by selecting **Export Code**. Select **Finish** to run the update of the cube.




Coalesce Incremental Data for a Cube

After a cube has been updated, you can choose to combine the aggregation racks (SPDE tables) that have been created during the cube update into a single SPDE table. In SAS OLAP Cube Studio, select a cube from the tree view. You can then right-click the cube and select **Incremental Update ► Coalesce Incremental Data**. You can also select the function from the **Actions** menu. This opens the Aggregations page of the Coalesce Incremental Data wizard.


On the Aggregations page, you can view the aggregations that will be coalesced from the **Selected aggregations** list. You can also observe the levels of the selected aggregations in the **Select aggregations to preview levels** list. The preview list is display only. Select **Next**.




On the Performance Options page, you can specify query performance settings for the cube aggregations. The **Default** tab enables you to change performance settings for all aggregations of the cube. Enter any needed setting values. Select the **Next** when you are finished.



On the Summary of Selections page, review the generated PROC OLAP code that is used to coalesce the aggregations. You can save the PROC OLAP code to a file by selecting **Export Code**. Select **Finish** to run the update of the cube.



A progress dialog box displays after you select **Finish**.




The SAS log that is generated contains status information for the coalesce process.

```
NOTE: Processing of all cube aggregations is completed.
NOTE: Aggregation "Default" was coalesced. Total records: 113149
NOTE: Cube "OrionStar_Schema_Gen3" was coalesced successfully.
NOTE: PROCEDURE OLAP used (Total process time):
      real time      11.42 seconds
      cpu time       3.21 seconds
```

Adding Calculated Members To A Cube


The Calculated Members function enables you to add new members for the Measures dimension. A calculated member is a definition (for a dimension member) that you create and store with the cube. The calculated member value is generated at a later point, during query time. You can also define a calculated member as a measure. You can access the Calculated Members function from the SAS OLAP Cube Studio **Actions** menu or from the context menu for a cube.

The following display shows the Calculations for Cube page of the Calculated Members function. The **Defined calculated members** list displays all calculated members that are defined in the metadata for the selected cube. On the Calculations for Cube page, you can add new calculated members or modify existing members for the selected cube. Select **Add** to launch the New Member wizard and define a new calculated member.




Creating a Simple Calculation

On the Calculation Type page of the New Member wizard, select the type of calculation that you want to create: a simple calculation, time analysis calculation, or a custom calculation. In this example, a simple calculation is created.



On the Simple Calculations page, select the mathematical calculation that you want to use in the calculated member from the list of radio buttons. Next, select the measures that you want to use with the calculation. In this example, the predicted and actual measures for the cube are used in the formula. It is the difference between these measures that the formula will calculate. Select **Next** when you are finished.



On the General page of the New Member wizard, enter the measure Name. In this example, the calculated measure is named PredAnDiff. Select the format and solve order for the measure. When you are finished, select **Next**.

New Member

General
Enter a name for Simple and Time calculations. For Custom calculations, the name may only be entered on the Custom window. You may also select a format and specify a solve order for the new member.

Calculation type: SIMPLE

Parent dimension: Measures

Parent member:

Name: PredAnDiff

Format: BEST15.

Solve Order: 0

< Back Next > Cancel Help

On the Finish page of the New Member wizard, a display panel shows the metadata that is generated for the new calculated member. Review the metadata. When you are finished, select **Finish** to complete the New Member Wizard.

New Member

The following metadata will be created:

A new calculated member for the cube FullyTestdonna will be defined with the following information:

Formula: [Measures].[predicted]-[Measures].[actual]

Parent Dimension: Measures

Parent Member:


Name: PredAnDiff

Format: BEST15.


Solve Order: 0

< Back Finish Cancel Help

The new calculated member PredAnDiff is now listed in the **Defined calculated members** list.



If needed, you can now make changes to the PredAnDiff calculated member. On the Calculations for Cube page, select the calculated member and click **Next**. This opens the Edit a Calculated Member page. You can then modify the measure name and formula. Select **Verify** to check the MDX code for errors and validate the code against the OLAP server. If you select **Build Formula**, the Build Formula dialog box opens.



Edit a Calculated Member ✕

Parent Dimension:


Parent Member:

Name:

Formula:


```
[Measures].[predicted]-[Measures].[actual],  
FORMAT_STRING="BEST15."
```

In the Build Formula dialog box, you can modify the MDX code with the help of an expression builder. You can also add MDX functions and Data Source elements to the expression. Select **OK** when you are finished editing the calculated member.




Creating a Time Analysis Calculation

You can also create a Time analysis calculated member with the New Member wizard. On the Calculations for Cube dialog box, select **Add** to launch the New Member wizard and define a new calculated member.



On the Calculation Type page of the New Member wizard, select **Time analysis calculation**. When you are finished, select **Next**.




On the Time Calculations page, select a time calculation. Next, select an existing measure from the **Formula** panel. The **Formula** panel will change depending on the time calculation radio button that you select. The **Existing measure** drop-down list is populated with all measures, including calculated measures. The **Time period** drop-down list is populated with members from the Time dimension. When you are finished, select **Next**.


Finish the New Member wizard by completing the **General** and **Finish** pages.

Creating a Custom Calculation


You can also create a custom calculated member with the New Member wizard. On the Calculations for Cube dialog box, select **Add** to launch the New Member wizard and define a new calculated member.




On the Calculation Type page of the New Member wizard, select **Custom calculation**. When you are finished, select **Next**.



On the Custom Calculations page, you can select the parent dimension and parent member for the measure. You can also enter the MDX formula for the measure. Enter the name for the calculated member and select the format and solve order.



If you select **Build Formula**, the Build Formula dialog box opens. In this dialog box you can modify the MDX code with the help of an expression builder. You can also add MDX functions and data source elements to the expression.




On the Custom Calculation page, select **Verify** to check any MDX code that you may have added for errors. This validates the code against the OLAP server. Select **OK** when you are finished editing the calculated member. Finish the New Member wizard by completing the **General** and **Finish** panels.

Setting Member Authorizations On A Dimension

SAS security enables you to set authorization permissions on a cube and the various components of a cube. You can apply member-specific filters to cube data by using MDX expression filters known as permission conditions. Permission conditions limit access to a cube dimension so that only designated portions of the data are visible to a user or group of users. With SAS OLAP data, permission conditions impose only explicit grants of the Read permission and can be specified only on dimension objects. You can add member authorizations to a cube dimension from either SAS OLAP Cube Studio or SAS Management Console. In this example, a retail company has a SAS OLAP cube that contains sales data. They need to apply Read restrictions to members of a dimension that contains data for the method of payment for merchandise.


To select a dimension from within SAS OLAP Cube Studio, select a cube in the tree view and drill down to a dimension. To select a dimension in SAS Management Console:

- 1 Select the **Authorization Manager** ► **By Type** ► **Dimension** and drill-down to a dimension.
- 2 Right-click the dimension and select **Properties**.
- 3 In the dimension's Properties dialog box, select the **Authorization** tab, as shown in the following display. Select (or add) the user or group whose Read access you want to limit. In this example, the **PUBLIC** group is restricted.
- 4 In the **Effective Permissions** list, add an explicit grant of the **Read** permission for that user or group. If the selected user or group does not already have a permission condition defined, the **Add Authorization** button is now enabled.
- 5 Click **Add Authorization** to open the Add Authorization dialog box.




On the Add Authorization dialog box, you can create an MDX expression by specifying members and permissions for the permission condition. You can choose to create a basic expression or an advanced expression for a member. By default, the option **Create a basic MDX expression by selecting one or more members and specifying permissions** is selected when you open the Add Authorization dialog box. In this example, the Terms dimension contains data for different methods of payment. The members AMEX and MC-VISA are restricted from view for the PUBLIC group. Drill down on the dimension in the list and select the AMEX member.


After you have selected the basic expression option and a member, the MDX code for the permission condition will appear in the **Member** text field. You can now select the option **Specify access to this member**. This enables you to apply a permission option to the currently selected member or a member and its descendants. You must select this option to activate one of the grant or deny permissions options. Click the **Deny read permissions to this member and its descendants** option. The MDX code for the permission condition is displayed in the **Expression** text box.



Repeat the process for the MC-VISA member. The MDX code for MC_VISA is appended to the MDX code for AMEX member in the **Expression** text box.




You can also choose to manually create an MDX expression filter. In the Add Authorization dialog box, select the option **Create an advanced expression using the expression builder**. The **Build Formula** button is now active. Click **Build Formula**. The Build Formula dialog box opens.




Note: If you have already specified a basic permission condition that denies access to a member, and you select the **Create an advanced expression using the expression builder** option, a message appears that asks: “Do you want your current expression copied to the expression builder?” See the following display. Select **Yes** or **No**. You can then click **Build Formula**.




In the Build Formula dialog box, you can create an MDX filter and observe the MDX expression as you build it. Use the logical operators to specify multiple clauses in your MDX expression in the **Expression Text** text box. Use the **Functions** tab to add MDX functions to your expression. Use the **Data Sources** tab to browse through the dimensions and hierarchies in your cube and select the members that require access control. You can use the **Add to expression** button to add your selections to the **Expression Text** text box. You can also check the accuracy of the expression you are building by selecting the **Validate Expression** button. When you are finished, select **OK**. You will return to the Add Authorization dialog box. Select **OK** again to save the permission condition and return to the Properties dialog box. △




After you have defined permission conditions for the dimension, the label for the **Add Authorization** button changes to **Edit Authorization**.



If you later select **Edit Authorization** to edit a permission condition, the Build Formula dialog box will open. You can then make any needed changes to the MDX code. Select **OK** to save the permission condition and return to the Properties dialog box.



After you have applied the permission conditions, you can validate the Read restrictions with the View Cube function in SAS OLAP Cube Studio. In SAS OLAP Cube Studio, select the SALES cube and select **View Cube** from the **Actions** menu. You can see in the following display that the members AMEX and MC-VISA are not available for view on the TERMS dimension.




You can also set permission conditions for a dimension by using PROC OLAP and MDX expressions. See the SECURITY_SUBSET option for the PROC OLAP statement and the SAS OLAP Security Totals and Permission Conditions topic in the *SAS OLAP Server: MDX Guide*.

See the topic “Setting Permission Conditions on Cube Dimensions” on page 73 for further information.


Setting Identity Driven Security

It is sometimes necessary to substitute identity values in a permission condition to further refine member-level security. Identity-specific values are dynamically derived according to the user ID with which a client is authenticated. Those values are then used to filter the target data. The identity-specific values are derived from identity-driven properties that are stored in the metadata repository for each user and group. You can set an identity driven authorization using the Member Authorization expression builder.


- 1 Select the **Authorization Manager ► By Type ► Dimension** and drill down to a dimension.
- 2 Right-click the dimension and select **Properties**.
- 3 In the dimension’s Properties dialog box, select the **Authorization** tab, as shown in the following display. Select (or add) the user or group whose Read access you want to limit. In this example, the **PUBLIC** group is restricted.
- 4 In the **Effective Permissions** list, add an explicit grant of the **Read** permission for that user or group. If the selected user or group does not already have a permission condition defined, the **Add Authorization** button is now enabled.
- 5 Click **Add Authorization** to open the Add Authorization dialog box.




In the Add Authorization dialog box, select the option **Create an advanced MDX expression using the expression builder** option. You can then click **Build Formula**. This opens the Build Formula dialog box.




In the Build Formula dialog box, you can create an MDX filter and observe the MDX expression as you build it. Use the logical operators to specify multiple clauses in your MDX expression in the **Expression Text** list. Use the **Functions** tab to add MDX functions to your expression. Use the **Insert** button to add your selections to the **Expression Text** list.



Use the **Data Sources** tab to browse through the dimensions and hierarchies in your cube and select the members that require access control. Use the **Add to Expression** button to add your selections to the **Expression Text** text field. You can also check the accuracy of the expression you are building by selecting the **Validate Expression** button.



To add identity values to the expression, click the **Identity Values** folder on the **Data Sources** tab. Select an identity value from the list. Use the **Add to Expression** button to add your selections to the **Expression Text** text field.



Here is a list of possible identity values:

SAS.ExternalIdentity

This property translates to optional, site-specific values such as Employee ID. Those values are not automatically stored in the metadata repository and need to be loaded and maintained.

SAS.IdentityGroupName

This property resolves to the name of the requesting group identity (for example, Portal Admins Group).

SAS.PersonName

This property resolves to the name of the requesting user identity (for example, SAS Demo User).

SAS.IdentityName

This property returns the name of either the requesting group identity or the requesting user identity, depending on whether the user ID is a group login or a personal login.

SAS.Userid

This property translates to the authenticated user ID, normalized to one of the uppercase formats USERID or USERID@DOMAIN (for example, SASDEMO@LXXXXX).

SAS.IdentityGroups

This property resolves to the names of the groups of which a user is a member.

When you are finished, select **OK**. You will return to the Add Authorization dialog box. Select **OK** again to save the permission condition and return to the Properties dialog box.

See the topics “Cube Security” on page 73 and specifically “Identity-Driven Security” on page 74 for further information.


Viewing a Cube in SAS OLAP Cube Studio

After you have built a cube in SAS OLAP Cube Studio, you can examine the contents of the cube with the View Cube function. The View Cube function is available from the **Actions** menu and from the cube context menu. The View Cube function is available for use only with physically built cubes. In addition, when you select the View Cube function, a connection to a SAS OLAP Server must be made.


When the View Cube function is selected, the View Cube dialog panel is loaded and a multidimensional view of the cube data is displayed in a table. By default, the first dimension (default hierarchy) of the cube is displayed on the row axis and the default measure is displayed on the column axis of the table. The dimensions that are not assigned to a row or column are automatically assigned to the slicer axis. The default member or All member is used to create the slicer filter. You can then drill into the dimensions for the cube and see figures for specific levels of the cube. In the View Cube dialog box you can choose to *add* level selections to the existing cube view, or you can select the check box **Replace current selection** to replace the existing levels in the display.

Adding Level Data to a Cube View


In this example, the option **Replace current selection** is not selected. A retailer is analyzing sales figures for outlets in various countries. In the initial cube view, the column axis displays the default measure for the **Sum of ACTUAL** sales. Displayed on the row axis is the default hierarchy for the **geo** (geography) dimension.




In the following display, the row axis view is filtered on the **products** dimension to the **FURNITURE** level. Furniture-specific sales numbers are added to the existing row axis. The column axis is not modified and is still filtered on the **Sum of ACTUAL** sales.




The following display shows continued filtering on the **FURNITURE** level down to the sales figures for sofas only.



You can also filter on the column values. In this display, the time dimension is filtered down to the sales figures for the month of November. In addition, the products dimension has been filtered on beds for the row axis.




In this display, the measure **Sum of PREDICTED** sales is added to the column axis. You can now compare the actual versus the predicted sales totals.




Replacing Level Data On A Cube View


In this example, the option **Replace current selection** is checked. A retailer is analyzing sales figures for outlets in various countries. In the initial cube view, the column axis displays the default measure for the **Sum of ACTUAL** sales. Displayed on the row axis is the default hierarchy for the **geo** (geography) dimension.



In the following display, the row axis view is filtered on the **products** dimension to the **FURNITURE** level. Because the **Replace current selection** check box is selected, the original, default row selections were replaced. The column axis is not modified and is still filtered on the **Sum of ACTUAL** sales.



You can also filter on the column values. In this display, the time dimension is filtered down to the sales figures for the month of November. Because the **Replace current selection** check box is selected, the original, default column selections were replaced.



Creating a Time Dimension in SAS OLAP Cube Studio


When you create a cube in SAS OLAP Cube Studio, the Cube Designer wizard enables you to create cube dimensions and add the needed levels and hierarchies to those dimensions. For time-specific dimensions, you can select from a list of supplied time hierarchies that will help you build the dimension and auto-populate the levels and some level properties for the cube.

Open the Cube Designer wizard and select the input data for the cube. Select **Add** on the Cube Designer – Dimensions page. This opens the Dimension Designer – General page. When you create a time dimension, you must select the **TIME** dimension type. See the following display.


Enter the remaining information for the dimension and select **Next**. This opens the Dimension Designer- Level page. On the Dimension Designer – Level page, the **Add** button becomes a drop-down list of options. The **Add levels** and **Add supplied time hierarchies** options are now available for selection.

Note: The Add (levels) function is always available for selection regardless of the type of dimension you are creating. When you specify a **TIME** type dimension, the **Add** button is converted to a drop-down list where you can select the Add (level) specific function that you want to use. △


Select the **Add supplied time hierarchies** option. On the Add Supplied page, the **Input column** drop-down list displays the available numeric values that have a date format applied to the column. This will default to the **Date** value if it is found in the data source list of columns. The **Format** and **Informat** values change, depending on the value selected in the **Input column** drop-down list. Select from the list of supplied time hierarchies to create the time levels. See the following display.



Selecting **OK** closes the dialog box and updates the Dimension Designer – Level table with new levels based on your selection. One level is created for each selected time period, even if it appears in multiple supplied hierarchies. You can then define properties such as the sort order and description for the levels that you have selected. See the following display.




The hierarchy or hierarchies that are selected on the Add Supplied page are listed in the **Hierarchies** panel on the Dimension Designer - Hierarchies page. If there is only one hierarchy, as with this example, the hierarchy name is changed to match the dimension name. See the following display.




Synchronizing Column Changes

The Synchronize Column Changes function enables you to synchronize a cube when the input table for an existing cube has encountered a column name change. This function finds the name differences between the cube and its input table and changes the hierarchy level names to match the input table column names. The Synchronize Levels function obtains the cube metadata from the metadata repository and compares the names between the input table and the cube hierarchy. If a discrepancy is found, a new cube file and definition are created with the new level name. The level name of the existing cube is then updated to reflect the new column name. The Synchronize Levels function is available for a cube if the cube physically exists and you have writemetadata permissions for that cube.

In SAS OLAP Cube Studio, select a cube in the tree view and right-click. On the context menu, select **Maintain ► Synchronize Column Changes**.



The Synchronize Column Changes message dialog box opens. Select **OK** to continue with the synchronize function. Select **cancel** to exit the function.



You can also use the SYNCHRONIZE_COLUMNS option if you are synchronizing the column changes with a PROC OLAP statement. See the SYNCHRONIZE_COLUMNS option on the “PROC OLAP Statement” on page 245 for more information.

Specifying an ESRI GIS Map For a Cube Dimension

When you create a cube in SAS OLAP Cube Studio, you can include the connection information for ESRI GIS mapping services. This GIS information can then be read by the SAS OLAP Server and returned during a cube query. In order to access ESRI

mapping data for a cube, you must have defined an ESRI map server in SAS Management Console and have installed the ESRI plug-in to SAS Management Console.

You can add ESRI mapping connections to a cube while you are creating the cube or at a later point when you are editing the cube. The Cube Designer wizard in SAS OLAP Cube Studio enables you to define a GIS specific dimension for a cube and specify the ESRI map server that you connect to. In this example, a dimension named US Geography is defined.


Note: For more information on defining ESRI GIS mapping connections for SAS OLAP cubes, see “Specifying GIS Map Information for a Dimension” on page 41. △

On the Dimension Designer – General page of the Cube Designer wizard, you can enter basic information for the dimension. For the dimension **Type**, select **GEO**. This will identify the dimension as a geography-specific dimension for ESRI mapping. You can have only one GEO-type dimension per cube. When you select **GEO**, the **Specify Map** button is enabled on the Cube Designer – Dimensions page.


The screenshot shows the 'Dimension Designer - General' dialog box. The title bar reads 'Dimension Designer' with a close button. Below the title bar, it says 'Dimension Designer - General' and 'Specify the information for the dimension.' The dialog contains several input fields and controls:

- Name:** A text box containing 'USGeography'.
- Caption:** An empty text box.
- Description:** An empty text box.
- Type:** A dropdown menu set to 'GEO'. To its right is a checked checkbox labeled 'Allow new members during incremental update'.
- Sort order:** A dropdown menu set to 'Ascending Uniform...'. To its right is an 'Advanced...' button.
- Star Schema Table:** A section with a checked checkbox 'Use the fact table' and a 'View Fact Table' button. Below this are four fields: 'Table', 'Key', 'Fact key', and 'Table column', each with a dropdown arrow. To the right of these fields are 'View Table' and 'View Date' buttons.
- Navigation:** At the bottom are four buttons: '< Back', 'Next >', 'Cancel', and 'Help'.


Continue defining the GEO-type dimension. Select the levels for the dimension with the Add Levels dialog box.




Specify any level properties on the Dimension Designer – Levels page.



Create the hierarchy for the dimension on the Cube Designer – Hierarchy page. Click **Add** to select the dimension levels for the hierarchy or select **Finish** to automatically generate a hierarchy that contains all selected levels for the dimension.




After the dimension is created, it will appear in the **Dimensions** list on the Cube Designer – Dimensions page. In addition, the **GIS Maps** button will be activated.



Click **GIS Maps** to open the GIS Maps dialog box. This dialog box enables you to assign ESRI spatial map information to levels of the GEO-type dimension. You can now specify the connection settings for the ESRI map service.


Select from the drop-down lists the **GIS Map Server** and **Map Service** that you need to connect to. You can now map individual levels to the map service. Select a level from the **Levels** list. Select the **Map Layer** and **Map Field ID** of the map service. Then select the corresponding **Field ID Column** for the cube input table.



You can also test the mapping of the Map Service IDs and the Cube Input Table ID. The following message is displayed. Select **OK** to continue with the testing.



The Test GIS Mappings dialog box is displayed when the testing is complete. Select **Close** to return to the GIS Maps dialog box.




When you are finished, select **OK** on the GIS Maps dialog box to return to the Cube Designer – Dimensions page. Finish creating or editing the cube. Afterwards, you can access the GIS Map functionality through the SAS Web OLAP Viewer and SAS Web OLAP Viewer for Java.

For further information see “Specifying GIS Map Information for a Dimension” on page 41.


Creating Multiple Hierarchies For a Cube

When you create a cube using the Cube Designer wizard in SAS OLAP Cube Studio, you can specify more than one hierarchy for a dimension. Hierarchies can be defined on the Dimension Designer wizard that is part of the Cube Designer wizard.


- 1 Create a cube in the Cube Designer wizard, selecting the input data for the cube and creating dimensions for the cube.
- 2 On the Dimension Designer - Hierarchy page, select **Add**. See the following display.




This opens the Dimension Designer – Define a Hierarchy page. See the following display.



- 3 On the Dimension Designer – Define a Hierarchy page, a hierarchy name is created for you by default that is based on the dimension name. Enter a different hierarchy name if needed. Enter a caption and a description if needed. You can then select the levels for the hierarchy from the **Available** list and add them to the **Selected** list. Select **OK** when you are finished. The hierarchy is created and listed on the Dimension Designer - Hierarchy page. See the following display.



- 4 Repeat the Dimension Designer – Define a Hierarchy function for each hierarchy that you need to create. As you create hierarchies they are listed in the **Hierarchies** panel of the Dimension Designer – Hierarchy page. You can then select a default hierarchy from the list of hierarchies that you have created. See the following display:



Note: If you are creating a time-specific dimension, use the Add Supplied dialog box to create the levels and hierarchies for the dimension. See the example “Creating a Time Dimension in SAS OLAP Cube Studio” on page 197. Δ

If you are creating the cube using PROC OLAP, you will include a HIERARCHY statement (see “HIERARCHY Statement” on page 267) for each hierarchy that you create for the dimension. See the following example PROC OLAP code:

```

DIMENSION Customers
  CAPTION          = 'Customers'
  SORT_ORDER      = ASCENDING
  HIERARCHIES     = (
    Customer CustomerUnit
  ) /* HIERARCHIES */;


HIERARCHY Customer
  ALL_MEMBER = 'All Customer'
  CAPTION    = 'Customer'
  LEVELS     = (
    Customer_Age_Group Customer_Gender
    Customer_Group Customer_Type
  ) /* LEVELS */
  DEFAULT;

HIERARCHY CustomerUnit
  ALL_MEMBER = 'All CustomerUnit'
  CAPTION    = 'CustomerUnit'
  LEVELS     = (
    Month_Num Tot_Cost Unit_Cost
  ) /* LEVELS */;


```

Set IGNORE_MISSING_DIMKEYS for a Star Schema

If you are building a cube with a star schema data source, you can define how the cube build process will respond to missing dimension keys. A missing dimension key is detected when the star schema fact table contains foreign key values that are not present in one of the corresponding dimension tables. You can set the option in SAS OLAP Cube Studio when you are creating or editing a star schema cube. On the Cube Designer – General page, select **Advanced**.



This opens the Advanced Cube Options dialog box. Select the **Star Schema Options** tab. You can now select one of the options for handling a missing dimension key if it is encountered.



Here are the options:

- | | |
|---|--|
| Print error and stop processing | Any missing dimension keys stop the build of the cube. This is the default option. |
| Print summary of missing keys and continue processing | The cube build continues and the fact table row with the missing key is ignored (it is not built into the cube). |
| Print details about missing keys and continue processing | The cube build continues and the fact table row with the missing key is ignored. In addition, the log receives additional detail and the missing keys are listed for each dimension table. |


You can also set the missing key option when you are updating a cube. See the example “Generating a New Cube” on page 164. If you are defining a cube in PROC OLAP you can set the option on the “PROC OLAP Statement” on page 245.

Implementing Drill-through to Detail Data in a SAS OLAP Cube


Many SAS OLAP applications give users the ability to select a cell or a range of cells and then view the input data that the cell data was summarized from. SAS OLAP enables you to assign a drill-through table to a cube. You can then access the underlying cell data.

Note: When you select a data table for drill-through, you may need to consider user access and security restrictions for that table. For further information see “Security for Drill-through Tables” on page 76. Δ


In this example a cube is assigned a drill-through table and then accessed in Microsoft Excel 2007. In SAS OLAP Cube Studio, open a cube in the Cube Designer wizard. In the Cube Designer – Drill-Through page, assign a drill-through table. This is often the same table that is used as the input table for the cube. Complete the Cube Designer wizard, saving the drill-through table assignment to the cube. See the following display:




In Microsoft Excel 2007, select **Data** ► **From Other sources** ► **From Data Connection Wizard**. This opens the Data Connection Wizard. Select the option **Other/Advanced**. Select **Next**.




On the Data Link Properties page, select the Provider tab. In the list of data provider options, select the option: **SAS OLAP Data Provider 9.2**. Select **Next**. This takes you to the **Connection** tab.




On the **Connection** tab, enter the **Data Source**, **User name** and **Password** for the SAS OLAP server that you are accessing.




At this point, you can choose to test the connection to the SAS OLAP Server. Select the **Test Connection** button. The following message is displayed if the connection is successful. You can then select **OK** on the Data Link Properties dialog box.




If the connection to the OLAP Server is successful, you can select the cube you want to view drill-through data for. On the Select Database and Table page, select the needed database and a cube. Select **Next**.



On the Save Data Connection File and Finish page, you can enter optional information about the cube so that you can reference and access the cube in a later session. You can enter a **Description**, **Friendly Name**, and **Search Keywords** for the cube. Select **Finish** to complete the Data Connection Wizard.



The Import Data dialog box opens. Select to view the cube data in a **PivotTable Report** or in a **PivotChart and a PivotTable Report**. Select **OK**.



The Data Link Properties open. Enter your password and select **OK**.





The detail data for the cell in the pivot table is displayed on a second sheet in Excel.

	ACTUAL	PRODUCT	COUNTRY	REGION	DIVISION	PROCTYPE	PRODUCT	QUARTER	YEAR	MONTH
4	8		975 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	1	1993	Jan
5	813		455 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	1	1993	Feb
6	779		799 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	1	1993	Mar
7	325		189 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	2	1993	Apr
8	567		147 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	2	1993	May
9	154		533 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	2	1993	Jun
10	247		381 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	3	1993	Jul
11	287		457 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	3	1993	Aug
12	867		441 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	3	1993	Sep
13	307		859 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	4	1993	Oct
14	573		177 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	4	1993	Nov
15	82		633 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	4	1993	Dec
16	208		382 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	1	1994	Jan
17	764		707 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	1	1994	Feb
18	842		568 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	1	1994	Mar
19	818		349 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	2	1994	Apr
20	817		138 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	2	1994	May
21	449		211 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	2	1994	Jun
22	895		254 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	3	1994	Jul
23	750		812 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	3	1994	Aug
24	718		362 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	3	1994	Sep
25	197		113 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	4	1994	Oct
26	278		86 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	4	1994	Nov
27	32		88 U.S.A.	EAST	EDUCATION	FURNITURE	SOFA	4	1994	Dec
28	523		204 U.S.A.	EAST	CONSUMER	FURNITURE	SOFA	1	1993	Jan
29	752		305 U.S.A.	EAST	CONSUMER	FURNITURE	SOFA	1	1993	Feb


Exporting a Cube From SAS OLAP Cube Studio

In SAS OLAP Cube Studio, you can select a cube to export a part of a SAS package. Select a cube in the tree view. Select **File** ► **Export SAS Package**. The Export SAS Package wizard opens.


On the Welcome to the Export SAS Package Wizard page, enter the name of the package you want to export. A default package named **Package1.spk** is automatically entered. You can also check the option **Include dependent objects when retrieving initial collection of objects**. This enables you to automatically include all dependent objects, recursively, in the export. Select **Next**.



On the Select Objects to Export page, check the objects that you want to include in the exported package. The cube job is automatically selected with the cube. Select **Next**.



On the Summary page, review the objects that are included in the exported package. Select **Next**.




On the Export Complete page, you can view the status of the export process. You can also select **View Log** to view the Export Log. Select **Finish** to complete the Export SAS Package wizard.




For detailed information on exporting SAS OLAP cubes, see “Export SAS Package and Import SAS Package” on page 235.

Importing a Cube Into SAS OLAP Cube Studio

In SAS OLAP Cube Studio, you can import a cube that has been previously exported as part of a SAS package. In the tree view, select the **Folders** tab. Select a folder to import a cube to. Select **File ► Import SAS Package**. The Import SAS Package wizard opens. On the Welcome to the Import SAS Package Wizard page, select the package that you want to import. On the **Import Options** panel, select the import options that you need. Select **Next**.




On the Select Objects to Import page, select the objects from the SAS package that you want to import. Select **Next**.




The About Metadata Connections page lists the metadata properties that you need to select values for. This list is determined from the objects that you selected to import. In this example, a SAS Application Server and a directory path must be selected in the following pages of the wizard. Select **Next**.




On the SAS Application Servers page, select the target SAS Application Server. Select **Next**.




On the OLAP Schemas page, select the OLAP schema on the target system that you are assigning the imported cube to. You must select an OLAP schema when importing SAS OLAP cubes. Select **Next**.



On the Directory Paths page, select the directory path, on the target system, that you want to import the cube files to. Select **Next**.




On the Summary page, review the objects that you want to import. Select **Next**.



On the Import Complete page, you can view the status of the import process. You can also select **View Log** to view the Import Log. Select **Finish** to complete the Import SAS Package wizard.



For detailed information on importing SAS OLAP cubes, see “Export SAS Package and Import SAS Package” on page 235.



CHAPTER

9

Using SAS OLAP Cubes

<i>Using a Cube with ADO MD</i>	225
<i>Using a Cube with OLE DB for OLAP</i>	225
<i>Using a Cube with Additional SAS Products</i>	226
<i>SAS Products That Use SAS OLAP Cubes</i>	226
<i>SAS Enterprise Guide</i>	226
<i>SAS Information Map Studio</i>	227
<i>SAS Web Report Studio</i>	227
<i>SAS Web OLAP Viewer for Java</i>	228
<i>SAS Add-In for Microsoft Office</i>	229
<i>Using a Cube with Third-Party Clients</i>	229
<i>Microsoft Excel Pivot Tables and Pivot Charts</i>	229

Using a Cube with ADO MD

Applications gain access to SAS OLAP cubes through ActiveX Data Objects – Multidimensional (ADO MD). ADO MD is an industry-standard programming interface to multidimensional data. It offers the same functionality as OLE DB for OLAP, but in a simpler programming model. Accessing SAS OLAP cubes through ADO MD requires the SAS OLAP Data Provider, which is a component of SAS Integration Technologies. The SAS OLAP Data Provider is installed with the SAS Integration Technologies Client for Windows. See the *SAS Providers for OLE DB: Cookbook* for more information about IOM data provider usage with ADO MD.

Using a Cube with OLE DB for OLAP

In addition to ADO MD, applications gain access to SAS OLAP cubes through *OLE DB for OLAP* (ODBO), an industry-standard set of programmable Component Object Model (COM) interfaces that expose multidimensional data. For SAS OLAP cubes, the OLE DB interfaces are exposed by the SAS OLAP Data Provider, a component of SAS Integration Technologies. The SAS OLAP Data Provider enables applications to perform data analysis by providing a means to view schema information, submit MDX queries, and retrieve results. The SAS OLAP Data Provider is installed with the SAS Integration Technologies Client for Windows. See the *SAS Providers for OLE DB: Cookbook* for more information about IOM data provider usage.

Using a Cube with Additional SAS Products

Listed below are the SAS products that enable you to access, view, query, and report on SAS OLAP cubes. Specific OLAP functionality within these products is also listed.

SAS Products That Use SAS OLAP Cubes

The following SAS products access SAS OLAP cubes:

- “SAS Enterprise Guide” on page 226
- “SAS Information Map Studio” on page 227
- “SAS Web Report Studio” on page 227
- “SAS Web OLAP Viewer for Java” on page 228
- “SAS Add-In for Microsoft Office” on page 229
- SQL Pass-Through Facility for OLAP. For more information, see “Accessing OLAP Cubes from SAS: SQL Pass-Through Facility for OLAP” on page 67.

Note: The SQL Pass-Through Facility for OLAP does not require additional licensing. \triangle

For information on how to access SAS OLAP cubes in these different products, see the individual product Help and documentation.

SAS Enterprise Guide

SAS Enterprise Guide provides a dedicated, intuitive and advanced interface is provided for analyzing business information stored in OLAP data cubes. SAS Enterprise Guide is an ODBO-compliant OLAP Viewer. When accessing OLAP cubes via ODBO, the SAS OLAP Server acts as an Open OLAP Provider, and SAS Enterprise Guide is an Open OLAP Consumer. SAS Enterprise Guide features drillable, interactive graphics and enables users to generate ad hoc queries for analytics based on cube data.

The *OLAP Analyzer* supports all of the functionality required to navigate through multidimensional data, add topic-specific business calculations, and extract information from multidimensional sources for further analysis with advanced statistical procedures or data mining. With the OLAP Analyzer and SAS Enterprise Guide, you can do the following:

- access SAS OLAP cubes directly from the metadata definition. You can preview each cube’s dimensions, measure its contents, and override the default view before accessing the cube.
- slice and dice data as needed to explore the information, as well as drill-through to underlying detailed data.

SAS Enterprise Guide provides the following features for OLAP cube access, visualization and manipulation:

- support for drilling, slicing and pivoting in a cube to explore the cube data. You can also drill-through to the underlying detailed data for a cube. With drill-through capability, you can navigate in your data from the most summarized levels to the most detailed levels. You can drill down on all members of a level or drill down on a specific member of the level. You can also drill in the graph view.
- calculation (calculated member or measure) support, including simple calculations, count analysis, relative contribution analysis and custom calculations (such as time series analysis).

- support for filtering that is based on a ranking function or on a range of values. You can now filter by member caption and member property. You can also create filters that can be used by multiple queries.
- support for totals, subtotals, and percent of totals calculations.
- multiple, independent views of a cube. Specific views on multidimensional information can be saved as bookmarks for easy reuse.
- surfacing of multidimensional information slices to other analytical SAS procedures for advanced analysis, including use in data mining procedures.
- surfacing of SAS OLAP data sources from the SAS OLAP Server or other third-party vendors supporting OLE DB for OLAP (for example, SAP BW and Microsoft Analysis Services).

For further information on SAS Enterprise Guide, see *Getting Started with SAS Enterprise Guide* and the SAS Enterprise Guide Help.

SAS Information Map Studio

SAS Information Maps can translate business questions into the necessary MDX code to access SAS OLAP structures. SAS Information Map Studio is a Java application that enables data modelers and data architects to create and manage SAS Information Maps, which are business metadata about your physical data. Information maps enable you to surface your warehouse data in business terms that typical business users understand, while storing key information that is needed to build appropriate queries.

Information maps can be built from various data structures, including star schemas, snowflake schemas, normalized data structures, or OLAP cubes, providing business users with easy access to enterprise data. The SAS OLAP Server is required to create information maps on top of multidimensional data sources.

For more information about SAS Information Map Studio, see *SAS Information Map Studio: Getting Started with SAS Information Maps*.

SAS Web Report Studio

SAS Web Report Studio is a Web-based application that enables you to create, view, and organize reports. You can use SAS Web Report Studio to perform these OLAP-specific tasks:

- drill and expand tables and graphs
- support ragged and unbalanced hierarchies
- pivot individual crosstabulation dimensions
- switch dimensions and measures with the Data Selection dialog box
- synchronize report components to display a common drill state or have them remain independent

In addition to OLAP-specific tasks, SAS Web Report Studio enables you to perform the following tasks:

Create reports	Beginning with a simple and intuitive view of your data provided by SAS Information Maps (created in SAS Information Map Studio), you can create reports based on either relational or multidimensional data sources.
View and manipulate reports	While viewing reports using a thin client (a Web browser), you can filter, sort, and rank the data that is shown in tables, crosstabulations, and graphs. With multidimensional data, you can

drill down on data in crosstabulations and graphs and drill-through to the underlying data.

For more information about using SAS Web Report Studio, see *SAS Web Report Studio: User's Guide*.

SAS Web OLAP Viewer for Java

The SAS Web OLAP Viewer is a Web application that provides a Web interface for viewing and exploring OLAP data. It enables business users to look at data from multiple angles, drill-through to detail data, and add linked graphs or maps to gain greater insights into trends, exceptions and opportunities.

The SAS Web OLAP Viewer surfaces SAS OLAP cubes, OLAP information maps, and OLAP-based data explorations. It provides an easy-to-use interface from which you can select a data source, view the data, and customize your view with features such as sorting and filtering. With the SAS Web OLAP Viewer, you can perform these tasks:

Explore OLAP data

You can select an OLAP data source and explore multidimensional data using a table or graph view. Specifically you can do the following:

- explore multidimensional data using a table or graph view (color-coded table, bar chart, bar-line chart, line plot, pie chart, scatter plot, tile chart, map)
- change business queries by selecting business items to be displayed from a side bar
- slice through multidimensional data by using a special slicer dimension
- dice multidimensional data by applying filters on any level of a hierarchy
- drill up/down through hierarchies or expand/collapse entire levels
- explore data following ragged or unbalanced hierarchies that model true hierarchies as they are experienced in business
- obtain detailed information for every cell (drill-through) with the option to export it to Microsoft Excel
- display and filter on properties associated with members inside the OLAP data source (member properties)

Analyze OLAP data

Multidimensional views can be adapted to business scenarios in many different ways, including applying filters and rankings, calculating new measures, sorting, and adding totals and subtotals. Specifically, you can do the following:

- rank multidimensional data to identify top performers or losers (ties can be handled)
- display totals and subtotals on cross tables
- add the percentage that a measure contributes to the total or a subtotal of a table with a single click on a dedicated button
- calculate new measures and add them to any view

Interactively explore geographical maps.

Maps from an ESRI ArcGIS Server can be used to display OLAP data just like other views of the data. Users can drill on map regions visualizing information from an OLAP data source in real time. Specifically, you can do the following:

- use maps from ESRI's ArcGIS Server to display OLAP data just like any other view on the data
- perform synchronized drill and display for map and table view

- drill on regions in the map visualizing information from an OLAP data source in real time, enabling a zoom down to the level of individual houses on a road
- drive the color coding of maps by data from the OLAP data source

Easily save and publish results

Views that reveal new business insights can be easily saved and shared among colleagues, saved as SAS Web Report Studio reports for further formatting, or saved as Microsoft Excel spreadsheets or Adobe PDF documents.

For more information, see the SAS Web OLAP Viewer for Java Help.

SAS Add-In for Microsoft Office

The SAS Add-In for Microsoft Office is a Component Object Model (COM) add-in that extends Microsoft Office by enabling you to harness the power of SAS analytics and access data directly from Microsoft Excel, Microsoft Word, and Microsoft PowerPoint. Specifically, you can do the following:

- access and view SAS OLAP Cubes or any data source that is available from your SAS server. There is no size limit on the SAS data sources that you can open.
- filter your data using an intuitive user interface or using an advanced SQL editor.
- sort your data by an unlimited number of variables.
- refresh your data to incorporate any changes that were made to a data source that is saved on a server.

For more information, see the *SAS Add-In for Microsoft Office: Getting Started with Data Analysis* and the SAS Add-In for Microsoft Office Help.

Using a Cube with Third-Party Clients


The SAS OLAP Server exposes multidimensional data through OLE DB for OLAP interfaces. Supporting these industry-standard interfaces enables the SAS OLAP Server to integrate with third-party clients. The user interfaces for these clients vary widely.

Microsoft Excel Pivot Tables and Pivot Charts

Microsoft Excel provides the ability to view and manipulate SAS OLAP cubes in Excel pivot tables and Excel pivot charts. You can connect to a SAS Metadata Server and access a SAS OLAP Cube in Excel. Once a cube is accessed, you can select which measures and dimension to place on the pivot table or pivot chart. Specifically, you can perform the following OLAP cube-related tasks:

- assign measures and dimensions to the columns and rows of a pivot table or pivot chart
- drill-through to the cube's detail data and save that data into a separate sheet
- generate pivot charts based on the cube data already selected in the pivot table
- perform various data analysis tasks such as filtering, sorting and subtotaling

For further information on working with SAS OLAP Cubes in Microsoft Excel, see the cube building example “Implementing Drill-through to Detail Data in a SAS OLAP Cube” on page 210.



CHAPTER

10

Importing and Exporting SAS OLAP Cubes

<i>Importing and Exporting SAS OLAP Cubes</i>	231
<i>Determining Which Tool to Use</i>	232
<i>ExportCubes and ImportCubes Batch Tools</i>	232
<i>ExportCubes Batch Tool</i>	232
<i>Overview</i>	232
<i>Syntax</i>	233
<i>ImportCubes Batch Tool</i>	233
<i>Overview</i>	233
<i>Cube Connection Points</i>	234
<i>Syntax</i>	234
<i>Export SAS Package and Import SAS Package</i>	235
<i>SAS Packages - Copying the Cube Metadata Registration for a SAS OLAP Cube</i>	235
<i>Export SAS Package</i>	236
<i>Import SAS Package</i>	236
<i>User Privilege and Permission Considerations</i>	237
<i>Creating Connection Points</i>	237
<i>File Naming Considerations</i>	238
<i>Multi-Language Cubes</i>	238
<i>Manually Copying Physical Files for a SAS OLAP Cube</i>	238
<i>SAS OLAP Cube Files</i>	238
<i>Creating Cube Directories</i>	239
<i>Manually Copying the Header Files for Cube Generations</i>	239
<i>Manually Copying the MOLAP Aggregation Tables</i>	239
<i>Manually Copying Rolap Files</i>	240
<i>Validating Data After It Is Moved</i>	240
<i>Cube Promotion and Migration Resources</i>	240

Importing and Exporting SAS OLAP Cubes

Companies that create and use SAS OLAP cubes often have more than one environment that they store cubes on. A company might have an environment for developing and building cubes, another environment for testing cubes, and finally a production environment where cubes are accessed and queried by end users. If you have multiple environments for SAS OLAP cubes, you need the ability to copy cubes and their supporting files and objects from one environment to another.

Determining Which Tool to Use

SAS 9.2 offers you different options to move your data from SAS 9.1.3 SP4. When you choose to copy or move SAS OLAP cubes between SAS environments, consider the following options for moving your cubes:

“ExportCubes and ImportCubes Batch Tools” on page 232	Cubes that were built with a version of SAS that is earlier than 9.2 can be moved to the SAS 9.2 environment by using the ExportCubes and ImportCubes batch tools.
“Export SAS Package and Import SAS Package” on page 235	Cubes built with SAS OLAP Cube Studio 4.2 can be copied or moved between SAS 9.2 systems with the Export SAS Package and Import SAS Package functions that are part of SAS OLAP Cube Studio 4.2.

ExportCubes and ImportCubes Batch Tools

The ExportCubes and ImportCubes batch tools enable you to move your SAS 9.1.3 SP4 cubes to the second maintenance release after SAS 9.2. To use the batch tools, you must go to the DOS command line on your operating system.

ExportCubes Batch Tool

Overview

The ExportCubes batch tool is a command-line tool that enables you to export SAS 9.1.3 OLAP cube metadata. This tool exports the metadata for your cubes into individual XML files that are stored in a designated directory. The ExportCubes batch tool exports all cubes on the metadata server. It is available for use on Windows and is executed from a DOS command line.

When exporting cubes, you designate the output directory that you want to store the XML files to and whether you want to generate a log file of the export process. You can also designate whether security should be exported with the cube metadata or not. When you export cube metadata, the name of the XML file is generated from a combination of the name of the cube’s OLAP schema (if there is one) and the name of the cube. If the OLAP schema or cube name contains characters that are not allowed for a filename on that particular operating system, then those characters are changed according to the native Java interface for that operating system.

When you export your cube metadata, the output directory is populated with an XML file for each extracted cube. All cubes in the Foundation and Custom repositories are exported. However, cubes in Project repositories are not exported and are ignored. A log file of the export execution is created in the output directory.

Note: The ExportCubes batch tool is supported in a hot fix that is applied to SAS 9.1.3 SP4. Cube metadata that is exported with the ExportCubes tool can then be imported in the second maintenance release after SAS 9.2 with the ImportCubes batch tool. In addition, cube metadata that is created with a version of SAS that is later than SAS 9.1.3 SP4 cannot be exported with the ExportCubes batch tool. If attempted, an error message is printed to the log. Δ

Syntax

To use the ExportCubes batch tool, go to the DOS command line on your operating system. The batch tool can be found in the following directory path:

SAS-Installation-directory\sasmanagementconsole\9.1\commands

You can enter the ExportCubes command with the needed options.

ExportCubes - <option(s)>

The following example shows the ExportCubes command:

```
ExportCubes -user "x2345\sasxxx" -password "Passwordxxx" -host "X2345.xy.abc.com"
-port 8561 -extractSecurity -outputDirectory c:\myextractedcubes
```

Option Name	Description
-extractSecurity	Specifies to extract security metadata for the cube. If not specified, security metadata is not extracted.
-host	Specifies the server to connect to. The metadata connection can be specified by a profile or the server, port, user, and password options.
-outputDirectory	Specifies the output directory. The cube metadata and log is written to this directory. If the directory exists, then a number (starting with 1) is appended to the end of the directory name until the name is unique. By default, the log file is named ExportCube.log.
-password "pw"	Specifies the password for the user. The metadata connection can be specified by a profile or the server, port, user, and password options.
-port number	Specifies the port for the server. The metadata connection can be specified by a profile or the server, port, user, and password options.
-profile filename	Specifies the workspace file to use for metadata connection information. The metadata connection can be defined by specifying the server, port, user, and password options or by specifying a metadata connection profile. This option can be used in place of the server, port, user, and password options.
-user "username"	Specifies the user to connect to the server with. The metadata connection can be specified by a profile or the server, port, user, and password options.

ImportCubes Batch Tool

Overview

The ImportCubes batch tool is a command-line tool that enables you to import SAS 9.1.3 OLAP cube metadata files that have been exported using the ExportCubes batch tool on your SAS 9.1.3 SP4 environment. The metadata for a cube is exported as an XML file. With the ImportCubes batch tool you can specify the location of multiple exported SAS 9.1.3 XML files and import them to your SAS 9.2 OLAP environment.

The SAS 9.1.3 SP4 cube metadata is imported and added to the SAS 9.2 Metadata Server. In addition, you can specify that a log file of the import execution be created. After the metadata files are imported, you can then rebuild your SAS 9.1.3 SP4 cubes on your SAS 9.2 system. You can use SAS OLAP Cube Studio to build the physical data for the cubes.

The ImportCubes batch tool is included with the standard SAS OLAP installation and is available for use on Windows. It is executed from a DOS command line. When you execute the ImportCubes command, all the cube XML extract files in a directory are imported. The cubes are imported into the repository from which they were exported. If that repository does not exist on the target metadata server, then the cubes are imported into the Foundation repository and this is noted in the log.

If needed, you can substitute file paths during the import of your cube metadata. To substitute file paths during the import, each XML file where you want to substitute file paths must be edited. To change file paths for a cube, you must edit the Substitutions section of the XML file. Specifically, you must change the string following the tag "DefaultValue=" to the directory path that you want. Do this for each XML file that you want to change, and then run the ImportCubes batch tool. The following example shows the Substitutions section of an XML file.

```
<Substitutions>
<Property Id="A5G2Y32Q.AI000EOC" TLObjN="1" Name="INDEXPATH"
  DefaultValue="c:\indexfilesAggrA"/>
<Property Id="A5G2Y32Q.AI000EOD" TLObjN="1" Name="DATAPATH"
  DefaultValue="c:\partitionsAggrA"/>
<Directory Id="A5G2Y32Q.B00000RU" TLObjN="1" Name="c:\v9cubes"
  DirectoryName="c:\v9cubes"/>
<Property Id="A5G2Y32Q.AI000EO9" TLObjN="1" Name="WORKPATH"
  DefaultValue="c:\v9cubes"/>
<Property Id="A5G2Y32Q.AI000EOA" TLObjN="1" Name="INDEXPATH"
  DefaultValue="c:\globalindex1"/>
<Property Id="A5G2Y32Q.AI000EOB" TLObjN="1" Name="DATAPATH"
  DefaultValue="c:\globalpartitions1"/>
</Substitutions>
```

Cube Connection Points

For each cube, the connection points that the ImportCubes batch tool expects to find must exist in the target repository. If a table, column, UniqueKey, or OLAP schema connection cannot be found in the target metadata server, the import of the cube fails and an error message is written to the ImportCubes log file. For the OLAP schema connection of a cube, if the OLAP schema does not exist on the target metadata server, then the cube is imported and not associated with any OLAP schema.

Note: The ImportCubes batch tool is supported in the second maintenance release of SAS 9.2. The ImportCubes batch tool is only compatible with cubes that were exported with the corresponding ExportCubes batch tool on your SAS 9.1.3 SP4 environment. The ImportCubes batch tool is not compatible with cubes that were exported from SAS 9.1.3 SP4 using the SAS OLAP Cube Studio Export Cube function. Δ

Syntax

To use the ImportCubes batch tool, go to the DOS command line on your operating system. The batch tool can be found in the following directory path:

SAS-Installation-directory\SASPlatformObjectFramework\9.2

You can enter the ImportCubes command with the needed options.

ImportCubes - <option(s)>

The following example shows the ImportCubes command:

```
ImportCubes -user "sasadm" -password "Password01" -host "X7891.xy.abc.com"
-port 8561 -inputDirectory c:\myextractedcubes
```

Option Name	Description
-host	Specifies the server to connect to. The metadata connection can be specified by a profile or the server, port, user, and password options.
-inputDirectory	Specifies the input directory that contains the XML cube extract files from the 9.1.3 ExportCubes batch tool. This option is required.
-log directory	Specifies the directory where the log file named ImportCubes.log is created. If this option is not specified, the log file is created in the user's home directory.
-password "pw"	Specifies the password for the user. The metadata connection can be specified by a profile or the server, port, user, and password options.
-port number	Specifies the port for the server. The metadata connection can be specified by a profile or the server, port, user, and password options.
-profile filename	Specifies the workspace file to use for metadata connection information. The metadata connection can be defined by specifying the server, port, user, and password options or by specifying a metadata connection profile. This option can be used in place of the server, port, user, and password options.
-user "username"	Specifies the user to connect to the server with. The metadata connection can be specified by a profile or the server, port, user, and password options.

Export SAS Package and Import SAS Package

A SAS OLAP cube contains a physical files component and a metadata component. Both of these components must be addressed and kept in-synch when copying or moving a cube. When copying or moving SAS 9.2 version cubes between SAS 9.2 systems, the physical files that compose a cube are copied with standard operating system functions and line commands. The cube metadata is then copied with Export SAS Package and Import SAS Package functions that are found in SAS OLAP Cube Studio and other SAS Intelligence Platform products.

SAS Packages - Copying the Cube Metadata Registration for a SAS OLAP Cube

Part of copying a cube includes copying the metadata registration for the cube. In SAS 9.2, OLAP cubes and their supporting objects can be exported and imported as a group in a SAS package (SPK) file. You can use the Export SAS Package function to

create an (SPK) file. This package file can then be imported into another system with the Import SAS Package function. The Export SAS Package and Import SAS Package functions are part of the SAS Intelligence Platform promotion tools. These tools enable you to promote individual metadata objects or groups of objects from one metadata server to another, or from one location to another on the same metadata server.

Export SAS Package

In SAS OLAP Cube Studio you can create a SAS package to export to another system. In the tree view, determine the object or objects that you want to export. Use the CTRL key to select multiple objects. These objects can include cubes, jobs, libraries, tables, folders, or OLAP schemas. You can also select multiples of the same object. For example, you can export two cubes in the same package.

Note: A cube must be exported with its corresponding cube job. △

After you select the needed objects, select **Export SAS Package** from the **File** menu or from the cube's context menu. This opens the Export SAS Package wizard. To export a SAS package, perform the following tasks:

- 1 On the Welcome page of the wizard, specify the name of the SAS package that you are exporting cube objects to. You can also select the option **Include dependent objects when retrieving initial collection of objects**. This signals the wizard to automatically include any dependent objects for the item(s) you initially selected in the tree view.
- 2 On the Select Objects to Export page of the wizard, select the objects to include in the exported SAS package.
- 3 On the Summary page of the wizard, confirm the objects to include in the SAS package. Select **Next**. The SAS package is created.
- 4 On the Export Complete page, you can view the export log and verify that the package was created.

Be aware that SAS OLAP cubes are dependent on a number of other objects. These include, in particular, cube jobs, tables, libraries, and OLAP schemas. For a successful import, all these objects need to be available on the target system. You must know whether these objects are already available in the target system. However, you can export the objects now and verify their availability later, when you import the package. If you are selecting a cube to export, always select the cube together with its job. The export does not work if the job is not present.

Import SAS Package

After you have exported the SAS package, you can import the package on the target system. Switch to your target system and make sure that the exported SAS package file and the copied cube files are available there.

At this point, you can open SAS OLAP Cube Studio and select a folder in the tree view. This selected folder is where you are importing the package to. If the OLAP schema that you need is on the target system, you can import your cube and related objects into the folder that contains your OLAP schema.

Note: The Import SAS Package function is available only when you select a folder in the tree view. △

Select **Import SAS Package** from the **File** menu or from the cube's context menu. The Import SAS Package wizard opens. To import a SAS package, perform the following tasks:

- 1 On the Welcome page of the wizard, specify the name of the SAS package that you are importing. On this page you can select to include the objects' access control templates. In addition, you can select to import all the objects in the package or only those that are new and do not exist in the destination folder.
- 2 On the Select Objects to Import page, you can select the objects from the package that you want to import. When you click an object in the list of objects to import, any available options for that object are available on the **Options** tab.

An important option to verify when importing a SAS OLAP cube is the cube option **The physical cube will need to be built after the import process**. This indicates that the physical cube files do exist on the target system. By default, this option is automatically selected. If you select this option, deselect any objects that are already available on the target system. Also, deselect any objects that you don't want to create in the target folder. The physical cube will be built after the cube is imported into the target system.

Note: A cube must be imported with its corresponding cube job. Δ

- 3 After you have selected the objects to import and build, the wizard identifies those objects that you must establish metadata definitions or connection points for. The About Metadata Connections page lists the metadata objects that you must define. This list is determined by the objects that you selected on the previous page. The Import SAS Package wizard enables you to specify values for the target location that correspond to values from the source location. Depending on what you import with the cube, this can include the following objects:
 - OLAP schema (This can have the same name as the original OLAP schema, but reside in a different folder.)
 - tables (These include fact tables, dimension tables, drill-through tables, or aggregation tables.)
 - libraries
 - SAS Application Server
 - directory paths (These include the cube path and, optionally, data paths, index paths, and library paths.)
- 4 Next, complete the wizard page for each metadata object identified on the About Metadata Connections page. When you have finished entering the connection point information, you can view the Summary page of the wizard. This page lists the objects that are being created and the import objects that are being mapped on the target system.

User Privilege and Permission Considerations

It is recommended that the person performing the cube export or import be the AdminUser and preferably the unrestricted user. This is because other users might have restrictions on parts of the cube that could result in a partial cube being exported and imported. It is also important that the identities, groups, and permissions for the imported data be set to correspond to the target environment security settings. Security settings can be created using the **Authorization** tab on the cube Properties dialog box.

Creating Connection Points

Certain metadata must exist on the target system prior to importing a cube. All metadata connection points for the cube must exist with the same names. If a metadata

connection object is not found, then the import fails. The following metadata connection points must exist:

Object	Connection Points	Details
Cube	Job(cube), OLAP schema, tables, directory path object	A cube must be imported with its corresponding job(cube).
Job(cube)	Cube	A cube must be imported with its corresponding job(cube).
OLAP Schema	SAS Application Server and its logical OLAP server	Only one OLAP schema can be associated with a SAS Application Server.
Library	SAS Application Server and directory path object	The target application server that you select determines the SAS libraries that are available.
Tables	Library	If there is a matching SAS library on the target system, then that SAS library is selected by default. This occurs if the target system library has the same library name and SAS Application Server name as the original SAS library.

File Naming Considerations

When importing cube metadata, if you have not selected an import file before, then the selected file shown by default is *user-default-directory\extractedcube.xml*. If you have selected an import file before, then the last selected file is displayed by default.

Multi-Language Cubes

You can export and import multi-language cubes. However, only the dimension tables for the server language (the first language in the UDT statement) are verified for registration. You must ensure that all the tables are present. If one of the tables is missing when the cube is imported, the import is still successful, but the cube might not rebuild correctly.

Manually Copying Physical Files for a SAS OLAP Cube

SAS OLAP Cube Files

When you copy a SAS OLAP cube between environments, you can choose to copy the physical files that are part of the cube. To do this, copy or move all of the files from the source cube subdirectory to the target cube subdirectory. You can use standard operating system functions to copy or move the files. It is important to check and verify that the target files have the same operating system permissions as your original files. In addition, when you are copying files between environments, note that you can copy SAS OLAP cube files only among systems that have the same data representation. You

cannot move between 32-bit and 64-bit systems, and you cannot move between Windows and UNIX. Manually copying SAS OLAP cube files involves the following components:

- cube directories
- cube header files
- MOLAP aggregation tables
- ROLAP data tables

Creating Cube Directories

Before you copy the files for a cube, you must create a directory for the cube on the target system. The path that you create the directory in is known as the *root* path for the cube and should have the same name as the cube. The cube's path is specified with either of the following methods:

In SAS OLAP Cube Studio When creating or editing a cube in the Cube Designer wizard, enter a file path in the **Path** field on the Cube Designer – General page.

Using PROC OLAP Include the **PATH=** option in the PROC OLAP statement.

For example, if you create the cube OrionStar in the cube path **myserver\testolap\testcube**, then the cube root path is **myserver\testolap\testcubes\OrionStar**.

When you copy the OrionStar cube to your target cube path, create a subdirectory with the cube's name. For example, if you want to copy or move the cube to **otherserver\prodolap\prodcubes**, you would create the directory **otherserver\prodolap\prodcubes\OrionStar**.

If the cube has aggregations or indexes stored in other directories, then these files must also be copied or moved to a subdirectory with the same name. You can use the **DATAPATH=** and **INDEXPATH=** options to specify what directories the aggregations are stored in.

Manually Copying the Header Files for Cube Generations

The header files for each generation of the cube are located in subdirectories in the cube root path. They are named **gen0000**, **gen0001**, and so on. For example, if the cube OrionStar is created in the cube path **myserver\testolap\testcubes**, then the cube header files are stored in the path **myserver\testolap\testcubes\OrionStar\gen0000**. If you have made incremental updates to the cube, you might have more than one **genNNNN** directory. Or your **genNNNN** directory might have a numeric suffix other than 0000. Copy or move all the **genNNNN** subdirectories with their contents to the target cube root directory.

Manually Copying the MOLAP Aggregation Tables

SAS OLAP cubes store aggregated data values in tables that have the same file structure as SAS SPD Engine tables. These tables consist of a number of different files and can be broken into partitions. By default, all the component files and partitions are stored in the cube root directory. However, you might have chosen to distribute the MOLAP aggregation storage over multiple locations. You can specify distributed file storage locations with either of the following methods:

In SAS OLAP Cube Studio

When creating or editing a cube in the Cube Designer wizard, you can define the storage location for one or more aggregations. On the Aggregations page of the

wizard, select the **Advanced** button to open the Performance Options dialog box. From here you can define the storage location for all aggregations or a single aggregation. On either the **Default** tab or the **Aggregations** tab, enter the file path in either of the following fields:

- Physical path to indexes**
- Physical path to aggregation tables**

Using PROC OLAP

Use the DATAPATH= and INDEXPATH= options in the PROC OLAP statement or in the individual AGGREGATION statements.

Copy your MOLAP aggregation tables by using operating system utilities into the target directories. Make sure to copy the files from all locations, if you have distributed your cube aggregation storage over multiple directories.

Manually Copying Rolap Files

ROLAP tables are used as data sources for SAS OLAP cubes and can be stored in many formats. These formats include SAS tables, SPD Engine tables, SPD Server tables, and tables in external RDBM systems. When you are building SAS OLAP cubes, ROLAP tables can be used as the following data sources:

- drill-through tables
- aggregation tables
- input data (if you used the NO_NWAY option)
- format catalogs (used by any of the above)

You can use PROC COPY to copy or move SAS tables, SPD Engine tables, and SPD Server tables. To copy or move tables in external RDBM systems, use standard operating system functions. When you copy or move any of these files, you should verify that your target SAS OLAP Server has access to those files. In addition, you need to change the LIBNAME specifications to point to the new file locations. Librefs for an OLAP server are allocated by using an AUTOEXEC file during the SAS OLAP Server invocation or by using preassigned libraries.

Validating Data After It Is Moved


After migration or promotion of your data is complete, the cubes must be rebuilt in SAS 9.2. In addition, other tasks must be completed for SAS OLAP Server and SAS OLAP Cube Studio. The following tasks must be performed:

- In SAS Management Console, ensure that all SAS OLAP Servers and OLAP schemas are available.
- In the Server Manager, validate the connections to the SAS OLAP Servers.
- Use SAS OLAP Cube Studio to build a test cube.
- Use the Cube Viewer in SAS OLAP Cube Studio to validate the structure of the test cube.
- Build cubes and validate their structure on all other SAS OLAP servers.

Cube Promotion and Migration Resources

In addition to the export and import cube functions discussed in this chapter, see the following resources for information about promoting and migrating your SAS data:

- “Promoting OLAP Cubes from SAS 9.1.3” in the *SAS Intelligence Platform: 9.1.3 to 9.2 Migration Guide*
- “Using the Promotion Tools” in the *SAS Intelligence Platform: Administration Guide*
- The SAS Migration Utility is used to upgrade the SAS content and configuration from an earlier SAS release to a later SAS release. For detailed information on the SAS Migration Utility, see the *SAS Intelligence Platform: 9.1.3 to 9.2 Migration Guide*



APPENDIX

1

The OLAP Procedure

<i>The OLAP Procedure</i>	244
<i>Syntax: OLAP Procedure</i>	244
<i>PROC OLAP Statement</i>	245
<i>Options</i>	245
<i>METASVR Statement</i>	256
<i>Required Argument</i>	256
<i>Options</i>	257
<i>DIMENSION Statement</i>	257
<i>Required Arguments</i>	258
<i>Options</i>	258
<i>LEVEL Statement</i>	263
<i>Required Arguments</i>	263
<i>Options</i>	263
<i>PROPERTY Statement</i>	266
<i>Required Arguments</i>	266
<i>Options</i>	266
<i>HIERARCHY Statement</i>	267
<i>Required Arguments</i>	268
<i>Options</i>	268
<i>MEASURE Statement</i>	269
<i>Required Arguments</i>	269
<i>Options</i>	272
<i>AGGREGATION Statement</i>	274
<i>Required Arguments</i>	274
<i>Options</i>	274
<i>DROP_AGGREGATION Statement</i>	276
<i>Required Arguments</i>	276
<i>DEFINE Statement</i>	276
<i>Required Arguments</i>	278
<i>UNDEFINE Statement</i>	278
<i>Required Arguments</i>	279
<i>USER_DEFINED_TRANSLATIONS Statement</i>	279
<i>Required Argument</i>	280
<i>SAS Servers and Character Encoding</i>	281
<i>REORGANIZE_LEVEL Statement</i>	281
<i>Loading Cubes</i>	281
<i>Loading Cubes from a Detail Table</i>	281
<i>Loading Cubes from a Star Schema</i>	283
<i>Loading Cubes Using Summarized Data</i>	285
<i>Maintaining Cubes</i>	286
<i>Building a Cube from an Existing Definition</i>	286

Adding Aggregations to an Existing Cube	287
Deleting Aggregations from an Existing Cube	287
Deleting Cubes	288
Specialized Options for PROC OLAP	288
Options for Managing Ragged Hierarchies	288
Options Used for Performance	289

The OLAP Procedure

The OLAP procedure is one of the SAS tools that you can use to create, update, and delete cubes. This includes adding and deleting cube aggregations.

Note: You can also use the Cube Designer wizard to maintain OLAP cubes. The Cube Designer wizard can be launched from SAS Data Integration Studio and SAS OLAP Cube Studio. Help on using the wizard to build cubes is available from within both applications.

Δ

In addition to the basic cube creation tasks, PROC OLAP also enables you to do the following:

- build cubes with ragged hierarchies
- rename cubes
- change nonstructural cube elements
- control resummarization behavior on cube security
- control options that can be used to optimize cube creation and query performance
- specify data set options on detail, fact, dimension, and drill-through tables
- create TIME dimensions
- design dimensions that have more than one hierarchy
- define global calculated members and named sets
- include SAS code when you submit PROC OLAP in batch mode
- read alternate locale data sets and create locale-specific metadata for use at query time

Syntax: OLAP Procedure

```

PROC OLAP <option(s)>;
METASVR OLAP_SCHEMA='schema-name' <option(s)>;
DIMENSION dim-name HIERARCHIES=(hier-nam ... hier-nameN) <option(s)>;
LEVEL level-name <option(s)>;
PROPERTY prop-name LEVEL=level-name <option(s)>;
HIERARCHY hier-name LEVELS=(level-name1 <level-name2 ...level-nameN>
    <option(s)>);
MEASURE measure-name STAT=statname <option(s)>;
AGGREGATION level-name <level-name2 level-name3 ...level-nameN> / <option(s)>;

```

```

DROP_AGGREGATION level-name1 < level-name2 ...level-nameN> /
    NAME=aggregation-name ;
DEFINE MEMBER | SET 'member-or-set-name' AS 'mdx-expression' ;
UNDEFINE MEMBER | SET 'member-or-set-name' ;
USER_DEFINED_TRANSLATIONS locale < locale2 ...localeN> ;
REORGANIZE_LEVEL | REORG_LEVEL

```

PROC OLAP Statement

The PROC OLAP statement specifies the input data source, cube name, and path. This statement can also be used to do the following:

- specify options that might improve query performance
- delete cubes
- specify global settings for handling missing hierarchy members in ragged and unbalanced hierarchies
- perform actions on existing cubes
- change nonstructural cube elements
- update cube data
- coalesce partition aggregations

```
PROC OLAP <option(s)>;
```

Options

Note: For a list of the options that can be used to optimize cube creation and query performance, see “Options Used for Performance” on page 289. △

ADD_DATA

specifies an incremental update to a cube. The ADD_DATA option enables you to add new members and data to a cube. You can update a cube in-place or create generations of the cube. The new data that the cube is updated with is specified with the DATA= option. Here is an example of the ADD_DATA option used with PROC OLAP:

```

proc olap data=mylib.newdata cube=cubeA add_data
    outcube=cubeB outschema=testSchema;
metasvr host="myhost" port=8561 repository=myrepository
    olap_schema=prodSchema;
run;

```

The ADD_DATA option requires that either UPDATE_IN_PLACE, OUTCUBE=, or OUTSCHEMA= also be specified. The DATA= or FACT= option is required to specify a table with the new input records.

Note: You cannot use ADD_DATA with cubes that have been created with the NONUPDATEABLE option. △

See Chapter 7, “Updating SAS OLAP Cubes,” on page 79 for further information.

ASYNCLIMIT= *n*

specifies a limit on the number of indices that will be created in parallel during the cube build process. By default, all indices for an aggregation in the cube are

created asynchronously. The number of indices built is based on the number of hierarchies present in the aggregation. Therefore, if your cube contains a large number of hierarchies and your machine has limited resources, you may find that limiting the indices created asynchronously can result in better cube build performance. Use a value greater than zero to set this limit. The default value is 0. There is no limit, as all indices for an aggregation are created asynchronously.

Note: The NWAY is the largest aggregation in the cube and contains data for ALL hierarchies. \triangle

COALESCE_AGGREGATIONS

specifies that all aggregations in a cube will be coalesced. When a cube is updated, multiple aggregation partitions (separate SPDE tables) are produced by one or more cube update (ADD_DATA) operations. The more partitions an aggregation has, the slower the access time will be. To prevent reduced query performance on a cube, these aggregation partitions can be coalesced (combined) back into a single aggregation table. Coalescing updated aggregations periodically keeps the number of aggregation partitions at a minimum. When an aggregation is coalesced, all existing partitions are grouped together to create a single SPDE table containing data from all partitions. This becomes the new aggregation table and the old partitions are deleted. Only aggregations with more than one partition are coalesced. All others are ignored (and a warning message is printed to the log).

Only MOLAP aggregations (those aggregations created by PROC OLAP) can be coalesced. No ROLAP aggregations are coalesced. If a ROLAP aggregation is specified in a COALESCE_AGGREGATION statement, that statement is ignored and a warning message is printed to the log.

COMPACT_NWAY

specifies that the cube build will include an additional summarization step that is designed to decrease the size of the NWAY aggregation and improve viewing performance. The amount of improvement depends on the nature of the data. The cubes that improve the most are those that have the largest number of rows that can be included in the additional summarization step.

Candidates for compaction are cubes that are built from star schema, where the cube does not define levels for all columns in all dimension tables. This can result in a fact table that contains many rows that belong to the same leaf member of a given hierarchy. These are the rows that are summarized to decrease the size of the NWAY aggregation.

For example, assume that a cube is built from a star schema that contains a Time dimension table. The Time table contains columns for year, quarter, month, and day, along with a primary key column. If the cube is defined so that the day column is not specified as a level of a time hierarchy, then there are up to 31 key values that refer to each unique combination of year, quarter, and month. Together, these key values define a unique leaf member of that hierarchy. These are the values that are summarized at build time.

The amount of compaction in the NWAY aggregation is determined by the number of source rows that can be summarized. The number of summarized rows depends on the number of unique key values in the fact table that refer to the same leaf member of a hierarchy. Another compaction factor is the number of rows in the fact table that contain unique combinations of keys; these rows are not compacted.

COMPRESS | NOCOMPRESS

specifies whether or not to store the aggregation tables in a compressed format on disk.

Note: This option applies to the automatically created NWAY and all aggregations that do not explicitly specify a COMPRESS option in the AGGREGATION statement. \triangle

Default: NOCOMPRESS

CONCURRENT=*n*

specifies the maximum number of aggregations to create in parallel. This option does not apply to the NWAY aggregation, which is always built first (unless the NO_NWAY option is set).

Default: 2, which is based on the results of a special algorithm that takes into consideration the number of aggregations that are being created and the number of processors that are available. The algorithm assumes that CPU resources should be reserved for creating aggregation indexes.

Tip: So that each built index has a fair share of the assigned INDEXSORTSIZE memory, INDEXSORTSIZE is divided by the CONCURRENT value. The value of INDEXSORTSIZE should give each concurrent index build enough memory to at least hold a table PARTSIZE. For best performance, INDEXSORTSIZE divided by CONCURRENT should be greater than PARTSIZE.

CUBE=*cube-name* | *folder-path/cube-name* <(cube)>

specifies a valid SAS name for the cube to be created, renamed, or updated. The folder path that is specified is a metadata folder path. It exists only in the metadata. In SAS 9.2, cubes are identified both by their location in an OLAP schema and by their association with a metadata folder. By default, if no folder-path is given, the cube will be associated with the folder that contains the OLAP schema. When you specify a metadata folder, you can optionally add the metadata type (CUBE) after the cube name. For naming guidelines, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30.

Note: The cube name must be unique within the OLAP schema *and* the metadata folder. The specified metadata folder path must exist. △

Here are some examples of the CUBE= option.

Example 1

```
proc olap cube='MYCUBE' /* other options */;
metasvr ... olap_schema='SASApp - OLAP Schema';
/* further statements */
run;
```

This example creates a cube named MYCUBE in the OLAP schema SASApp - OLAP Schema and also surfaces the cube in that schema's metadata folder. It is located by default in **/Shared Data/SASApp - OLAP Schema:**

Example2

```
proc olap cube='/Shared Data/Cubes/MYCUBE(cube)' /* other options */;
metasvr ... olap_schema='SASApp - OLAP Schema';
/* further statements */
run;
```

This example creates a cube named MYCUBE in the OLAP schema SASApp - OLAP Schema and also surfaces the cube in the metadata folder **/Shared Data/Cubes.**

DATA | FACT=*dsname*

specifies the data source for the cube. The unsummarized data source can be any SAS data file, including files that are supported by SAS/ACCESS software engines. If you load the cube from a star schema, then the *dsname* is the name of the fact table that contains the analysis variables from which to derive the measures for the cube. The fact table must also contain fact keys that correspond to dimension tables in the star schema.

You can also provide data set options along with DATA | FACT=. Options are stored within the cube and reapplied when the data is accessed at run time. For more information, see “Data Set Options” in *SAS Language Reference: Concepts*.

Note: This option is not required if you want to define the cube by using input data from a fully summarized external data source (a crossing of all dimensions known as an NWAY). In that case, you specify the data source for the cube by using the TABLE= option in the AGGREGATION statement. \triangle

Interaction: If you load the cube from a star schema, then you must use the DIMENSION statement to do the following:

- specify the dimension table name (the DIMTBL= option)
- specify the dimension (primary) key column (the DIMKEY= option)
- specify the column (foreign key) in the fact table that corresponds to the dimension key column (the FACTKEY= option)

DATAPATH=(*'pathname' ...'pathnameN'*)

specifies the location of one or more partitions in which to place aggregation table data. The data is distributed by cycling through each partition location according to the partition size (set using the PARTSIZE= option). For example, if you specify **DATAPATH=**(**'C:\data1' 'D:\data2'**), then PROC OLAP places the first partition of each aggregation table into directory **C:\data1**, the second partition of each table into directory **D:\data2**, the third partition of each table into **C:\data1**, and so on. It is also possible to have aggregation tables that use fewer than the specified number of partitions. For example, your cube might contain an aggregation table that fits entirely into **C:\data1**.

Note: This option applies to the automatically created NWAY and all aggregations that do not explicitly specify a DATAPATH= option in the AGGREGATION statement. \triangle

Default: The cube subdirectory of the location that is specified by the PATH= option in the PROC OLAP statement

DELETE

deletes the physical cube that is specified with the CUBE= option. It also deletes the cube's definition, which is stored in the metadata repository.

If either the physical cube, its registration, or both are not present, then the DELETE option behaves as explained in the following table:

Note: The use of the DELETE option will remove all information about a cube, including security information and information maps. \triangle

Table A1.1 How the DELETE Option Behaves If the Physical Cube or Its Registration Is Not Present

Physical Cube Exists	Registration eExists	Behavior of DELETE Option
No	Yes	The physical cube is not deleted. The registration is deleted. If there is a registration, and you use the DELETE option, the registration is always deleted and you cannot recreate the cube from the registration. You can only recreate the cube from the registration when you use the DELETE_PHYSICAL option.
No	No	Fails because there is nothing to delete.
Yes	No	Fails because the cube cannot be located without its registration information. You must manually delete the cube files.

DELETE_PHYSICAL

deletes the physical cube that is specified with the CUBE= option but leaves the cube definition intact. This enables you to build a new cube based on the saved cube definition. With the DELETE_PHYSICAL option, the cube registration is maintained in the metadata. As a result, you can change the definition of the cube without changing the cube name or file path. The physical cube must be removed before that cube can be rebuilt.

Note: The use of the DELETE_PHYSICAL option is preferable when rebuilding a cube, because it preserves information about a cube, including security information and information maps. △

If either the physical cube, its registration, or both are not present, then the DELETE_PHYSICAL option behaves as explained in the following table:

Table A1.2 How the DELETE_PHYSICAL Option Behaves If the Physical Cube or Its registration Is Not Present

Physical Cube Exists	Registration Exists	Behavior of DELETE_PHYSICAL Option
No	Yes	A warning message is given that there is no physical cube to delete.
No	No	Fails because there is nothing to delete.
Yes	No	Fails because the cube cannot be located without its registration information. You must manually delete the physical cube files.

DESC | DESCRIPTION='cube-description'

specifies up to 200 characters to be stored as descriptive text. If the text includes blank spaces or any characters that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Note: Cubes that are built with a cube description that is longer than 200 characters cannot be fully registered in the SAS Metadata Repository. The cube description will be truncated to 200 characters. △

DRILLTHROUGH_TABLE | DT_TABLE | DT_TBL=table-name

specifies an optional drill-through table. Drill-through tables can be used by client applications to provide a view from processed data into the underlying data source.

You can specify the DATA | FACT= table or a different table that includes the necessary data and columns.

You can also specify data set options with this option. Options are stored within the cube and reapplied when the data is accessed at run time. For more information, see “Data Set Options” in *SAS Language Reference: Concepts*.

EMPTY_CHAR=*'string'*

specifies the quoted text string that identifies members of character levels that are to be skipped or disregarded. Members are skipped in order to create ragged or unbalanced hierarchies, as described in “Defining Ragged and Unbalanced Hierarchies for a Dimension” on page 44. The maximum length of the quoted string is 256 characters.

To be skipped, a member in a character level must have a caption whose value matches the value of the EMPTY_CHAR= option. For example, if a member in a character level is skipped, and if the caption of that member is **Empty**, then the EMPTY_CHAR= option is specified as follows:

```
empty_char='Empty'
```

Interaction: When specified in the PROC OLAP statement, the EMPTY_CHAR= option can be overridden by the EMPTY_CHAR= or IGNORE_EMPTY options in a HIERARCHY statement or by the EMPTY= or IGNORE_EMPTY options in a LEVEL statement.

To skip members in numeric levels, use the EMPTY_NUM= option.

EMPTY_NUM=*'string'*

specifies the quoted text string that identifies members of numeric levels that are to be skipped or disregarded. Members are skipped in order to create ragged or unbalanced hierarchies, as described in “Defining Ragged and Unbalanced Hierarchies for a Dimension” on page 44. The maximum length of the quoted string is 256 characters.

To be skipped, a member in a numeric level must have a caption whose value matches the value of the EMPTY_NUM= option. For example, if a member in a numeric level is skipped, and if the caption of that member is **Empty**, then the EMPTY_NUM= option is specified as follows:

```
empty_num='.'
```

Interaction: When specified in the PROC OLAP statement, the EMPTY_NUM= option can be overridden by the EMPTY_NUM= or IGNORE_EMPTY options in a HIERARCHY statement or by the EMPTY= or IGNORE_EMPTY options in a LEVEL statement.

Note: If there is no format that is associated with the member value, then BEST12 is used as the format. \triangle

To skip members in character levels, use the EMPTY_CHAR= option.

ESRI_MAP_SERVER=*MapServerName*

specifies the ESRI map server to which this cube should be linked. The map server must already be defined on the metadata server. The ESRI map server must be an ArcGIS server defined on the metadata server. The name will be unique across all ArcGIS servers defined on the metadata server. There can only be one map server per cube.

ESRI_REPLACE

indicates that any existing associations to ESRI metadata should be replaced with those specified in this PROC OLAP session. If no ESRI options are provided, all existing linkages are removed. If ESRI_REPLACE is not specified and there is a conflict between existing associations and those specified in this session, an error results and no associations are added or changed.

IGNORE_MISSING_DIMKEYS=TERSE | VERBOSE

when specified while building a cube from a star schema, causes SAS to ignore an error condition, log the error, and continue building the cube. The error condition is detected when the fact table contains foreign key values that are not present in one of the contributing dimension tables. By default, and when this option is not specified, any missing dimension keys stop the build of the cube. When **IGNORE_MISSING_DIMKEYS=TERSE** is specified, the cube build continues and the fact table row with the missing key is ignored (it is not built into the cube). The SAS log receives an entry that lists the total number of key values that are missing from each dimension table. Specifying a value of **VERBOSE** produces the same behavior, except that the log receives additional details; the missing keys are listed for each dimension table.

INDEX | NOINDEX

specifies whether or not to create the aggregations with indexes. For faster cube creation and adding and deleting aggregations, you can set this option to **NOINDEX**. However, the lack of indexes will adversely affect query performance.

Note: This option applies to the automatically created **NCW** and all aggregations that do not explicitly specify an **INDEX** option in the **AGGREGATION** statement. △

Default: **INDEX**

INDEXPATH=('pathname' ...'pathnameN')

specifies the locations of the index component files that correspond to each aggregation table partition as specified by the **DATAPATH=** option.

Note: This option applies to the automatically created **NCW** and all aggregations that do not explicitly specify an **INDEXPATH=** option in the **AGGREGATION** statement. △

Note: Indexes are not created for aggregations that have fewer than 1,024 records. △

Default: The cube subdirectory of the location that is specified by the **PATH=** option in the **PROC OLAP** statement

INDEXSORTSIZE=n

specifies the amount of memory in megabytes that is available when aggregations are being created.

Default: The system's available memory

Tip: So that each built index has a fair share of the assigned **INDEXSORTSIZE** memory, **INDEXSORTSIZE** is divided by the **CONCURRENT** value. The value of **INDEXSORTSIZE** should give each concurrent index build enough memory to at least hold a table **PARTSIZE**. For best performance, **INDEXSORTSIZE** divided by **CONCURRENT** should be greater than **PARTSIZE**.

MAX_RETRIES =

specifies the number of times **proc OLAP** will attempt to reconnect to the metadata server after a connection is lost. The value must be between 0 and 25. The default is 1, which indicates a single attempt to reconnect. A value of 0 means that **PROC OLAP** will not be attempted to reconnect.

MAX_RETRY_WAIT=

specifies the maximum number of seconds to wait before **PROC OLAP** attempts to reconnect to the metadata server. The time to wait (initially **MIN_RETRY_WAIT**) will be doubled each time a reconnect fails until this limit is reached. The value must be between 1 and 3600. The default value is the greater of **MIN_RETRY_WAIT** and 30.

MAXTHREADS=*n*

specifies the maximum number of threads that are used to asynchronously create the aggregation indexes. The processing engine calculates how many threads are needed based on the number of indexes that are being created and the INDEXSORTSIZE= value. This option sets a limit on the number of threads regardless of the number that is calculated by the processing engine. However, if the processing engine determines that fewer than the maximum number of threads is needed, then only the calculated number of threads are used.

Default: The value of the SAS system option SPDEMAXTHREADS or 0. If the value is 0, then the processing engine determines the number of threads based on the number of indexes that are created plus the available memory. The maximum value is 65,536 threads.

MIN_RETRY_WAIT=

specifies the initial number of seconds to wait before PROC OLAP attempts to reconnect to the metadata server. This delay applies to a second attempt to reconnect to the metadata server. There is no delay for the first attempt to reconnect. The value must be between 1 and 3600 (one hour). The default value is the lesser of MAX_RETRY_WAIT and 30.

NO_NWAY

prevents PROC OLAP from automatically creating an NWAY aggregation (the crossing of all dimension levels) for the new cube. The automatically created NWAY is usually the largest in the cube and most resembles the content of the unsummarized data source.

Interaction: If you use this option, then the input data source that is specified with the DATA= or FACT= option must be available at run time; otherwise, queries that are not covered by other aggregations will fail.

NONUPDATEABLE

specifies that the dimension (or dimensions) for a cube should be built with the minimum amount of disk space to represent the members available when the cube is created. By default, new dimensions are built to allow for new members to be added in future updates. NONUPDATEABLE is valid only when the cube is first created.

If this option is set, no new members can be added to the dimension in future updates of the cube. This option can be specified either on the PROC OLAP statement or on the individual DIMENSION statements. To make individual dimensions non-updateable, use this option on the DIMENSION statement instead.

OUTCUBE=*cube-name2* | *folder-path/cube-name* <(cube)>

Specifies a valid SAS name for the new cube registration. For naming guidelines, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30. The folder path that is specified is a metadata folder path. It exists only in the metadata. The OUTCUBE= option can be used in combination with either the RENAME= options or the ADD_DATA option.

If you are updating a cube, you can use OUTCUBE= option with the ADD_DATA option to specify the location of the new cube registration. If you are renaming an existing cube, you can use the OUTCUBE= with the RENAME option.

OUTCUBE= does not create a copy of existing cubes. In SAS 9.2, cubes are identified both by their location in an OLAP schema and by their association with a metadata folder. By default, if no folder-path is given, the cube will be associated with the folder of the original cube’s OLAP schema, or, if specified, with the folder of the OLAP schema specified in the OUTSCHEMA= option. When you specify a metadata folder, you can optionally add the metadata type (CUBE) after the cube name.

Note: The cube name must be unique within the OLAP schema and the metadata folder. The specified metadata folder path must exist. △

Here are some examples of the OUTCUBE= option.

Example 1

```
proc olap rename cube=MYCUBE_A outcube='/Shared Data/Cubes/MYCUBE_B';
metasvr ... olap_schema='SASApp - OLAP Schema';
run;
```

This example creates a cube registration named MYCUBE_B in the schema **SASApp - OLAP Schema** and the folder **Shared Data/Cubes**.

Example 2

```
proc olap add_data cube=MYCUBE_A outcube='MYCUBE_B' /* more options */;
metasvr ... olap_schema='SASApp - OLAP Schema';
/* more statements */
run;
```

This example creates a registration named MYCUBE_B for the updated version of the cube MYCUBE_A. The new registration is located in the same OLAP schema and folder as MYCUBE_A.

Note: See Chapter 7, “Updating SAS OLAP Cubes,” on page 79 for further information. △

OUTSCHEMA=*olap-schema-name* |

specifies the name of the OLAP schema into which a new cube should be placed. The OLAP schema must already exist within the SAS Metadata Server and the metadata (or a repository on which it depends) specified in the METASVR statement. Changing the OLAP schema for the cube is strictly a change to the metadata and does not result in a physical move.

You can use OUTSCHEMA= in combination with the RENAME option to move an existing cube registration into a different OLAP schema. You can simultaneously use the OUTCUBE= option to also change the cube’s name. In addition, OUTSCHEMA= can be used in combination with the ADD_DATA option to specify the target OLAP schema for the registration of the updated cube.

Note: See Chapter 7, “Updating SAS OLAP Cubes,” on page 79 for further information. △

PARTSIZE=*partition-size*

specifies the partition size in megabytes of the aggregation table partitions and their corresponding index components.

Note: This option applies to the automatically created NWAY and all aggregations that do not explicitly specify a PARTSIZE= option in the AGGREGATION statement. △

Default: 128 megabytes. The minimum value is 16 megabytes.

PATH=*'pathname'*

specifies the physical path to the location of a new cube. Within the specified path, the cube is stored in a directory that uses the name of the cube. However, if the cube folder already exists, a unique subdirectory is generated. When a cube is renamed with the OUTCUBE= option, either with the RENAME or ADD_DATA option, the physical path name does not change. Here is an example where this might cause a unique subdirectory to be generated:

- 1 Create a cube named MRKTDATA.
- 2 Rename the cube MRKTDATA as NEW_MARKET_DATA (or update MRKTDATA and specify OUTCUBE=NEW_MARKET_DATA).

3 Create another cube named MRKTDATA.

The new cube MRKTDATA would have a cube folder of **C:\v9cubes\MRKTDATA1** because the cube folder **C:\v9cubes\MRKTDATA** still exists for cube NEW_MARKET_DATA.

REGISTER_ONLY

specifies that metadata for a cube is to be registered, but the cube is not to be physically built. All of the metadata for the cube is added to the SAS Metadata Repository. The physical cube can be built later using the existing metadata definition, with either the short form of PROC OLAP or in SAS OLAP Cube Studio. Note that all data sets must physically exist at registration time. The data sets can be empty—they do not need to contain data. Complete data sets are required when the cube is physically built.

RENAME

indicates that the cube should be renamed. Renaming a cube updates the metadata for the cube but does not change the file structure or physical location of the cube. This process requires an exclusive lock on the cube. If the cube is being queried by SAS OLAP Server sessions, the cube will need to be disabled (see the OLAPOPERATE procedure) before it is renamed.

Here is the syntax usage for the RENAME option:

```
PROC OLAP RENAME CUBE=cube-name1 <OUTCUBE=cube-name2>
  <OUTSCHEMA=olap-schema-name>; METASVR options; RUN;
```

Note: The cube that is being renamed must be defined in the metadata but does not need to physically exist. Δ

REORGANIZE_LEVELS | REORG_LEVELS

enables you to reorganize all eligible levels for a cube. It is used when you have performed multiple updates on a cube such that the level values must be reorganized. If this option is used, only those levels that have previously been updated with new captions during prior instances of cube update are eligible to be reorganized.

Here is an example of the REORGANIZE_LEVELS option:

```
proc olap cube=test REORG_LEVELS;
metasvr ... etc;;
run
```

SECURITY_SUBSET= YES|NO

controls how permission conditions are interpreted at query time. With SECURITY_SUBSET= YES, cell values are recalculated at query time based on the security subset defined by the active permission conditions for the given user. SECURITY_SUBSET=NO does not recalculate the cell values. The default value (NO) includes all members within a total.

Note: This option can also be found in SAS OLAP Cube Studio. See the option **Include secured member values in presummarized computations** on the Cube Designer - General page of the Cube Designer wizard. Δ

SEGSIZE=*rows-per-segment*

specifies the number of observations (table rows) in the file segment of the index component. The value is expressed in multiples of 1,024. The minimum value is 1 (1,024 rows). The segmented indexes are used to optimize the processing of WHERE expressions. Each parallel thread is given a segment of the table to evaluate that is equal to the value of the SEGSIZE= option multiplied by 1,024.

Note: This option applies to the NWAY aggregation and all aggregations that do not explicitly specify a SEGSIZE= option in the AGGREGATION statement. Δ

Default: 8 (8 x 1,024 = 8,192 rows)

SYNCHRONIZE_COLUMNS

synchronizes the column names stored in the cube's internal metadata with the input column names for a cube. This is necessary when dataset column names have changed since the cube was last built. This applies to cube levels, measures, and properties.

Note: If column names for a cube have changed and the cube was not synchronized, you will not be able to update the cube. If the cube has presummarized aggregations, you will not be able to query the cube. △

UPDATE_DISPLAY_NAMES

enables you to update the captions for dimensions, hierarchies, levels, member properties, and measures. This option specifies that captions or descriptions on the cube are going to be modified. The UPDATE_DISPLAY_NAMES option is allowed with the ADD_DATA option or by itself on the PROC OLAP statement. It can also be used to change the description or caption on the DIMENSION, HIERARCHY, LEVEL, MEASURE and PROPERTY statements. Here is the syntax usage for the UPDATE_DISPLAY_NAMES option:

```
PROC OLAP CUBE=cubename DESCRIPTION="new description"
UPDATE_DISPLAY_NAMES DT=newdtname; METASVR options; DIMENSION
TIME DESCRIPTION="New TimeDescription" CAPTION="New TimeCaptions";
```

Specific considerations for the UPDATE_DISPLAY_NAMES option include the following:

Drill-through Table	can be updated with the UPDATE_DISPLAY_NAMES option. When you rename a cube, only the cube description or drill-through table name can be updated via the UPDATE_DISPLAY_NAMES option.
NUNIQUE Measures	When the UPDATE_DISPLAY_NAMES option is specified, the NUNIQUE measure's INCLUDE_CALCULATED_MEMBERS option can be toggled ON or OFF (with the addition of the NOINCLUDE_CALCULATED_MEMBERS option).
RENAME	The only description that can change during a RENAME of a cube is the cube description (the drill-through table name can also be changed during a RENAME). Other descriptions and captions must be altered in a separate step after the RENAME is executed.

UPDATE_IN_PLACE

enables you to control how the updated cube is made available to users. An update of the existing cube will occur. UPDATE_IN_PLACE can only be used when the ADD_DATA option is specified. It cannot be used in combination with the OUTCUBE= or OUTSCHEMA= options. See Chapter 7, "Updating SAS OLAP Cubes," on page 79 for further information.

WORKPATH=(*'pathname1' ... 'pathnameN'*)

specifies one or more locations for temporary work files.

Default: For all operating environments except z/OS and VMS, if the WORKPATH= option is not specified, PROC OLAP uses the SPDEUTILLOC= system option. If SPDEUTILLOC= is not specified, PROC OLAP uses the UTILLOC= system option. If UTILLOC= is not specified, or if you do not have write access to the specified path, the following message is generated:

```
ERROR: Cannot create temporary index for proc olap.
NOTE: The SAS System stopped processing this step
because of errors.
```

For z/OS and VMS, PROC OLAP uses the SPDEUTILLOC= system option only.

Note: The SPD Engine options are honored by PROC OLAP only if the REGISTER_ONLY option is set on the PROC OLAP statement, and only if the long form of the procedure is used to build a cube. The long form of the procedure is used when you run a SAS program that contains PROC OLAP code. The short form of the procedure is used by SAS OLAP Cube Studio. Δ

METASVR Statement

The METASVR statement identifies the SAS metadata repository in which existing cube metadata information exists or in which metadata about a new cube is stored.

METASVR OLAP_SCHEMA=*'schema-name'* < option(s)>;

The METASVR statement options can be used to override the metadata repository connection values that are specified through SAS start-up options.

Note: During an interactive SAS session, if connection information is not available either through start-up settings or through a METASVR statement, then the user is prompted for the missing information. For more information about SAS start-up options, see *SAS Language Reference: Dictionary*. Δ

Following is an example of a METASVR statement with all of its options set:

```
metasvr olap_schema='Banking Schema'
  repository='financial repository'
  host='misdept.us.mar.com'
  port=9999
  userid=jjones
  pw='my password';
```

Required Argument

OLAP_SCHEMA=*'schema-name'*

is a string that specifies the name of the schema that has been defined in a SAS metadata repository. The name can be a maximum of 32 characters. The OLAP schema specifies which group of cubes that a SAS OLAP Server can access. Each OLAP schema can be accessed by multiple SAS OLAP Servers. However, each SAS OLAP Server has access to only one OLAP schema. When using embedded blanks or special characters in the schema name, enclose the name in quotation marks.

Options

HOST=*'metadata-server-host-name'*

is a string that specifies the IP address of the metadata repository host. An example is **'misdept.us.mar.com'**. The address can be a maximum of 256 characters. When using lowercase letters, embedded blanks, or special characters in the host name, enclose the name in quotation marks.

PORT=*port-number*

specifies the numeric value of the port on which the metadata repository resides.

PW=*'password'*

is a string that specifies the password for the user identified with the USERID= option. The password can be a maximum of 512 characters. When using lowercase letters, embedded blanks, or special characters in the password, enclose the password in quotation marks.

REPOSITORY=*'repos-name'*

is a string that specifies the name of a SAS metadata repository in which existing cube metadata information exists or in which metadata about a new cube is stored. The name can be a maximum of 60 characters. When using lowercase letters, embedded blanks, or special characters in the repository name, enclose the name in quotation marks.

USERID=*'userid'*

is a string that specifies the user's identification for the specified metadata repository. The identification can be a maximum of 256 characters. When using lowercase letters, embedded blanks, or special characters in the user ID, enclose the user ID in quotation marks.

DIMENSION Statement

The DIMENSION statement defines the logical and hierarchical relationships between the variables in the input data.

DIMENSION *dim-name* HIERARCHIES=(*hier-nam ... hier-nameN*) <*option(s)*>;

At least one DIMENSION statement must be specified when the cube is created. The DIMENSION statement is not used when adding or deleting aggregations from cubes.

The maximum number of dimensions that can be defined in a cube is determined by combining the number of dimensions with the number of multiple hierarchies that are defined in those dimensions. The maximum value of that sum is 128. Mathematically, the sum is expressed as follows:

$$\text{MaxDims} = \text{NumDims} + \text{NumMultipleHeirarchies} = 128$$

All hierarchies other than the first hierarchy in each dimension apply to the total.

Here are some examples of cubes that are defined with the maximum number of dimensions:

128 dimensions, each dimension has 1 hierarchy

127 dimensions, 1 dimension has 2 hierarchies

126 dimensions, 1 dimension has 3 hierarchies

126 dimensions, 2 dimensions have 2 hierarchies

The DIMENSION statement does *not* create aggregations. To create aggregations, use the AGGREGATION statement.

A DIMENSION statement must include the name of at least one hierarchy in its HIERARCHIES= option. In addition, a HIERARCHY statement must include the name of at least one level in its LEVELS= option. Note that you cannot use the same level in more than one dimension.

You can use LEVEL statements to specify a time period for each level in a TIME dimension. You can also use LEVEL statements to supply information such as a level-specific sort order or a level description.

The following example uses one DIMENSION statement, two HIERARCHY statements, and three optional LEVEL statements to define a fully specified dimension. In the example, the same levels are being used in different ways.

```
dimension time
  hierarchies=(Year_Months Year_Quarters)
  ;
  hierarchy Year_Months
    levels=(year month day)
  ;
  hierarchy Year_Quarters
    levels=(year quarter day)
  ;
  level year
    type=year
    caption='Year'
  ;
  level quarter
    type=quarters
    caption='Quarter'
  ;
  level month
    type=months
    caption='Month'
  ;
  level day
    type=days
    caption='Day'
  ;
```

Required Arguments

dim-name

names a dimension by using a valid SAS name up to 32 characters. For naming guidelines, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30.

HIERARCHIES=(*hier-name...hier-nameN*)

specifies the name of one or more hierarchies as defined by HIERARCHY statements.

Options

CAPTION=*string*

specifies a maximum of 200 characters that can be used to create a meaningful description of the dimension. Third-party applications that report on cube data might display this description. If the text includes blank spaces or any characters

that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Note: Cubes that are built with captions that are longer than 200 characters cannot be fully registered in the SAS Metadata Repository. Captions will be truncated to 200 characters. △

Default: *dim-name*

DESC | DESCRIPTION=*'string'*

specifies any number of characters that can be used to create a meaningful description of the dimension. Third-party applications that report on cube data might display this description. If the text includes blank spaces or any characters that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Default: *dim-name*

DIMKEY=*dimension-table-column*

specifies the name of a column in the dimension table that is specified in the DIMTBL= option. That column must contain values that correspond to fact key values in the fact table and be a value that corresponds to a unique combination of level values in the fact table.

Note: The corresponding fact key is specified with the FACTKEY= option. The fact table is specified with the FACT= option in the PROC OLAP statement. △

For example, for a dimension that is composed of three levels—NAME, ADDRESS, and INCOME—a dimension key named CUSTOMER_ID might exist. In this dimension, each unique value of CUSTOMER_ID corresponds to a unique combination of NAME, ADDRESS, and INCOME.

Table A1.3 Sample Dimension Data That Illustrates How Unique DIMKEY Values Correspond to Unique Combinations of Level Values

CUSTOMER_ID	NAME	ADDRESS	INCOME
1	Juan	hostel	2000
2	Shelly	apartment	2000
3	Paul	house	25000
4	Makoto	castle	250000000

DIMTBL=*libname.memname*

specifies the two-level SAS name or a dimension table that matches the fact table that is specified with the FACT= option in the PROC OLAP statement. The dimension table must contain one column for each dimension level and each level property and one column for the dimension key. However, if the dimension key is also a level, then the dimension table needs to have only as many columns as there are levels in the dimension. Member metadata for the dimension is derived from the information in the level columns of the dimension table.

You can also specify data set options with DIMTBL=. Options are stored within the cube and reapplied when the data is accessed at run time. For more information, see “Data Set Options” in *SAS Language Reference: Concepts*.

Note: The same dimension tables can be used to load cubes that have some, but not all, dimensions in common. This means that it is possible for multiple cubes to share the same dimension *data*. △

Note: If you are building a cube that will contain multiple national languages, then replace the DIMTBL= option with DIMTABLELIBREF= and

DIMTABLEMEMMPREF= options. In addition, you must create a USER_DEFINED_TRANSLATIONS statement. △

DIMTABLELIBREF=

specifies the library for the data sets that exist, for this dimension, in each language that is specified by the USER_DEFINED_TRANSLATIONS statement. The library is associated with the dimension and not the language. You cannot put different languages in different libraries, but you can put different dimensions in different libraries. This option is required if you are using the Multiple Language Support capabilities of the SAS OLAP Server. It is also used in conjunction with the DIMTABLEMEMMPREF= option.

Note: If you are building a cube that will contain multiple national languages, then DIMTABLELIBREF= and DIMTABLEMEMMPREF= are required instead of DIMTBL=. △

DIMTABLEMEMMPREF=

specifies the member prefix for the translated dimension tables. The member prefix is the prefix of the data set name. The suffix of the name is provided by the USER_DEFINED_TRANSLATIONS statement. For example, if the member prefix is `dealdim_` and the suffix is `da_DK`, then PROC OLAP looks for a data set named `dealdim_da_DK.sas7bdat` in the library that is specified by the DIMTABLELIBREF= option. DIMTABLEMEMMPREF= is required if you are using the Multiple Language Support capabilities of the SAS OLAP Server. It is used in conjunction with the DIMTABLELIBREF= option and the USER_DEFINED_TRANSLATIONS statement.

Note: If you are building a cube that will contain multiple national languages, then DIMTABLELIBREF= and DIMTABLEMEMMPREF= are required instead of DIMTBL=. △

EMPTY_CHAR='string'

specifies the text string that identifies members of character levels that are to be skipped or disregarded. Members are skipped in order to create ragged or unbalanced hierarchies, as described in “Defining Cube Hierarchies” on page 42.

To be skipped, a member in a character level must have a caption whose value matches the value of the EMPTY_CHAR= option. For example, if a member in a character level is skipped, and if the caption of that member is **Empty**, then the EMPTY_CHAR= option is specified as follows:

```
empty_char='Empty'
```

The maximum length of the quoted string is 256 characters.

Interaction: When specified in the HIERARCHY statement, the EMPTY_CHAR= option overrides (for that hierarchy) any specification of the EMPTY_CHAR= option in the PROC OLAP statement. In turn, the EMPTY_CHAR= option in the HIERARCHY statement is overridden by the EMPTY= or IGNORE_EMPTY options in the LEVEL statements in that hierarchy.

To skip members in numeric levels, use the EMPTY_NUM= option.

EMPTY_NUM='string'

specifies the text string that identifies members of numeric levels that are to be skipped or disregarded. Members are skipped in order to create ragged or unbalanced hierarchies, as described in “Defining Cube Hierarchies” on page 42.

To be skipped, a member in a numeric level must have a caption whose value matches the value of the EMPTY_NUM= option. For example, if a member in a numeric level is skipped, and if the caption of that member is **Empty**, then the EMPTY_NUM= option is specified as follows:

`empty_num='Empty'`

The maximum length of the quoted string is 256 characters.

Interaction: When specified in the HIERARCHY statement, the EMPTY_NUM= option overrides (for that hierarchy) any specification of the EMPTY_NUM= option in the PROC OLAP statement. In turn, the EMPTY_NUM= option in the HIERARCHY statement is overridden by the EMPTY= or IGNORE_EMPTY options in the LEVEL statements in that hierarchy.

Note: If there is no format that is associated with the member value, then BEST12 is used as the format. △

To skip members in character levels, use the EMPTY_CHAR= option.

FACTKEY=*fact-table-column*

specifies the name of the column in the fact table that corresponds to the dimension table column that is specified with the DIMKEY= option. The name does not have to match the DIMKEY name. Referring back to the previously discussed example, the FACTKEY name could be CUST_NO even though the DIMKEY name is CUSTOMER_ID. However, even if the names are different, the underlying data must match. For example, you must match numeric columns with numeric columns and character columns with character columns. In addition, if the FACTKEY is a character column, then it must be the same length as the DIMKEY column. If the FACTKEY is a numeric column, then it is handled as a decimal precision number (rather than as an integer).

IGNORE_EMPTY

specifies that, for this hierarchy, any values that were specified for the EMPTY_CHAR= and EMPTY_NUM= options in the PROC OLAP statement are to be ignored. This option can be overridden by specifications of EMPTY_CHAR= and EMPTY_NUM= in the same HIERARCHY statement. The IGNORE_EMPTY option can also be overridden in subsequent LEVEL statements by using the EMPTY= option. For further information, see “Defining Cube Hierarchies” on page 42.

MAP_SERVICE=*MapServiceName*

specifies the map service to which this cube should be linked. The dimension must have a dimension type of GEO. The map service must be a map service defined with the ESRI server specified above. There can only be one map service per GEO dimension.

NONUPDATEABLE

specifies that the dimension should be built with the minimum amount of disk space to represent the members available when the cube is created. By default, new dimensions are built to allow for new members to be added in future updates. NONUPDATEABLE is valid only when the cube is first created.

If this option is set, no new members can be added to the dimension in future updates of the cube. If you want to make all dimensions of a cube non-updateable, you can use the NONUPDATEABLE option on the PROC OLAP statement instead.

**SORT_ORDER=ASCENDING | DESCENDING | ASCFORMATTED |
DESFORMATTED | DSORDER**

specifies a sort order for all levels in the dimension. Values that are returned from queries display in this order by default.

Default: ASCENDING

Interaction: This setting is overridden if sort order is set in a LEVEL statement.

Tip: To specify a sort order for each level within a dimension, set the SORT_ORDER= option in each LEVEL statement. Values that are returned from queries display in this order.

Note: The sort order can be changed at query time using the MDX ORDER functions.

TYPE=TIME | GEO

identifies the dimension as a TIME or GEO dimension. The GEO type is used when defining ESRI map information for a cube. There can only be one GEO and one TIME dimension for a cube.

Requirement: You must set this option for a TIME or GEO dimension. TIME and GEO are the only valid values for this option.

Interaction: You can use LEVEL statements to specify the time period of each level in the TIME dimension. Specifying TYPE=TIME also allows you to use the MDX time series functions during data query.

Note: In order to add geographic information to an existing cube, you must use SAS OLAP Cube Studio. The OLAP procedure does not support adding the GEO type to an existing dimension. \triangle

UPDATE_DIMENSION

enables you to add new members to a cube and update member properties for existing members of the dimension. UPDATE_DIMENSION can only be used on a dimension without the NONUPDATEABLE designation. It can be used with or without the ADD_DATA option on the PROC statement. It can be used with one of the following parameters:

MEMBERS

This means that the dimension table currently associated with the dimension should be read and processed for new members of every hierarchy in that dimension. If ADD_DATA option is used, this is the default behavior for cubes loaded from a detail table or a star schema. If the ADD_DATA option is not used, it can be applied to individual dimensions for cubes loaded from star schemas only. It is not valid for a cube loaded from a detail table. Thus this option provides a way to update individual dimensions without adding data to the cube.

MEMBERS_AND_PROPERTIES

This means that the dimension table currently associated with the dimension should be read and processed for new members PLUS member properties for existing members should be changed with the new values in the dimension table. If the ADD_DATA option is used, this option can be used for cubes loaded from a detail table or a star schema. If the ADD_DATA option is not used it can be applied to individual dimensions, for cubes loaded from star schemas only. It is not valid for a cube loaded from a detail table. Thus this option provides a way to update individual dimensions without adding data to the cube.

OFF

This means that the dimension table currently associated with the dimension should NOT be read at all. It applies only to cubes loaded from a star schema and never to a cube loaded from a single detail table. If the ADD_DATA option is used on the PROC, this option must be used on any UPDATEABLE dimensions whose dimensions tables should not be processed, since the default behavior is UPDATE_DIMENSIONS= MEMBERS. If the ADD_DATA option is not used on the PROC, this is the DEFAULT option for all dimensions in the cube.

LEVEL Statement

The LEVEL statement provides additional information about a level specified with the LEVELS= option in a HIERARCHY statement, and enables you to set options for ragged hierarchies. Each LEVEL statement must correspond to a HIERARCHY statement. You cannot have a level that does not belong to a hierarchy.

LEVEL *level-name* <*option(s)*>;

For TIME dimensions, you can use LEVEL statements to specify a time period for each level in the dimension. However, if you specify the time period for one level, then you must specify the time period for all levels. You also use LEVEL statements to supply information such as a level description or a level-specific sort order. You can have a maximum of 256 levels per cube and a maximum of 19 levels per hierarchy.

Note: Levels that are shared between hierarchies share the values of the options EMPTY_CHAR=, EMPTY_NUM=, EMPTY=, and IGNORE_EMPTY. These options are used to create ragged or unbalanced hierarchies, as described in “Defining Cube Hierarchies” on page 42. △

Note: When you rebuild a cube that has been physically deleted, the rebuilt cube still uses the formats that were originally used to build the cube and were saved in the cube’s metadata. This means that the rebuilt cube does not automatically include any formatting changes that you might have made in the input data source. To manually specify the new formats, edit and rebuild the cube by using SAS OLAP Cube Studio. △

Required Arguments

level-name

specifies a valid SAS name. For naming guidelines, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30. This is the same name that is used in the LEVELS= option in the HIERARCHY statement. Level names must be unique within a cube.

Interaction: Level names should not conflict with the ALL member name in the HIERARCHY statement. You should not use *ALL* for the level name .

Options

CAPTION=*'string'*

specifies a maximum of 200 characters that can be used to create a meaningful description of the level. Third-party applications that report on cube data might display this description. If the text includes blank spaces or special characters that are not permitted in a valid SAS name, then enclose the caption within quotation marks.

Note: Cubes that are built with captions that are longer than 200 characters cannot be fully registered in the SAS Metadata Repository. Captions will be truncated to 200 characters. △

Default: The column’s label in the input data source. If there is no label available, the default is the level name.

COLUMN=column-name

specifies the name of a column from the input data source. Use this option if the column name is not the same as the LEVEL name. This option is useful in the following scenarios:

- You want your level name to be different from the input column name.
- Different dimension tables for the cube each have a column with the same name
- You want to use the same dimension table and columns in more than one dimension of your cube.
- You load multiple levels from the same input column (in combination with the FORMAT= option).

You can use a column as a level even if it is also being used as a measure.

DESC | DESCRIPTION='string'

specifies any number of characters that can be used to create a meaningful description of the level. Third-party applications that report on cube data might display this description. If the text includes blank spaces or any characters that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Default: The value of the CAPTION= option if one exists; otherwise, the column's label. If there is no label available, the default is the level name.

EMPTY='string'

specifies the text string that identifies members that are to be skipped or disregarded. Members are skipped in order to create ragged or unbalanced hierarchies, as described in “Defining Cube Hierarchies” on page 42.

To be skipped, a member must have a caption whose value matches the value of the EMPTY= option. For example, if a member is skipped, and if the caption of that member is **Empty**, then the EMPTY= option is specified as follows:

```
empty=' Empty'
```

The maximum length of the quoted string is 256 characters.

Interaction: The EMPTY= option overrides for that level any specification of EMPTY_CHAR=, EMPTY_NUM=, or IGNORE_EMPTY that might have been specified in the respective HIERARCHY or PROC OLAP statement.

ESRI_MAP_LAYER=MapLayerName

specifies the ESRI map layer which should be associated with this level. It must be a map layer defined in the specified map service. The level must be a level contained in the dimension with dimension type GEO. Only one map layer may be defined per level. The same map layer can be specified for different levels in the dimension. However, it is generally with a different map field specified.

FORMAT=sas-format-name

specifies the SAS format to be used when reading in the member caption data from the input data source for the LEVEL. By default, the column format specified in the input table is used. However, you can override that format by using the SAS FORMAT statement in the PROC OLAP step or the FORMAT= option on the LEVEL statement.

This option can be particularly useful if you want to load multiple levels from the same column. A typical use is loading the levels of a time dimension. For example, if your input data contain a column with a SAS date values, you could load the time dimension levels, as shown in the following example.

```
dimension Time
  hierarchies=(Time)
```

```

type=time
;
hierarchy Time
  levels=(Year Quarter Month Day)
;
level Year
  type=year
  column=date_id
  format=year.
;
level Quarter
  type=quarters
  column=date_id
format=qtr.
;
level Month
  type=months
  column=date_id
  format=monname.
;
level Day
  type=days
  column=date_id
  format=weekdate.
;

```

IGNORE_EMPTY

specifies that any value of the `EMPTY_CHAR=` option (for character levels) or `EMPTY_NUM=` option (for numeric levels) that was specified in the respective `HIERARCHY` or `PROC OLAP` statement is to be ignored. The level is not to be skipped in the cube build. For further information, see “Defining Cube Hierarchies” on page 42

SORT_ORDER=ASCENDING | DESCENDING | ASCFORMATTED | DESFORMATTED | DSORDER

specifies a sort order for a level within a dimension. Values that are returned from queries display in this order.

Default: If a sort order is not specified in the `DIMENSION` statement or in the `LEVEL` statement, then the default order of `ASCENDING` is applied.

Interaction: This setting overrides the `SORT_ORDER=` setting in the `DIMENSION` statement.

TYPE=YEAR | HALF_YEARS | QUARTERS | MONTHS | WEEKS | DAYS | HOURS | MINUTES | SECONDS

specifies the time period for the dimension levels when you specify the `TYPE=TIME` option in the `DIMENSION` statement.

Requirement: If you specify a time period for one level in the `TIME` dimension, then you must specify the time period for all levels in the dimension. With regard to drill path, identify the levels from the most general time period to the most specific.

PROPERTY Statement

The PROPERTY statement assigns properties to specific levels within specified hierarchies.

PROPERTY *prop-name* LEVEL=*level-name* <*option(s)*>;

Each level can have more than one property assigned to it by using multiple PROPERTY statements. Property names must match the name of a column in the input data source, or you must use the COLUMN= option to specify the column name.

In the following example, the COLUMN= option is used in the first two PROPERTY statements because the column name is different from the property name. In this way, the property named **Population** can be assigned to both the **country** level and the **state** level in the **geo** hierarchy. The level **state** has two properties: **Population** and **West_of_Miss**.

```
property Population
  column=p_country
  hierarchy=geo
  level=country
;
property Population
  column=p_state
  hierarchy=geo
  level=state
;
property West_of_Miss
  hierarchy=geo
  level=state
;
```

Required Arguments

prop-name

specifies a valid SAS name for the property. Usually this is the name of a column in the input data source. If it is not the name of a column, then you must include the COLUMN= option to specify the column name. For naming guidelines, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30.

LEVEL=*level-name*

specifies the name of the level that you are assigning the property to.

Options

CAPTION=*'string'*

specifies a maximum of 200 characters that can be used to create a meaningful description of the level. Third-party applications that report on cube data might display this description. If the caption includes blank spaces or special characters that are not permitted in a valid SAS name, then enclose the caption within quotation marks.

Note: Cubes that are built with captions that are longer than 200 characters cannot be fully registered in the SAS Metadata Repository. Captions will be truncated to 200 characters. Δ

Default: The column's label if it exists, otherwise the property name.

COLUMN=column-name

specifies the name of a column from the input data source. You must use this option if the column name is not the same as the property name.

DESC | DESCRIPTION='string'

specifies any number of characters that can be used to create a meaningful description of the level. Third-party applications that report on cube data might display this description. If the description includes blank spaces or any characters that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Default: The value of the **CAPTION=** option if one exists; otherwise, the column's label.

ESRI_MAP_FIELD=MapFieldName

specifies the field of the map layer to be associated with the level. It must be a map field associated with the **LEVEL** statement *MapLayerName*. The OLAP property must be named "SAS_SPATIAL_ID" and must be a property of the level associated with the map layer. It must be associated with all hierarchies in the dimension. If the **HIERARCHY=** (or **HIERARCHIES=**) option is specified, all hierarchies in the dimension must be specified.

Note: There only can be one map field for each SAS_SPATIAL_ID property. In addition, there can only be one SAS_SPATIAL_ID property per level. △

HIERARCHY=(hier-name ... hier-nameN)

specifies the name of one or more hierarchies that contain the level. If you do not include the **HIERARCHY** option, then the property is automatically assigned to all occurrences of the level in all of the hierarchies in which it appears. Otherwise, the property is assigned to the level only in the specified hierarchies.

SAS_SPATIAL_ID

specifies the ESRI mapping properties for the cube levels that you are mapping to. The **COLUMN**, **HIERARCHY**, and **LEVEL** property options are included.

To associate the values in your OLAP data with the regions of an ESRI map, your cube must contain member properties whose values match the field values in the ESRI map service. For each level of a hierarchy that you want to use with the ESRI map component, you must define a member property named **SAS_SPATIAL_ID** to identify the map region. Here is an example:

```
PROPERTY SAS_SPATIAL_ID
          LEVEL=Level2_ProvinceID
          HIERARCHY= (geography )
          COLUMN=Level2_ProvinceID_Char
          CAPTION='SAS_SPATIAL_ID'
```

HIERARCHY Statement

The **HIERARCHY** statement specifies the navigational order of the levels in a dimension.

HIERARCHY *hier-name* LEVELS=(*level-name1* <*level-name2* ...*level-nameN*> <*option(s)*>);

You must define at least one hierarchy for each dimension. Specifically, each **DIMENSION** statement must identify at least one unique **HIERARCHY** statement.

The maximum number of hierarchies that can be defined in a cube is 128. Mathematically, the sum is expressed as follows:

$$\text{MaxHiers} = \text{NumMultHiers} + \text{NumDimensions} = 128$$

All hierarchies other than the first hierarchy in each dimension apply to the total. Here are some examples of cubes that meet the maximum number of hierarchies:

128 dimensions, each dimension has 1 hierarchy

127 dimensions, 1 dimension has 2 hierarchies

126 dimensions, 1 dimension has 3 hierarchies

126 dimensions, 2 dimensions have 2 hierarchies

Levels in the same dimension can be shared between hierarchies. You can have a maximum of 19 levels per hierarchy. There is no limit to the number of hierarchies per dimension.

Following is an example of a HIERARCHY statement that specifies three levels:

```
hierarchy Geography
  levels=(country region division);
```

Required Arguments

hier-name

specifies a valid SAS name for the hierarchy. This name is also used in the HIERARCHIES= option in the DIMENSION statement. The *hier-name* cannot be the same as any of its level names. Hierarchy names must be unique within the cube. If the hierarchy that you are defining is the only one in the dimension, then the hierarchy name must match the dimension name. For other naming guidelines, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30.

LEVELS=(*level-name1* <*level-name2* ...*level-nameN*>)

specifies a valid SAS name for at least one level. These names correspond to the level names used in any optional LEVEL statements. Level names must be unique within a cube and cannot be the same as the *hier-name*. (You can use a column as a level even if it is also being used as a measure.) Enter one or more names, separated by a space. Enter the level names in the order in which you want them to be used, beginning with the top level. For naming guidelines, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30.

Requirement: If the hierarchy is part of a TIME dimension, then the levels must be listed in order from most general to least general based on their assigned TYPE. For example, a TYPE=YEAR level must be listed before a TYPE=QUARTER level.

Interaction: Level names should not conflict with the ALL member name in the HIERARCHY statement. You should not use *ALL* for the level name .

Options

ALL_MEMBER=*'string'*

specifies the caption for the ALL member of the hierarchy.

Interaction: The ALL member name should not conflict with Level names in the LEVEL or HIERARCHY statement. When selecting the ALL member name, follow these guidelines:

- The all member caption should not be *all*.
- A level name within the hierarchy should not be named *all*.

- The all member caption should not be a level name that is within the hierarchy.

When a cube is built, a warning message is displayed in the SAS log and in SAS OLAP Cube Studio if a conflict between the ALL member name and a level name is detected.

CAPTION=*'string'*

specifies a maximum of 200 characters that can be used to create a meaningful description of the hierarchy. Third-party applications that report on cube data might display this description. If the text includes blank spaces or any characters that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Note: Cubes that are built with captions that are longer than 200 characters cannot be fully registered in the SAS Metadata Repository. Captions will be truncated to 200 characters. △

Default: *hier-name*

DEFAULT

identifies a hierarchy as the default hierarchy for the dimension that is defined by the DIMENSION statement.

Default: The first hierarchy listed for the dimension

DESC | DESCRIPTION=*'string'*

specifies any number of characters that can be used to create a meaningful description of the hierarchy. Third-party applications that report on cube data might display this description. If the text includes blank spaces or any characters that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Default: The hierarchy caption, which may be the default, *hier-name*.

MEASURE Statement

The MEASURE statement defines the cube's measures and indicates how they map to the input data.

MEASURE *measure-name* STAT=*statname* <*option(s)*>;

Include one MEASURE statement for each measure in the cube. Each cube must have at least one measure. Measure names must be unique. You can have a maximum of 1,024 measures per cube.

Note: All cube aggregations have identical measures. △

Required Arguments

measure-name

specifies a valid SAS name for the measure. The name must be unique. For naming guidelines, see “Naming Guidelines and Rules for the SAS OLAP Server” on page 30.

STAT= *statname*

specifies the statistic for the measure. The following base statistics are available: N, NMISS, NUNIQUE, SUM, MAX, MIN, or USS. In addition, these derived

statistics are also available: AVG, RANGE, CSS, VAR, STD, STDERR, CV, T, PRT, LCLM, or UCLM.

Note: At least one non-NUNIQUE measure must be defined. \triangle

New cubes that are based on a data source that contains *existing summarized data* (where such data has been indicated in at least one AGGREGATION statement via the TABLE= option), must include MEASURE statements for the stored statistics required for each derived statistic that you want to create for the new cube. For example, if you want to calculate AVG, you must create measures for N and SUM, as well as AVG.

Note: MOLAP aggregations do not require the N and SUM. \triangle

The following table indicates which stored statistics are required for each derived statistic:

Table A1.4 Stored Statistics Required for Each Derived Statistic

Derived Statistics	Required Stored Statistics
AVG	N, SUM
CSS	N, SUM, USS
RANGE	MIN, MAX
VAR, STD, STDERR, CV, T, PRT, LCLM, UCLM	N, SUM, USS

Note: For information about statistic formulas, see “Keywords and Formulas” in *Base SAS Procedures Guide*. \triangle

For cubes that are *not* loaded from a fully summarized data source (that is, you specified a data source by using the DATA | FACT= option), some statistics use formats taken from the input data source. Specifically, if the statistic is SUM, MIN, MAX, RANGE, AVG, STD, STDERR, LCLM, or UCLM, then PROC OLAP uses the format that is assigned to the column specified by the COLUMN | ANALYSIS= option. The following table lists the formats used for the other supported statistics:

Table A1.5 Default Formats Used for Statistics

Statistic	Format Used
CSS	BEST.
CV	8.2
N	12.0
NMISS	10.0
PRT	6.4
T	7.3
USS	BEST.
VAR	BEST.

For cubes that *are* loaded from a fully summarized data source (that is, you specified the data source by using the AGGREGATION statement), the default format is BEST12.

To override the default formats, you can either set the FORMAT= option or use a SAS FORMAT statement.

Note: The `FORMAT=` option also overrides a `FORMAT` statement. △

Note: When you rebuild a cube that has been physically deleted, the rebuilt cube still uses the formats originally saved in the cube's metadata. This means that the rebuilt cube does not automatically include any formatting changes that you might have made in the input data source. To manually specify the new formats, edit and rebuild the cube by using SAS OLAP Cube Studio. △

COLUMN | ANALYSIS=*anlvar*

specifies the name of a numeric column that is contained in the cube's input data source. (You can use a column as a measure even if it is also being used as a level.)

If the cube is based on an *unsummarized* data source, then *anlvar* is the name of the column in that data source from which the measure will be calculated. Use `COLUMN=` to specify the column.

If the cube is based on a *summarized* data source, then *anlvar* can be the name of the numeric column in the data source that was used as the analysis variable for the pre-calculated measure. Use `ANALYSIS=` to specify the column. It can also be a name that identifies a logical association between measures with the same *anlvar* name.

For example, if your cube has three measures, `N`, `SUM`, and `AVERAGE`, and if those measures were derived from the same analysis variable, then you could specify `ANALYSIS=Sales` to logically link the three measures through their shared analysis variable. You would also identify the analysis variable in the `AGGR_COLUMN=` option.

As a further illustration, assume that you were building a cube with an `NWAY` aggregation that was specified using a summarized SAS dataset. The dataset contains the columns `Country`, `Region`, `Division`, `Year`, `Quarter`, `Month`, `SumOfSales`, and `NumOfSales`. You would use two `MEASURES` statements, one for `SumOfSales` and another for `NumOfSales`, as follows.

```
measure Sales_Sum
  stat=sum
  aggr_column="SumOfSales"
  analysis="Sales"
  desc='Sum of Sales'
  units='Dollars'
  format=dollar10.2
;
measure Sales_N
  stat=n
  aggr_column="NumOfSales"
  analysis="Sales"
  desc='Number of Sales'
  units='Dollars'
  format=dollar10.2
;
```

The `Sales` column becomes logically linked with the physical columns `SumOfSales` and `NumOfSales`.

If the cube consists of a combination of summarized and unsummarized data sources, then *anlvar* refers to both a physical and a logical entity. For example, you might have a cube that requires a physical analysis variable to create a crossing, but that same cube already contains other, higher-level aggregations. In this case, the analysis variable is also used to logically link the measures in the pre-existing aggregations that were derived from the same input column. You would also identify the analysis variable in the `AGGR_COLUMN=` option.

Default: *measure-name*

Interaction: An unsummarized data source is specified with the DATA | FACT= option in the PROC OLAP statement. A summarized data source is specified with the TABLE= option in an AGGREGATION statement.

Note: The COLUMN argument is not required for the NUNIQUE statistic and will be ignored for the NUNIQUE statistic if specified. Δ

Options

AGGR_COLUMN=*input-column*

specifies the name of the numeric column in the summarized input data that contains the values for the measure. The source of the summarized input data is specified in the AGGREGATION statement. This option is valid only for stored statistics.

CAPTION=*'string'*

specifies a maximum of 200 characters that can be used to create a meaningful description of the measure. Third-party applications that report on cube data might display this description. If the text includes blank spaces or any characters that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Note: Cubes that are built with captions that are longer than 200 characters cannot be fully registered in the SAS Metadata Repository. Captions will be truncated to 200 characters. Δ

Default: The default is based on the statistic and the COLUMN= value, as shown in the following table. For example, if the statistic is **SUM** and the COLUMN= value is **Sales**, then the default caption is **Sum of Sales**.

Table A1.6 Defaults for the CAPTION= Option If No Caption Is Specified

Statistic Used for Measure	Default Caption
AVG	Average <i>measure-column-name</i>
CSS	Corrected Sum of Squares of <i>measure-column-name</i>
CV	<i>Measure-column-name</i> Coefficient of Variation
LCLM	<i>Measure-column-name</i> Lower Confidence Limit
MAX	Maximum <i>measure-column-name</i>
MIN	Minimum <i>measure-column-name</i>
N	Number of Values for <i>measure-column-name</i>
NMISS	Number of Missing Values for <i>measure-column-name</i>
NUNIQUE	Number of Unique Values for level-name in hierarchy-name
PRT	Probability of Greater Absolute Value for <i>measure-column-name</i>
RANGE	<i>Measure-column-name</i> Range

Statistic Used for Measure	Default Caption
STD	<i>Measure-column-name</i> Standard Deviation
STDERR	<i>Measure-column-name</i> Standard Error of Mean
SUM	Sum of <i>measure-column-name</i>
T	<i>Measure-column-name</i> T Value
UCLM	<i>Measure-column-name</i> Upper Confidence Limit
USS	<i>Measure-column-name</i> Uncorrected Sum of Squares
VAR	<i>Measure-column-name</i> Variance

DEFAULT

identifies a measure as the default measure for the cube.

Default: The measure defined in the first MEASURE statement

DESC | DESCRIPTION='string'

specifies any number of characters that can be used to create a meaningful description of the measure. Third-party applications that report on cube data might display this description. If the text includes blank spaces or any characters that are not permitted in a valid SAS name, then enclose the text within quotation marks.

Default: *measure-name*

FORMAT=sas-format-name

specifies the SAS format to be used to display the value of the measure. This format overrides the default format (see STAT= for more information) and any format that is specified in a SAS FORMAT statement. Both SAS-supplied formats and user-defined formats are supported.

Note: When you rebuild a cube that has been physically deleted, the rebuilt cube still uses the formats that were originally saved in the cube's metadata. This means that the rebuilt cube does not automatically include any formatting changes that you might have made in the input data source. To manually specify the new formats, edit and rebuild the cube by using SAS OLAP Cube Studio. △

HIERARCHY='string'

specifies the hierarchy in which the level resides. This option is used only with the NUNIQUE statistic. If there is only one hierarchy, then the option can be omitted.

Note: The HIERARCHY= option will be ignored for non-NUNIQUE statistics if specified. △

INCLUDE_CALCULATED_MEMBER | INCLUDE_CALC

specifies that calculated members are included in the NUNIQUE count for the measure statement. This option applies to the STAT= NUNIQUE statistic option only.

LEVEL='string'

specifies the level for which a unique count is determined. This option is used only with the NUNIQUE statistic. The default is the Measure name.

Note: The LEVEL= option is ignored for non-NUNIQUE statistics if specified. △

NOINCLUDE_CALCULATED_MEMBER | NOINCLUDE_CALC

specifies that calculated members are *not* included in the MEASURE statement.

UNITS=*string*'

specifies a maximum of 256 characters that can be used to create a meaningful description of the measure's units (for example, "pounds sterling"). Third-party applications that report on cube data might display this description. If the text includes blank spaces, mixed-case letters, special characters, then enclose the text within quotation marks.

AGGREGATION Statement

The AGGREGATION statement defines an aggregation of the cube based on level information that you provide.

AGGREGATION *level-name* <*level-name2 level-name3 ...level-nameN*> / <*option(s)*>;

You can specify level names that are associated with an unsummarized data source, or you can specify level names that match columns in a table that contains existing aggregated data. The levels can exist in more than one dimension. You do not need to include dimension names, because level names must be unique across dimensions.

Here is an example of an AGGREGATION statement that specifies three levels and uses the NAME= option. The slash character (/) is required to separate level names from option specifications.

```
aggregation country prodtype year /
    name='Product Types by Country';
```

Required Arguments

level-name

is the level that is to be used to create the aggregation. Additional level names are optional. Names are separated by spaces. Names are separated from option specifications with a required slash character (/). You do not have to include all levels that are specified in all HIERARCHY statements, but the names that you do specify must match the names that are used in the HIERARCHY statements. You can include a TABLE= option to identify a table that contains existing aggregated information for your specified levels. The levels that you specify must match columns in the input table.

Restriction: Levels must be listed in drill-path order. You cannot specify an aggregation that contains a summary level that could never be requested. For example, if your TIME hierarchy contains the levels **Year**, **Month**, and **Day**, you could specify **Year** and **Month** as an aggregation, but not **Month** by itself.

Options

Note: For a list of the options that can be used to optimize cube creation and query performance, see "Options Used for Performance" on page 289. Δ

COMPRESS | NOCOMPRESS

specifies whether or not to store the aggregation table in a compressed format on disk.

Default: NOCOMPRESS

DATAPATH=(*pathname* ...*pathnameN*)

specifies the location of one or more partitions in which to place aggregation table data. The data is distributed by cycling through each partition location according to the partition size. This is set by using the PARTSIZE= option. For example, if you specify **DATAPATH**=(**C:\data1** **D:\data2**), then PROC OLAP places the first partition of the aggregation table into directory **C:\data1**, the second partition of the table into directory **D:\data2**, the third partition of the table into **C:\data1**, and so on. It is also possible to have aggregation tables that use fewer than the specified number of partitions. For example, your aggregation table might fit entirely into **C:\data1**.

Default: The cube subdirectory of the location that is specified by the **PATH=** option in the PROC OLAP statement

INDEX | **NOINDEX**

specifies whether or not to create the specified aggregation with indexes. For faster cube creation and updates, you can set this option to **NOINDEX**. However, the lack of indexes might adversely affect query performance.

Note: Indexes are not created for aggregations that have fewer than 1,024 records. △

Default: **INDEX**

INDEXPATH=(*pathname* ...*pathnameN*)

specifies the locations of the index component files that correspond to each aggregation table partition as specified by the **DATAPATH=** option.

Default: The cube subdirectory of the location that is specified by the **PATH=** option in the PROC OLAP statement

NAME=*aggregation-name*

specifies a maximum of 256 characters as the name of the aggregation. If the name includes blank spaces or any characters that are not permitted in a valid SAS name, then the name must be enclosed within quotation marks. The name is stored with the cube's metadata.

Default: A name assigned by SAS, such as **AGGR1**.

PARTSIZE=*partition-size*

specifies the partition size in megabytes of the aggregation table partitions and their corresponding index components.

Default: 128 megabytes. The minimum value is 16.

SEGSIZE=*rows-per-segment*

specifies the number of observations (table rows) in the file segment of the index component. The value is expressed in multiples of 1,024. The minimum value is 1 (1,024 rows). The segmented indexes are used to optimize the processing of **WHERE** expressions. Each parallel thread is given a segment of the table to evaluate that is equal to the value of the **SEGSIZE=** option multiplied by 1,024.

Default: 8 (8 x 1,024 = 8,192 rows)

Interaction: The value of this option overrides for the current aggregation any such value that was specified for all aggregations in the PROC OLAP statement.

TABLE=*libname.dataset*

specifies the name of a SAS data set or data view that contains the data for one aggregation. Every level that is listed in the **AGGREGATION** statement must match a column that contains aggregation information in the specified table. Place this option after the list of level names.

Analysis columns in the table are mapped to the numeric columns that are specified with the **AGGR_COLUMN=** option in **MEASURE** statements.

You can also set data set options with TABLE=. Options are stored within the cube and reapplied when the data is accessed at run time. For more information, see “Data Set Options” in *SAS Language Reference: Concepts*.

Restriction: You cannot use the TABLE= option in an AGGREGATION statement that is used to add an aggregation to an existing cube.

DROP_AGGREGATION Statement

The DROP_AGGREGATION statement removes an aggregation from the specified cube.

```
DROP_AGGREGATION level-name1 < level-name2 ...level-nameN> /
  NAME=aggregation-name;
```

You can specify the levels that are in the aggregation, or the name of the aggregation, or both the levels and the name. The slash character (/) is required to separate level names from option specifications.

Required Arguments

At least one of the following arguments is required for a DROP_AGGREGATION statement:

level-name1

specifies the names of the level that is in the aggregation that you want to drop. Additional levels can be specified using blank spaces to separate the level names, as shown in the following example.

```
drop_aggregation Year Month Product /
  name=Sales ;
```

NAME=*'aggregation-name'*

specifies the name of the aggregation that you want to drop. If the name includes blank spaces or any characters that are not permitted in a valid SAS name, then the name must be enclosed within quotation marks.

DEFINE Statement

The DEFINE statement defines a global calculated member or a named set for any cube that is registered in a SAS Metadata Repository.

```
DEFINE MEMBER | SET 'member-or-set-name' AS 'mdx-expression' ;
```

A calculated member is a dimension member that has been calculated from the member values in the input table. Only the definition of the member is stored; the value is calculated when a query is submitted. A named set is an alias for a specified MDX expression. Named sets are often used to make complex MDX queries easier to read and maintain.

The defined calculated members and named sets are available to any session that creates a query in the context of the SAS OLAP Server and the schema defined in the METASRV statement of the PROC OLAP code that is used to create the global member or set.

DEFINE statements can apply to more than one cube, so the CUBE= option is not required to use this statement. The METASVR statement verifies that the cube definition exists in the metadata repository.

The DEFINE statement can be used alone as shown in this example, which defines two calculated members and one named set. The METASVR is the only other required statement. To define multiple sets or calculated members, separate option values with a comma.

```
proc olap;
metasvr olap_schema='Services Schema'
  repository='services'
  host='misdept.us.mar.com'
  port=9999
  userid=jjones
  pw='my password'
;
define member '[mddbcars].[Measures].[avg]' as
  '[Measures].[sales_sum]/[Measures].[sales_n]',
member '[sales].[Measures].[stat1]' as
  '[Measures].[qty] +1',
set '[campaign].[myset]' as
  '[campaign_dates].[All campaign_dates].children'
;
run;
```

The DEFINE statement can also be used with a PROC OLAP program that creates a cube or with a program that adds aggregations to or deletes aggregations from an existing cube. Cube builds, additions, and deletions occur before the DEFINE statement is processed, so the DEFINE statement is not processed if those statements fail.

```
proc olap data=olapsio.cars
  cube=mddbcars
  path='d:\services\'
;
metasvr olap_schema='Services Schema'
  repository='cars'
  host='misdept.us.mar.com'
  port=9999
  userid=jjones
  pw='my password'
;
dimension date
  hierarchies=(date)
  sort_order=scending
;
hierarchy date
  LEVELS=(dte)
;
level dte
;
dimension cars
  hierarchies=(cars
  sort_order=ascending)
;
hierarchy cars
  levels=(car color)
```

```

;
dimension dealers
  hierarchies=(dealers)
  sort_order=ascending
;
hierarchy dealers
  levels=(dealer dest)
;
measure sales_sum
  column=sales
  stat=sum
  format=dollar15.2
;
measure sales_n
  column=sales
  stat=n
  format=12.0
;
define member '[mddbcars].[Measures].[avg]' as
  '[Measures].[sales_sum] / [Measures].[sales_n]'
;
run;

```

Required Arguments

MEMBER | SET

indicates whether you are creating a calculated member or a named set.

'member-or-set-name'

specifies the name of the member or set that you are creating. If you are creating a calculated member, then this value specifies a name for the member that will be calculated by the MDX expression. If you are creating a named set, then this value is the alias for the specified MDX expression.

AS *'mdx-expression'*

specifies the MDX expression.

UNDEFINE Statement

The UNDEFINE statement deletes from a SAS metadata repository one or more global calculated members or named sets.

UNDEFINE MEMBER | SET *'member-or-set-name'* ;

To delete multiple calculated members or named sets in a single UNDEFINE statement, use commas to separate instances of MEMBER | SET *'member-or-set-name'*.

For additional information on calculated members and named sets, see “DEFINE Statement” on page 276.

The following example shows how a single UNDEFINE statement can be used to delete from a metadata repository two calculated members and one named set.

```

proc olap1;
metasvr olpa_schema='Services Schema'
  repository='services'

```

```

host='misdept.us.mar.com'
port=9999
userid=jjones
pw='my password'
;
undefine member '[carsCube].[Measures].[avg]',
member '[sales].[Measures].[stat1]',
set '[campaign].[myset]'
;
run;

```

Required Arguments

MEMBER | SET

indicates whether you are deleting a global calculated member or a named set.

'member-or-set-name'

specifies the name of the member or set that is to be deleted from the metadata repository. Cube names and dimension names are required for each *member-or-set-name*. Square brackets ([]) are optional inside the quotation marks of *member-or-set-name*, as shown in the preceding example.

USER_DEFINED_TRANSLATIONS Statement

The USER_DEFINED_TRANSLATIONS statement is required to use the Multiple Language Support capabilities of the SAS OLAP Server. This statement specifies the locales that are associated with the data sets that you specify in the DIMENSION statement.

USER_DEFINED_TRANSLATIONS *locale* <*locale2* ...*localeN*> ;

Note: Alternative statement names are UDT and USER_DEFINED_TRANSLATION. △

PROC OLAP uses the UDT statement information, along with DIMENSION statement options, to read your alternate locale data sets and create locale-specific metadata for use at query time. Query results are returned in the language of the requested locale. The Multiple Language Support feature is available only for cubes that are loaded from a star schema. The alternate locale data set names consist of a prefix, which indicates the member, and a suffix, which indicates the language. The DEFINE statement supplies the suffix. The DIMTABLEMEMMPREF= option in the DIMENSION statement specifies the member prefix. For example, if the member prefix is **dealdim_** and the suffix is **p1_PL**, then PROC OLAP looks for a data set named **dealdim_p1_PL.sas7bdat** in the library that is specified by the DIMTABLELIBREF= option.

The following sample code looks for these dimension data sets in the **mylib** library. The default locale is the first locale specified in the UDT statement. Additionally, the default locale does not use the suffix that is defined by the UDT statement. In this example, Polish is the default locale, so the suffix is not used.

```

dimension date
hierarchies=(date)
sort_order=ascending
dimtablelibref=mylib

```

```

    dimtablemempref=ctimedim_
    factkey=dte
    dimkey=dte
    ;
hierarchy date levels=(dte)
    ;
level dte
    ;
dimension cars
    hierarchies=(cars)
    sort_order=ascending
    dimtablelibref=mylib
    dimtablemempref=cardim_
    factkey=carkey
    dimkey=carkey
    ;
dimension cars
    levels=(car color)
    ;
dimension dealers
    hierarchies=(dealers)
    sort_order=ascendng
    dimtablelibref=mylib
    dimtablemempref=dealdim_
    factkey=dealerkey
    dimkey=dealerkey
    ;
hierarchy dealers
    levels=(dealer dest)
    ;
user_defined_translations
    pl_PL /* Polish as used in Poland */
    en_US /* English as used in the United States */
    ja_JP /* Japanese as used in Japan */
    ;

```

Table A1.7 Locales and Associated Data Set Names

Locale	Dimension Data Sets		
English	ctimedim_en_US	cardim_en_US	dealdim_en_US
Japanese	ctimedim_ja_JP	cardim_ja_JP	dealdim_ja_JP
Polish	ctimedim_	cardim_	dealdim_

Required Argument

locale

specifies the locales that correspond to the data sets contained in the library that is specified by the DIMTABLELIBREF= option in the DIMENSION statement. Separate locales with a space.

SAS Servers and Character Encoding

If your server metadata contains characters other than those typically found in English, then you must be careful to start your server with an `ENCODING=` or `LOCALE=` system option that accommodates those characters. For example, a SAS server started with the default US English locale cannot read metadata that contains Japanese characters. SAS will fail to start and will log a message indicating a transcoding failure.

In general, different SAS jobs or servers can run different encodings (such as ASCII/EBCDIC or various Asian DBCS encodings) as long as the encoding that is used by the particular job or server can represent all the characters of the data being processed. In the context of server start up, this requires that you review the characters used in the metadata describing your server (as indicated by the `SERVER=` object server parameter) to ensure that SAS runs under an encoding that supports those characters.

REORGANIZE_LEVEL Statement

REORGANIZE_LEVEL | REORG_LEVEL

This option enables you to reorganize individual, eligible levels for a cube. It is used when you have performed multiple cube updates on a cube such that the levels values must be reorganized. Any level used in this statement must have been updated with new captions during prior instances of cube update.

Here is an example of the `REORGANIZE_LEVEL` statement:

```
proc olap cube=test ;
  metasvr... etc;;

  REORG_LEVEL CUSTOMER_NAMES;

  REORG_LEVEL PRODUCT;
;
run ;
```

Loading Cubes

Loading Cubes from a Detail Table

The following table lists the PROC OLAP statements and options that you use to load a cube from a detail table. The detail table has a column for each level and at least one numeric analysis column from which one or more measures can be generated. The PROC OLAP code that is used for a cube that is built from a detail table will use the `DATA=` option to specify the data source for the cube.

Table A1.8 Statements and Options Used to Load Cubes from a Detail Table

Statements	Options	Required or Optional?
PROC OLAP	DATA=	Required
	CUBE=	Required

Statements	Options	Required or Optional?
	PATH=	Required
	DESC=	Optional
	NO_NWAY	Optional
	WORKPATH=	Optional
	DT_TABLE	Optional
METASVR	OLAP_SCHEMA=	Required
	REPOSITORY=	Optional
	HOST=	Optional
	PORT=	Optional
	PROTOCOL=	Optional
	USERID=	Optional
	PW=	Optional
DIMENSION	HIERARCHIES=	Required
	DESC=	Optional
	CAPTION=	Optional
	TYPE=TIME	Required only for TIME dimensions
	SORT_ORDER=	Optional
LEVEL		The LEVEL statement is optional unless you want to specify time periods for each level in a TIME dimension. If you specify a time period for one level, then you must specify a time period for all levels. To specify a time period, you use the TYPE= option.
HIERARCHY	LEVELS=	Required
	DESC=	Optional
	CAPTION=	Optional
MEASURE	STAT=	Required
	COLUMN ANALYSIS=	Required
	AGGR_COLUMN=	Required if you use the AGGREGATION statement with the /TABLE= option
	DESC=	Optional
	CAPTION=	Optional
	UNITS=	Optional
	FORMAT=	Optional
	DEFAULT=	Optional

Statements	Options	Required or Optional?
PROPERTY	prop-name	Required
	LEVEL=	Required
AGGREGATION		The AGGREGATION statement is optional unless you are creating additional aggregations. In that case, you must specify the names of the contiguous levels to be used to create the aggregation. Use the TABLE= option for cubes that contain aggregated data from tables other than the input data source.

Loading Cubes from a Star Schema

The following table lists the PROC OLAP statements and options that you use to load a cube from a star schema. A star schema is a set of input tables that are defined in a repository. The set of tables includes a single fact table and one or more dimension tables. The fact table must contain at least one numeric analysis column for each set of measures that will be generated. To specify the data source for a star schema you must use the FACT=, DIMTBL=, DIMKEY=, and FACTKEY= options

Table A1.9 Statements and Options Used to Load Cubes from a Star Schema

Statements	Options	Required or Optional?
PROC OLAP	FACT=	Required
	CUBE=	Required
	PATH=	Required
	DESC=	Optional
	NO_NWAY	Optional
	WORKPATH=	Optional
	DT_TABLE	Optional
METASVR	OLAP_SCHEMA=	Required
	REPOSITORY=	Optional
	HOST=	Optional
	PORT=	Optional
	PROTOCOL=	Optional
	USERID=	Optional
	PW=	Optional
DIMENSION	HIERARCHIES=	Required

Statements	Options	Required or Optional?
	DIMTBL=	Required for cubes that support one locale. If the cube will contain multiple national languages, replace this option with DIMTABLELIBREF= and DIMTABLEMEMMPREF=.
	FACTKEY=	Required
	DIMKEY=	Required
	DIMTABLELIBREF	Required if you build a cube that will contain multiple national languages. Replaces DIMTBL=.
	DIMTABLEMEMMPREF	Required if you build a cube that will contain multiple national languages. Replaces DIMTBL=.
	DESC=	Optional
	CAPTION=	Optional
	TYPE=TIME	Required only for TIME dimensions
	SORT_ORDER=	Optional
LEVEL		The LEVEL statement is optional unless you want to specify time periods for each level in a TIME dimension. If you specify a time period for one level, then you must specify a time period for all levels. To specify a time period, you use the TYPE= option.
HIERARCHY	LEVELS=	Required
	DESC=	Optional
	CAPTION=	Optional
MEASURE	STAT=	Required
	COLUMN ANALYSIS=	Required
	AGGR_COLUMN=	Required if you use the AGGREGATION statement with the /TABLE= option
	DESC=	Optional
	CAPTION=	Optional
	UNITS=	Optional
	FORMAT=	Optional
	DEFAULT=	Optional
PROPERTY	prop-name	Required

Statements	Options	Required or Optional?
	LEVEL=	Required
AGGREGATION		The AGGREGATION statement is optional unless you are creating additional aggregations. In that case, you must specify the names of the contiguous levels to be used to create the aggregation. Use the /TABLE= option for cubes that contain aggregated data from tables other than the input data source.

Loading Cubes Using Summarized Data

The following table lists the PROC OLAP statements and options that you use to load cubes from a fully summarized data source (a crossing of all dimensions also known as an NWAY). Unlike a detail table or star schema, a fully summarized cube does not use either the DATA= or FACT= option to specify the data that is used to build the cube. Instead the TABLE= option is used on the AGGREGATION statement.

Table A1.10 Statements and Options Used to Load Cubes from Fully Summarized Data

Statements	Options	Required or Optional?
PROC OLAP	CUBE=	Required
	PATH=	Required
	DESC=	Optional
	NO_NWAY	Optional
	WORKPATH=	Optional
	DT_TABLE	Optional
METASVR	OLAP_SCHEMA=	Required
	REPOSITORY=	Optional
	HOST=	Optional
	PORT=	Optional
	PROTOCOL=	Optional
	USERID=	Optional
DIMENSION	PW=	Optional
	HIERARCHIES=	Required
	DESC=	Optional
	CAPTION=	Optional
	TYPE=TIME	Required only for TIME dimensions
	SORT_ORDER=	Optional

Statements	Options	Required or Optional?
LEVEL		The LEVEL statement is optional unless you want to specify time periods for each level in a TIME dimension. If you specify a time period for one level, then you must specify a time period for all levels. To specify a time period, you use the TYPE= option.
HIERARCHY	LEVELS=	Required
	DESC=	Optional
	CAPTION=	Optional
MEASURE	STAT=	Required
	COLUMN ANALYSIS=	Required
	AGGR_COLUMN=	Required
	DESC=	Optional
	CAPTION=	Optional
	UNITS=	Optional
	FORMAT=	Optional
	DEFAULT=	Optional
PROPERTY	prop-name	Required
	LEVEL=	Required
AGGREGATION	names of the contiguous levels to be used to create the aggregation	Required (additional AGGREGATION statements without the TABLE= option can be used to create aggregations other than the automatically defined NWAY).
	TABLE=	Required

Maintaining Cubes

Building a Cube from an Existing Definition

It is possible to have cube definitions in the SAS metadata that do not have associated physical cubes. For example, you can use the DELETE_PHYSICAL= option in the PROC OLAP statement to delete a cube but leave its definition intact. You can also use SAS OLAP Cube Studio to save only the definition of a new cube.

The following table lists the PROC OLAP statements and options that you use to build a cube from an existing metadata definition:

Table A1.11 Statements and Options Used to Build a Cube from an Existing Definition

Statements	Options
PROC OLAP	CUBE=
METASVR	OLAP_SCHEMA=

Adding Aggregations to an Existing Cube

The following table lists the PROC OLAP statements and options that you use to add aggregations to an existing cube.

Table A1.12 Statements and Options Used to Add Aggregations to an Existing Cube

Statements	Options	Required or Optional?
PROC OLAP	CUBE=	Required
METASVR	OLAP_SCHEMA=	Required
AGGREGATION	Names of the contiguous levels to be used to create the aggregation	Required
	NAME=	Optional
	DATAPATH=	Optional
	INDEXPATH=	Optional
	COMPRESS NOCOMPRESS	Optional
	INDEX NOINDEX	Optional
	PARTSIZE=	Optional
	SEGSIZE=	Optional

Note: You can add and delete aggregations in the same PROC OLAP script. △

Note: You cannot add aggregations to a cube that contains aggregated data from a source other than the input data source. △

Deleting Aggregations from an Existing Cube

The following table lists the PROC OLAP statements and options that you use to delete aggregations from an existing cube.

Table A1.13 Statements and Options Used to Drop Aggregations from an Existing Cube

Statements	Options
DROP_AGGREGATION	Specify one or more level names that correspond to the aggregations that you want to remove, or use the aggregation name to specify the aggregation that you want to remove.
METASVR	OLAP_SCHEMA=
PROC OLAP	CUBE=

Note: You can add and delete aggregations in the same PROC OLAP script. △

Note: You cannot delete aggregations from a cube that contains aggregated data from a source other than the input data source. △

Deleting Cubes

The following table lists the PROC OLAP statements and options that you use to delete existing cubes.

If you use the DELETE option, then both the physical cube and its definition, which is stored in the metadata server, are deleted.

Table A1.14 Statements and Options Used to Delete a Cube and Its Metadata

Statements	Options
METASVR	OLAP SCHEMA=
PROC OLAP	CUBE= DELETE

If you use the DELETE_PHYSICAL option, then only the physical cube is deleted; the definition remains intact.

Table A1.15 Statements and Options Used to Delete a Cube but Retain Its Metadata

Statements	Options
METASVR	OLAP_SCHEMA=
PROC OLAP	CUBE= DELETE_PHYSICAL

Specialized Options for PROC OLAP

Options for Managing Ragged Hierarchies

If a hierarchy is balanced, then all of its branches descend to the same level, and each member has a parent level that is positioned immediately above it. However, hierarchies are not always balanced, and sometimes they contain missing hierarchy members. To manage missing hierarchy members, you can use these four options, which were created specifically for ragged hierarchies:

Table A1.16 Options That Can Be Set to Manage Missing Hierarchy Members in Ragged Hierarchies

Options	Statements Where Option is Available
EMPTY_CHAR=	PROC OLAP and HIERARCHY
EMPTY_NUM=	PROC OLAP and HIERARCHY

Options	Statements Where Option is Available
EMPTY=	LEVEL
IGNORE_EMPTY	HIERARCHY and LEVEL

Options Used for Performance


When you create a cube, you can set some options that can be used to optimize cube creation and query performance. If you set the options in the PROC OLAP statement, then the settings are applied to all aggregations in the cube. If you set the options in the AGGREGATION statement, then the options apply to that specific aggregation. Options set for individual aggregations override any options set in the PROC OLAP statement.

Here are the options:

- ASYNCINDEXLIMIT=
- COMPACT_NWAY
- COMPRESS | NOCOMPRESS
- CONCURRENT=
- DATAPATH=
- INDEXPATH=
- INDEXSORTSIZE=
- MAXTHREADS=
- INDEX | NOINDEX
- PARTSIZE=
- SEGSIZE=

Note: ASYNCINDEXLIMIT=, CONCURRENT=, INDEXSORTSIZE=, and MAXTHREADS= are available only in the PROC OLAP statement. △

For an explanation of these options, see the PROC OLAP statement and the AGGREGATION statement.



APPENDIX

2

The OLAPOPERATE Procedure

<i>The OLAPOPERATE Procedure</i>	291
<i>Syntax: OLAPOPERATE Procedure</i>	292
<i>Connection Options</i>	292
<i>Disconnecting from an OLAP Server</i>	293
<i>Listing Active Sessions on an OLAP Server</i>	293
<i>Listing Session Queries on an OLAP Server</i>	294
<i>Listing Session Rowsets on an OLAP Server</i>	294
<i>Closing a Session</i>	295
<i>Cancelling a Query Result Set</i>	296
<i>Disabling a Cube</i>	296
<i>Enabling a Cube</i>	296
<i>Enabling and Disabling Cubes for a Cube Update</i>	296
<i>Refreshing a Cube</i>	296
<i>Stopping an OLAP Server</i>	297

The OLAPOPERATE Procedure

The OLAPOPERATE procedure enables users to manage an OLAP server, its cubes, its sessions, and their queries. It is intended as a batch interface for many of the functions provided in the OLAP Server Monitor plug-in for SAS Management Console. More specifically, you can use PROC OLAPOPERATE to connect to an OLAP server and do the following:

- list sessions and queries for an active OLAP server
- close individual active sessions
- cancel or close individual open result sets
- disable a cube or all cubes
- enable a cube or all cubes
- refresh a cube or all cubes
- stop the OLAP server

You can establish a connection to a running SAS OLAP Server by providing connection information on the PROC OLAPOPERATE statement or by running the CONNECT statement. All other OLAPOPERATE statements require a connection to an OLAP server. Once you have connected, all subsequent statements apply to that server until a DISCONNECT or STOP SERVER statement is executed.

If a connection is made to a secured OLAP server, you must have administrative privileges defined with your login information in the metadata in order to execute these tasks. The OLAPOPERATE procedure is an interactive procedure and its statements are executed as they are encountered. The OLAPOPERATE procedure is terminated by

a QUIT or a RUN statement. For example, you can list sessions for a server and then close certain sessions based on the output of the LIST SESSIONS statement, without having to reinvoke the procedure (and reconnect to the OLAP server).

The OLAPOPERATE procedure can be used on 9.1 and 9.2 OLAP servers. For 9.1 OLAP servers, some of the statements are not available.

Syntax: OLAPOPERATE Procedure

```
PROC OLAPOPERATE [<connection options>];
  CONNECT <connection options>;
  DISCONNECT;
  LIST SESSIONS [CUBE=cube-name ][USER=user-id];
  LIST QUERIES [SESSION=owner];
  LIST ROWSETS [SESSION=owner];
  CLOSE SESSION id | _ALL_;
  CANCEL QUERY id | _ALL_;
  DISABLE CUBE cube-name | _ALL_;
  ENABLE CUBE cube-name | _ALL_;
  REFRESH CUBE cube-name | _ALL_;
  STOP SERVER;
```

Connection Options

The CONNECT statement enables you to connect to a SAS OLAP Server.

```
CONNECT <connection options>;
```

The following server connection options on the OLAPOPERATE statement establish communication with a SAS OLAP Server.

HOST=	specifies the host name or network IP address of the computer hosting the SAS OLAP Server (for example, HOST=d1234.na.sas.com). The value "localhost" can be used if the SAS session is connecting to a server on the same computer.
PORT=	specifies the TCP port to which the SAS OLAP Server listens for connections (for example PORT= 5451).
USERID=	specifies a user ID to be used to connect to the SAS OLAP Server. The user must have Administer privileges for the OLAP Server in order to execute statements of the OLAPOPERATE procedure.
PASSWORD= PW=	specifies the password that corresponds to the user ID.
SSPI	(Security Support Provider Interface) specifies that the identity of the user running the SAS session will be used to establish the connection. The USERID and PASSWORD options are unnecessary and will be ignored when SSPI is specified. For further information on SSPI, see "Integrated Windows Authentication and Single Sign-On" on page 75 and the SSPI topic in <i>Communications Access Methods for SAS/CONNECT and SAS/SHARE</i> .

SPN= specifies the Service Principal Name of the connection destination. This option is optional even if SSPI is specified. When SPN is not specified, one is created using the HOST and PORT values.

Disconnecting from an OLAP Server

You can disconnect from the current SAS OLAP server with the DISCONNECT statement. After you have disconnected from the server, the only other statement you can execute is the CONNECT statement to connect to another (or the same) server.

When the OLAPOPERATE procedure is terminated with the QUIT or the RUN statement, an automatic server disconnect is executed. It is not necessary to explicitly disconnect from the server before terminating the procedure.

DISCONNECT;

Listing Active Sessions on an OLAP Server

You can list information about active sessions on the OLAP Server with the LIST SESSIONS statement.

LIST SESSIONS [CUBE=*cube-name*] [USER=*user-id*];

If the CUBE option is specified, the output will contain information about all sessions currently accessing the specified cube. If the USER option is specified, the output will contain information about all sessions for the specified user. If neither option is specified, then all active sessions are displayed.

The LIST SESSIONS statement generates the following statistics:

Session ID is the unique session identifier.
Session Owner is the user ID of the session owner.
Seconds Inactive is the number of seconds since the last query for this session.
Results Sets Open is the number of open results sets for the session.

Here is an output example.

```
Session ID: 390C78CD-EAD0-4FEF-9069-2E151AA8B524
Session owner: sasdemo@t2333
Session name:
Session description:
Session seconds inactive = 324 and open results = 0
```

For clients using the 9.2 MDX IOM interface, a confirmation message listing the cubes that are active for the session is also given. For example:

```
Cube ORIONSTAR is active in this session.
```

The following statistics are also generated by the LIST SESSIONS statement:

Number of Active Cubes is the number of cubes being used by this session.
Cube Names is the names of the cubes being used by this session.

Listing Session Queries on an OLAP Server

You can list information about session queries with the LIST QUERIES statement.

LIST QUERIES [SESSION=*owner*];

The LIST QUERIES statement generates the following statistics:

MDX String contains the MDX string for the query. The cube referenced in the query is easily determined because the FROM clause is placed at the front of the query. For example, an MDX query such as

```
select measures.members on columns from mddbcars
```

is displayed as *from mddbcars select measures.members on columns*.

Results Set Type indicates the type of the query's results set. Valid values for this statistic are *Multidimensional*, *Flattened*, or *Schema Rowset* (not used).

Size indicates the approximate memory used in the query.

Note: For SAS versions earlier than 9.2, this statistic indicates the number of cells returned in the query. Δ

Here is an example output for a 9.2 server.

```
Query Statement "FROM STAR SELECT"
Query type is Multidimensional
Query size in cells = 1
The query is running.
```

For 9.2 OLAP Servers, the following statistics are also available:

Last Update Time indicates the last time the query was referenced.

Query ID is the query ID.

Read Cells indicates the number of cells read.

Total Cells indicates the total number of cells.

Note: For SAS versions earlier than 9.2, this statistic is equivalent to the **Size** statistic. Δ

Here is an example output for a 9.2 OLAP server.

```
Query Statement "FROM SALES SELECT"
Query ID: E8900912-3A26-4A06-9315-34A111AF57D0
Query type is Flattened
Query size in MB = 693464
Total cells = 1 and cells read = 1
The time of last data set access is 12Nov2007:16:51:21
```

Listing Session Rowsets on an OLAP Server

You can list information about session queries with the LIST ROWSETS statement.

If the SESSION option is specified, then the rowsets associated with a particular user are listed. Otherwise, all queries are listed. The LIST ROWSETS statement is

only available for SAS 9.2 OLAP servers. It is not available for OLAP servers prior to SAS 9.2.

LIST ROWSETS [SESSION=*owner*];

The LIST ROWSETS statement generates the following statistics:

Cube names	lists the names of the cubes opened in this rowset.
Last Update Time	indicates the last time the rowset was referenced.
Number of open cubes	indicates the number of cubes opened in this rowset.
Rowset ID	is the rowset ID.
Rowset name	indicates the type name of the rowset. The following are possible values for Rowset name statistic.
	ACTIONSET
	CATALOGSET
	CUBESET
	DIMENSIONSET
	FUNCTIONSET
	HIERARCHYSET
	LEVELSET
	MEMBERSET
	MEASURESET
	PROPERTYSET
	SCHEMASET
	SETSET

Here are some examples of session rowset output from an OLAP server.

```
Rowset name "CUBESET"
Rowset ID: 09FDD38D-237C-42AD-BD92-1126ADEC7046
The time of last rowset access is 13Nov2007:14:03:40
```

```
Rowset name "DIMENSIONSET"
Rowset ID: E9B538AA-D80E-40E5-9FD2-8C374EE1B37F
The time of last rowset access is 13Nov2007:14:03:41
Cube MAILORDERINFO is being accessed by this rowset.
```

Closing a Session

You can close an OLAP session using the CLOSE SESSION statement.

CLOSE SESSION *id* | **_ALL_**;

The *id* is the unique session identifier returned from the LIST SESSIONS statement. The session will be closed on the server. If the **_ALL_** keyword is specified, then all open sessions are closed. Any queries that are running for an OLAP session will be cancelled when that session is closed.

Cancelling a Query Result Set

You can cancel (close) an OLAP query result set with the CANCEL QUERY statement.

```
CANCEL QUERY id |_ALL_;
```

The *id* is the unique result set identifier returned from the LIST QUERIES statement. Cancelling a query will stop the query processing and close the result set on the server. If the *_ALL_* keyword is specified, then all running queries will be cancelled.

Disabling a Cube

You can disable a cube by using the DISABLE CUBE statement. When you disable a cube, the server will not accept any *new* queries against that cube. Existing queries will continue to run to completion. In addition, the cube will not show up in the list of available cubes in the client session.

```
DISABLE CUBE cube-name |_ALL_;
```

The statement disallows new MDX queries which reference the specified cube (or all cubes if *_ALL_* is specified). It does not terminate currently executing queries. Those queries will either finish on their own, or the sessions can be terminated using the CLOSE SESSION command. If a query is executed on a disabled cube, the following message is displayed.

```
ERROR: The cube cannot be opened
```

Enabling a Cube

You can enable access to a disabled cube by using the ENABLE CUBE statement. The statement allows queries to be run against the cube again.

```
ENABLE CUBE cube-name |_ALL_;
```

If the *_ALL_* keyword is specified, all cubes are enabled for access by OLAP sessions.

Enabling and Disabling Cubes for a Cube Update

After a cube has been updated, the updates can be seen only after all open sessions have finished using the cube. After that, the next open session will display the new updates. In a busy environment, with numerous sessions, this could take considerable time. Executing the DISABLE and then ENABLE statements forces new sessions to immediately access the new data. In addition, if a cube is actively being queried, it must be disabled before you can use the DELETE or DELETE_PHYSICAL PROC OLAP options. For more information on updating cubes, see Chapter 7, “Updating SAS OLAP Cubes,” on page 79 and the ADD_DATAPROC OLAP option.

Refreshing a Cube

You can refresh a cube on an OLAP Server by using the REFRESH CUBE statement.

REFRESH CUBE *cube-name* |_ALL_;

The statement allows the server to access global calculated members and named sets that were defined or deleted since the cube header was loaded into the servers cube cache. This includes removing access to any global definitions which have been deleted since the server was started. If the `_ALL_` keyword is specified, all cubes on the OLAP server are refreshed.


Note: The REFRESH CUBE statement will not show changes to a cube that were made with the UPDATE_IN_PLACE or the ADD_DATA PROC OLAP option. You must disable and then enable a cube to see cube updates. △

Stopping an OLAP Server

You can stop a running OLAP Server by using the STOP SERVER statement.

STOP SERVER;

The STOP SERVER statement provides a clean shutdown of the current OLAP server. After the STOP SERVER statement is complete, the only statement which can be executed is a CONNECT statement to connect to another server.



APPENDIX

3

Resources and Reference Topics for SAS OLAP Cubes

<i>OLAP Server Resources</i>	299
<i>OLAP Schema Resources</i>	300
<i>Cube Resources</i>	300
<i>Table Resources</i>	300
<i>Dimension Resources</i>	301
<i>Level Resources</i>	302
<i>Hierarchy Resources</i>	302
<i>Measure Resources</i>	303
<i>Member and Member Properties Resources</i>	303
<i>Calculated Members and Measures Resources</i>	303
<i>Aggregation Resources</i>	304
<i>Incremental Cube Update Resources</i>	304
<i>Drill-Through to Detail Resources</i>	305
<i>GIS Mapping Resources</i>	305

OLAP Server Resources

The following are various links and references to documentation and Help topics that are related to the SAS OLAP Server.

- *SAS OLAP Server: User's Guide* Topics
 - "SAS OLAP Server" on page 9
- SAS OLAP Cube Studio Help Topics:
 - SAS Servers
 - Defining a SAS OLAP Server
 - Securing a SAS OLAP Server
 - OLAP Schema Properties
- Other Resources:
 - "Administering SAS OLAP Servers" in the *SAS Intelligence Platform: Application Server Administration Guide*
 - "Enabling ARM Logging in SAS OLAP Servers" in the *SAS Intelligence Platform: System Administration Guide*

OLAP Schema Resources

The following are various links and references to documentation and Help topics that are related to SAS OLAP schemas.

- *SAS OLAP Server: User's Guide* Topics
 - "OLAP Schemas" on page 14
- SAS OLAP Cube Studio Help Topics
 - Editing an OLAP Schema
 - Changing the OLAP Schema for a Cube
 - Deleting an OLAP Schema
 - OLAP Schema Properties
- OLAP Procedure
 - OLAP_SCHEMA argument on the "METASVR Statement" on page 256

Cube Resources

The following are various links and references to documentation and Help topics that are related to SAS OLAP cubes.

- *SAS OLAP Server: User's Guide* Topics
 - "What Is a Cube?" on page 3
 - "Understanding the Cube Structure" on page 4
 - "Why You Should Use Cubes" on page 5
 - "Cube Designer Wizard" on page 38
 - "Deleting Cubes and Cube Objects" on page 57
 - "Disabling and Enabling Cubes" on page 72
- SAS OLAP Cube Studio Help Topics
 - SAS OLAP and Cubes Overview
 - Creating Cubes – Cube Designer Wizard
 - Editing a Cube in SAS OLAP Cube Studio
 - SAS OLAP Cube Studio – View Cube
 - Cube Properties – Structure Tab
- Cube Examples
 - "Building a Cube from a Detail Table" on page 96
 - "Building a Cube from a Star Schema" on page 114
 - "Building a Cube from a Summary Table" on page 133
 - "Viewing a Cube in SAS OLAP Cube Studio" on page 192
- OLAP Procedure
 - "PROC OLAP Statement" on page 245

Table Resources

The following are various links and references to documentation and Help topics that are related to data tables.

- *SAS OLAP Server: User's Guide* Topics
 - “Data Tables” on page 11
 - “Data Tables Used to Define SAS OLAP Cubes” on page 24
 - “Multiple Language Support and Dimension Table Translations” on page 65
- SAS OLAP Cube Studio Help Topics
 - Data Tables Used to Define Cubes
 - Defining Tables Used to Build Cubes
 - Viewing Table Data in SAS OLAP Cube Studio
 - Renaming Tables
 - Deleting Tables
 - Cube Designer – Input
 - Cube Designer – Dimension Tables
 - Multiple Language Support for Cubes
 - Cube Designer – Drill-Through
 - Cube Designer – Aggregation Tables
 - Cube Designer – Table Options
 - Editing the Display Settings for the View Data panel
- Cube Examples
 - “Building a Cube from a Detail Table” on page 96
 - “Building a Cube from a Star Schema” on page 114
 - “Building a Cube from a Summary Table” on page 133
- OLAP Procedure
 - The following options are used on the “PROC OLAP Statement” on page 245:
 - DATA | FACT=
 - DATAPATH=
 - DRILLTHROUGH_TABLE | DT_TABLE | DT_TBL=
 - The following options are used on the “DIMENSION Statement” on page 257:
 - DIMKEY=
 - DIMTBL=
 - DIMTABLELIBREF=
 - DIMTABLEMEMMPREF=
 - FACTKEY=
 - TABLE= option on the “AGGREGATION Statement” on page 274

Dimension Resources

The following are various links and references to documentation and Help topics that are related to cube dimensions.

- *SAS OLAP Server: User's Guide* Topics
 - “Defining Dimensions and Levels” on page 40
 - “Specifying the TIME Dimension” on page 40
 - “Specifying GIS Map Information for a Dimension” on page 41

- SAS OLAP Cube Studio Help Topics
 - Defining the Number of Dimensions and Hierarchies For a Cube
 - Cube Designer – Dimensions
 - Dimension Designer – General
 - Specifying a Time Dimension For a Cube
 - Specifying GIS Map Information for a Dimension
- Cube Examples
 - “Setting Member Authorizations On A Dimension” on page 180
 - “Creating a Time Dimension in SAS OLAP Cube Studio” on page 197
 - “Specifying an ESRI GIS Map For a Cube Dimension” on page 200
- OLAP Procedure
 - “DIMENSION Statement” on page 257

Level Resources

The following are various links and references to documentation and Help topics that are related to cube levels.

- *SAS OLAP Server: User’s Guide* Topics
 - “Defining Dimensions and Levels” on page 40
- SAS OLAP Cube Studio Help Topics
 - Dimension Designer – Add Levels
 - Dimension Designer – Level
- Cube Examples
 - “Building a Cube from a Detail Table” on page 96
 - “Creating a Time Dimension in SAS OLAP Cube Studio” on page 197
- OLAP Procedure
 - “LEVEL Statement” on page 263

Hierarchy Resources

The following are various links and references to documentation and Help topics that are related to cube hierarchies.

- *SAS OLAP Server: User’s Guide* Topics
 - “Defining Cube Hierarchies” on page 42
- SAS OLAP Cube Studio Help Topics
 - Defining the Number of Dimensions and Hierarchies For a Cube
 - Dimension Designer – Hierarchy
 - Ragged and Unbalanced Hierarchies in SAS OLAP Cubes
 - Specifying Multiple Hierarchies For a Cube
- Cube Examples
 - “Creating Multiple Hierarchies For a Cube” on page 205
 - “Building a Cube from a Detail Table” on page 96

- OLAP Procedure
 - “HIERARCHY Statement” on page 267

Measure Resources

The following are various links and references to documentation and Help topics that are related to cube measures.

- *SAS OLAP Server: User’s Guide* Topics
 - “Defining Measures For A Cube” on page 47
- SAS OLAP Cube Studio Help Topics
 - Defining Measures For a Cube
 - Cube Designer – Select Measures
 - Cube Designer – Measure Details
 - Statistics Available for Measures
 - SAS Formats Available for Measures
- Cube Examples
 - “Building a Cube from a Detail Table” on page 96
 - “Building a Cube from a Summary Table” on page 133
- OLAP Procedure
 - “MEASURE Statement” on page 269

Member and Member Properties Resources

The following are various links and references to documentation and Help topics that are related to cube members and member properties.

- *SAS OLAP Server: User’s Guide* Topics
 - “Defining Member Properties” on page 49
- SAS OLAP Cube Studio Help Topics
 - Defining Member Properties For a Cube
 - Cube Designer – Member Properties
- Cube Examples
 - “Building a Cube from a Detail Table” on page 96
- OLAP procedure
 - “PROPERTY Statement” on page 266
 - UPDATE_DIMENSION option on the “DIMENSION Statement” on page 257

Calculated Members and Measures Resources

The following are various links and references to documentation and Help topics that are related to calculated members and measures for cubes.

- *SAS OLAP Server: User’s Guide* Topics
 - “Specifying Calculated Members and Measures” on page 63

- SAS OLAP Cube Studio Help Topics
 - Specifying Calculated Members and Measures
- Cube Examples
 - “Adding Calculated Members To A Cube” on page 170
- OLAP Procedure
 - “DEFINE Statement” on page 276
 - “UNDEFINE Statement” on page 278

Aggregation Resources

The following are various links and references to documentation and Help topics that are related to cube aggregations.

- *SAS OLAP Server: User’s Guide* Topics
 - “Aggregation Design” on page 26
 - “Aggregation Storage” on page 28
 - “Defining Aggregations While Building a Cube” on page 50
 - “Tuning Cube Aggregations” on page 58
 - “Specifying Tuning and Performance Options in Cube Aggregations” on page 61
 - “Coalescing Cube Aggregations” on page 82
- SAS OLAP Cube Studio Help Topics
 - Tuning SAS OLAP Cube Aggregations
 - Aggregation Tuning
 - Defining Aggregations While Building a Cube
 - Coalesce Incremental Data - Aggregations
 - Performance Options – Default
 - Performance Options – Single Aggregation
- Cube Examples
 - “ Tuning Aggregations For a Cube” on page 151
- OLAP Procedure
 - “AGGREGATION Statement” on page 274
 - COALESCE_AGGREGATIONS
 - “DROP_AGGREGATION Statement” on page 276

Incremental Cube Update Resources

The following are various links and references to documentation and Help topics that are related to incremental cube updates.

- *SAS OLAP Server: User’s Guide* Topics
 - Chapter 7, “Updating SAS OLAP Cubes,” on page 79
 - Appendix 2, “The OLAPOPERATE Procedure,” on page 291
- SAS OLAP Cube Studio Help Topics
 - Incrementally Updating SAS OLAP Cubes

- Coalesce Incremental Data – Aggregations
 - Cube Properties – Generations tab
- Cube Examples
 - “Adding Data to a Cube with Cube Update” on page 160
- OLAP Procedure
 - The following options are used on the “PROC OLAP Statement” on page 245:
 - ADD_DATA
 - COALESCE_AGGREGATIONS
 - OUTSCHEMA=
 - OUTCUBE=
 - UPDATE_DISPLAY_NAMES
 - UPDATE_INPLACE
 - NONUPDATEABLE
 - The following options are used on the “DIMENSION Statement” on page 257
 - UPDATE_DIMENSION
 - NONUPDATEABLE

Drill-Through to Detail Resources


The following are various links and references to documentation and Help topics that are related to drill-through capabilities for a cube.

- *SAS OLAP Server: User’s Guide* Topics
 - “Defining Drill-Through Tables” on page 39
 - The following products enable you to drill through to the detail data for a cube:
 - “SAS Enterprise Guide” on page 226
 - “SAS Web Report Studio” on page 227
 - “SAS Web OLAP Viewer for Java” on page 228
 - “Microsoft Excel Pivot Tables and Pivot Charts” on page 229
- SAS OLAP Cube Studio Help Topics
 - Data Tables Used to Define Cubes
 - Cube Designer – Drill-Through
- Other Resources
 - “Display Detail Data for a Large Cube” in the *SAS Intelligence Platform: Data Administration Guide*
- Cube Examples
 - “Implementing Drill-through to Detail Data in a SAS OLAP Cube” on page 210
- OLAP Procedure
 - DRILLTHROUGH_TABLE= option on the “PROC OLAP Statement” on page 245

GIS Mapping Resources

The following are various links and references to documentation and Help topics that are related to the GIS mapping of a cube dimension.

- *SAS OLAP Server: User's Guide* Topics
 - “Specifying GIS Map Information for a Dimension” on page 41
- SAS OLAP Cube Studio Help Topics
 - Specifying GIS Map Information for a Dimension
 - GIS Maps
 - Test GIS Mappings
- Other Resources
 - “Configuring the ESRI Map Component” in the *SAS Intelligence Platform: Web Application Administration Guide*
- Cube Examples
 - “Specifying an ESRI GIS Map For a Cube Dimension” on page 200
- OLAP Procedure
 - MAP_SERVICE= option on the “DIMENSION Statement” on page 257



APPENDIX

4

SAS OLAP Cube Studio Accessibility Features

SAS OLAP Cube Studio Accessibility Features 307

SAS OLAP Cube Studio Accessibility Features

SAS OLAP Cube Studio includes the following accessibility and compatibility features that improve usability of the product for users with disabilities. These features are related to accessibility standards for electronic information technology that were adopted by the U.S. Government under Section 508 of the U.S. Rehabilitation Act of 1973, as amended.

For further information on accessibility features in SAS, see *Accessibility Features in SAS under Windows* and *Default Key Settings for Interactive SAS Sessions under Windows* in SAS OnlineDoc.

Note: If you have questions or concerns about the accessibility of SAS products, send e-mail to accessibility@sas.com. Δ


Table A4.1 Window Manipulation Keys

Window Manipulation Key	Action
ALT+F4	Closes active applications on the desktop. ALT+F4 Closes SAS OLAP Cube Studio.
ALT+F6	Switches to the next window between modeless secondary windows and their primary window.
ALT+SPACEBAR	Opens the title bar icon and displays the program menu of the leftmost icon on the title bar.
ALT+TAB	Provides a pop-up menu of active applications, identified by their product graphics.
ALT+TAB+SHIFT	Reverses direction through the pop-up menu.
ALT+HYPHEN	Displays the shortcut menu for the active MDI window.
ALT+ESC	Navigates active applications on the desktop and then desktop shortcut bar items. Releasing the ESC key selects an application.

Window Manipulation Key	Action
ALT+ESC+SHIFT	Reverses the navigation direction of active applications on the desktop and the desktop shortcut bar. Navigates windows in reverse order.
ALT+ENTER	Displays the properties of a selected item while you are working in a window.
PRINT SCREEN	Copies a screen image to the clipboard.
ALT+ PRINT SCREEN	Copies an active window image to the Windows clipboard.

Table A4.2 Actions Performed with the ALT Key.

Action	Procedure
Move the application window	In SAS OLAP Cube Studio, select the ALT key and perform the following steps: <ol style="list-style-type: none"> 1 Choose Move from title bar menu. 2 Use the arrow keys to move the window. 3 Press ENTER to accept or ESC to cancel.
Size the application window	In SAS OLAP Cube Studio, select the ALT key and perform the following steps: <ol style="list-style-type: none"> 1 Choose Size from title bar menu. 2 Use an arrow key to choose which window border to move. 3 Resize the window with the appropriate arrow keys. 4 Press ENTER to accept or ESC to cancel.
Minimize/Maximize the application window	In SAS OLAP Cube Studio, perform the following steps: <ol style="list-style-type: none"> 1 Choose Minimize from the title bar menu. 2 Choose Maximize from the title bar menu.
Restore the application window	In SAS OLAP Cube Studio, select the ALT key and select Restore from title bar menu. The application is returned to its original state.



APPENDIX

5

Recommended Reading

Recommended Reading 309

Recommended Reading

Below is the recommended reading list for SAS OLAP Server:

- SAS OLAP Server: MDX Guide*
- SAS Providers for OLE DB: Cookbook*
- SAS Language Reference: Concepts*
- SAS Language Reference: Dictionary*
- SAS Intelligence Platform: Overview*
- SAS Intelligence Platform: Administration Guide*
- SAS Intelligence Platform: Data Administration Guide*
- SAS Intelligence Platform: System Administration Guide*
- SAS Intelligence Platform: Application Server Administration Guide*
- SAS Intelligence Platform: Security Administration Guide*
- SAS Intelligence Platform: Web Application Administration Guide*
- SAS Intelligence Platform: 9.1.3 to 9.2 Migration Guide*
- SAS Intelligence Platform: Desktop Application Administration Guide*
- Getting Started with SAS Enterprise Guide*
- SAS Information Map Studio: Getting Started with SAS Information Maps*
- SAS Web Report Studio: User's Guide*
- SAS Add-In for Microsoft Office: Getting Started with Data Analysis*
- SAS 9.2 Web OLAP Viewer for Java: Help*
- SAS 9.2 OLAP Server Monitor: Help*
- SAS Management Console main application: Help*
- SAS Companion that is specific to your operating environment

For a complete list of SAS publications, see the current *SAS Publishing Catalog*. To order the most current publications or to receive a free copy of the catalog, contact a SAS representative at

SAS Publishing Sales
 SAS Campus Drive
 Cary, NC 27513

Telephone: (800) 727-3228*

Fax: (919) 677-8166

E-mail: **sasbook@sas.com**

Web address: **support.sas.com/pubs**

* For other SAS Institute business, call (919) 677-8000.

Customers outside the United States should contact their local SAS office.

Glossary

aggregation

a summary of detail data that is stored with or referred to by a cube. Aggregations support rapid and efficient answers to business questions.

aggregation table

a table that contains pre-calculated totals. Aggregation tables can be referred to by cubes, reducing the amount of time that is required for building the cubes.

Application Response Measurement

the name of an application programming interface that was developed by an industry partnership and which is used to monitor the availability and performance of software applications. ARM monitors the application tasks that are important to a particular business. Short form: ARM.

ARM

See Application Response Measurement.

base table

a table that contains detail data that is used for building cubes or aggregation tables.

calculated member

in a dimension, a member whose value is derived from the values of other members.

child

within a dimension hierarchy, a descendant in level n-1 of a member that is at level n. For example, if a Geography dimension includes the levels Country and City, then Bangkok would be a child of Thailand, and Hamburg would be a child of Germany.

connection profile

a client-side definition of where a metadata server is located. The definition includes a computer name and a port number. In addition, the connection profile can also contain user connection information.

cube

a logical set of data that is organized and structured in a hierarchical, multidimensional arrangement. A cube is a directory structure, not a single file. A cube includes measures, and it can have numerous dimensions and levels of data.

data cleansing

the process of eliminating inaccuracies, irregularities, and discrepancies from data.

data scrubbing

another term for data cleansing. See also data cleansing.

data sparsity

a characteristic of a multidimensional data source in which there is a relatively high proportion of empty cells (which indicate missing data values) to filled cells.

data warehouse

a collection of data that is extracted from one or more sources for the purpose of querying and analysis.

descendant

in a dimension hierarchy, a member that resides at a lower level in relation to other members in the hierarchy. For example, if a Geography dimension includes the levels Country, State, and City, then California and Los Angeles would be descendants of USA.

detail data

nonsummarized (or partially summarized) factual information that pertains to a single area of interest, such as sales figures, inventory data, or human-resource data.

dimension

a group of closely related hierarchies. Hierarchies within a dimension typically represent different groupings of information that pertains to a single concept. For example, a Time dimension might consist of two hierarchies: (1) Year, Month, Date, and (2) Year, Week, Day. See also hierarchy.

dimension table

in a star schema, a table that contains the data for one of the dimensions. The dimension table is connected to the star schema's fact table by a primary key. The dimension table contains fields for each level of each hierarchy that is included in the dimension.

drill down

in a view of an OLAP cube, to start at one level of a dimension hierarchy and to click through one or more lower levels until you reach the data that you are interested in.

drill up

in a view of an OLAP cube, to start at one level of a dimension hierarchy and to click through one or more higher levels until you reach the level of summarized data that you are interested in.

drill-through table

a view, data set, or other data file that contains data that is used to define a cube. Drill-through tables can be used by client applications to provide a view from processed data into the underlying data source.

fact

a single piece of factual information in a data table. For example, a fact can be an employee name, a customer's phone number, or a sales amount. It can also be a derived value such as the percentage by which total revenues increased or decreased from one year to the next.

fact table

the central table in a star schema. The fact table contains the individual facts that are being stored in the database as well as the keys that connect each fact to the appropriate value in each dimension.

foreign key

a column or combination of columns in one table that references the corresponding primary key in another table. A foreign key must have the same attributes as the primary key that it references.

granularity

the relative level of detail that a data item represents. From the top of a dimension to the bottom, granularity increases. For example, in a Time dimension that consists of a Year-Month-Day hierarchy, Month is more granular than Year, and Day is more granular than Month.

hierarchy

an arrangement of members of a dimension into levels that are based on parent-child relationships. Members of a hierarchy are arranged from more general to more specific. For example, in a Time dimension, a hierarchy might consist of the members Year, Quarter, Month, and Day. In a Geography dimension, a hierarchy might consist of the members Country, State or Province, and City. More than one hierarchy can be defined for a dimension. Each hierarchy provides a navigational path that enables users to drill down to increasing levels of detail. See also member and level.

HOLAP

See hybrid online analytical processing.

hybrid online analytical processing

a type of OLAP in which relational OLAP (ROLAP) and multidimensional OLAP (MOLAP) are combined. In HOLAP, the source data is usually stored using a ROLAP strategy, and aggregations are stored using a MOLAP strategy. This combination usually results in the smallest amount of storage space. In HOLAP, aggregates can be pre-calculated and can be linked into a hybrid storage model. Short form: HOLAP.

informat

a pattern or set of instructions that SAS uses to determine how data values in an input file should be interpreted. SAS provides a set of standard informats and also enables you to define your own informats.

leaf member

the lowest-level member of a hierarchy. Leaf members do not have any child members.

level

an element of a dimension hierarchy. Levels describe the dimension from the highest (most summarized) level to the lowest (most detailed) level. For example, possible levels for a Geography dimension are Country, Region, State or Province, and City.

logical server

in the SAS Metadata Server, the second-level object in the metadata for SAS servers. A logical server specifies one or more of a particular type of server component, such as one or more SAS Workspace Servers.

MDDB

See multidimensional database.

MDX language

See multidimensional expressions language.

measure

a special dimension that contains summarized numeric data values that are analyzed. Total Sales and Average Revenue are examples of measures. For example, you might drill down within the Clothing hierarchy of the Product dimension to see the value of the Total Sales measure for the Shirts member.

member

a name that represents a particular data element within a dimension. For example, September 1996 might be a member of the Time dimension. A member can be either unique or non-unique. For example, 1997 and 1998 represent unique members in the

Year level of a Time dimension. January represents non-unique members in the Month level, because there can be more than one January in the Time dimension if the Time dimension contains data for more than one year.

metadata server

a server that provides metadata management services to one or more client applications. A SAS Metadata Server is an example.

MOLAP

See multidimensional online analytical processing.

multidimensional database

another term for cube. Short form: MDDB. See also cube.

multidimensional expressions language

a standardized, high-level language that is used to query multidimensional data sources. The MDX language is the multidimensional equivalent of SQL (Structured Query Language). Short form: MDX language.

multidimensional online analytical processing

a type of OLAP that stores aggregates in multidimensional database structures. Short form: MOLAP.

multi-threading

See threading.

navigate

to purposefully move from one view of the data in a table (or in some other data structure, such as a cube) to another. Drilling down and drilling up are two examples of navigation.

NWAY aggregation

the aggregation that has the minimum set of dimension levels that is required for answering any business question. The NWAY aggregation is the aggregation that has the finest granularity. See also granularity.

ODBO

See OLE DB for OLAP.

OLAP

See online analytical processing.

OLAP schema

a group of cubes. A cube is assigned to an OLAP schema when it is created, and an OLAP schema is assigned to a SAS OLAP Server when the server is defined in the metadata. A SAS OLAP Server can access only the cubes that are in its assigned OLAP schema.

OLE DB for OLAP

an OLAP API that is used to link OLAP clients and servers by means of a multidimensional expressions (MDX) language. Short form: ODBO. See also multidimensional expressions language.

online analytical processing

a software technology that enables users to dynamically analyze data that is stored in multidimensional database (MDDB) tables. Short form: OLAP.

parallel I/O

a method of input and output that takes advantage of multiple CPUs and multiple controllers, with multiple disks per controller to read or write data in independent threads.

parallel processing

a method of processing that divides a large job into several smaller jobs that can be executed in parallel on multiple CPUs.

parent

within a dimension hierarchy, the ancestor in level *n* of a member in level *n-1*. For example, if a Geography dimension includes the levels Country and City, then Thailand would be the parent of Bangkok, and Germany would be the parent of Hamburg. The parent value is usually a consolidation of all of its children's values.

primary key

a column or combination of columns that uniquely identifies a row in a table.

reach-through

the act of retrieving and displaying to a user the (unsummarized) detail data from which the summarized data in a multidimensional database is derived, when that detail data is stored in a separate data repository.

relational online analytical processing

a type of OLAP in which the multidimensional data is stored in a relational database. Short form: ROLAP.

ROLAP

See relational online analytical processing.

roll up

to summarize (or apply some other type of calculation or formula to) data values at one level of a dimension hierarchy in order to derive values for a parent level. For example, sales figures for January can be rolled up to Quarter1, and employee data for one department can be rolled up to the division level.

SAS application server

a server that provides SAS services to a client. In the SAS Open Metadata Architecture, the metadata for a SAS application server specifies one or more server components that provide SAS services to a client.

SAS ARM interface

an interface that can be used to monitor the performance of SAS applications. In the SAS ARM interface, the ARM API is implemented as an ARM agent. In addition, SAS supplies ARM macros, which generate calls to the ARM API function calls, and ARM system options, which enable you to manage the ARM environment and to log internal SAS processing transactions. See also ARM (Application Response Measurement).

SAS format

a pattern or set of instructions that SAS uses to determine how the values of a variable (or column) should be written or displayed. SAS provides a set of standard formats and also enables you to define your own formats.

SAS informat

a pattern or set of instructions that SAS uses to determine how data values in an input file should be interpreted. SAS provides a set of standard informats and also enables you to define your own informats.

SAS Management Console

a Java application that provides a single user interface for performing SAS administrative tasks.

SAS Metadata Server

a multi-user server that enables users to read metadata from or write metadata to one or more SAS Metadata Repositories. The SAS Metadata Server uses the

Integrated Object Model (IOM), which is provided with SAS Integration Technologies, to communicate with clients and with other servers.

SAS OLAP Cube Studio

a Java interface for defining and building OLAP cubes in SAS System 9 or later. Its main feature is the Cube Designer wizard, which guides you through the process of registering and creating cubes.

SAS OLAP Server

a SAS server that provides access to multidimensional data. The data is queried using the multidimensional expressions (MDX) language.

SAS Open Metadata Architecture

a general-purpose metadata management facility that provides metadata services to SAS applications. The SAS Open Metadata Architecture enables applications to exchange metadata, which makes it easier for these applications to work together.

Scalable Performance Data Engine

a SAS engine that is able to deliver data to applications rapidly because it organizes the data into a streamlined file format. The SPD Engine divides a problem (such as a WHERE clause) into smaller problems that can be processed in parallel. Short form: SPD Engine. See also parallel processing.

schema

a map or model of the overall data structure of a database. An OLAP schema specifies which group of cubes an OLAP server can access.

scrubbing

another term for data cleansing. See also data cleansing.

slice

a subset of data from a cube, where the data in the slice pertains to one or more members of one or more dimensions. For example, from a cube that contains data about customer feedback, one slice might pertain to feedback on one particular product (one member of the Product dimension). Another slice might pertain to feedback on that product from customers residing in particular geographic areas who submitted their feedback during a certain time period (one member of the Product dimension, multiple members of the Geography dimension, one or more members of the Time dimension).

SMP

See symmetric multiprocessing.

sparsity

See data sparsity.

SPD Engine

See Scalable Performance Data Engine.

SQL

See Structured Query Language.

star schema

tables in a database in which a single fact table is connected to multiple dimension tables. This is visually represented in a star pattern. SAS OLAP cubes can be created from a star schema.

stored statistics

statistics that are stored in a cube. Stored statistics can be used to derive higher-level statistics. Examples include sum, minimum, and maximum.

Structured Query Language

a standardized, high-level query language that is used in relational database management systems to create and manipulate database management system objects. Short form: SQL.

symmetric multiprocessing

a hardware and software architecture that can improve the speed of I/O and processing. An SMP machine has multiple CPUs and a thread-enabled operating system. An SMP machine is usually configured with multiple controllers and with multiple disk drives per controller. Short form: SMP.

thread

a single path of execution of a process in a single CPU, or a basic unit of program execution in a thread-enabled operating system. In an SMP environment, which uses multiple CPUs, multiple threads can be spawned and processed simultaneously. Regardless of whether there is one CPU or many, each thread is an independent flow of control that is scheduled by the operating system. See also symmetric multiprocessing, thread-enabled operating system, and threading.

thread-enabled operating system

an operating system that can coordinate symmetric access by multiple CPUs to a shared main memory space. This coordinated access enables threads from the same process to share data very efficiently.

threading

a high-performance method of data I/O or data processing in which the I/O or processing is divided into multiple threads that are executed in parallel. In the boss-worker model of threading, the same code for the I/O or calculation process is executed simultaneously in separate threads on multiple CPUs. In the pipeline model, a process is divided into steps, which are then executed simultaneously in separate threads on multiple CPUs. See also parallel I/O, parallel processing, and SMP (symmetric multiprocessing).

Time dimension

a dimension that divides time into levels such as Year, Quarter, Month, and Day.

tuple

a data object that contains two or more components. In OLAP, a tuple is a slice of data from a cube. It is a selection of members (or cells) across dimensions in a cube. It can also be viewed as a cross-section of member data in a cube. For example, ([time].[all time].[2003], [geography].[all geography].[u.s.a.], [measures].[actualsum]) is a tuple that contains data from the Time, Geography, and Measures dimensions.

wizard

an interactive utility program that consists of a series of dialog boxes, windows, or pages. Users supply information in each dialog box, window, or page, and the wizard uses that information to perform a task.

Index

A

ADO MD
 cubes with 225
 AGGR_COLUMN= option
 MEASURE statement (OLAP) 272
 AGGREGATION statement
 OLAP procedure 274
 aggregations 2
 adding to cubes 287
 deleting from cubes 276, 287
 ANALYSIS argument
 MEASURE statement (OLAP) 271

C

calculated members
 adding to cubes 63
 Calculated Members plug-in 63
 calculation types 64
 custom calculations 65
 simple calculations 64
 time analysis calculations 64
 CAPTION= option
 DIMENSION statement (OLAP) 258
 HIERARCHY statement (OLAP) 269
 LEVEL statement (OLAP) 263
 MEASURE statement (OLAP) 272
 PROPERTY statement (OLAP) 266
 character encoding
 SAS servers and 66, 281
 COLUMN argument
 MEASURE statement (OLAP) 271
 COLUMN= option
 PROPERTY statement (OLAP) 267
 COMPACT_NWAY option
 PROC OLAP statement 246
 COMPRESS option
 AGGREGATION statement (OLAP) 274
 OLAP procedure 62
 PROC OLAP statement 246
 CONCURRENT= option
 OLAP procedure 62
 PROC OLAP statement 247
 cube aggregations
 performance options 61
 tuning options 61
 Cube Designer
 building cubes from detail tables 96
 defining member properties 50

defining multiple hierarchies 43
 defining ragged and unbalanced hierarchies 44
 dimension table translations 66
 setting tuning and performance options 61
 cube metadata
 refreshing 57
 CUBE= option
 PROC OLAP statement 247
 cube structure 4
 cubes 3
 accessing from SQL Pass-Through facility 67
 adding aggregations to 287
 adding calculated members 63
 adding system options to 66
 ADO MD with 225
 building from detail tables 96
 building from existing definition 286
 building from star schema 114
 building from summary tables 133
 data management 6
 data storage with 5
 default hierarchy 42
 defining member properties 49
 deleting 288
 deleting aggregations from 276, 287
 loading from detail tables 281
 loading from star schema 283
 loading with summarized data 285
 maintenance 6
 multi-threading and 5
 multiple hierarchies 42
 OLE DB for OLAP with 225
 ragged and unbalanced hierarchies 44
 SAS products with 226
 saving OLAP procedure code 108
 setup 6
 third-party clients with 229
 updating 56

D

data access 1
 data management 6
 DATA= option
 PROC OLAP statement 247
 data storage 1
 cubes for 5
 data warehouse 2

DATAPATH= option
 AGGREGATION statement (OLAP) 275
 OLAP procedure 62
 PROC OLAP statement 248
 default hierarchy 42
 DEFAULT option
 HIERARCHY statement (OLAP) 269
 MEASURE statement (OLAP) 273
 DEFINE statement
 OLAP procedure 276
 DELETE option
 PROC OLAP statement 248
 DELETE_PHYSICAL option
 PROC OLAP statement 249
 DESC= option
 DIMENSION statement (OLAP) 259
 HIERARCHY statement (OLAP) 269
 LEVEL statement (OLAP) 264
 MEASURE statement (OLAP) 273
 PROC OLAP statement 249
 PROPERTY statement (OLAP) 267
 detail tables
 building cubes from 96
 loading cubes from 281
 DIMENSION statement
 OLAP procedure 257
 dimension table translations 65
 dimensions
 GIS map information for 41
 multiple hierarchies for 42
 ragged and unbalanced hierarchies for 44
 DIMKEY= option
 DIMENSION statement (OLAP) 259
 DIMTABLELIBREF= option
 DIMENSION statement (OLAP) 260
 DIMTABLEMEMMPREF= option
 DIMENSION statement (OLAP) 260
 DIMITBL= option
 DIMENSION statement (OLAP) 259
 DRILLTHROUGH_TABLE= option
 PROC OLAP statement 249
 DROP_AGGREGATION statement
 OLAP procedure 61, 276

E

EMPTY= option
 LEVEL statement (OLAP) 264
 EMPTY_CHAR= option
 HIERARCHY statement (OLAP) 260
 PROC OLAP statement 250
 EMPTY_NUM= option
 HIERARCHY statement (OLAP) 260
 PROC OLAP statement 250

F

FACT= option
 PROC OLAP statement 247
 FACTKEY= option
 DIMENSION statement (OLAP) 261
 FORMAT= option
 MEASURE statement (OLAP) 273

G

GIS map information 41

H

hierarchies
 default hierarchy 42
 multiple hierarchies for a dimension 42
 ragged and unbalanced 44
 HIERARCHIES= argument
 DIMENSION statement (OLAP) 258
 HIERARCHY= option
 PROPERTY statement (OLAP) 267
 HIERARCHY statement
 OLAP procedure 43, 267
 HOST= option
 METASVR statement (OLAP) 257
 hyper-cubes
 See cubes

I

IGNORE_EMPTY option
 HIERARCHY statement (OLAP) 261
 LEVEL statement (OLAP) 265
 INDEX option
 AGGREGATION statement (OLAP) 275
 OLAP procedure 63
 PROC OLAP statement 251
 INDEXPATH= option
 AGGREGATION statement (OLAP) 275
 OLAP procedure 63
 PROC OLAP statement 251
 INDEXSORTSIZE= option
 OLAP procedure 63
 PROC OLAP statement 251

L

language support 65
 LEVEL= argument
 PROPERTY statement (OLAP) 266
 LEVEL statement
 OLAP procedure 263
 LEVELS= argument
 HIERARCHY statement (OLAP) 268
 loading cubes
 from detail tables 281
 from star schema 283
 with summarized data 285

M

MAXTHREADS= option
 OLAP procedure 63
 PROC OLAP statement 252
 MEASURE statement
 OLAP procedure 269
 MEMBER argument
 DEFINE statement (OLAP) 278
 UNDEFINE statement 279
 member properties 49
 members
 unique names 46

- metadata
 - refreshing cube metadata 57
- METASVR statement
 - OLAP procedure 256
- multi-cubes
 - See cubes
- multi-threading 5
- multiple hierarchies 42
- multiple language support 65

N

- NAME= argument
 - DROP_AGGREGATION statement (OLAP) 276
- NAME= option
 - AGGREGATION statement (OLAP) 275
- NO_NWAY option
 - PROC OLAP statement 252

O

- OLAP 1
 - benefits of 2
 - data storage and access 1
 - variations of 3
- OLAP procedure 244
 - AGGREGATION statement 274
 - building cubes 110
 - COMPRESS option 62
 - CONCURRENT= option 62
 - DATAPATH= option 62
 - DEFINE statement 276
 - defining member properties 50
 - defining multiple hierarchies 43
 - defining ragged and unbalanced hierarchies 45
 - DIMENSION statement 257
 - dimension table translations 66
 - DROP_AGGREGATION statement 61, 276
 - HIERARCHY statement 43, 267
 - INDEX option 63
 - INDEXPATH= option 63
 - INDEXSORTSIZE= option 63
 - LEVEL statement 263
 - MAXTHREADS= option 63
 - MEASURE statement 269
 - METASVR statement 256
 - overview 244
 - PARTSIZE= option 63
 - performance options 289
 - PROC OLAP statement 245
 - PROPERTY statement 50, 266
 - ragged hierarchy options 288
 - REFRESH statement 58
 - refreshing cube metadata 58
 - saving code for cubes 108
 - SEGSIZE= option 63
 - setting tuning and performance options 62
 - syntax 244
 - tuning cube aggregations 61
 - UNDEFINE statement 278
 - USER_DEFINED_TRANSLATIONS statement 66, 279
 - WORKPATH= option 63
- OLAP_SCHEMA= argument
 - METASVR statement (OLAP) 256
- OLE DB for OLAP
 - cubes with 225

- Online Analytical Processing
 - See OLAP

P

- PARTSIZE= option
 - AGGREGATION statement (OLAP) 275
 - OLAP procedure 63
 - PROC OLAP statement 253
- PATH= option
 - PROC OLAP statement 253
- performance
 - cube aggregations 61
 - OLAP procedure options for 289
- PORT= option
 - METASVR statement (OLAP) 257
- PROC OLAP statement 245
- PROPERTY statement
 - OLAP procedure 50, 266
- PW= option
 - METASVR statement (OLAP) 257

R

- ragged hierarchies 44
 - OLAP procedure options for 288
 - unique member names and 46
- REFRESH statement
 - OLAP procedure 58
- refreshing cube metadata 57
- REGISTER_ONLY option
 - PROC OLAP statement 254
- REPOSITORY= option
 - METASVR statement (OLAP) 257

S

- SAS Enterprise Guide 226
- SAS Information Map Studio 227
- SAS OLAP Cube Studio
 - building cubes 96
- SAS OLAP Server
 - monitoring performance 61
- SAS servers
 - character encoding and 66, 281
- SAS Web OLAP Viewer 228
- SAS Web Report Studio 227
- SEGSIZE= option
 - AGGREGATION statement (OLAP) 275
 - OLAP procedure 63
 - PROC OLAP statement 254
- SET argument
 - DEFINE statement (OLAP) 278
 - UNDEFINE statement 279
- SORT_ORDER= option
 - DIMENSION statement (OLAP) 261
 - LEVEL statement (OLAP) 265
- SQL Pass-Through facility
 - accessing cubes from 67
 - conversion issues 67
 - example 69
 - SQL procedure syntax 69
- SQL procedure 69
- star schemas
 - building cubes from 114
 - loading cubes from 283

- STAT= argument
 - MEASURE statement (OLAP) 269
- summarized data
 - loading cubes with 285
- summary tables
 - building cubes from 133
- system options
 - adding to cubes 66

T

- TABLE= option
 - AGGREGATION statement (OLAP) 275
- third-party clients
 - cubes with 229
- time analysis calculations 64
- tuning cube aggregations
 - DROP_AGGREGATION statement for 61
 - monitoring OLAP server performance 61
 - specifying tuning options 61
- TYPE= option
 - DIMENSION statement (OLAP) 262
 - LEVEL statement (OLAP) 265

U

- unbalanced hierarchies 44
- UNDEFINE statement
 - OLAP procedure 278
- unique member names 46
- UNITS= option
 - MEASURE statement (OLAP) 274
- USER_DEFINED_TRANSLATIONS statement
 - OLAP procedure 66, 279
- USERID= option
 - METASVR statement (OLAP) 257

V

- VALIDVARNAME= system option 68

W

- WORKPATH= option
 - OLAP procedure 63
 - PROC OLAP statement 255

Your Turn

We welcome your feedback.

- If you have comments about this book, please send them to yourturn@sas.com. Include the full title and page numbers (if applicable).
- If you have comments about the software, please send them to suggest@sas.com.

SAS® Publishing Delivers!

Whether you are new to the work force or an experienced professional, you need to distinguish yourself in this rapidly changing and competitive job market. SAS® Publishing provides you with a wide range of resources to help you set yourself apart. Visit us online at support.sas.com/bookstore.

SAS® Press

Need to learn the basics? Struggling with a programming problem? You'll find the expert answers that you need in example-rich books from SAS Press. Written by experienced SAS professionals from around the world, SAS Press books deliver real-world insights on a broad range of topics for all skill levels.

support.sas.com/saspress

SAS® Documentation

To successfully implement applications using SAS software, companies in every industry and on every continent all turn to the one source for accurate, timely, and reliable information: SAS documentation. We currently produce the following types of reference documentation to improve your work experience:

- Online help that is built into the software.
- Tutorials that are integrated into the product.
- Reference documentation delivered in HTML and PDF – free on the Web.
- Hard-copy books.

support.sas.com/publishing

SAS® Publishing News

Subscribe to SAS Publishing News to receive up-to-date information about all new SAS titles, author podcasts, and new Web site features via e-mail. Complete instructions on how to subscribe, as well as access to past issues, are available at our Web site.

support.sas.com/spn



THE
POWER
TO KNOW®

