



THE
POWER
TO KNOW.

SAS/ACCESS[®] 9.1.3

for Relational Databases

Reference

Fifth Edition

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2007. *SAS/ACCESS® 9.1.3 for Relational Databases: Reference, Fifth Edition*. Cary, NC: SAS Institute Inc.

SAS/ACCESS® 9.1.3 for Relational Databases: Reference, Fifth Edition

Copyright © 2007, SAS Institute Inc., Cary, NC, USA

ISBN 978-1-59994-584-2

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 printing, December 2007

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/pubs 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>	vii
Overview	vii
All Supported SAS/ACCESS Interfaces to Relational Databases	vii
SAS/ACCESS Interface to DB2 under UNIX and PC Hosts	viii
SAS/ACCESS Interface to DB2 under z/OS	ix
SAS/ACCESS Interface to HP Neoview	ix
SAS/ACCESS Interface to Informix	ix
SAS/ACCESS Interface to Microsoft SQL Server	x
SAS/ACCESS Interface to MySQL	x
SAS/ACCESS Interface to Netezza	x
SAS/ACCESS Interface to ODBC	x
SAS/ACCESS Interface to OLE DB	xi
SAS/ACCESS Interface to Oracle	xi
SAS/ACCESS Interface to Sybase	xi
SAS/ACCESS Interface to Teradata	xii
Documentation Enhancements	xii

PART 1 **Concepts 1**

Chapter 1 \triangle **Overview of the SAS/ACCESS Interface to Relational Databases 3**

About This Document	3
Methods for Accessing Relational Database Data	3
Selecting a SAS/ACCESS Method	4
SAS Views of DBMS Data	6

Chapter 2 \triangle **SAS Names and Support for DBMS Names 7**

Introduction to SAS/ACCESS Naming	7
SAS Naming Conventions	8
SAS/ACCESS Default Naming Behaviors	9
Renaming DBMS Data	10
Options That Affect SAS/ACCESS Naming Behavior	11
Naming Behavior When Retrieving DBMS Data	12
Naming Behavior When Creating DBMS Objects	12
SAS/ACCESS Naming Examples	13

Chapter 3 \triangle **Data Integrity and Security 21**

Introduction to Data Integrity and Security	21
DBMS Security	21
SAS Security	22
Potential Result Set Differences When Processing Null Data	27

Chapter 4 \triangle **Performance Considerations 31**

Increasing Throughput of the SAS Server	31
Limiting Retrieval	31
Repeatedly Accessing Data	33
Sorting DBMS Data	33
Temporary Table Support for SAS/ACCESS	34
Chapter 5 \triangle Optimizing Your SQL Usage	37
Overview of Optimizing Your SQL Usage	37
Passing Functions to the DBMS Using PROC SQL	38
Passing Joins to the DBMS	38
Passing the DELETE Statement to Empty a Table	40
When Passing Joins to the DBMS Will Fail	40
Passing DISTINCT and UNION Processing to the DBMS	42
Optimizing the Passing of WHERE Clauses to the DBMS	42
Using the DBINDEX=, DBKEY=, and MULTI_DATASRC_OPT= Options	43
Chapter 6 \triangle Threaded Reads	47
Overview of Threaded Reads in SAS/ACCESS	47
Underlying Technology of Threaded Reads	48
SAS/ACCESS Interfaces and Threaded Reads	48
Scope of Threaded Reads	49
Options That Affect Threaded Reads	49
Generating Trace Information for Threaded Reads	50
Performance Impact of Threaded Reads	53
Autopartitioning Techniques in SAS/ACCESS	53
Data Ordering in SAS/ACCESS	54
Two-Pass Processing for SAS Threaded Applications	54
When Threaded Reads Will Not Occur	55
Summary of Threaded Reads	55
Chapter 7 \triangle How SAS/ACCESS Works	57
Introduction to How SAS/ACCESS Works	57
How the SAS/ACCESS LIBNAME Statement Works	58
How the Pass-Through Facility Works	59
How the ACCESS Procedure Works	60
How the DBLOAD Procedure Works	62
PART 2 General Reference	63
Chapter 8 \triangle SAS/ACCESS Features by Host	65
Introduction	65
SAS/ACCESS Interface to DB2 under UNIX and PC Hosts: Supported Features	65
SAS/ACCESS Interface to DB2 under z/OS: Supported Features	66
SAS/ACCESS Interface to HP Neoview: Supported Features	66
SAS/ACCESS Interface to Informix: Supported Features	67
SAS/ACCESS Interface to Microsoft SQL Server: Supported Features	67

SAS/ACCESS Interface to MySQL: Supported Features	68
SAS/ACCESS Interface to Netezza: Supported Features	68
SAS/ACCESS Interface to ODBC: Supported Features	69
SAS/ACCESS Interface to OLE DB: Supported Features	70
SAS/ACCESS Interface to Oracle: Supported Features	70
SAS/ACCESS Interface to Sybase: Supported Features	71
SAS/ACCESS Interface to Teradata: Supported Features	72
Chapter 9 \triangle The LIBNAME Statement for Relational Databases	73
Overview of the LIBNAME Statement for Relational Databases	73
Assigning a Libref Interactively	74
LIBNAME Options for Relational Databases	78
Chapter 10 \triangle Data Set Options for Relational Databases	157
Overview of Data Set Options for Relational Databases	157
Chapter 11 \triangle Macro Variables and System Options for Relational Databases	269
Introduction to Macro Variables and System Options for Relational Databases	269
Macro Variables for Relational Databases	269
System Options for Relational Databases	271
Chapter 12 \triangle The Pass-Through Facility for Relational Databases	285
Overview of the SQL Procedure's Interactions with SAS/ACCESS	285
Syntax for the Pass-Through Facility for Relational Databases	286

PART 3 **Converting SAS/ACCESS Descriptors to PROC SQL Views** **297**

Chapter 13 \triangle The CV2VIEW Procedure	299
Overview of the CV2VIEW Procedure	299
Procedure Syntax	300
CV2VIEW Procedure Examples	304

PART 4 **Appendixes** **309**

Appendix 1 \triangle The ACCESS Procedure for Relational Databases	311
Overview of the ACCESS Procedure for Relational Databases	311
Procedure Syntax	313
Using Descriptors with the ACCESS Procedure	325
Examples of Using the ACCESS Procedure	327
Appendix 2 \triangle The DBLOAD Procedure for Relational Databases	329
Overview of the DBLOAD Procedure for Relational Databases	329
Procedure Syntax	331
Example of Using the DBLOAD Procedure	342
Appendix 3 \triangle Recommended Reading	343
Recommended Reading	343

Glossary **345**

Index **351**

What's New

Overview

SAS/ACCESS Interfaces for Relational Databases has the following changes and enhancements:

- Beginning with SAS 9.1.3 Service Pack 4, SAS/ACCESS Interface for HP Neoview and SAS/ACCESS Interface for Netezza are new, including documentation enhancements.
- Beginning with SAS 9.1.3, SAS/ACCESS supports additional hosts for existing DBMSs.
- Beginning with SAS 9.1.2, SAS/ACCESS Interface for MySQL databases is new.
- Beginning with SAS 9.0, you can use threaded reads to complete jobs in substantially less time than if each task is handled sequentially.
- The new CV2VIEW procedure converts SAS/ACCESS view descriptors into SQL views.

All Supported SAS/ACCESS Interfaces to Relational Databases

- You can use the CHANGE statement to rename SAS/ACCESS tables“Renaming SAS/ACCESS Tables” on page 10.
- The CV2VIEW procedure converts SAS/ACCESS view and access descriptors to the SAS 9.0 format. It can also convert a view descriptor to a SAS 9.0 SQL view. As SAS/ACCESS moves forward with LIBNAME enhancements and tighter integration with the SAS Open Metadata Repository, SAS/ACCESS views no longer is the method of choice.
- The DBMASTER= data set option designates which table is the master table when you are processing a join that involves tables from two different types of databases.
- DBMS metadata can now be accurately maintained within the SAS Open Metadata Repository.
- You can now encode the DBMS password that appears in SAS source code so that it does not appear as text in SAS programs.

- You can use the `DIRECT_EXE= LIBNAME` option to pass an SQL statement directly to a database by using explicit pass-through when you use `PROC SQL` with a libref.
- The new SQL `MULTI_DATASRC_OPT= LIBNAME` option and the `DBMASTER=` data set option optimize the performance of the SQL procedure. For more detailed information, see the passing joins to the DBMS, determining when joins will fail, and optimizing `WHERE` clauses topics.
- You can use the `MULTI_DATASRC_OPT= LIBNAME` option instead of the `DBKEY=` data set option to improve performance when you process a join between two data sources.
- The `SASTRACE=` system option now provides improved debugging capabilities.
- With temporary table support, DBMS temporary tables can persist from one SAS step to the next. It involves establishing a SAS connection to the DBMS that persists across SAS procedures and `DATA` steps.
- Threaded reads divide resource-intensive tasks into multiple independent units of work and execute those units in parallel.
- SAS/ACCESS supports these hosts:
 - AIX (RS/6000) for MySQL (beginning with SAS 9.1.3 Service Pack 2)
 - HP for Itanium for Sybase (beginning with SAS 9.1.3 Service Pack 2) and Teradata (beginning with SAS 9.1.3 Service Pack 3)
 - Linux for Intel for MySQL (beginning with SAS 9.1.3 Service Pack 1) and Teradata (beginning with SAS 9.1.3)
 - Linux for Itanium for DB2, Informix, Microsoft SQL Server, MySQL, ODBC, Oracle, and Sybase, beginning with SAS 9.1.3 Service Pack 1
 - 64-bit Solaris systems are supported for Teradata (beginning with SAS 9.1.3 Service Pack 4)
 - 64-bit Windows for Oracle and DB2—in addition to 64-bit UNIX, which was provided in SAS 8.2
- Beginning with SAS 9.0, support for these items is discontinued:
 - SAS/ACCESS Interface to CA-OpenIngres
 - SAS/ACCESS Interface to Oracle Rdb under OpenVMS Alpha
 - OS/2, OpenVMS VAX, MIPS ABI, Intel ABI, UNIX MP-RAS, and CMS operating environments
 - CV2ODBC procedure

SAS/ACCESS Interface to DB2 under UNIX and PC Hosts

The following options are new:

- The `DBSLICE=` data set option specifies user-supplied `WHERE` clauses to partition a DBMS query into component queries for threaded reads.
- The data set and `LIBNAME` options control the scope of DBMS-threaded reads and the number of threads.
- The `IGNORE_READ_ONLY_COLUMNS=` data set and `LIBNAME` options specify whether to ignore or include columns whose data types are read-only when generating an SQL statement for inserts or updates.
- The `SQL_FUNCTIONS= LIBNAME` option specifies that the SQL functions that match the functions that SAS supports are passed to the DBMS for processing.

SAS/ACCESS Interface to DB2 under z/OS

Note: z/OS is the successor to the OS/390 operating system. SAS/ACCESS 9.1 (and later) for z/OS is supported on both OS/390 and z/OS operating systems. Throughout this document, any reference to z/OS also applies to OS/390, unless otherwise stated. Δ

SAS/ACCESS Interface to DB2 under z/OS features stored procedure support that includes passing input parameters, retrieving output parameters into SAS macro variables, and retrieving result sets into SAS tables.

The following options are new:

- The BL_DB2CURSOR= data set option specifies a string that contains a valid DB2 SELECT statement that points to either local or remote objects (tables or views). After your database administrator populates the communication database with the appropriate entries, you can select data from a remote location to load DB2 tables directly from other DB2 and non-DB2 objects.
- The BL_DB2LDCT3= data set option specifies a string in the LOAD utility control statement, following the field specification.
- The DBSLICE= data set option specifies user-supplied WHERE clauses to partition a DBMS query into component queries for threaded reads.
- The DBSLICEPARAM= data set and LIBNAME options control the scope of DBMS threaded reads and the number of threads.
- The DEGREE= LIBNAME option determines whether DB2 uses parallelism.
- The REMOTE_DBTYPE= LIBNAME option ensures that the SQL that some SAS procedures use to access the DB2 catalog tables is generated properly, based on the database server type.
- The TRAP151= data set option removes columns that cannot be updated from a FOR UPDATE OF clause so that column update can continue.

SAS/ACCESS Interface to HP Neoview

SAS/ACCESS Interface to HP Neoview is a new database engine that runs on specific UNIX and Windows platforms. SAS/ACCESS Interface to HP Neoview: Supported Features” on page 66 provides direct, transparent access to HP Neoview databases through LIBNAME statements or the Pass-Through Facility. You can use various LIBNAME statement options and data set options that the LIBNAME engine supports to control the data that is returned to SAS.

SAS/ACCESS Interface to Informix

The following options are new:

- The DBSLICE= data set option specifies user-supplied WHERE clauses to partition a DBMS query into component queries for threaded reads.
- The DBSLICEPARAM= data set and LIBNAME options control the scope of DBMS threaded reads and the number of threads.

SAS/ACCESS Interface to Microsoft SQL Server

The following options are new:

- The DBSLICE= data set option specifies user-supplied WHERE clauses to partition a DBMS query into component queries for threaded reads.
- The DBSLICEPARM= data set and LIBNAME options control the scope of DBMS threaded reads and the number of threads.
- The ERRLIMIT= LIBNAME option specifies the number of errors that are allowed while using the Fastload utility before SAS stops loading data to Teradata.
- The IGNORE_READ_ONLY_COLUMNS= data set and LIBNAME options specify whether to ignore or include columns where data types are read-only when generating an SQL statement for inserts or updates.

SAS/ACCESS Interface to MySQL

Beginning with SAS 9.1.3 Service Pack 4, you can use the INSERTBUFF= data set and LIBNAME options to specify the number of rows in a single insert operation.

Beginning with SAS 9.1.2, SAS/ACCESS Interface to MySQL is new. MySQL software is an open-source SQL database server that runs on 32-bit Windows systems, 64-bit HP systems, and 64-bit Solaris systems. SAS/ACCESS Interface to MySQL provides direct transparent access to MySQL databases through LIBNAME statements or the Pass-Through Facility. To control data that is returned to SAS, you can use various LIBNAME and data set options that the LIBNAME engine supports.

SAS/ACCESS Interface to Netezza

Beginning with SAS 9.1.3 Service Pack 4, SAS/ACCESS Interface to Netezza is a new database engine that runs on specific UNIX and Windows“SAS/ACCESS Interface to Netezza: Supported Features” on page 68 platforms. SAS/ACCESS Interface to Netezza provides direct, transparent access to Netezza databases through LIBNAME statements or the Pass-Through Facility. To control data that is returned to SAS, you can use various LIBNAME and data set options that the LIBNAME engine supports.

SAS/ACCESS Interface to ODBC

The following feature and options are new:

- ODBC 3.x standard API is supported.
- The DBSLICE= data set option specifies user-supplied WHERE clauses to partition a DBMS query into component queries for threaded reads.
- The DBSLICEPARM= data set and LIBNAME options control the scope of DBMS threaded reads and the number of threads.
- The IGNORE_READ_ONLY_COLUMNS= data set and LIBNAME options specify whether to ignore or include columns whose data types are read-only when generating an SQL statement for inserts or updates.
- The SQL_FUNCTIONS= LIBNAME option specifies that the SQL functions that match the functions that SAS supports by SAS are passed to the DBMS.

SAS/ACCESS Interface to OLE DB

The following options are new:

- The IGNORE_READ_ONLY_COLUMNS= data set and LIBNAME options specify whether to ignore or include columns whose data types are read-only when generating an SQL statement for inserts or updates.
- The INSERTBUFF= LIBNAME option specifies the number of rows in a single insert operation.
- The SQL_FUNCTIONS= LIBNAME option specifies that the SQL functions that match the functions that SAS supports are passed to the DBMS.

SAS/ACCESS Interface to Oracle

The following options are new:

- You can use the BL_INDEX_OPTIONS= data set option to specify SQL*Loader Index options with bulk loading. You can boost performance by specifying the SORTED INDEXES index option.
- The BL_PARFILE= data set option creates a file that contains the SQL*Loader command line options.
- You can use the BL_RECOVERABLE= data set option to specify whether the load process is recoverable. It enhances the performance of the bulk load.
- To improve performance, you can use the BL_SUPPRESS_NULLIF= data set option to specify whether to suppress the NULLIF clause for specified columns when a table is created.
- The DBSLICE= data set option specifies user-supplied WHERE clauses to partition a DBMS query into component queries for threaded reads.
- The DBSLICEPARAM= data set and LIBNAME options control the scope of DBMS threaded reads and the number of threads.
- You can use the OR_PARTITION= data set option to read, update, and delete data from a specific partition in a partitioned table, as well as insert and bulk-load data into a specific partition in a partitioned table. It also boosts performance.
- In addition to being a LIBNAME option, OR_UPD_NOWHERE= is now also a data set option.

SAS/ACCESS Interface to Sybase

The following options are new:

- The DBSLICE= data set option specifies user-supplied WHERE clauses to partition a DBMS query into component queries for threaded reads.
- The DBSLICEPARAM= data set and LIBNAME options control the scope of DBMS threaded reads and the number of threads.

SAS/ACCESS Interface to Teradata

The following features and options are new:

- The DBSLICE= data set option specifies user-supplied WHERE clauses to partition a DBMS query into component queries for threaded reads.
- The DBSLICEPARM= data set and LIBNAME options control the scope of DBMS threaded reads and the number of threads.
- The ERRLIMIT= LIBNAME option specifies the number of errors that are allowed while using the Fastload utility before SAS stops loading data to Teradata.
- The LOGDB= LIBNAME option specifies the name of an alternate database where the restart log tables from the Teradata FastExport utility should be created.
- You can extract data faster with the Teradata FastExport utility.
- SAS/ACCESS Interface to Teradata supports the MultiLoad bulk-load facility. MultiLoad greatly accelerates insertion of data into both empty and existing Teradata tables. The SAS/ACCESS MultiLoad facility is similar to the native Teradata MultiLoad. You invoke the MultiLoad facility with the MULTILOAD= data set option. MULTILOAD= works together with several other new MultiLoad data set options.
- Beginning with SAS 9.1.3 Service Pack 4, when processing WHERE statements that contain *literal* values for TIME or TIMESTAMP, the SAS engine passes the values to Teradata exactly as they were entered, without rounding or truncation.
- Beginning with SAS 9.1, specifying OBS=*n* causes SAS/ACCESS Interface to Teradata to append SAMPLE *n* to the SQL that Teradata generates.


Documentation Enhancements

Two new supplements are available:

- *SAS/ACCESS Supplement for Netezza (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for HP Neoview (SAS/ACCESS for Relational Databases)*


These supplements explain how SAS/ACCESS software is customized to accommodate the particular requirements and capabilities of each respective DBMS. When you use them with the separately available *SAS/ACCESS for Relational Databases: Reference*, they provide comprehensive guides for using SAS/ACCESS features and options for the DBMSs. These titles are available from SAS OnlineDoc 9.1.3 on the Web and are available for purchase from the SAS Publications Catalog.

In addition, the DISTRIBUTE_ON= data set option is new in *SAS/ACCESS Supplement for Netezza (SAS/ACCESS for Relational Databases)* and *SAS/ACCESS for Relational Databases: Reference*.



Concepts

<i>Chapter 1</i>	Overview of the SAS/ACCESS Interface to Relational Databases	3
<i>Chapter 2</i>	SAS Names and Support for DBMS Names	7
<i>Chapter 3</i>	Data Integrity and Security	21
<i>Chapter 4</i>	Performance Considerations	31
<i>Chapter 5</i>	Optimizing Your SQL Usage	37
<i>Chapter 6</i>	Threaded Reads	47
<i>Chapter 7</i>	How SAS/ACCESS Works	57



CHAPTER

1

Overview of the SAS/ACCESS Interface to Relational Databases

<i>About This Document</i>	3
<i>Methods for Accessing Relational Database Data</i>	3
<i>Selecting a SAS/ACCESS Method</i>	4
<i>Methods for Accessing DBMS Tables and Views</i>	4
<i>SAS/ACCESS LIBNAME Statement Advantages</i>	4
<i>Pass-Through Facility Advantages</i>	5
<i>SAS/ACCESS Features for Common Tasks</i>	5
<i>SAS Views of DBMS Data</i>	6

About This Document

This document provides conceptual, reference, and usage information for the SAS/ACCESS Interface to relational database management systems (DBMSs). The information in this document applies generally to all of the relational database management systems that are supported by SAS/ACCESS software. *Because the availability and behavior of SAS/ACCESS features vary from one interface to another, you should use this general document in conjunction with the documentation for your SAS/ACCESS interface.* There is an individual document for each supported DBMS, and those documents are sold separately.

This document is intended for applications programmers and end users who meet the following conditions:

- familiar with the basics of their DBMS and its SQL (Structured Query Language)
- know how to use their operating environment
- can use basic SAS commands and statements.

Database administrators might also want to read this document to understand how the interface is implemented and administered.

Methods for Accessing Relational Database Data

The SAS/ACCESS interface to relational databases is a family of interfaces (each of which is licensed separately) that enable you to interact with data in other vendors' databases from within SAS. SAS/ACCESS provides the following methods for accessing relational DBMS data:

- The LIBNAME statement enables you to assign SAS librefs to DBMS objects such as schemas and databases. After a database is associated with a libref, you can

use a SAS two-level name to specify any table or view in the database and then work with the table or view as you would with a SAS data set.

- The Pass-Through Facility enables you to interact with a data source using its native SQL syntax without leaving your SAS session. The SQL statements are passed directly to the data source for processing.
- The ACCESS and DBLOAD procedures support indirect access to DBMS data. These procedures are no longer the recommended method for accessing DBMS data, but they continue to be supported for the database systems and environments on which they were available for SAS Version 6.

See “Selecting a SAS/ACCESS Method” on page 4 for information about when to use each method.

Note: Not all SAS/ACCESS interfaces support all of these features. To determine the features that are available in your environment, see the about features by host section. \triangle

Selecting a SAS/ACCESS Method

Methods for Accessing DBMS Tables and Views

In SAS/ACCESS, there are often several ways to complete a task. For example, you can access DBMS tables and views by using the LIBNAME statement or the Pass-Through Facility. The advantages and limitations of these features are described below. Before processing complex or data-intensive operations, you might want to test several of these features to determine the most efficient feature for your particular task.

SAS/ACCESS LIBNAME Statement Advantages

It is generally recommended that you use the SAS/ACCESS LIBNAME statement to access your DBMS data because this is usually the fastest and most direct method. An exception to this is when you need to use non-ANSI standard SQL. ANSI standard SQL is required when you use the SAS/ACCESS library engine in the SQL procedure. The Pass-Through Facility, however, accepts all the extensions to SQL that are provided by your DBMS.

The SAS/ACCESS LIBNAME statement has the following advantages:

- Significantly fewer lines of SAS code are required to perform operations on your DBMS. For example, a single LIBNAME statement establishes a connection to your DBMS, enables you to specify how your data is processed, and enables you to easily view your DBMS tables in SAS.
- You do not need to know the SQL language of your DBMS in order to access and manipulate data on your DBMS. You can use SAS procedures, such as PROC SQL, or DATA step programming on any libref that references DBMS data. You can read, insert, update, delete, and append data, as well as create and drop DBMS tables by using SAS syntax.
- The LIBNAME statement provides more control over DBMS operations such as locking, spooling, and data type conversion through the use of LIBNAME and data set options.

- The engine can optimize the processing of joins and WHERE clauses by passing these operations directly to the DBMS. This takes advantage of your DBMS's indexing and other processing capabilities. For more information, see “Overview of Optimizing Your SQL Usage” on page 37.
- The engine can pass some functions directly to the DBMS for processing.

Pass-Through Facility Advantages

The Pass-Through Facility has the following advantages:

- Pass-Through Facility statements enable the DBMS to optimize queries, particularly when you join tables. The DBMS optimizer can take advantage of indexes on DBMS columns to process a query more quickly and efficiently.
- Pass-Through Facility statements enable the DBMS to optimize queries when the queries have summary functions (such as AVG and COUNT), GROUP BY clauses, or columns created by expressions (such as the COMPUTED function). The DBMS optimizer can use indexes on DBMS columns to process the queries more quickly.
- On some DBMSs, you can use Pass-Through Facility statements with SAS/AF applications to handle the transaction processing of the DBMS data. Using a SAS/AF application gives you complete control of COMMIT and ROLLBACK transactions. Pass-Through Facility statements give you better access to DBMS return codes.
- The Pass-Through Facility accepts all the extensions to ANSI SQL that are provided by your DBMS.

SAS/ACCESS Features for Common Tasks

The following table contains a list of tasks and the features that you can use to accomplish them.

Table 1.1 SAS/ACCESS Features for Common Tasks

Task	SAS/ACCESS Features
Read DBMS tables or views	LIBNAME statement*
	Pass-Through Facility
	View descriptors**
Create DBMS objects, such as tables	LIBNAME statement*
	DBLOAD procedure
	Pass-Through Facility's EXECUTE statement
Update, delete, or insert rows into DBMS tables	LIBNAME statement*
	View descriptors**
	Pass-Through Facility's EXECUTE statement
Append data to DBMS tables	DBLOAD procedure with APPEND option
	LIBNAME statement and APPEND procedure*
	Pass-Through Facility's EXECUTE statement

Task	SAS/ACCESS Features
List DBMS tables	LIBNAME statement and SAS Explorer window*
	LIBNAME statement and DATASETS procedure*
	LIBNAME statement and CONTENTS procedure*
	LIBNAME statement and SQL procedure dictionary tables*
Delete DBMS tables or views	LIBNAME statement and SQL procedure's DROP TABLE statement*
	LIBNAME statement and DATASETS procedure's DELETE statement*
	DBLOAD procedure with SQL DROP TABLE statement
	Pass-Through Facility's EXECUTE statement

* LIBNAME statement refers to the SAS/ACCESS LIBNAME statement.

** View descriptors refer to view descriptors that are created in the ACCESS procedure.

SAS Views of DBMS Data

SAS/ACCESS enables you to create a SAS view of data that exists in a relational database management system. A *SAS data view* defines a virtual data set that is named and stored for later use. A view contains no data, but rather describes data that is stored elsewhere. There are three types of SAS data views:

- *DATA step views* are stored, compiled DATA step programs.
- *SQL views* are stored query expressions that read data values from their underlying files, which can include SAS data files, SAS/ACCESS views, DATA step views, other SQL views, or relational database data.
- *SAS/ACCESS views* (also called view descriptors) describe data that is stored in DBMS tables. This is no longer a recommended method for accessing relational DBMS data. Use the CV2VIEW procedure to convert existing view descriptors into SQL views.


You can use all types of views as inputs into DATA steps and procedures. You can specify views in queries as if they were tables. A view derives its data from the tables or views that are listed in its FROM clause. The data accessed by a view is a subset or superset of the data in its underlying table(s) or view(s).

SQL views and SAS/ACCESS views can be used to update their underlying data if the view is based on only one DBMS table or is based on a DBMS view that is based on only one DBMS table, and if the view has no calculated fields. DATA step views cannot be used to update their underlying data; they can only read the data.

Your options for creating a SAS view of DBMS data are determined by the SAS/ACCESS feature that you are using to access the DBMS data. The following table lists the recommended methods for creating SAS views.

Table 1.2 Creating SAS Views

Feature You Use to Access DBMS Data	SAS View Technology You Can Use
SAS/ACCESS LIBNAME statement	SQL view or DATA step view of the DBMS table
SQL Procedure Pass-Through Facility	SQL view with CONNECTION TO component



CHAPTER

2

SAS Names and Support for DBMS Names

<i>Introduction to SAS/ACCESS Naming</i>	7
<i>SAS Naming Conventions</i>	8
<i>Length of Name</i>	8
<i>Case-sensitivity</i>	8
<i>SAS Name Literals</i>	9
<i>SAS/ACCESS Default Naming Behaviors</i>	9
<i>Modification and Truncation</i>	9
<i>ACCESS Procedure</i>	9
<i>DBLOAD Procedure</i>	10
<i>Renaming DBMS Data</i>	10
<i>Renaming SAS/ACCESS Tables</i>	10
<i>Renaming SAS/ACCESS Columns</i>	11
<i>Renaming SAS/ACCESS Variables</i>	11
<i>Options That Affect SAS/ACCESS Naming Behavior</i>	11
<i>Naming Behavior When Retrieving DBMS Data</i>	12
<i>Naming Behavior When Creating DBMS Objects</i>	12
<i>SAS/ACCESS Naming Examples</i>	13
<i>Replacing Unsupported Characters</i>	13
<i>Preserving Column Names</i>	14
<i>Preserving Table Names</i>	15
<i>Using DQUOTE=ANSI</i>	17
<i>Using Name Literals</i>	18
<i>Using DBMS Data to Create a DBMS Table</i>	19
<i>Using a SAS Data Set to Create a DBMS Table</i>	20

Introduction to SAS/ACCESS Naming

Because some DBMSs allow case-sensitive names and names with special characters, you must show special consideration when you use the names of DBMS objects (such as tables and columns) with SAS/ACCESS features. This section presents SAS/ACCESS naming conventions, default naming behaviors, options that can modify naming behavior, and usage examples. See the documentation for your SAS/ACCESS interface for information about how SAS handles your DBMS names.

SAS Naming Conventions

Length of Name

SAS naming conventions allow long names for SAS data sets and SAS variables. For example, MYDB.TEMP_EMPLOYEES_QTR4_2000 is a valid two-level SAS name for a data set.

The names of the following SAS language elements can be up to 32 characters in length:

- members of SAS libraries, including SAS data sets, data views, catalogs, catalog entries, and indexes
- variables in a SAS data set
- macros and macro variables.

The following SAS language elements have a maximum length of eight characters:

- librefs and filerefs
- SAS engine names
- names of SAS/ACCESS access descriptors and view descriptors
- variable names in SAS/ACCESS access descriptors and view descriptors.

For a complete description of SAS naming conventions, see the *SAS Language Reference: Dictionary*.

Case-sensitivity

When SAS encounters mixed-case or case-sensitive names in SAS code, it stores and displays the names as they are specified. If the SAS variables `Flight` and `dates` are defined in mixed case, for example,

```
input Flight $3. +3 dates date9.;
```

then SAS displays the variable names as defined. Note how the column headings appear as defined:

Output 2.1 Mixed-Case Names Displayed in Output

SAS System		
Obs	Flight	dates
1	114	01MAR2000
2	202	01MAR2000
3	204	01MAR2000

Although SAS *stores* variable names as they are defined, it *recognizes* variables for processing without regard to case. For example, SAS processes these variables as `FLIGHT` and `DATES`. Likewise, renaming the `Flight` variable to "flight" or "FLIGHT" would result in the same processing.

SAS Name Literals

A SAS *name literal* is a name token that is expressed as a quoted string, followed by the letter **n**. Name literals enable you to use special characters or blanks that are not otherwise allowed in SAS names when you specify a SAS data set or variable. Name literals are especially useful for expressing database column and tables names that contain special characters.

Examples of name literals are

```
data mydblib.'My Staff Table'n;
```

and

```
data Budget_for_1999;
```

```
input '$ Amount Budgeted'n 'Amount Spent'n;
```

Name literals are subject to certain restrictions:

- You can use a name literal only for SAS variable and data set names, statement labels, and DBMS column and table names.
- You can use name literals only in a DATA step or in the SQL procedure.
- If a name literal contains any characters that are not allowed when VALIDVARNAME=V7, then you must set the system option to VALIDVARNAME=ANY. For details about using the VALIDVARNAME= system option, see “VALIDVARNAME= System Option” on page 282.

SAS/ACCESS Default Naming Behaviors

Modification and Truncation

When SAS/ACCESS reads DBMS column names that contain characters that are not standard in SAS names, the default behavior is to replace an unsupported character with an underscore (_). For example, the DBMS column name Amount Budgeted\$ becomes the SAS variable name Amount_Budgeted_.

Note: Nonstandard names include those with blank spaces or special characters (such as @, #, %) that are not allowed in SAS names. Δ

When SAS/ACCESS encounters a DBMS name that exceeds 32 characters, it truncates the name.

After it has modified or truncated a DBMS column name, SAS appends a number to the variable name, if necessary, to preserve uniqueness. For example, DBMS column names MY\$DEPT, My\$Dept, and my\$dept become SAS variable names MY_DEPT, MY_Dept0, and my_dept1.

ACCESS Procedure

If you attempt to use long names in the ACCESS procedure, you get an error message advising you that long names are not supported. Long member names, such as access descriptor and view descriptor names, are truncated to eight characters. Long DBMS column names are truncated to 8-character SAS variable names within the SAS access descriptor. You can use the RENAME statement to specify 8-character SAS variable names, or you can accept the default truncated SAS variable names that are assigned by the ACCESS procedure.

The ACCESS procedure converts DBMS object names to uppercase characters unless they are enclosed in quotation marks. Any DBMS objects that are given lowercase names when they are created, or whose names contain special or national characters, must be enclosed in quotation marks.

DBLOAD Procedure

You can use long member names, such as the name of a SAS data set that you want to load into a DBMS table, in the DBLOAD procedure DATA= option. However, if you attempt to use long SAS variable names, you get an error message advising you that long variable names are not supported in the DBLOAD procedure. You can use the RENAME statement to rename the 8-character SAS variable names to long DBMS column names when you load the data into a DBMS table. You can also use the SAS data set option RENAME to rename the columns after they are loaded into the DBMS.

Most DBLOAD procedure statements convert lowercase characters in user-specified values and default values to uppercase. If your host or database is case sensitive and you want to specify a value that includes lowercase alphabetic characters (for example, a user ID or password), enclose the entire value in quotation marks. You must also put quotation marks around any value that contains special characters or national characters.

The only exception is the DBLOAD SQL statement. The DBLOAD SQL statement is passed to the DBMS exactly as you type it, with case preserved.

Renaming DBMS Data

Renaming SAS/ACCESS Tables

You can rename DBMS tables and views using the CHANGE statement, as shown in the following example:

```
proc datasets lib=x;
  change oldtable=newtable;
quit;
```

You can rename tables using this method for the following engines:

DB2 UNIX/PC

DB2 z/OS

HP Neoview

Informix

MySQL

Microsoft SQL
Server

Netezza

ODBC

OLE DB

Oracle

Sybase

Teradata

Note: If you change a table name, any view dependent on that table will no longer work, unless the view references the new table name. Δ

Renaming SAS/ACCESS Columns

You can use the RENAME statement to rename the 8-character default SAS variable names to long DBMS column names when you load the data into a DBMS table. You can also use the SAS data set option RENAME= to rename the columns after they are loaded into the DBMS.

Renaming SAS/ACCESS Variables

You can use the RENAME statement to specify 8-character SAS variable names such as access descriptors and view descriptors.

Options That Affect SAS/ACCESS Naming Behavior

To change how SAS handles case-sensitive or nonstandard DBMS table and column names, specify one or more of the following options.

PRESERVE_COL_NAMES=YES

is a SAS/ACCESS LIBNAME and data set option that applies *only* to creating DBMS tables. When set to YES, this option preserves spaces, special characters, and mixed case in DBMS column names. See “PRESERVE_COL_NAMES= LIBNAME Option” on page 130 for more information about this option.

PRESERVE_TAB_NAMES=YES

is a SAS/ACCESS LIBNAME option. When set to YES, this option preserves blank spaces, special characters, and mixed case in DBMS table names. See “PRESERVE_TAB_NAMES= LIBNAME Option” on page 132 for more information about this option.

Note: Specify the alias PRESERVE_NAMES=YES | NO if you plan to specify both the PRESERVE_COL_NAMES= and PRESERVE_TAB_NAMES= options in your LIBNAME statement. Using this alias saves time when you are coding. Δ

DQUOTE=ANSI

is a PROC SQL option. This option specifies whether PROC SQL treats values within *double* quotation marks as a character string or as a column name or table name. When you specify DQUOTE=ANSI, your SAS code can refer to DBMS names that contain characters and spaces that are not allowed by SAS naming conventions. Specifying DQUOTE=ANSI enables you to preserve special characters in table and column names in your SQL statements by enclosing the names in double quotation marks.

To preserve table names, you must also specify PRESERVE_TAB_NAMES=YES. To preserve column names when you create a table, you must also specify PRESERVE_COL_NAMES=YES.

VALIDVARNAME=ANY

is a global system option that can override the SAS naming conventions. See “VALIDVARNAME= System Option” on page 282 for information about this option.

The availability of these options and their default settings are DBMS-specific, so consult the SAS/ACCESS documentation for your DBMS to learn how the SAS/ACCESS engine for your DBMS processes names.

Naming Behavior When Retrieving DBMS Data

The following two tables illustrate how SAS/ACCESS processes DBMS names when retrieving data from a DBMS. This information applies generally to all the interfaces. In some cases, however, it is not necessary to specify these options because the option default values are DBMS-specific. See the documentation for your SAS/ACCESS interface for details.

Table 2.1 DBMS Column Names to SAS Variable Names When Reading DBMS Data

DBMS Column Name	Desired SAS Variable Name	Options
Case-sensitive DBMS column name, such as Flight	Case-sensitive SAS variable name, such as Flight	No options are necessary
DBMS column name with characters that are not valid in SAS names, such as My\$Flight	Case-sensitive SAS variable name where an underscore replaces the invalid characters, such as My_Flight	No options are necessary
DBMS column name with characters that are not valid in SAS names, such as My\$Flight	Nonstandard, case-sensitive SAS variable name, such as My\$Flight	PROC SQL DQUOTE=ANSI or, in a DATA or PROC step, use a SAS name literal such as 'My\$Flight'n and VALIDVARNAME=ANY

Table 2.2 DBMS Table Names to SAS Data Set Names When Reading DBMS Data

DBMS Table Name	Desired SAS Data Set Name	Options
Default DBMS table name, such as STAFF	Default SAS data set or member name (uppercase), such as STAFF	PRESERVE_TAB_NAMES=NO
Case-sensitive DBMS table name, such as Staff	Case-sensitive SAS data set, such as Staff	PRESERVE_TAB_NAMES=YES
DBMS table name with characters that are not valid in SAS names, such as All\$Staff	Nonstandard, case-sensitive SAS data set name, such as All\$Staff	PROC SQLDQUOTE=ANSI and PRESERVE_TAB_NAMES=YES or, in a DATA step or PROC, use a SAS name literal such as 'All\$Staff'n and PRESERVE_TAB_NAMES=YES

Naming Behavior When Creating DBMS Objects

The following two tables illustrate how SAS/ACCESS handles variable names when creating DBMS objects such as tables and views. This information applies generally to

all the interfaces. In some cases, however, it is not necessary to specify these options because the option default values are DBMS-specific. See the documentation for your DBMS for details.

Table 2.3 SAS Variable Names to DBMS Column Names When Creating Tables

SAS Variable Name as Input	Desired DBMS Column Name	Options
Any SAS variable name, such as Miles	Default DBMS column name (normalized to follow the DBMS's naming conventions), such as MILES	PRESERVE_COL_NAMES=NO
A case-sensitive SAS variable name, such as Miles	Case-sensitive DBMS column name, such as Miles	PRESERVE_COL_NAMES=YES
A SAS variable name with characters that are not valid in a normalized SAS name, such as Miles-to-Go	Case-sensitive DBMS column name that matches the SAS name, such as Miles-to-Go	PROC SQL DQUOTE=ANSI and PRESERVE_COL_NAMES=YES or, in a DATA or PROC step, use a SAS name literal and PRESERVE_COL_NAMES=YES and VALIDVARNAME=ANY

Table 2.4 SAS Data Set Names to DBMS Table Names

SAS Data Set Name as Input	Desired DBMS Table Name	Options
Any SAS data set name, such as Payroll	Default DBMS table name (normalized to follow the DBMS's naming conventions), such as PAYROLL	PRESERVE_TAB_NAMES=NO
A case-sensitive SAS data set name, such as Payroll	Case-sensitive DBMS table name, such as Payroll	PRESERVE_TAB_NAMES=YES
A case-sensitive SAS data set name with characters that are not valid in a normalized SAS name, such as Payroll-for-QC	Case-sensitive DBMS table name that matches the SAS name, such as Payroll-for-QC	PROC SQL DQUOTE=ANSI and PRESERVE_TAB_NAMES=YES or, in a DATA or PROC step, use a SAS name literal and PRESERVE_TAB_NAMES=YES

SAS/ACCESS Naming Examples

Replacing Unsupported Characters

In the following example, a view, myview, is created from the Oracle table mytable.

```
proc sql;
connect to oracle (user=testuser password=testpass);
create view myview as
  select * from connection to oracle
    (select "Amount Budgeted$", "Amount Spent$"
      from mytable);
quit;
```

```
proc contents data=myview;
run;
```

In the output produced by PROC CONTENTS, the Oracle column names (that were processed by the SQL view of MYTABLE) are renamed to different SAS variable names: Amount Budgeted\$ becomes Amount_Budgeted_ and Amount Spent\$ becomes Amount_Spent_.

Preserving Column Names

The following example uses the Oracle table PAYROLL to create a new Oracle table, PAY1, and then prints the table. Both the PRESERVE_COL_NAMES=YES and the PROC SQL DQUOTE=ANSI options are used to preserve the case and nonstandard characters in the column names. You do not need to quote the column aliases in order to preserve the mixed case. You only need double quotation marks when the column name has nonstandard characters or blanks.

By default, most SAS/ACCESS interfaces use DBMS-specific rules to set the case of table and column names. Therefore, even though the new Oracle table name pay1 is created in lowercase in this example, Oracle stores the name in uppercase as PAY1. If you want the table name to be stored as "pay1", you must set PRESERVE_TAB_NAMES=NO.

```
options linesize=120 pagesize=60 nodate;

libname mydblib oracle user=testuser password=testpass path='ora8_servr'
        schema=hrdept preserve_col_names=yes;

proc sql dquote=ansi;
create table mydblib.pay1 as
    select idnum as "ID #", sex, jobcode, salary,
           birth as BirthDate, hired as HiredDate
    from mydblib.payroll
    order by birth;

title "Payroll Table with Revised Column Names";
select * from mydblib.pay1;
quit;
```

SAS recognizes the JOBCODE, SEX, and SALARY column names, whether you specify them in your SAS code as lowercase, mixed case, or uppercase. In the Oracle table PAYROLL, the SEX, JOBCODE, and SALARY columns were created in uppercase; therefore, they retain this case in the new table (unless you rename them). A partial output from the example is shown:

Output 2.2 DBMS Table Created with Nonstandard and Standard Column Names

Payroll Table with Revised Column Names					
ID #	SEX	JOBCODE	SALARY	BirthDate	HiredDate
1118	M	PT3	11379	16JAN1944:00:00:00	18DEC1980:00:00:00
1065	M	ME2	35090	26JAN1944:00:00:00	07JAN1987:00:00:00
1409	M	ME3	41551	19APR1950:00:00:00	22OCT1981:00:00:00
1401	M	TA3	38822	13DEC1950:00:00:00	17NOV1985:00:00:00
1890	M	PT2	91908	20JUL1951:00:00:00	25NOV1979:00:00:00

Preserving Table Names

The following example uses PROC PRINT to print the DBMS table PAYROLL. The DBMS table was created in uppercase and since PRESERVE_TAB_NAMES=YES, the table name must be specified in uppercase. (If you set the PRESERVE_TAB_NAMES=NO, you can specify the DBMS table name in lowercase.) A partial output follows the example.

```
options nodate linesize=64;
libname mydblib oracle user=testuser password=testpass
        path='ora8_servr' preserve_tab_names=yes;

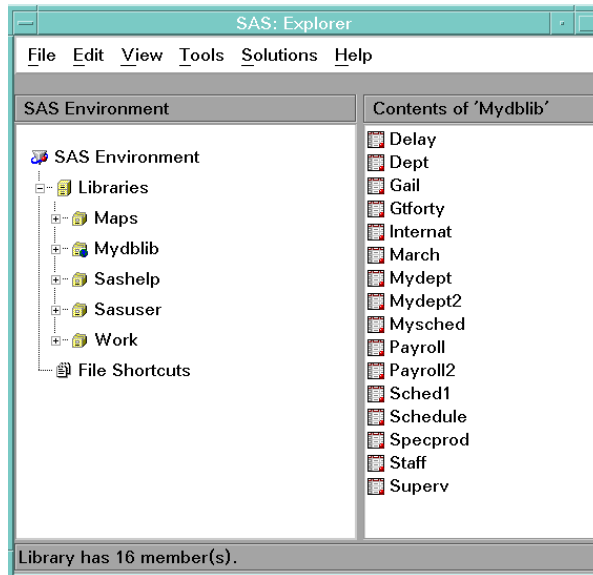
proc print data=mydblib.PAYROLL;
    title 'PAYROLL Table';
run;
```

Output 2.3 DBMS Table with a Case-sensitive Name

PAYROLL Table					
Obs	IDNUM	SEX	JOBCODE	SALARY	BIRTH
1	1919	M	TA2	34376	12SEP1960:00:00:00
2	1653	F	ME2	35108	15OCT1964:00:00:00
3	1400	M	ME1	29769	05NOV1967:00:00:00
4	1350	F	FA3	32886	31AUG1965:00:00:00
5	1401	M	TA3	38822	13DEC1950:00:00:00

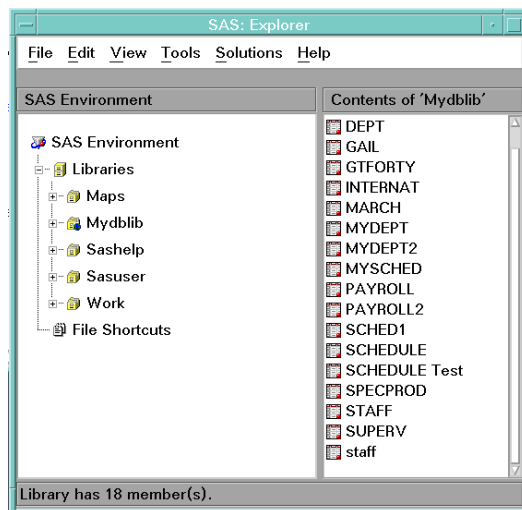
The following example submits a SAS/ACCESS LIBNAME statement and then opens the SAS Explorer window, which lists the Oracle tables and views that are referenced by the MYDBLIB libref. Notice that 16 members are listed and that all of the member names are in the case (initial capitalization) that is set by the Explorer window. The table names are capitalized because PRESERVE_TAB_NAMES= defaulted to NO.

```
libname mydblib oracle user=testuser pass=testpass;
```

Display 2.1 SAS Explorer Window Listing DBMS Objects

If you submit a SAS/ACCESS LIBNAME statement with `PRESERVE_TAB_NAMES=YES` and then open the SAS Explorer window, you see a different listing of the Oracle tables and views referenced by the MYDBLIB libref.

```
libname mydblib oracle user=testuser password=testpass
        preserve_tab_names=yes;
```

Display 2.2 SAS Explorer Window Listing Case-Sensitive DBMS Objects

Notice that there are 18 members listed, including one that is in lowercase and one that has a name separated by a blank space. Because `PRESERVE_TAB_NAMES=YES`, SAS displays the tables names in the exact case in which they were created.

Using DQUOTE=ANSI

The following example creates a DBMS table with a blank space in its name. Double quotation marks are used to specify the table name, International Delays. Both of the preserve names LIBNAME options are also set by using the alias PRESERVE_NAMES=. Because PRESERVE_NAMES=YES, the schema airport is now case-sensitive for Oracle.

```
options linesize=64 nodate;

libname mydblib oracle user=testuser password=testpass path='airdata'
        schema=airport preserve_names=yes;

proc sql dquote=ansi;
create table mydblib."International Delays" as
    select int.flight as "FLIGHT NUMBER", int.dates,
           del.orig as ORIGIN,
           int.dest as DESTINATION, del.delay
    from mydblib.INTERNAT as int,
         mydblib.DELAY as del
    where int.dest=del.dest and int.dest='LON';
quit;

proc sql dquote=ansi outobs=10;
    title "International Delays";
select * from mydblib."International Delays";
```

Notice that you use single quotation marks to specify the data value for London (`int.dest='LON'`) in the WHERE clause. Because of the preserve name LIBNAME options, using double quotation marks would cause SAS to interpret this data value as a column name.

Output 2.4 DBMS Table with Nonstandard Column Names

International Delays				
FLIGHT NUMBER	DATES	ORIGIN	DESTINATION	DELAY
219	01MAR1998:00:00:00	LGA	LON	18
219	02MAR1998:00:00:00	LGA	LON	18
219	03MAR1998:00:00:00	LGA	LON	18
219	04MAR1998:00:00:00	LGA	LON	18
219	05MAR1998:00:00:00	LGA	LON	18
219	06MAR1998:00:00:00	LGA	LON	18
219	07MAR1998:00:00:00	LGA	LON	18
219	01MAR1998:00:00:00	LGA	LON	18
219	02MAR1998:00:00:00	LGA	LON	18
219	03MAR1998:00:00:00	LGA	LON	18

If you query a DBMS table and use a label to change the FLIGHT NUMBER column name to a standard SAS name (Flight_Number), a label (enclosed in single quotation marks) changes the name only in the output. Because this column name and the table name, International Delays, each have a space in their names, you have to enclose the names in double quotation marks. A partial output follows the example.

```

options linesize=64 nodate;

libname mydblib oracle user=testuser password=testpass path='airdata'
        schema=airport preserve_names=yes;

proc sql dquote=ansi outobs=5;
    title "Query from International Delays";
select "FLIGHT NUMBER" label='Flight_Number', dates, delay
    from mydblib."International Delays";

```

Output 2.5 Query Renaming a Nonstandard Column to a Standard SAS Name

Query from International Delays		
Flight_	DATES	DELAY
Number		

219	01MAR1998:00:00:00	18
219	02MAR1998:00:00:00	18
219	03MAR1998:00:00:00	18
219	04MAR1998:00:00:00	18
219	05MAR1998:00:00:00	18

You can preserve special characters by specifying DQUOTE=ANSI and using double quotation marks around the SAS names in your SELECT statement.

```

proc sql dquote=ansi;
    connect to oracle (user=testuser password=testpass);
    create view myview as
        select "Amount Budgeted$", "Amount Spent$"
        from connection to oracle
            (select "Amount Budgeted$", "Amount Spent$"
            from mytable);
quit;
proc contents data=myview;
run;

```

Output from this example would show that Amount Budgeted\$ remains Amount Budgeted\$ and Amount Spent\$ remains Amount Spent\$.

Using Name Literals

The following example creates a table using name literals. You must specify the SAS option VALIDVARNAME=ANY in order to use name literals. Use PROC SQL to print the new DBMS table because name literals work only with PROC SQL and the DATA step. PRESERVE_COLUMN_NAMES=YES is required *only* because the table is being created with nonstandard SAS column names.

```

options ls=64 validvarname=any nodate;

libname mydblib oracle user=testuser password=testpass path='ora8servr'
        preserve_col_names=yes preserve_tab_names=yes ;

data mydblib.'Sample Table'n;
    'EmpID# 'n=12345;
    Lname='Chen';

```

```

Salary in $'n=63000;

proc sql;
  title "Sample Table";
  select * from mydblib.'Sample Table'n;

```

Output 2.6 DBMS Table to Test Column Names

Sample Table		
EmpID#	Lname	Salary in \$
12345	Chen	63000

Using DBMS Data to Create a DBMS Table

The following example uses PROC SQL to create a DBMS table based on data from other DBMS tables. You preserve the case-sensitivity of the aliased column names by using PRESERVE_COL_NAMES=YES. A partial output is displayed after the code.

```

libname mydblib oracle user=testuser password=testpass
  path='hrdata99' schema=personnel preserve_col_names=yes;

proc sql;
create table mydblib.gtforthy as
  select lname as LAST_NAME,
         fname as FIRST_NAME,
         salary as ANNUAL_SALARY
  from mydblib.staff a,
       mydblib.payroll b
  where (a.idnum eq b.idnum) and
        (salary gt 40000)
  order by lname;

proc print noobs;
  title 'Employees with Salaries over $40,000';
run;

```

Output 2.7 Updating DBMS Data

Employees with Salaries over \$40,000		
LAST_NAME	FIRST_NAME	ANNUAL_SALARY
BANADYGA	JUSTIN	88606
BAREFOOT	JOSEPH	43025
BRADY	CHRISTINE	68767
BRANCACCIO	JOSEPH	66517
CARTER-COHEN	KAREN	40260
CASTON	FRANKLIN	41690
COHEN	LEE	91376
FERNANDEZ	KATRINA	51081

Using a SAS Data Set to Create a DBMS Table

The following example uses a SAS data step to create a DBMS table, College-Hires-1999, from a temporary SAS data set that has case-sensitive names. It creates the temporary data set and then defines the LIBNAME statement. Because it uses a DATA step to create the DBMS table, it must specify the table name as a name literal and specify the PRESERVE_TAB_NAMES= and PRESERVE_COL_NAMES= options (in this case, by using the alias PRESERVE_NAMES=).

```
options validvarname=any nodate;

data College_Hires_1999;
  input IDnum $4. +3 Lastname $11. +2
        Firstname $10. +2 City $15. +2
        State $2.;
  datalines;
3413   Schwartz      Robert      New Canaan   CT
3523   Janssen       Heike       Stamford     CT
3565   Gomez         Luis        Darien       CT
;


libname mydblib oracle user=testuser password=testpass
        path='hrdata99' schema=hrdept preserve_names=yes;

data mydblib.'College-Hires-1999'n;
  set College_Hires_1999;

proc print;
  title 'College Hires in 1999';
run;
```

Output 2.8 DBMS Table with Case-Sensitive Table and Column Names

College Hires in 1999					
Obs	IDnum	Lastname	Firstname	City	State
1	3413	Schwartz	Robert	New Canaan	CT
2	3523	Janssen	Heike	Stamford	CT
3	3565	Gomez	Luis	Darien	CT



CHAPTER

3

Data Integrity and Security

<i>Introduction to Data Integrity and Security</i>	21
<i>DBMS Security</i>	21
<i>Privileges</i>	21
<i>Triggers</i>	22
<i>SAS Security</i>	22
<i>Securing Data</i>	22
<i>Assigning SAS Passwords</i>	22
<i>Protecting Connection Information</i>	24
<i>Extracting DBMS Data to a SAS Data Set</i>	24
<i>Defining Views and Schemas</i>	25
<i>Controlling DBMS Connections</i>	25
<i>Locking, Transactions, and Currency Control</i>	26
<i>Customizing DBMS Connect and Disconnect Exits</i>	27
<i>Potential Result Set Differences When Processing Null Data</i>	27

Introduction to Data Integrity and Security

This section briefly describes DBMS security issues and then presents measures you can take on the SAS side of the interface to help protect DBMS data from accidental update or deletion. This section also provides information about how SAS handles null values that help you achieve consistent results.

DBMS Security

Privileges

The database administrator (DBA) controls who has privileges to access or update DBMS objects. The DBA also controls who can create objects, and creators of the objects control who can access the objects. A user cannot use DBMS facilities to access DBMS objects through SAS/ACCESS software unless the user has the appropriate DBMS privileges or authority on those objects. You can grant privileges on the DBMS side by using the Pass-Through Facility to EXECUTE an SQL statement, or by issuing a GRANT statement from the DBLOAD procedure SQL statement.

On the DBMS, you should give users only the privileges they must have. Privileges are granted on whole tables or views. A user must explicitly be granted privileges on the DBMS tables or views that underlie a view in order to use that view.

Note: See your DBMS documentation for more information about ensuring security on the DBMS side of the interface. \triangle

Triggers

If your DBMS supports triggers, you can use them to enforce security authorizations or business-specific security considerations. When and how triggers are executed is determined by when the SQL statement is executed and how often the trigger is executed. Triggers can be executed before an SQL statement is executed, after an SQL statement is executed, or for each row of an SQL statement. Also, triggers can be defined for DELETE, INSERT, and UPDATE statement execution.

Enabling triggers can provide more specific security for delete, insert, and update operations. SAS/ACCESS abides by all constraints and actions that are specified by a trigger. For more information, see the documentation for your DBMS.

SAS Security

Securing Data

SAS preserves the data security provided by your DBMS and operating system; SAS/ACCESS does not override the security of your DBMS. To secure DBMS data from accidental update or deletion, you can take steps on the SAS side of the interface such as the following:

- specifying the SAS/ACCESS LIBNAME option DBPROMPT= to avoid saving connection information in your code
- creating SQL views and protecting them from unauthorized access by applying passwords.

These and other approaches are discussed in detail in the following sections.

Assigning SAS Passwords

By using SAS passwords, you can protect SQL views, SAS data sets, and descriptor files from unauthorized access. The following table summarizes the levels of protection that SAS passwords provide. Note that you can assign multiple levels of protection.

Table 3.1 Password Protection Levels and Their Effects

File Type	READ=	WRITE=	ALTER=
PROC SQL view of DBMS data	Protects the underlying data from being read or updated through the view; does not protect against replacement of the view	Protects the underlying data from being updated through the view; does not protect against replacement of the view	Protects the view from being modified, deleted, or replaced
Access descriptor	No effect on descriptor	No effect on descriptor	Protects the descriptor from being read or edited
View descriptor	Protects the underlying data from being read or updated through the view	Protects the underlying data from being updated through the view	Protects the descriptor from being read or edited

You can use the following methods to assign, change, or delete a SAS password:

- the global SETPASSWORD command, which opens a dialog box
- the DATASETS procedure's MODIFY statement.

The syntax for using PROC DATASETS to assign a password to an access descriptor, a view descriptor, or a SAS data file is as follows:

```
PROC DATASETS LIBRARY=libref MEMTYPE=member-type;  
    MODIFY member-name (password-level = password-modification);  
RUN;
```

The *password-level* argument can have one or more of the following values: READ=, WRITE=, ALTER=, or PW=. PW= assigns read, write, and alter privileges to a descriptor or data file. The *password-modification* argument enables you to assign a new password or to change or delete an existing password. For example, this PROC DATASETS statement assigns the password MONEY with the ALTER level of protection to the access descriptor ADLIB.SALARIES:

```
proc datasets library=adlib memtype=access;  
    modify salaries (alter=money);  
run;
```

In this case, users are prompted for the password whenever they try to browse or update the access descriptor or try to create view descriptors that are based on ADLIB.SALARIES.

In the next example, the PROC DATASETS statement assigns the passwords MYPW and MYDEPT with READ and ALTER levels of protection to the view descriptor VLIB.JOBC204:

```
proc datasets library=vlib memtype=view;  
    modify jobc204 (read=mypw alter=mydept);  
run;
```

In this case, users are prompted for the SAS password when they try to read the DBMS data or try to browse or update the view descriptor VLIB.JOBC204. You need both levels to protect the data and descriptor from being read. However, a user could still update the data accessed by VLIB.JOBC204, for example, by using a PROC SQL UPDATE. Assign a WRITE level of protection to prevent data updates.

Note: When you assign multiple levels of passwords, use a different password for each level to ensure that you grant only the access privileges that you intend. Δ

To delete a password, put a slash after the password:

```
proc datasets library=vlib memtype=view;
    modify jobc204 (read=myspw/ alter=mydept/);
run;
```

Protecting Connection Information

In addition to directly controlling access to data, you can protect the data indirectly by protecting the connection information that SAS/ACCESS uses to reach the DBMS. Generally, this is achieved by not saving connection information in your code.

One way to protect connection information is by storing user name, password, and other connection options in a local environment variable. Access to the DBMS is denied unless the correct user and password information is stored in a local environment variable. See the documentation for your DBMS to determine whether this alternative is supported.

Another way to protect connection information is by requiring users to manually enter it at connection time. When you specify DBPROMPT=YES in a SAS/ACCESS LIBNAME statement, each user has to provide DBMS connection information in a dynamic, interactive manner. This is demonstrated in the following statement, which causes a dialog box to prompt the user to enter connection information, such as a username and password:

```
libname myoralib oracle dbprompt=yes defer=no;
```

The dialog box that appears contains the DBMS connection options that are valid for the SAS/ACCESS engine that is being used; in this case, Oracle.

Using the DBPROMPT= option on the LIBNAME statement offers several advantages. DBMS account passwords are protected because they do not need to be stored in a SAS program or descriptor file. Also, when a password or username changes, the SAS program does not need to be modified. Another advantage is that the same SAS program can be used by any valid username and password combination that is specified during execution. You can also use connection options in this interactive manner when you want to run a program on a production server instead of testing a server without making modifications to your code. By using the prompt window, the new server name can be specified dynamically.

Note: The DBPROMPT= option is not available in the SAS/ACCESS interface to DB2 under z/OS. \triangle

Extracting DBMS Data to a SAS Data Set

If you are the owner of a DBMS table and do not want anyone else to read the data, you can extract the data (or a subset of the data) and not distribute information about either the access descriptor or view descriptor.

Note: You might need to take additional steps to restrict LIBNAME or Pass-Through access to the extracted data set. \triangle

If you extract data from a view that has a SAS password assigned to it, the new SAS data file is automatically assigned the same password. If a view does not have a password, you can assign a password to the extracted SAS data file by using the MODIFY statement in the DATASETS procedure. See the *Base SAS Procedures Guide* for more information.

Defining Views and Schemas

If you want to provide access to some but not all fields in a DBMS table, create a SAS view that prohibits access to the sensitive data by specifying that particular columns be dropped. Columns that are dropped from views do not affect the underlying DBMS table and can be reselected for later use.

Some SAS/ACCESS engines support LIBNAME options that restrict or qualify the scope, or schema, of the tables in the libref. For example, the DB2 engine supports the AUTHID= and LOCATION= options, and the Oracle engine supports the SCHEMA= and DBLINK= options. See the SAS/ACCESS documentation for your DBMS to determine which options are available to you.

The following example uses the SAS/ACCESS interface to Oracle:

```
libname myoralib oracle user=testuser password=testpass
      path='myoraserver' schema=testgroup;

proc datasets lib=myoralib;
run;
```

In this example, the MYORALIB libref is associated with the Oracle schema named TESTGROUP. The DATASETS procedure lists only the tables and views that are accessible to the TESTGROUP schema. Any reference to a table that uses the libref MYORALIB is passed to the Oracle server as a qualified table name; for example, if the SAS program reads a table by specifying the SAS data set MYORALIB.TESTTABLE, the SAS/ACCESS engine passes the following query to the server:

```
select * from "testgroup.testtable"
```

Controlling DBMS Connections

Because the overhead of executing a connection to a DBMS server can be resource-intensive, SAS/ACCESS supports the CONNECTION= and DEFER= options to control when a DBMS connection is made, and how many connections are executed within the context of your SAS/ACCESS application. For most SAS/ACCESS engines, a connection to a DBMS begins one transaction, or work unit, and all statements issued in the connection execute within the context of the active transaction.

The CONNECTION= LIBNAME option enables you to specify how many connections are executed when the library is used and which operations on tables are shared within a connection. By default, the value is CONNECTION=SHAREDREAD, which means that a SAS/ACCESS engine executes a *shared read* DBMS connection when the library is assigned. Every time a table in the library is read, the read-only connection is used. However, if an application attempts to update data using the libref, a separate connection is issued, and the update occurs in the new connection. As a result, there is one connection for read-only transactions and a separate connection for each update transaction.

In the following example, the SAS/ACCESS engine issues a connection to the DBMS when the libref is assigned. The PRINT procedure reads the table by using the first connection. When the PROC SQL updates the table, the update is performed with a second connection to the DBMS.

```
libname myoralib oracle user=testuser password=testpass
      path='myoraserver';

proc print data=myoralib.mytable;
run;
```

```
proc sql;
  update myoralib.mytable set acctnum=123
    where acctnum=456;
quit;
```

The following example uses the SAS/ACCESS interface to DB2 under z/OS. The LIBNAME statement executes a connection by way of the DB2 Call Attach Facility to the DB2 DBMS server:

```
libname mydb2lib db2 authid=testuser;
```

If you want to assign more than one SAS libref to your DBMS server and if you do not plan to update the DBMS tables, SAS/ACCESS enables you to optimize how the engine performs connections. Your SAS librefs can share a single read-only connection to the DBMS if you use the CONNECTION=GLOBALREAD option. The following example shows you how to use the CONNECTION= option with the ACCESS= option to control your connection and to specify read-only data access.

```
libname mydblib1 db2 authid=testuser
  connection=globalread access=readonly;
```

If you do not want the connection to occur when the library is assigned, you can delay the connection to the DBMS by using the DEFER= option. When you specify DEFER=YES on the LIBNAME statement, for example,

```
libname mydb2lib db2 authid=testuser defer=yes;
```

the SAS/ACCESS engine connects to the DBMS the first time a DBMS object is referenced in a SAS program.

Note: If you use DEFER=YES to assign librefs to your DBMS tables and views in an AUTOEXEC program, the processing of the AUTOEXEC file is faster because the connections to the DBMS are not made every time SAS is invoked. Δ

Locking, Transactions, and Currency Control

SAS/ACCESS provides options that enable you to control some of the row-, page-, or table-locking operations that are performed by the DBMS and the SAS/ACCESS engine as your programs are executed. For example, by default, the SAS/ACCESS Oracle engine does not lock any data when it reads rows from Oracle tables. However, you can override this behavior by using the locking options that are supported in the SAS/ACCESS interface to Oracle.

If you want to lock the data pages of a table while SAS reads the data to prevent other processes from updating the table, you can use the READLOCK_TYPE= option, as in the following example.

```
libname myoralib oracle user=testuser pass=testpass
  path='myoraserver' readlock_type=table;

data work.mydata;
  set myoralib.mytable(where=(colnum > 123));
run;
```

In this example, the SAS/ACCESS Oracle engine obtains a TABLE SHARE lock on the table so that the data cannot be updated by other processes while your SAS program is reading it.

In the following example, Oracle acquires row-level locks on rows read for update in the tables in the libref.

```
libname myoralib oracle user=testuser password=testpass
  path='myoraserver' updatelock_type=row;
```

Note: Each SAS/ACCESS interface supports specific options; see the SAS/ACCESS documentation for your DBMS to determine which options it supports. Δ

Customizing DBMS Connect and Disconnect Exits

You can specify DBMS commands or stored procedures to be executed immediately after a DBMS connection or before a DBMS disconnect by using the DBCONINIT= and DBCONTERM= LIBNAME options. See the following example:

```
libname myoralib oracle user=testuser password=testpass
  path='myoraserver' dbconinit="EXEC MY_PROCEDURE";

proc sql;
  update myoralib.mytable set acctnum=123
    where acctnum=567;
quit;
```

When the libref is assigned, the SAS/ACCESS engine connects to the DBMS and passes a command to the DBMS to execute the stored procedure MY_PROCEDURE. By default, a new connection to the DBMS is made for every table that is opened for updating. This means that MY_PROCEDURE is executed again after a connection is made to update the table MYTABLE.

To execute a DBMS command or stored procedure *only* after the first connection in a library assignment, you can use the DBLIBINIT= option. Similarly, the DBLIBTERM= option enables you to specify a command to be executed before the disconnection of only the first library connection, as in the following example:

```
libname myoralib oracle user=testuser password=testpass
  dblibinit="EXEC MY_INIT" dblibterm="EXEC MY_TERM";
```

Potential Result Set Differences When Processing Null Data

When your data contains null values or when internal processing generates intermediate data sets that contain null values, you might get different result sets depending on whether the processing is done by SAS or by the DBMS. Although in many cases this does not present a problem, it is important to understand how these differences occur.

Most relational database systems have a special value called null, which means an absence of information and is analogous to a SAS missing value. SAS/ACCESS translates SAS missing values to DBMS null values when creating DBMS tables from within SAS and, conversely, translates DBMS null values to SAS missing values when reading DBMS data into SAS.

There is, however, an important difference in the behavior of DBMS null values and SAS missing values:

- A DBMS null value is interpreted as the absence of data, so you cannot sort a DBMS null value or evaluate it with standard comparison operators.
- A SAS missing value is interpreted as its internal floating-point representation because SAS supports 28 missing values (where a period (.) is the most common missing value). Because SAS supports multiple missing values, you can sort a SAS missing value and evaluate it with standard comparison operators.

This means that SAS and the DBMS interpret null values differently, which has significant implications when SAS/ACCESS passes queries to a DBMS for processing. This can be an issue in the following situations:

- when filtering data (for example, in a WHERE clause, a HAVING clause, or an outer join ON clause). SAS interprets null values as missing; many DBMS exclude null values from consideration. For example, if you have null values in a DBMS column that is used in a WHERE clause, your results might differ depending on whether the WHERE clause is processed in SAS or is passed to the DBMS for processing. This is because the DBMS removes null values from consideration in a WHERE clause, but SAS does not.
- when using certain functions. For example, if you use the MIN aggregate function on a DBMS column that contains null values, the DBMS does not consider the null values. However, SAS interprets the null values as missing and this affects the result.
- when submitting outer joins where internal processing generates nulls for intermediate result sets.
- when sorting data. SAS sorts null values low; most DBMSs sort null values high. (See “Sorting DBMS Data” on page 33 for more information.)

For example, create a simple data set that consists of one observation and one variable.

```
libname myoralib oracle user=testuser password=testpass;
data myoralib.table;
x=.;          /*create a missing value */
run;
```

Then, print the data set using a WHERE clause, which SAS/ACCESS passes to the DBMS for processing.

```
proc print data=myoralib.table;
  where x<0;
run;
```

The log indicates that no observations were selected by the WHERE clause, because Oracle interprets the missing value as the absence of data, and does not evaluate it with the less-than (<) comparison operator.

When there is the potential for inconsistency, consider using one of the following strategies:

- Use the LIBNAME option DIRECT_SQL= to control whether the processing is done by SAS or by the DBMS.
- Use the Pass-Through Facility to ensure that the processing is done by the DBMS.
- Add the "is not null" expression to WHERE clauses and ON clauses to ensure that you get the same result regardless of whether SAS or the DBMS does the processing.

Note: Use the data set option NULLCHAR= to specify how the DBMS interprets missing SAS character values when updating DBMS data or inserting rows into a DBMS table. Δ


You can use the first of these strategies to force SAS to process the data in the example below:

```
libname myoralib oracle user=testuser password=testpass
      direct_sql=nowhere; /* forces SAS to process WHERE clauses */
data myoralib.table;
x=.;          /*create a missing value */
run;
```

Then, print the data set using a WHERE clause:

```
proc print data=myoralib.table;
      where x<0;
run;
```

This time, the log indicates that one observation was read from the data set, because SAS evaluates the missing value as satisfying the less-than-zero condition in the WHERE clause.



CHAPTER

4

Performance Considerations

<i>Increasing Throughput of the SAS Server</i>	31
<i>Limiting Retrieval</i>	31
<i>Column Selection</i>	31
<i>The KEEP= and DROP= Options</i>	32
<i>Repeatedly Accessing Data</i>	33
<i>Sorting DBMS Data</i>	33
<i>Temporary Table Support for SAS/ACCESS</i>	34
<i>Overview</i>	34
<i>General Temporary Table Use</i>	36
<i>Pushing Heterogeneous Joins</i>	36
<i>Pushing Updates</i>	36

Increasing Throughput of the SAS Server

When you invoke SAS as a server that responds to multiple clients, you can use the DBSRVTP= system option to improve the performance of the clients. The DBSRVTP= option tells the SAS server whether to put a hold (or block) on the originating client while making performance-critical calls to the database. By holding or blocking the originating client, the SAS/ACCESS server remains available for other clients; they do not have to wait for the originating client to complete its call to the database.

Limiting Retrieval

Column Selection

Limiting the number of rows that are returned to SAS from the DBMS is an important performance consideration. Because the less data that is requested by the SAS job, the faster the job runs.

Where possible, specify selection criteria in order to limit the number of rows that the DBMS returns to SAS. Use the SAS WHERE clause to retrieve a subset of the DBMS data.

Likewise, select only the DBMS columns that your program needs. Selecting unnecessary columns slows your job.

The KEEP= and DROP= Options

Just as with a SAS data set you can use the DROP= and KEEP= data set options to prevent retrieving unneeded columns from your DBMS table.

In this example, the KEEP= data set option causes the SAS/ACCESS engine to select only the SALARY and DEPT columns when it reads the MYDBLIB.EMPLOYEES table.

```
libname mydblib db2 user=testid password=testpass database=testdb;

proc sql;
select *
    from mydblib.employees(keep=salary dept)
    where dept='ACC024';
quit;
```

The generated SQL that is processed by the DBMS is similar to the following:

```
SELECT "SALARY", "DEPT" FROM EMPLOYEES
    WHERE(DEPT="ACC024")
```

Without the KEEP option, the SQL processed by the DBMS would be similar to the following:

```
SELECT * FROM EMPLOYEES    WHERE(DEPT="ACC024")
```

This would result in all of the columns from the EMPLOYEES table being read in to SAS.

The DROP= data set option is a parallel option that specifies columns to omit from the output table. Keep in mind that the DROP= and KEEP= data set options are not interchangeable with the DROP and KEEP statements. Use of the DROP and KEEP statements when selecting data from a DBMS can result in retrieval of all column into SAS, which can seriously impact performance.

For example, the following would result in all of the columns from the EMPLOYEES table being retrieved into SAS. The KEEP statement would be applied when creating the output data set.

```
libname mydblib db2 user=testid password=testpass database=testdb;

data temp;
    set mydblib.employees;
    keep salary;
run;
```

The following is an example of how to use the KEEP data set option to retrieve only the SALARY column:

```
data temp;
    set mydblib.employees(keep=salary);
run;
```

Repeatedly Accessing Data

CAUTION:

If you need to access the most current DBMS data, then access it directly from the database every time. Do not follow the extraction suggestions that are presented in this section. △

It is sometimes more efficient to extract (copy) DBMS data to a SAS data file than to repeatedly read the data by using a SAS view. SAS data files are organized to provide optimal performance with PROC and DATA steps. Programs that use SAS data files are often more efficient than SAS programs that read DBMS data directly.

Consider extracting data when you are working with a large DBMS table and you plan to use the same DBMS data in several procedures or DATA steps during the same SAS session.

Note: You can extract DBMS data to a SAS data file by using the OUT= option, a DATA step, or ACCESS procedures. △

Sorting DBMS Data

Sorting DBMS data can be resource-intensive, whether you use the SORT procedure, a BY statement, or an ORDER BY clause on a DBMS data source; or, in the SQL procedure's SELECT statement. Sort data only when it is needed for your program. The following list contains guidelines for sorting data:

- If you specify a BY statement in a DATA or PROC step that references a DBMS data source, you associate the BY variable (in a DATA or PROC step) with an indexed DBMS column. This is recommended for performance reasons. If you reference DBMS data in a SAS program and the program includes a BY statement for a variable that corresponds to a column in the DBMS table, the SAS/ACCESS LIBNAME engine automatically generates an ORDER BY clause for that variable. The ORDER BY clause causes the DBMS to sort the data before the DATA or PROC step uses the data in a SAS program. If the DBMS table is very large, this sorting can adversely affect your performance. Use a BY variable that is based on an indexed DBMS column in order to reduce this negative impact.
- The outermost BY or ORDER BY clause overrides any embedded BY or ORDER BY clauses, including those specified by the DBCONDITION= option, those specified in a WHERE clause, and those in the selection criteria in a view descriptor. In the following example, the EXEC_EMPLOYEES data set includes a BY statement that sorts the data by the variable SENIORITY. However, when that data set is used in the following PROC SQL query, the data is ordered by the SALARY column and not by SENIORITY.

```
libname mydblib oracle user=testuser password=testpass;
data exec_employees;
  set mydblib.staff (keep=lname fname idnum);
  by seniority;
  where salary >= 150000;
run;

proc sql;
select * from exec_employees
  order by salary;
```

- Do not use PROC SORT to sort data from SAS back into the DBMS, because this impedes performance and has no effect on the order of the data.
- The database does not guarantee sort stability when you use PROC SORT. Sort stability means that the ordering of the observations in the BY statement is exactly the same every time the sort is run against static data. If you absolutely require sort stability, you must place your database data into a SAS data set, and then use PROC SORT.
- When you use PROC SORT, be aware that the sort rules for SAS and for your DBMS might be different. Use the Base SAS system option SORTPGM to specify which rules (host, SAS, or DBMS) are applied:

`SORTPGM=BEST`

sorts data according to the DBMS sort rules, then the host sort rules, and then the SAS sort rules. (Sorting uses the first available and pertinent sorting algorithm in this list.) This is the default.

`SORTPGM=HOST`

sorts data according to host rules and then SAS rules. (Sorting uses the first available and pertinent sorting algorithm in this list.)

`SORTPGM=SAS`

sorts data by SAS rules.

Temporary Table Support for SAS/ACCESS

Overview

DBMS temporary table support in SAS consists of the ability to retain DBMS temporary tables from one SAS step to the next. This ability is a result of establishing a SAS connection to the DBMS that persists across multiple SAS procedures and DATA steps.

Temporary table support for SAS 9.1 is available on the following DBMSs:

DB2 UNIX/PC

DB2 z/OS

HP Neoview

Informix

Netezza

ODBC

OLE DB

Oracle

Sybase

Teradata

The value of DBMS temporary table support in SAS is increased performance potential. By pushing processing to the DBMS in certain situations, an overall performance gain can be achieved. The following processes outline, in general, how to use DBMS temporary tables in SAS 9.1.

General Temporary Table Use

To use temporary tables on the DBMS, complete the following steps:

- 1 Establish a global connection to the DBMS that persists across SAS procedure and DATA step boundaries.
- 2 Create a DBMS temporary table and load it with data.
- 3 Use the DBMS temporary table with SAS.

Note: Closing the global connection causes the DBMS temporary table to close as well. \triangle

Pushing Heterogeneous Joins

To push heterogeneous joins to the DBMS, complete the following steps:


- 1 Establish a global connection to the DBMS that persists across SAS procedure and DATA step boundaries.
- 2 Create a DBMS temporary table and load it with data.
- 3 Perform a join on the DBMS using the DBMS temporary and DBMS permanent tables.
- 4 Process the result of the join with SAS.

Pushing Updates

To push updates (process transactions) to the DBMS, complete the following steps:

- 1 Establish a global connection to the DBMS that will persist across SAS procedure and DATA step boundaries.
- 2 Create a DBMS temporary table and load it with data.
- 3 Issue SQL that uses values in the temporary table to process against the production table.
- 4 Process the updated DBMS tables with SAS.

Note: These processing scenarios are purposefully generic; however they apply to each DBMS that supports temporary tables. Refer to the SAS/ACCESS section for your database for complete details. \triangle



CHAPTER

5

Optimizing Your SQL Usage

<i>Overview of Optimizing Your SQL Usage</i>	37
<i>Passing Functions to the DBMS Using PROC SQL</i>	38
<i>Passing Joins to the DBMS</i>	38
<i>Passing the DELETE Statement to Empty a Table</i>	40
<i>When Passing Joins to the DBMS Will Fail</i>	40
<i>Passing DISTINCT and UNION Processing to the DBMS</i>	42
<i>Optimizing the Passing of WHERE Clauses to the DBMS</i>	42
<i>Passing Functions to the DBMS Using WHERE Clauses</i>	43
<i>Using the DBINDEX=, DBKEY=, and MULTI_DATASRC_OPT= Options</i>	43

Overview of Optimizing Your SQL Usage

SAS/ACCESS takes advantage of DBMS capabilities by passing certain SQL operations to the DBMS whenever possible. This can reduce data movement, which can improve performance. The performance impact can be significant when you access large DBMS tables and the SQL that is passed to the DBMS subsets the table to reduce the amount of rows. SAS/ACCESS sends operations to the DBMS for processing in the following situations:

- When operations use the Pass-Through Facility. When you use the Pass-Through Facility, you submit DBMS-specific SQL statements that are sent directly to the DBMS for execution. For example, when you submit Transact-SQL statements to be passed to a Sybase database.
- When SAS/ACCESS can translate the operations into the SQL of the DBMS. When you use the SAS/ACCESS LIBNAME statement and PROC SQL, you submit SAS statements that SAS/ACCESS can translate into the SQL of the DBMS and then pass to the DBMS for processing.

By using the automatic translation abilities, you can often achieve the performance benefits of the Pass-Through Facility without needing to write DBMS-specific SQL code. The following sections describe the SAS SQL operations that SAS/ACCESS can pass to the DBMS for processing. See “Optimizing the Passing of WHERE Clauses to the DBMS” on page 42 for information about passing WHERE clauses to the DBMS.

Note: There are certain conditions that prevent operations from being passed to the DBMS. For example, when you use an INTO clause or any data set option, operations are processed in SAS instead of being passed to the DBMS. Re-merges, union joins, and truncated comparisons also prevent operations from being passed to the DBMS.

Additionally, it is important to note that when you join tables across multiple tables, implicit pass-through utilizes the first connection. Consequently, LIBNAME options from subsequent connections are ignored.

You can use the `SASTRACE=` system option to determine whether an operation is processed by SAS or is passed to the DBMS for processing. Δ

To prevent operations from being passed to the DBMS, use the `LIBNAME` option `DIRECT_SQL=`.

Passing Functions to the DBMS Using PROC SQL

When you use the `SAS/ACCESS LIBNAME` statement, the SAS SQL aggregate functions `MIN`, `MAX`, `AVG`, `MEAN`, `FREQ`, `N`, `SUM`, and `COUNT` are passed to the DBMS because they are SQL ANSI-defined aggregate functions.

For example, the following query of the Oracle table `EMP` is passed to the DBMS for processing:

```
libname myoralib oracle user=testuser password=testpass;
proc sql;
  select count(*) from myoralib.emp;
quit;
```

This code causes Oracle to process the following query:

```
select COUNT(*) from EMP
```

`SAS/ACCESS` can also translate other SAS functions into DBMS-specific functions so they can be passed to the DBMS.

In the following example, the SAS function `UPCASE` is translated into the Oracle function `UPPER`:

```
libname myoralib oracle user=testuser password=testpass;
proc sql;
  select customer from myoralib.customers
  where upcase(country)="USA";
quit;
```

The translated query that is processed in Oracle is

```
select customer from customers where upper(country)='USA'
```

The functions that are passed are different for each DBMS. See the documentation for your `SAS/ACCESS` interface to determine which functions it translates.

Passing Joins to the DBMS

When you perform a join across tables in a single DBMS, `SAS/ACCESS` can often pass the join to the DBMS for processing. Before implementing a join, `PROC SQL` checks to see whether the DBMS can do the join. A comparison is made using the `SAS/ACCESS LIBNAME` statement for the tables. Certain criteria must be met for the join to proceed.

If it can, `PROC SQL` passes the join to the DBMS, which performs the join and returns only the results to SAS. If the DBMS cannot do the join, `PROC SQL` processes the join.

The following types of joins are eligible for passing to the DBMS:

- for all DBMSs, inner joins between two or more tables.
- for DBMSs that support ANSI outer join syntax, outer joins between two or more DBMS tables.
- for ODBC and Microsoft SQL Server, outer joins between two or more tables. However, the outer joins must not be mixed with inner joins in a query.
- For DBMSs that support nonstandard outer join syntax (ORACLE, Sybase and INFORMIX), outer joins between two or more tables, with the following restrictions:

Full outer joins are not supported.

Only a comparison operator is allowed in an ON clause. For Sybase, the only valid comparison operator is '='.

For ORACLE and Sybase, both operands in an ON clause must reference a column name. If an operand is a literal, it cannot be passed to the DBMS. Since these DBMSs do not support this, all ON clauses are transformed into WHERE clauses before attempting to pass the join to the DBMS. This can result in queries not being passed to the DBMS if they include additional WHERE clauses or contain complex join conditions.

For INFORMIX, the outer joins cannot consist of more than two tables and cannot contain a WHERE clause.

For Sybase: If there are multiple joins or joins with additional WHERE clauses, it might be advisable to have the join processed internally by PROC SQL rather than passing it to the DBMS. This is because Sybase evaluates multi-joins with WHERE clauses differently than SAS. To allow PROC SQL to process the join internally use the SAS/ACCESS DIRECT_SQL= option.

Note: If PROC SQL cannot successfully pass down a complete query to the DBMS, it might re-attempt passing down a subquery. The SQL that is passed to the DBMS can be analyzed by turning on SAS tracing options. The SAS trace information displays the exact queries being passed to the DBMS for processing.

Δ

In the following example, two large DBMS tables (TABLE1 and TABLE2) each have a column named DeptNo. Retrieve the rows from an inner join of these tables where the DeptNo value in TABLE1 equals the DeptNo value in TABLE2. The join between two tables in the DBLIB library (which references an Oracle database) is detected by PROC SQL and passed by SAS/ACCESS directly to the DBMS. The DBMS processes the inner join between the two tables and returns only the resulting rows to SAS.

```
libname dblib oracle user=testuser password=testpass;
proc sql;
  select tab1.deptno, tab1.dname from
    dblib.table1 tab1,
    dblib.table2 tab2
  where tab1.deptno = tab2.deptno;
quit;
```

The query is passed to the DBMS, generating the following Oracle code:

```
select table1."deptno", table1."dname" from TABLE1, TABLE2
  where TABLE1."deptno" = TABLE2."deptno"
```

In the following example, an outer join between two Oracle tables, TABLE1 and TABLE2, is passed to the DBMS for processing.

```
libname myoralib oracle user=testuser password=testpass;
proc sql;
  select * from myoralib.table1 right join myoralib.table2
    on table1.x = table2.x
  where table2.x > 1;
quit;
```

The query is passed to the DBMS, generating the following Oracle code:

```
select table1."X", table2."X" from TABLE1, TABLE2
  where TABLE1."X" (+)= TABLE2."X"
  and (TABLE2."X" > 1)
```

Passing the DELETE Statement to Empty a Table

When you use the SAS/ACCESS LIBNAME statement with the DIRECT_EXE option set to DELETE, the SAS SQL DELETE statement gets passed to the DBMS for execution as long as it does not contain a WHERE clause. The DBMS deletes all the rows, but does not delete the table itself.

The following example illustrates how a DELETE statement gets passed to Oracle for the table EMP to be emptied:

```
libname myoralib oracle user=testuser password=testpass direct_exe=delete;
proc sql;
  delete from myoralib.emp;
quit;
```

This code causes Oracle to execute the following:

```
delete from emp
```

When Passing Joins to the DBMS Will Fail

SAS/ACCESS will, by default, attempt to pass certain types of SQL statements directly to the DBMS for processing. Most notable are SQL join statements that otherwise would be processed as individual queries to each data source that belonged to the join. In that instance, the join would then be performed internally by PROC SQL. Passing the join to the DBMS for direct processing can result in significant performance gains.

However, there are several reasons why a join statement under PROC SQL might not be passed to the DBMS for processing. In general, the success of the join depends upon the nature of the SQL that was coded and the DBMS's acceptance of the generated syntax. It is also greatly influenced by the use of option settings. The following are the primary reasons why join statements might fail to be passed:

- The generated SQL syntax is not accepted by the DBMS.

PROC SQL attempts to pass the SQL join query directly to the DBMS for processing. The DBMS can reject the syntax for any number of reasons. In this event, PROC SQL attempts to open both tables individually and perform the join internally.

- The SQL query involves multiple librefs that do not share connection characteristics.

If the librefs are specified using different servers, user IDs, or any other connection options, PROC SQL does not attempt to pass the statement to the DBMS for direct processing.
- The use of data set options in the query.

The specification of any data set option on a table that is referenced in the SQL query prohibits the statement from passing to the DBMS for direct processing.
- The use of certain LIBNAME options.

The specification of LIBNAME options that request member level controls, such as table locks (“READ_LOCK_TYPE= LIBNAME Option” on page 138 or “UPDATE_LOCK_TYPE= LIBNAME Option” on page 151), will prohibit the statement from successfully passing to the DBMS for direct processing.
- The “DIRECT_SQL= LIBNAME Option” on page 114 option setting.

The DIRECT_SQL= option default setting is YES. PROC SQL attempts to pass SQL joins directly to the DBMS for processing. Other settings for the DIRECT_SQL= option influence the nature of the SQL statements that PROC SQL tries to pass down to the DBMS, or even if it tries to pass anything at all.

DIRECT_SQL=YES

PROC SQL automatically attempts to pass the SQL join query to the DBMS. This is the default setting for this option. The join attempt could fail due to a DBMS return code. If this happens, PROC SQL attempts to open both tables individually and perform the join internally.

DIRECT_SQL=NO

PROC SQL does not attempt to pass SQL join queries to the DBMS. Other SQL statements can be passed, however. If the “MULTI_DATASRC_OPT= LIBNAME Option” on page 126 is in effect, the generated SQL can also be passed.

DIRECT_SQL=NONE

PROC SQL does not attempt to pass any SQL directly to the DBMS for processing.

DIRECT_SQL=NOWHERE

PROC SQL attempts to pass SQL to the DBMS including SQL joins. However, it does not pass any WHERE clauses associated with the SQL statement. This causes any join that is attempted with direct processing to fail.

DIRECT_SQL=NOFUNCTIONS

PROC SQL does not pass any statements in which any function is present to the DBMS. Normally PROC SQL attempts to pass down any functions coded in the SQL to the DBMS, provided the DBMS supports the given function.

DIRECT_SQL=NOGENSQL

PROC SQL does not attempt to pass SQL join queries to the DBMS. Other SQL statements can be passed down, however. If the MULTI_DATASRC_OPT= option is in effect, the generated SQL can be passed.

DIRECT_SQL=NOMULTOUTJOINS

PROC SQL does not attempt to pass any multiple outer joins to the DBMS for direct processing. Other SQL statements can be passed, however, including portions of a multiple outer join.

- The use of SAS functions on the SELECT clause can prevent joins from being passed.

Passing DISTINCT and UNION Processing to the DBMS

When you use the SAS/ACCESS LIBNAME statement to access DBMS data, the DISTINCT and UNION operators are processed in the DBMS rather than in SAS. For example, when PROC SQL detects a DISTINCT operator, it passes the operator to the DBMS to check for duplicate rows. The DBMS then returns only the unique rows to SAS.

In the following example, the Oracle table CUSTBASE is queried for unique values in the STATE column.

```
libname myoralib oracle user=testuser password=testpass;
proc sql;
  select distinct state from myoralib.custbase;
quit;
```

The DISTINCT operator is passed to Oracle, generating the following Oracle code.

```
select distinct custbase."STATE" from CUSTBASE
```

Oracle passes the results from this query back to SAS.

Optimizing the Passing of WHERE Clauses to the DBMS

Use the following general guidelines for writing efficient WHERE clauses:

- Avoid the NOT operator if you can use an equivalent form.
 - Inefficient: **where zipcode not>8000**
 - Efficient: **where zipcode<=8000**
- Avoid the >= and <= operators if you can use the BETWEEN predicate.
 - Inefficient: **where ZIPCODE>=70000 and ZIPCODE<=80000**
 - Efficient: **where ZIPCODE between 70000 and 80000**
- Avoid LIKE predicates that begin with % or _ .
 - Inefficient: **where COUNTRY like '%INA'**
 - Efficient: **where COUNTRY like 'A%INA'**
- Avoid arithmetic expressions in a predicate.
 - Inefficient: **where SALARY>12*4000.00**
 - Efficient: **where SALARY>48000.00**
- Use DBKEY=, DBINDEX=, and MULTI_DATASRC_OPT= when appropriate. See “Using the DBINDEX=, DBKEY=, and MULTI_DATASRC_OPT= Options” on page 43 for details about these options.

Whenever possible, SAS/ACCESS passes WHERE clauses to the DBMS, because the DBMS processes them more efficiently than SAS does. SAS translates the WHERE clauses into generated SQL code. The performance impact can be particularly significant when you are accessing large DBMS tables. The following section describes how and when functions are passed to the DBMS. For information about passing processing to the DBMS when you are using PROC SQL, see “Overview of Optimizing Your SQL Usage” on page 37.

Note: If you have NULL values in a DBMS column that is used in a WHERE clause, be aware that your results might differ depending on whether the WHERE clause is processed in SAS or is passed to the DBMS for processing. This is because DBMSs tend to remove NULL values from consideration in a WHERE clause, while SAS does not. Δ

To prevent WHERE clauses from being passed to the DBMS, use the LIBNAME option DIRECT_SQL= NOWHERE.

Passing Functions to the DBMS Using WHERE Clauses

When you use the SAS/ACCESS LIBNAME statement, SAS/ACCESS translates several SAS functions in WHERE clauses into DBMS-specific functions so they can be passed to the DBMS.

In the following SAS code, SAS can translate the FLOOR function into a DBMS function and pass the WHERE clause to the DBMS.

```
libname myoralib oracle user=testuser password=testpass;
proc print data=myoralib.personnel;
  where floor(hourlywage)+floor(tips)<10;
run;
```

The generated SQL that the DBMS processes would be similar to this code:

```
SELECT "HOURLYWAGE", "TIPS" FROM PERSONNEL
  WHERE ((FLOOR("HOURLYWAGE") + FLOOR("TIPS")) < 10)
```

If the WHERE clause contains a function that SAS cannot translate into a DBMS function, SAS retrieves all the rows from the DBMS and then applies the WHERE clause.

The functions that are passed are different for each DBMS. See the documentation for your SAS/ACCESS interface to determine which functions it translates.

Using the DBINDEX=, DBKEY=, and MULTI_DATASRC_OPT= Options

When you code a join operation in SAS and the join cannot be passed directly to a DBMS for processing, the join is performed by SAS. Normally, this processing involves individual queries to each data source that belonged to the join and the join being performed internally by SAS. When you join a large DBMS table and a small SAS data set or DBMS table, using the DBKEY=, DBINDEX=, and MULTI_DATASRC_OPT= options might enhance performance. You can use them to retrieve a subset of the DBMS data into SAS for the join.

When MULTI_DATASRC_OPT=IN_CLAUSE is specified for DBMS data sources in a PROC SQL join operation, the procedure retrieves the unique values of the join column from the smaller table to construct an IN clause. This IN clause is used when SAS retrieves the data from the larger DBMS table. The join is performed in SAS. If a SAS data set is used, no matter how large, it is always in the IN_CLAUSE. For better performance, it is recommended that the SAS data set be smaller than the DBMS table, otherwise processing can be extremely slow.

MULTI_DATASRC_OPT= generates a SELECT COUNT to determine the size of data sets that are not SAS data sets. If you know the size of your data set, you can use DBMASTER to designate the larger table.

MULTI_DATASRC_OPT= might provide performance improvements over DBKEY=. If both options are specified, DBKEY= overrides MULTI_DATASRC_OPT=.

MULTI_DATASRC_OPT= is used only when SAS processes a join with PROC SQL. It is not used for SAS datastep processing. For certain joins operations, such as those involving additional subsetting applying to the query, PROC SQL might determine that it is more efficient to process the join internally. In these situations, it does not use the MULTI_DATASRC_OPT= optimization even when specified. If PROC SQL determines

it can pass the join directly to the DBMS, it does not also use this option even though it is specified.

In this example, the MULTI_DATASRC_OPT= option is used to improve the performance of an SQL join statement. MULTI_DATASRC_OPT= instructs PROC SQL to pass the WHERE clause to the SAS/ACCESS engine with an IN clause built from the SAS table. The engine then passes this optimized query to the DBMS server. The IN clause is built from the unique values of the SAS DeptNo variable. As a result, only rows that match the WHERE clause are retrieved from the DBMS. Without this option, PROC SQL retrieves all the rows from the Dept table and applies the WHERE clause during PROC SQL processing in SAS. Processing can be both CPU-intensive and I/O-intensive if the Oracle Dept table is large.

```
data keyvalues;
  deptno=30;
  output;
  deptno=10;
  output;
run;

libname dblib oracle user=testuser password=testpass
  path='myorapath' multi_datasrc_opt=in_clause;

proc sql;
  select bigtab.deptno, bigtab.loc
  from dblib.dept bigtab,
       keyvalues smalllds
  where bigtab.deptno=smalllds.deptno;
quit;
```

The SQL statement that is created by SAS/ACCESS and passed to the DBMS is similar to the following

```
SELECT "DEPTNO", "LOC" FROM DEPT WHERE (("DEPTNO" IN (10,30)))
```

Using DBKEY or DBINDEX decreases performance when the SAS data set is too large. These options cause each value in the transaction data set to generate a new result set (or open cursor) from the DBMS table. For example, if your SAS data set has 100 observations with unique key values, you request 100 result sets from the DBMS, which might be very expensive. You must determine whether use of these options is appropriate, or whether you can achieve better performance by reading the entire DBMS table (or by creating a subset of the table).

DBINDEX= and DBKEY= are mutually exclusive. If you specify them together, DBKEY= overrides DBINDEX=. Both of these options are ignored if you specify the SAS/ACCESS data set option DBCONDITION= or the SAS data set option WHERE=.

DBKEY= does not require that any database indexes be defined; nor does it check the DBMS system tables. This option instructs SAS to use the specified DBMS column name or names in the WHERE clause that is passed to the DBMS in the join.

The DBKEY= option can also be used in a SAS DATA step, with the KEY= option in the SET statement, to improve the performance of joins. You specify a value of KEY=DBKEY in this situation. The following DATA step creates a new data file by joining the data file KEYVALUES with the DBMS table MYTABLE. The variable DEPTNO is used with the DBKEY= option to cause a WHERE clause to be issued by SAS/ACCESS.

```
data sasuser.new;
  set sasuser.keyvalues;
  set dblink.mytable(dbkey=deptno) key=dbkey;
run;
```

Note: When you use DBKEY= with the DATA step MODIFY statement, there is no implied ordering of the data that is returned from the database. If the master DBMS table contains records with duplicate key values, using DBKEY= can alter the outcome of the DATA step. Because SAS regenerates result sets (open cursors) during transaction processing, the changes you make during processing have an impact on the results of subsequent queries. Before you use DBKEY= in this context, determine whether your master DBMS file has duplicate values for keys. Remember that the REPLACE, OUTPUT, and REMOVE statements can cause duplicate values to appear in the master table. Δ

The DBKEY= option does not require or check for the existence of indexes created on the DBMS table. Therefore, the DBMS system tables are not accessed when you use this option. The DBKEY= option is preferred over the DBINDEX= option for this reason. If you perform a join and use PROC SQL, you *must* ensure that the columns that are specified through the DBKEY= option match the columns that are specified in the SAS data set.


CAUTION:

Before you use the DBINDEX= option, take extreme care to evaluate some characteristics of the DBMS data. The number of rows in the table, the number of rows returned in the query, and the distribution of index values in the table are among the factors to consider. Some experimentation might be necessary to discover the optimum settings. Δ

You can use the DBINDEX= option instead of the DBKEY= option if you know that the DBMS table has one or more indexes that use the column(s) that the join is performed on. Use DBINDEX=*index-name* if you know the name of the index, or use DBINDEX=YES if you do not know the name of the index. Use this option as a data set option, and *not* a LIBNAME option, because index lookup can potentially be an expensive operation.

DBINDEX= requires that the join table *must* have a database index that is defined on the columns involved in the join. If there is no index, then all processing of the join takes place in SAS, where all rows from each table are read into SAS and SAS performs the join.

Note: The data set options NULLCHAR= and NULLCHARVAL= determine how SAS missing character values are handled during DBINDEX= and DBKEY= processing. Δ



CHAPTER

6

Threaded Reads

<i>Overview of Threaded Reads in SAS/ACCESS</i>	47
<i>Underlying Technology of Threaded Reads</i>	48
<i>SAS/ACCESS Interfaces and Threaded Reads</i>	48
<i>Scope of Threaded Reads</i>	49
<i>Options That Affect Threaded Reads</i>	49
<i>Generating Trace Information for Threaded Reads</i>	50
<i>Performance Impact of Threaded Reads</i>	53
<i>Autopartitioning Techniques in SAS/ACCESS</i>	53
<i>Data Ordering in SAS/ACCESS</i>	54
<i>Two-Pass Processing for SAS Threaded Applications</i>	54
<i>When Threaded Reads Will Not Occur</i>	55
<i>Summary of Threaded Reads</i>	55

Overview of Threaded Reads in SAS/ACCESS

In Version 8 and earlier, SAS opened a single connection to the DBMS to read a table. SAS statements requesting data were converted to an SQL statement and passed to the DBMS. The DBMS processed the SQL statement, produced a result set consisting of table rows and columns, and transferred the result set back to SAS on the single connection.

With a threaded read, the table read time can be reduced by retrieving the result set on multiple connections between SAS and the DBMS. SAS is able to create multiple threads, and a read connection is established between the DBMS and each SAS thread. The result set is partitioned across the connections, and rows are passed to SAS simultaneously (in parallel) across the connections, improving performance.

Underlying Technology of Threaded Reads

To perform a threaded read, SAS first creates threads, which are standard operating system tasks controlled by SAS, within the SAS session. Next, SAS establishes a DBMS connection on each thread. SAS then causes the DBMS to partition the result set and reads one partition per thread. To cause the partitioning, SAS appends a WHERE clause to the SQL so that a single SQL statement becomes multiple SQL statements, one for each thread. For example:

```
proc reg SIMPLE
  data=dblib.salesdata (keep=salesnumber maxsales);

  var ALL;
run;
```

Previous versions of SAS opened a single connection and issued:

```
SELECT salesnumber,maxsales FROM SALESDATA;
```

Assuming that SalesData has an integer column EmployeeNum, SAS 9.1, *might* open two connections, issuing the following statements:

```
SELECT salesnumber,maxsales FROM salesdata WHERE (EMPLOYEEENUM mod 2)=0;
```

and

```
SELECT salesnumber,maxsales FROM SALESDATA WHERE (EMPLOYEEENUM mod 2)=1;
```

For more information about MOD, see “Autopartitioning Techniques in SAS/ACCESS” on page 53

Note: *Might* is an important word here. Most, but not all, SAS/ACCESS interfaces support threaded reads in SAS 9.1. The partitioning WHERE clauses that SAS generates vary. SAS cannot always generate partitioning WHERE clauses. In that case, the SAS user can supply them. There also might be other means used to partition the data, instead of the WHERE clauses. \triangle

SAS/ACCESS Interfaces and Threaded Reads

Here are the SAS/ACCESS interfaces that support threaded reads for SAS 9.1. More interfaces are expected to support threaded reads in future releases.

DB2 under UNIX
and PC Hosts

DB2 under z/OS

Informix

Microsoft SQL
Server

ODBC

Oracle*

Sybase

Teradata**

*Threaded reads for the Oracle engine are not supported under MVS (z/OS).

**Teradata on OS/390, z/OS, and UNIX supports only FastExport threaded reads; see the Teradata specific documentation for details.

Threaded reads work across all UNIX and Windows platforms where you run SAS. There are special considerations for z/OS; see the Oracle-specific and Teradata-specific sections about threaded reads for details.

Scope of Threaded Reads

SAS steps called threaded applications are automatically eligible for a threaded read. Threaded applications are bottom-to-top fully threaded SAS procedures that perform data reads, numerical algorithms, and data analysis in threads. Only some SAS procedures are threaded applications. Here is a basic example of PROC REG, a SAS threaded application:

```
libname lib oracle user=scott password=tiger;
proc reg simple
data=lib.salesdata (keep=salesnumber maxsales);
var _all_;
run;
```

For DBMSs, many more SAS steps can become eligible for a threaded read, specifically, steps with a read-only table. A libref has the form Lib.DbTable, where Lib is a SAS libref that "points" to DBMS data, and DbTable is a DBMS table. Here are sample read-only tables for which threaded reads can be turned on:

```
libname lib oracle user=scott password=tiger;
proc print data=lib.dbtable;
run;

data local;
set lib.families;
where gender="F";
run;
```

An eligible SAS step can require user assistance in order to actually perform threaded reads. If SAS is unable to automatically generate a partitioning WHERE clause or to otherwise perform threaded reads, then the user can code an option that supplies partitioning. To determine whether SAS can automatically generate a partitioning WHERE clause, use SASTRACE= system option "SASTRACE= System Option" on page 272 and SASTRACELOC= system option "SASTRACELOC= System Option" on page 281.

Threaded reads can be turned off altogether. This eliminates additional DBMS activity associated with SAS threaded reads, such as additional DBMS connections and multiple SQL statements.

Threaded reads are not supported for the Pass-Through Facility, in which you code your own DBMS-specific SQL that is passed directly to the DBMS for processing.

Options That Affect Threaded Reads

SAS/ACCESS provides two options precisely for threaded reads from DBMSs: "DBSLICE= Data Set Option" on page 217 and "DBSLICEPARM= Data Set Option" on page 219.

DBSLICE= is a data set option, applicable only to a table reference. It permits you to code your own WHERE clauses to partition table data across threads, and is useful when you are familiar with your table data. For instance, if your DBMS table has a CHAR(1) column Gender, and your clients are approximately half female, Gender equally partitions the table into two parts. Therefore, here is a possible DBSLICE= example:

```
proc print data=lib.dhtable (dbslice=("gender='f'" "gender='m'"));
  where dbcol>1000;
run;
```

SAS creates two threads and about half of the data is delivered in parallel on each connection.

When applying DBSLICEPARAM=ALL instead of DBSLICE=, SAS attempts to "autopartition" the table for you. With the default DBSLICEPARAM=THREADED_APPS setting, threaded reads are automatically attempted only for SAS threaded applications (SAS procedures that thread I/O and numeric operations). DBSLICEPARAM=ALL extends threaded reads to more SAS procedures, specifically steps that only read tables. Or, DBSLICEPARAM=NONE turns it off entirely. It can be specified as a data set option, LIBNAME option, or as a global SAS option.

The first argument to DBSLICEPARAM= is required and extends or restricts threaded reads. The second, optional argument is not commonly used and limits the number of DBMS connections. The following examples demonstrate the different uses of DBSLICEPARAM=:

- UNIX or Windows SAS invocation option that turns on threaded reads for all read-only libref.

```
--dbsliceparm ALL
```

- Global SAS option that turns off threaded reads.

```
option dbsliceparm=NONE;
```

- LIBNAME option that restricts threaded reads to just SAS threaded applications.

```
libname lib oracle user=scott password=tiger dbsliceparm=THREADED_APPS;
```

- Table option that turns on threaded reads (with a maximum of three connections, in the example below)

```
proc print data=lib.dhtable(dbsliceparm=(ALL,3));
  where dbcol>1000;
run;
```

DBSLICE= and DBSLICEPARAM= apply only to DBMS table reads. THREADS= and CPUCOUNT= are additional SAS options that apply to threaded applications. For more information about these options, see the *SAS Language Reference: Dictionary*.

Generating Trace Information for Threaded Reads

A threaded read is a complex feature. A SAS step can be eligible for a threaded read, but not have it applied. Performance effect is not always easy to predict. Use the SASTRACE option to see if threaded reads occurred and to help assess performance. The following examples demonstrate usage scenarios with SAS/ACCESS to Oracle. Keep in mind that trace output is in English only and changes from release to release.

```

/*Turn on SAS tracing */
options sastrace='',t,' sastraceloc=saslog nostsuffix;

```

```

/* Run a SAS job */

```

```

data work.locemp;
set trlib.MYEMPS(DBSLICEPARM=(ALL,3));
where STATE in ('GA', 'SC', 'NC') and ISTENURE=0;
run;

```

The above job produces the following trace messages:

```

406 data work.locemp;
407 set trlib.MYEMPS(DBSLICEPARM=(ALL, 3));
408 where STATE in ('GA', 'SC', 'NC') and ISTENURE=0;
409 run;

```

```

ORACLE: DBSLICEPARM option set and 3 threads were requested
ORACLE: No application input on number of threads.

```

```

ORACLE: Thread 2 contains 47619 obs.
ORACLE: Thread 3 contains 47619 obs.
ORACLE: Thread 1 contains 47619 obs.
ORACLE: Threaded read enabled. Number of threads created: 3

```

If you want to see the SQL that is executed during the threaded read, you can set tracing to `sastrace=',t,d'` and run the job again. This time, the output contains the threading information as well as all of the SQL statements processed by Oracle:

```

ORACLE_9: Prepared:
SELECT * FROM MYEMPS 418 data work.locemp;

```

```

419 set trlib.MYEMPS(DBSLICEPARM=(ALL, 3));
420 where STATE in ('GA', 'SC', 'NC') and ISTENURE=0;
421 run;

```

```

ORACLE: DBSLICEPARM option set and 3 threads were requested
ORACLE: No application input on number of threads.

```

```

ORACLE_10: Executed:
SELECT "HIREDATE", "SALARY", "GENDER", "ISTENURE", "STATE", "EMPNUM", "NUMCLASSES"
FROM MYEMPS WHERE ( ( "STATE" IN ( 'GA' , 'NC' , 'SC' ) ) ) AND
("ISTENURE" = 0 ) ) AND ABS(MOD("EMPNUM",3))=0

```

```

ORACLE_11: Executed:
SELECT "HIREDATE", "SALARY", "GENDER", "ISTENURE", "STATE", "EMPNUM", "NUMCLASSES"
FROM MYEMPS WHERE ( ( "STATE" IN ( 'GA' , 'NC' , 'SC' ) ) ) AND
("ISTENURE" = 0 ) ) AND ABS(MOD("EMPNUM",3))=1

```

```

ORACLE_12: Executed:
SELECT "HIREDATE", "SALARY", "GENDER", "ISTENURE", "STATE", "EMPNUM", "NUMCLASSES"
FROM MYEMPS WHERE ( ( "STATE" IN ( 'GA' , 'NC' , 'SC' ) ) ) AND
("ISTENURE" = 0 ) ) AND (ABS(MOD("EMPNUM",3))=2 OR "EMPNUM" IS NULL)

```

```

ORACLE: Thread 2 contains 47619 obs.
ORACLE: Thread 1 contains 47619 obs.
ORACLE: Thread 3 contains 47619 obs.
ORACLE: Threaded read enabled. Number of threads created: 3

```

Notice that the Oracle engine used the EMPNUM column as a partitioning column.

If a threaded read cannot be done either because all of the candidates for autopartitioning are in the WHERE clause or because the table does not contain a column that fits the criteria, you see a warning in your log. For example, the data set below uses a WHERE clause that contains all of the possible autopartitioning columns:

```

data work.locemp;
set trlib.MYEMPS (DBLSICEPARM=ALL);
where EMPNUM<=30 and ISTEURE=0 and SALARY<=35000 and NUMCLASSES>2;
run;

```

You receive the following warnings:

```

ORACLE: No application input on number of threads.
ORACLE: WARNING: Unable to find a partition column for use w/ MOD()
ORACLE: The engine cannot automatically generate the partitioning WHERE clauses.
ORACLE: Using only one read connection.
ORACLE: Threading is disabled due to an error. Application reverts to non-threading
I/O's.

```

If the SAS job contains any options that are invalid when the engine tries to perform threading, you also receive a warning.

```

libname trlib oracle user=orauser pw=orapw path=oraserver DBSLICEPARM=(ALL);

proc print data=trlib.MYEMPS (OBS=10);
where EMPNUM<=30;
run;

```

This produces the following message:

```

ORACLE: Threading is disabled due to the ORDER BY clause or the FIRSTOBS/OBS option.
ORACLE: Using only one read connection.

```

To produce timing information, add an 's' in the last slot of sastrace. For example:

```

options sastrace=',,t,s' sastraceloc=saslog nostsuffix;

data work.locemp;
set trlib.MYEMPS (DBSLICEPARM=ALL);
where EMPNUM<=10000;
run;

```

This produces the following:

```

ORACLE: No application input on number of threads.
ORACLE: Thread 1 contains 5000 obs.
ORACLE: Thread 2 contains 5000 obs.

```

```

Thread 0 fetched 5000 rows
DBMS Threaded Read Total Time: 1234 mS
DBMS Threaded Read User CPU: 46 mS
DBMS Threaded Read System CPU: 0 mS

```



```

Thread 1 fetched 5000 rows
DBMS Threaded Read Total Time: 469 mS
DBMS Threaded Read User CPU: 15 mS
DBMS Threaded Read System CPU: 15 mS
ORACLE: Threaded read enabled. Number of threads created: 2
NOTE: There were 10000 observations read from the data set TRLIB.MYEMPS.
      WHERE EMPNUM<=10000;

```

```

Summary Statistics for ORACLE are: Total SQL prepare seconds were: 0.001675
Total seconds used by the ORACLE ACCESS engine were 7.545805

```

For more information regarding tracing, please see the SASTRACE documentation.

Performance Impact of Threaded Reads

Threaded reads only increase performance when the DBMS result set is large. Performance is optimal when the partitions are similar in size. Threaded reads should reduce the elapsed time of your SAS step, but unusual cases can slow the SAS step. They generally increase the workload on your DBMS.

For instance, threaded reads for DB2 under z/OS involve a tradeoff, generally reducing job elapsed time but increasing DB2 workload and CPU utilization. See the auto partitioning documentation for DB2 under z/OS for details.

SAS automatically tries to autopartition table references for SAS in threaded applications. To determine whether autopartitioning is occurring and to assess its performance, complete the following tasks:

- Turn on SAS tracing to determine whether SAS is autopartitioning and to view the SQL associated with each thread.
- Know your DBMS algorithm for autopartitioning.
- Turn threaded reads on and off, and compare the elapsed times.

To optimally tune threaded reads, follow these guidelines:

- Use it only when pulling large result sets into SAS from the DBMS.
- Use DBSLICE= to partition if SAS autopartitioning does not occur.
- Override autopartitioning with DBSLICE= if you can manually provide substantially better partitioning. The best partitioning equally distributes the result set across the threads.
- Consult the DBMS-specific section of this documentation for information and tips concerning your specific DBMS.

Threaded reads are most effective on new, faster computer hardware running SAS, and with a powerful parallel edition of the DBMS. For example, if SAS runs on a fast uniprocessor or on a multiprocessor machine and your DBMS runs on a high-end SMP server, you receive substantial performance gains. However you receive minimal gains or even performance degradation when running SAS on an old desktop model with a nonparallel DBMS edition running on old hardware.

Autopartitioning Techniques in SAS/ACCESS

SAS/ACCESS products share an autopartitioning scheme based on the MOD function. Some products support additional techniques. For example, if your Oracle tables are physically partitioned in the DBMS, the SAS/ACCESS interface to Oracle

automatically partitions in accordance with Oracle physical partitions rather than using MOD. The SAS/ACCESS interface to Teradata uses FastExporting, if available, which enables the FastExport utility to direct the partitioning.

MOD is a mathematical function that produces the remainder of a division operation. Your DBMS table must contain a column to which SAS can apply the MOD function — a numeric column constrained to integral values. DBMS integer and small integer columns suit this purpose. Integral decimal (numeric) type columns can work as well. On each thread, SAS appends a WHERE clause to your SQL that uses the MOD function with the numeric column to create a subset of the result set. Combined, these subsets add up to exactly the result set for your original single SQL statement.

For example, assume that your original SAS-produced SQL is **SELECT CHR1, CHR2 FROM DBTAB** and that table Ddbtab contains integer column IntCol. SAS creates two threads and issues:

```
SELECT CHR1, CHR2 FROM DBTAB WHERE (MOD(INTCOL,2)=0)
```

and

```
SELECT CHR1, CHR2 FROM DBTAB WHERE (MOD(INTCOL,2)=1)
```

Rows with an even value for IntCol are retrieved by the first thread. Rows with an odd value for IntCol are retrieved by the second thread. Distribution of rows across the two threads is optimal if IntCol has a 50/50 distribution of even and odd values.

SAS modifies the SQL for columns containing negative integers, for nullable columns, and to combine SAS WHERE clauses with the partitioning WHERE clauses. SAS can also run more than two threads. You use the second parameter of the DBSLICEPARM= option to increase the number of threads.

The success of this technique depends on the distribution of the values in the chosen integral column. Without knowledge of the distribution, your SAS/ACCESS product attempts to pick the best possible column. For example, indexed columns are given preference for some DBMSs. However, column selection is more or less a guess, and the SAS guess might cause poor distribution of the result set across the threads. If no suitable numeric column is found, MOD cannot be used at all, and threaded reads do not occur if your SAS/ACCESS product has no other partitioning technique. For these reasons, you should explore autopartitioning particulars for your DBMS and judiciously utilize DBSLICE= to augment autopartitioning. For details about autopartitioning, see the documentation for your DBMS.

Data Ordering in SAS/ACCESS

The order in which table rows are delivered to SAS varies each time a step is rerun with threaded reads. Most DBMS editions, especially increasingly popular parallel editions, do not guarantee consistent ordering.

Two-Pass Processing for SAS Threaded Applications

Two-pass processing occurs when a SAS threaded application “Scope of Threaded Reads” on page 49 requests that data be made available for multiple pass reading (that is, more than one pass through the data set). In the context of DBMS engines, this requires that as the data is read from the database, temporary spool files are written containing the read data. There is one temporary spool file per thread, and each spool

file contains all the data read on that thread. If three threads are specified for threaded reads, then three temporary spool files are written.

As the application requests subsequent passes of the data, the data is read from the temporary spool files, not re-read from the database. The temporary spool files can be written on different disks, reducing any disk read contention, and enhancing performance. To accomplish this, the SAS option UTILLOC= is used to define different disk devices and directory paths when creating temporary spool files. There are several different ways to specify this option:

- In the SAS config file, add the line:

```
--utilloc("C:\path" "D:\path" "E:\path")
```

- Specify the UTILLOC= option on the SAS command line:

on Windows:

```
sas --utilloc(c:\path d:\path e:\path)
```

on UNIX:

```
sas --utilloc '(\path \path2 \path3)'
```

For more information about the UTILLOC= SAS option, see the *SAS Language Reference: Dictionary*.

When Threaded Reads Will Not Occur

Threading does not occur


- when a BY statement is used in a PROC or DATA step
- when the OBS or the FIRSTOBS option is in a PROC or DATA step
- when the KEY or the DBKEY option is used PROC or DATA step
- if no column in the table exists to which SAS can apply the MOD function. For more information, see Autopartitioning Techniques in SAS“Autopartitioning Techniques in SAS/ACCESS” on page 53.
- if all columns within a table to which SAS can apply the MOD function are in WHERE clauses.

For more information, see Autopartitioning Techniques in SAS“Autopartitioning Techniques in SAS/ACCESS” on page 53.

- if the NOTTHREADS system option is set
- if DBSLICEPARAM=NONE

Summary of Threaded Reads

For large reads of table data, SAS 9.1 threaded reads can speed up SAS jobs. They are particularly useful when you understand the autopartitioning technique specific to your DBMS and use DBSLICE= to manually partition only when appropriate. Look for enhancements in future SAS releases.



CHAPTER

7

How SAS/ACCESS Works

<i>Introduction to How SAS/ACCESS Works</i>	57
<i>Installation Requirements</i>	57
<i>SAS/ACCESS Interfaces</i>	58
<i>How the SAS/ACCESS LIBNAME Statement Works</i>	58
<i>Accessing Data From a DBMS Object</i>	58
<i>Processing Queries, Joins, and Data Functions</i>	59
<i>How the Pass-Through Facility Works</i>	59
<i>How the ACCESS Procedure Works</i>	60
<i>Overview of the ACCESS Procedure</i>	60
<i>Reading Data</i>	60
<i>Updating Data</i>	61
<i>How the DBLOAD Procedure Works</i>	62

Introduction to How SAS/ACCESS Works

Note: Not all features are supported by all SAS/ACCESS interfaces. See the documentation for your SAS/ACCESS interface to determine which features are supported in your environment. △

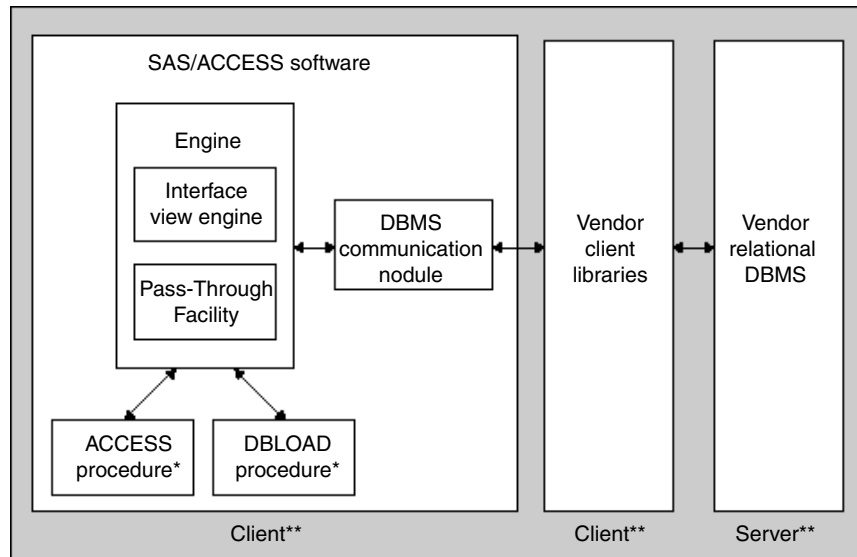
Installation Requirements

Before you use any SAS/ACCESS features, you must install Base SAS, the SAS/ACCESS interface for the DBMS that you are accessing, and any required DBMS client software. See the SAS installation instructions and DBMS client installation instructions for more information.

SAS/ACCESS Interfaces

Each SAS/ACCESS interface consists of one or more data access engines that translate read and write requests from SAS into appropriate calls for a specific DBMS. The following image depicts the relationship between a SAS/ACCESS interface and a relational DBMS.

Figure 7.1 How SAS Connects to the DBMS



* The ACCESS procedure and the DBLOAD procedure are not supported by all SAS/ACCESS interfaces.

** In some cases, both client and server software can reside on the same machine.

You can invoke a SAS/ACCESS relational DBMS interface by using either a LIBNAME statement or a PROC SQL statement. (You can also use the ACCESS and DBLOAD procedures with some of the SAS/ACCESS relational interfaces. However, these procedures are no longer the recommended way to access relational database data.)

How the SAS/ACCESS LIBNAME Statement Works

Accessing Data From a DBMS Object

SAS/ACCESS enables you to read, update, insert, and delete data from a DBMS object as if it were a SAS data set. The process is as follows:

- 1 You invoke a SAS/ACCESS interface by specifying a DBMS engine name and the appropriate connection options in a LIBNAME statement.
- 2 You enter SAS requests as you would when accessing a SAS data set.
- 3 SAS/ACCESS generates DBMS-specific SQL statements that are equivalent to the SAS requests that you enter.
- 4 SAS/ACCESS submits the generated SQL to the DBMS.

The SAS/ACCESS engine defines which operations are supported on a table and calls code that translates database operations such as open, get, put, or delete into DBMS-specific SQL syntax. SAS/ACCESS engines use an established set of routines with calls that are tailored to each DBMS.

Processing Queries, Joins, and Data Functions

To enhance performance, SAS/ACCESS can also transparently pass a growing number of queries, joins, and data functions to the DBMS for processing (instead of retrieving the data from the DBMS and then doing the processing in SAS). For example, an important use of this feature is the handling of PROC SQL queries that access DBMS data. Here's how it works:

- 1 PROC SQL examines each query to determine whether it might be profitable to send all or part of the query to the DBMS for processing.
- 2 A special query textualizer in PROC SQL translates queries (or query fragments) into DBMS-specific SQL syntax.
- 3 The query textualizer submits the translated query to the SAS/ACCESS engine for approval.
- 4 If SAS/ACCESS approves the translation, it sends an approval message to PROC SQL and the query (or query fragment) gets processed by the DBMS, which returns the results to SAS. Any queries or query fragments that can not be passed to the DBMS are processed in SAS.

See the chapter on performance considerations for detailed information about tasks that SAS/ACCESS can pass to the DBMS.

How the Pass-Through Facility Works

When you read and update DBMS data with the Pass-Through Facility, SAS/ACCESS passes SQL statements directly to the DBMS for processing. Here are the steps:

- 1 Invoke PROC SQL and submit a PROC SQL CONNECT statement that includes a DBMS name and the appropriate connection options to establish a connection with a specified database.
- 2 Use a CONNECTION TO component in a PROC SQL SELECT statement to read data from a DBMS table or view. In the SELECT statement (that is, the PROC SQL query) that you write, use the SQL that is native to your DBMS. SAS/ACCESS passes the SQL statements directly to the DBMS for processing. If the SQL syntax that you enter is correct, the DBMS processes the statement and returns any results to SAS. If the DBMS does not recognize the syntax that you enter, it returns an error that appears in the SAS log. The SELECT statement can be stored as a PROC SQL view. For example:

```
proc sql;
connect to oracle (user=scott password=tiger);
create view budget2000 as select amount_b,amount_s
    from connection to oracle
    (select Budgeted, Spent from annual_budget);
quit;
```

- 3 Use a PROC SQL EXECUTE statement to pass any dynamic, non-query SQL statements (such as INSERT, DELETE, and UPDATE) to the database. As with the CONNECTION TO component, all EXECUTE statements are passed to the DBMS exactly as you submit them. INSERT statements must contain literal values. For example:

```
proc sql;
  connect to oracle(user=scott password=tiger);
  execute (create view whotookorders as select ordernum, takenby,
    firstname, lastname, phone from orders, employees
    where orders.takenby=employees.empid) by oracle;
  execute (grant select on whotookorders to testuser) by oracle;
  disconnect from oracle;
quit;
```

- 4 Terminate the connection with the DISCONNECT statement.

See Chapter 12, “The Pass-Through Facility for Relational Databases,” on page 285 for more details.

How the ACCESS Procedure Works

Overview of the ACCESS Procedure

When you use the ACCESS procedure to create an access descriptor, the SAS/ACCESS interface view engine requests the DBMS to execute a SQL SELECT statement to the data dictionary tables in your DBMS dynamically (by using DBMS-specific call routines or interface software). The ACCESS procedure then issues the equivalent of a DESCRIBE statement to gather information about the columns in the specified table. The access descriptor’s information about the table and its columns is then copied into the view descriptor when it is created. Therefore, it is not necessary for SAS to call the DBMS when it creates a view descriptor.

The process is as follows:

- 1 When you supply the connection information to PROC ACCESS, the SAS/ACCESS interface calls the DBMS to connect to the database.
- 2 SAS constructs a SELECT * FROM *table-name* statement and passes it to the DBMS to retrieve information about the table from the DBMS data dictionary. This SELECT statement is based on the information you supplied to PROC ACCESS. It enables SAS to determine whether the table exists and can be accessed.
- 3 The SAS/ACCESS interface calls the DBMS to get table description information, such as the column names, data types (including width, precision, and scale), and whether the columns accept null values.
- 4 SAS closes the connection with the DBMS.

Reading Data

When you use a view descriptor in a DATA step or procedure to read DBMS data, the SAS/ACCESS interface view engine requests the DBMS to execute a SQL SELECT statement. The interface view engine follows these steps:

- 1 Using the connection information that is contained in the created view descriptor, the SAS/ACCESS interface calls the DBMS to connect to the database.

- 2 SAS constructs a SELECT statement that is based on the information stored in the view descriptor (table name and selected columns and their characteristics) and passes this information to the DBMS.
- 3 SAS retrieves the data from the DBMS table and passes it back to the SAS procedures as if it were observations in a SAS data set.
- 4 SAS closes the connection with the DBMS.

For example, if you execute the following SAS program using a view descriptor, the previous steps are executed once for the PRINT procedure and then a second time for the GCHART procedure. (The data used for the two procedures is not necessarily the same because the table might have been updated by another user between procedure executions.)

```
proc print data=vlib.allemp;
run;

proc gchart data=vlib.allemp;
  vbar jobcode;
run;
```

Updating Data

You use a view descriptor, DATA step, or procedure to update DBMS data in much the same way as when reading data. Any of the following steps might also occur:


- Using the connection information that is contained in the specified access descriptor, the SAS/ACCESS interface calls the DBMS to connect to the database.
- When rows are added to a table, SAS constructs a SQL INSERT statement and passes it to the DBMS. When you reference a view descriptor, you can use the ADD command in FSEDIT and FSVIEW, the APPEND procedure, or an INSERT statement in PROC SQL to add data to a DBMS table. (You can also use the Pass-Through Facility's EXECUTE statement to add, delete, or modify DBMS data directly. Literal values must be used when inserting data with the Pass-Through Facility.)
- When rows are deleted from a DBMS table, SAS constructs a SQL DELETE statement and passes it to the DBMS. When you reference a view descriptor, you can use the DELETE command in FSEDIT and FSVIEW or a DELETE statement in PROC SQL to delete rows from a DBMS table.
- When data in the rows is modified, SAS constructs a SQL UPDATE statement and passes it to the DBMS. When you reference a view descriptor, you can use FSEDIT, the MODIFY command in FSVIEW, or an INSERT statement in PROC SQL to update data in a DBMS table. You can also reference a view descriptor in the DATA step's UPDATE, MODIFY, and REPLACE statements.
- SAS closes the connection with the DBMS.

How the DBLOAD Procedure Works

When you use the DBLOAD procedure to create a DBMS table, the procedure issues dynamic SQL statements to create the table and insert data from a SAS data file, DATA step view, PROC SQL view, or view descriptor into the table.


The SAS/ACCESS interface view engine completes the following steps:

- 1 When you supply the connection information to PROC DBLOAD, the SAS/ACCESS interface calls the DBMS to connect to the database.
- 2 SAS uses the information that is provided by the DBLOAD procedure to construct a `SELECT * FROM table-name` statement, and passes the information to the DBMS to determine if the table already exists. PROC DBLOAD continues only if a table with that name does not exist, unless you use the DBLOAD APPEND option.
- 3 SAS uses the information that is provided by the DBLOAD procedure to construct a `SQL CREATE TABLE` statement and passes it to the DBMS.
- 4 SAS constructs a `SQL INSERT` statement for the current observation and passes it to the DBMS. New INSERT statements are constructed and then executed repeatedly until all of the observations from the input SAS data set are passed to the DBMS. Some DBMSs have a bulkcopy capability that allows a group of observations to be inserted at once. See your DBMS documentation to determine if your DBMS has this capability.
- 5 Additional nonquery SQL statements specified in the DBLOAD procedure are executed as submitted by the user. The DBMS returns an error message if a statement does not execute successfully.
- 6 SAS closes the connection with the DBMS.



General Reference

- Chapter 8*. **SAS/ACCESS Features by Host** 65
- Chapter 9*. **The LIBNAME Statement for Relational Databases** 73
- Chapter 10*. **Data Set Options for Relational Databases** 157
- Chapter 11*. **Macro Variables and System Options for Relational Databases** 269
- Chapter 12*. **The Pass-Through Facility for Relational Databases** 285



CHAPTER

8

SAS/ACCESS Features by Host

<i>Introduction</i>	65
<i>SAS/ACCESS Interface to DB2 under UNIX and PC Hosts: Supported Features</i>	65
<i>SAS/ACCESS Interface to DB2 under z/OS: Supported Features</i>	66
<i>SAS/ACCESS Interface to HP Neoview: Supported Features</i>	66
<i>SAS/ACCESS Interface to Informix: Supported Features</i>	67
<i>SAS/ACCESS Interface to Microsoft SQL Server: Supported Features</i>	67
<i>SAS/ACCESS Interface to MySQL: Supported Features</i>	68
<i>SAS/ACCESS Interface to Netezza: Supported Features</i>	68
<i>SAS/ACCESS Interface to ODBC: Supported Features</i>	69
<i>SAS/ACCESS Interface to OLE DB: Supported Features</i>	70
<i>SAS/ACCESS Interface to Oracle: Supported Features</i>	70
<i>SAS/ACCESS Interface to Sybase: Supported Features</i>	71
<i>SAS/ACCESS Interface to Teradata: Supported Features</i>	72

Introduction

This section lists by host environment the features that are supported in each SAS/ACCESS relational interface.

SAS/ACCESS Interface to DB2 under UNIX and PC Hosts: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to DB2 under UNIX and PC hosts. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.1 Features by Host Environment for DB2 under UNIX and PC Hosts

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
AIX (RS/ 6000)	X	X		X	X
HP-UX	X	X		X	X

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
Microsoft Windows NT & 2000	X	X		X	X
Linux for Intel	X	X		X	X
Linux for Itanium	X	X		X	X
Solaris for SPARC (64-bit)	X	X		X	X
Solaris for x64	X	X		X	X

SAS/ACCESS Interface to DB2 under z/OS: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to DB2 under z/OS. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.2 Features by Host Environment for DB2 under z/OS

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
z/OS	X	X	X	X	X

SAS/ACCESS Interface to HP Neoview: Supported Features

The following table lists the features that SAS/ACCESS Interface to HP Neoview supports. To find out which versions of your DBMS are supported, see your System Requirements manual.

Table 8.3 Features by Host Environment for HP Neoview

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
HP-UX for Itanium	X	X			
Linux for Intel	X	X			

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
Microsoft Windows for Intel	X	X			
Solaris on SPARC (64-bit)	X	X			

SAS/ACCESS Interface to Informix: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to Informix. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.4 Features by Host Environment for Informix

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
AIX (RS/ 6000)	X	X			
HP-UX	X	X			
HP-UX for Itanium	X	X			
Linux for Itanium	X	X			
Solaris for SPARC (64-bit)	X	X			

SAS/ACCESS Interface to Microsoft SQL Server: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to Microsoft SQL Server. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.5 Features by Host Environment for Microsoft SQL Server

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
AIX (RS/ 6000)	X	X		X	
HP-UX	X	X		X	

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
Linux for Itanium	X	X		X	
Solaris for SPARC (64-bit)	X	X		X	

SAS/ACCESS Interface to MySQL: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to MySQL. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.6 Features by Host Environment for MySQL

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
AIX (RS/ 6000)	X	X			
HP-UX	X	X			
Linux for Intel	X	X			
Linux for Itanium	X	X			
Microsoft Windows NT and 2000	X	X			
Solaris for SPARC (64-bit)	X	X			
Solaris for x64	X	X			

SAS/ACCESS Interface to Netezza: Supported Features

The following table lists the features that SAS/ACCESS Interface to Netezza supports. To find out which versions of your DBMS are supported, see your System Requirements manual.

Table 8.7 Features by Host Environment for Netezza

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
AIX (RS/ 6000)	X	X			X
HP-UX	X	X			X
HP-UX for Itanium	X	X			X
Linux for Intel	X	X			X
Microsoft Windows for Intel	X	X			X
Solaris on SPARC (64-bit)	X	X			X

For information about these features, see “Methods for Accessing Relational Database Data” on page 3 and the LIBNAME option BULKLOAD=“BULKLOAD=LIBNAME Option” on page 84.

SAS/ACCESS Interface to ODBC: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to ODBC. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.8 Features by Host Environment for ODBC

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
AIX (RS/ 6000)	X	X		X	
HP-UX	X	X		X	
HP-UX for Itanium	X	X		X	
HP-UX	X	X		X	
Linux for Intel	X	X		X	
Linux for Itanium	X	X		X	
Microsoft Windows NT & 2000	X	X		X	X*

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
Solaris for SPARC (64-bit)	X	X		X	
Solaris for x64	X	X		X	

* Bulk load support is available only with the SQL Server driver on Windows platforms.

SAS/ACCESS Interface to OLE DB: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to OLE DB. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.9 Features by Host Environment for OLE DB

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
Microsoft Windows NT & 2000	X	X			X

SAS/ACCESS Interface to Oracle: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to Oracle. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.10 Features by Host Environment for Oracle

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
HP-UX	X	X	X	X	X
Linux for Intel	X	X	X	X	X
Linux for Itanium	X	X	X	X	X
Microsoft Windows NT & 2000	X	X	X	X	X
OpenVMS Alpha	X	X	X	X	X

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
Solaris for SPARC (64-bit)	X	X	X	X	X
Solaris for x64	X	X	X	X	X
z/OS	X	X	X	X	X

SAS/ACCESS Interface to Sybase: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to Sybase. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.11 Features by Host Environment for Sybase


	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
AIX (RS/ 6000)	X	X	X	X	X
HP-UX	X	X	X	X	X
HP-UX for Itanium	X	X	X	X	X
Microsoft Windows NT & 2000	X	X	X	X	X
Linux for Intel	X	X	X	X	X
Linux for Itanium	X	X	X	X	X
Solaris for SPARC (64-bit)	X	X	X	X	X
Solaris for x64	X	X	X	X	X

SAS/ACCESS Interface to Teradata: Supported Features

The following table lists the features that are supported in the SAS/ACCESS Interface to Teradata. To find out which versions of your DBMS are supported, refer to your System Requirements manual.

Table 8.12 Features by Host Environment for Teradata

	SAS/ACCESS LIBNAME Statement	Pass-Through Facility	ACCESS Procedure	DBLOAD Procedure	Bulk-Load Support
AIX (RS/6000)	X	X			X
HP-UX	X	X			X
HP-UX for Itanium	X	X			X
Linux for Intel	X	X			X
Microsoft Windows NT & 2000	X	X			X
Solaris for SPARC (64-bit)	X	X			X
Solaris for x64	X	X			X
z/OS	X	X			X



CHAPTER

9

The LIBNAME Statement for Relational Databases

<i>Overview of the LIBNAME Statement for Relational Databases</i>	73
<i>Assigning Librefs</i>	73
<i>Sorting Data</i>	73
<i>Using SAS Functions</i>	74
<i>Assigning a Libref Interactively</i>	74
<i>LIBNAME Options for Relational Databases</i>	78

Overview of the LIBNAME Statement for Relational Databases

Assigning Librefs

The SAS/ACCESS LIBNAME statement extends the SAS global LIBNAME statement to enable you to assign a libref to a relational DBMS. This feature lets you reference a DBMS object directly in a DATA step or SAS procedure. You can use this direct reference to read from and write to a DBMS object as if it were a SAS data set. You can associate a SAS libref with a relational DBMS database, schema, server, or group of tables and views.

Sorting Data

When you use the SAS/ACCESS LIBNAME statement to associate a libref with relational DBMS data, you might observe some behavior that differs from that of normal SAS librefs. Because these librefs refer to database objects such as tables and views, they are stored in the format of your DBMS. This format differs from the format of normal SAS data sets. This is helpful to remember when you access and work with DBMS data.

For example, you can sort the observations in a normal SAS data set and store the output to another data set. However, in a relational DBMS, sorting data often has no effect on how it is stored. Because you cannot depend on your data to be sorted in the DBMS, you must sort the data at the time of query. Furthermore, when you sort DBMS data, the results can vary depending on whether your DBMS places data with NULL values at the beginning or the end of the result set. The NULL values are translated in SAS to missing values.

Using SAS Functions

When you use librefs that refer to DBMS data with SAS functions, some functions might return a value that differs from what is returned when you use functions with normal SAS data sets. For example, the `PATHNAME` function might return a blank value. For a normal SAS libref, a blank value means that the libref is not valid. However, for a libref associated with a DBMS object, a blank value means only that there is no pathname associated with the libref.

Usage of some functions might also vary. For example, the `LIBNAME` function can accept an optional *SAS-data-library* argument. When you use the `LIBNAME` function to assign or de-assign a libref that refers to DBMS data, you omit this argument. For full details about how to use SAS functions, see the *SAS Language Reference: Dictionary*.

Assigning a Libref Interactively

An easy way to associate a libref with a relational DBMS is to use the New Library window. One method to open this window is to issue the `DMLIBASSIGN` command from your SAS session's command box or command line. The window can also be opened by clicking the file cabinet icon in the SAS Explorer toolbar. In the display below, the user Samantha assigns a libref `MYORADB` to an Oracle database referred to by the `SQL*Net` alias `ORahrdept`. The `LIBNAME` option, `SCHEMA=`, enables Samantha to access database objects that are owned by another user.

Display 9.1 New Library Window

The following list describes how to use the features of the New Library window:

- **Name:** enter the libref that you want to assign to a SAS library or a relational DBMS.
- **Engine:** click the down arrow to select a name from the pull-down listing.
- **Enable at startup:** click this if you want the specified libref to be assigned automatically when you open a SAS session.
- **Library Information:** these fields represent the SAS/ACCESS connection options and vary according to the SAS/ACCESS engine that you specify. Enter the appropriate information for your site's DBMS. You can use the **Options** field to enter SAS/ACCESS `LIBNAME` options. Use blanks to separate multiple options.
- **OK:** click this button to assign the libref, or click **Cancel** to exit the window without assigning a libref.

LIBNAME Statement Syntax for Relational Databases

Associates a SAS libref with a DBMS database, schema, server, or group of tables and views

Valid: Anywhere

Syntax

- ❶ **LIBNAME** *libref* *engine-name*
 <SAS/ACCESS-connection-options>
 <SAS/ACCESS-LIBNAME-options>;
- ❷ **LIBNAME** *libref* CLEAR | _ALL_ CLEAR;
- ❸ **LIBNAME** *libref* LIST | _ALL_ LIST;

Arguments

The SAS/ACCESS LIBNAME statement takes the following arguments:

libref

is any SAS name that serves as an alias to associate SAS with a database, schema, server, or group of tables and views. Like the global SAS LIBNAME statement, the SAS/ACCESS LIBNAME statement creates shortcuts or nicknames for data storage locations. While a SAS libref is an alias for a virtual or physical directory, a SAS/ACCESS libref is an alias for the DBMS database, schema, or server where your tables and views are stored.

engine-name

is the SAS/ACCESS engine name for your DBMS, such as **oracle** or **db2**. The engine name is required. Because the SAS/ACCESS LIBNAME statement associates a libref with a SAS/ACCESS engine that supports connections to a particular DBMS, it requires a DBMS-specific engine name.

See the documentation for your SAS/ACCESS interface to find your engine's name.

SAS/ACCESS-connection-options

provide connection information and control how SAS manages the timing and concurrence of the connection to the DBMS; these arguments are different for each database. Here is an example of connecting to an Oracle database. Your connection options are USER=, PASSWORD=, and PATH=:

```
libname myoralib oracle user=testuser password=testpass path='voyager';
```

If the connection options contain characters that are not allowed in SAS names, enclose the values of the arguments in quotation marks. On some DBMSs, if you specify the appropriate system options or environment variables for your database, you can omit the connection options.

See the documentation for your SAS/ACCESS interface for detailed information about your connection options.

SAS/ACCESS-LIBNAME-options

define how DBMS objects are processed by SAS. Some LIBNAME options can enhance performance; others determine locking or naming behavior. For example, you can use the PRESERVE_COL_NAMES= option to specify whether to preserve spaces, special characters, and mixed case in DBMS column names when creating tables. The availability and default behavior of many of these options are DBMS-specific.

See the documentation for your SAS/ACCESS interface for a list of the LIBNAME options that are available for your DBMS.

See “LIBNAME Options for Relational Databases” on page 78 for detailed information about all of the LIBNAME options.

CLEAR

disassociates one or more currently assigned librefs.

Specify *libref* to disassociate a single libref. Specify `_ALL_` to disassociate all currently assigned librefs.

`_ALL_`

specifies that the CLEAR or LIST argument applies to all currently-assigned librefs.

LIST

writes the attributes of one or more SAS/ACCESS libraries or SAS data libraries to the SAS log.

Specify *libref* to list the attributes of a single SAS/ACCESS library or SAS library. Specify `_ALL_` to list the attributes of all libraries that have librefs in your current session.

Details

1 Using Data from a DBMS You can use a LIBNAME statement to read from and write to a DBMS table or view as if it were a SAS data set.

For example, in MYDBLIB.EMPLOYEES_Q2, MYDBLIB is a SAS libref that points to a particular group of DBMS objects, and EMPLOYEES_Q2 is a DBMS table name. When you specify MYDBLIB.EMPLOYEES_Q2 in a DATA step or procedure, you dynamically access the DBMS table. SAS supports reading, updating, creating, and deleting DBMS tables dynamically.

2 Disassociating a Libref from a SAS Library To disassociate or clear a libref from a DBMS, use a LIBNAME statement. Specify the libref (for example, MYDBLIB) and the CLEAR option as follows:

```
libname mydblib CLEAR;
```

You can clear a single specified libref or all current librefs.

The database engine disconnects from the database and closes any free threads or resources that are associated with that libref's connection.

3 Writing SAS Library Attributes to the SAS Log Use a LIBNAME statement to write the attributes of one or more SAS/ACCESS libraries or SAS data libraries to the SAS log. Specify *libref* to list the attributes of a single SAS/ACCESS library or SAS library, as follows:

```
libname mydblib LIST;
```

Specify `_ALL_` to list the attributes of all libraries that have librefs in your current session.

SQL Views with Embedded LIBNAME Statements

With SAS software, you can embed LIBNAME statements in the definition of an SQL view. This means that you can store, in an SQL view, a LIBNAME statement that contains all the information required to connect to a DBMS. Whenever the SQL view is read, PROC SQL uses the embedded LIBNAME statement to assign a libref. After the view has been processed, PROC SQL de-assigns the libref.

In the following example, an SQL view of the Oracle table DEPT is created. Whenever this view is used in a SAS program, the library ORALIB is assigned. The library ORALIB uses the connection information (user name, password, and data source) that is provided in the embedded LIBNAME statement.

```
proc sql;
  create view sasuser.myview as
    select dname from oralib.dept
    using libname oralib oracle
    user=scott pw=tiger datasrc=orsrv;
quit;
```

Note: The USING LIBNAME syntax is used to embed LIBNAME statements in SQL views. For more information about the USING LIBNAME syntax, see the PROC SQL topic in the *Base SAS Procedures Guide*. Δ

Assigning a Libref with a SAS/ACCESS LIBNAME Statement

The following statement creates a libref, MYDBLIB, that uses the SAS/ACCESS interface for DB2:

```
libname mydblib db2 ssid=db2a authid=testid server=os390svr;
```

The following statement associates the SAS libref MYDBLIB with an Oracle database that uses the SQL*Net alias AIRDB_REMOTE. You specify the SCHEMA= option on the SAS/ACCESS LIBNAME statement to connect to the Oracle schema in which the database resides. In this example, Oracle schemas reside in a database.

```
libname mydblib oracle user=testuser password=testpass
  path=airdb_remote schema=hrdept;
```

The AIRDB_REMOTE database contains a number of DBMS objects, including several tables, such as STAFF. After you assign the libref, you can reference the Oracle table like a SAS data set and use it as a data source in any DATA step or SAS procedure. In the following SQL procedure statement, MYDBLIB.STAFF is the two-level SAS name for the STAFF table in the Oracle database AIRDB_REMOTE:

```
proc sql;
  select idnum, lname
    from mydblib.staff
   where state='NY'
   order by lname;
quit;
```

You can use the DBMS data to create a SAS data set:

```
data newds;
  set mydblib.staff(keep=idnum lname fname);
run;
```

You can also use the libref and data set with any other SAS procedure. This statement prints the information in the STAFF table:

```
proc print data=mydblib.staff;
run;
```

This statement lists the database objects in the MYDBLIB library:

```
proc datasets library=mydblib;
quit;
```

Using the Prompting Window When Specifying LIBNAME Options

The following statement uses the DBPROMPT= option to cause the DBMS connection prompting window to appear and prompt you for connection information:

```
libname mydblib oracle dbprompt=yes;
```

When you use this option, you enter connection information into the fields in the prompting window rather than on the LIBNAME statement.

You can add the DEFER=NO option to make the prompting window appear at the time that the libref is assigned rather than when the table is opened:

```
libname mydblib oracle dbprompt=yes defer=no;
```

Assigning a Libref to a Remote DBMS

SAS/CONNECT (single-user) and SAS/SHARE (multiple user) software give you access to data by means of *remote library services* (RLS). You can use RLS to access your data on a remote machine as if it were local. For example, it permits a graphical interface to reside on the local machine while the data remains on the remote machine.

This access is given to data stored in many kinds of SAS files, such as external databases (through the SAS/ACCESS LIBNAME statement and views created with it) and SAS data views (views created with PROC SQL, the DATA step, and SAS/ACCESS software). You can use RLS to access SAS data sets, SAS views, and relational DBMS data that are defined by SAS/ACCESS LIBNAME statements. For more information, see the discussion about remote library services in the *SAS/SHARE User's Guide*.

You can use RLS to update relational DBMS tables that are referenced with the SAS/ACCESS LIBNAME statement.

In the following example, the SAS/ACCESS LIBNAME statement makes a connection to a DB2 database that resides on the remote SAS/SHARE server REMOS390. This LIBNAME statement is submitted in a local SAS session. The SAS/ACCESS engine name is specified in the remote option RENGINE=. The DB2 connection option and any LIBNAME options are specified in the remote option ROPTIONS=. Options are separated by a blank space. RLSDB2.EMPLOYEES is a SAS data set that references the DB2 table EMPLOYEES.

```
libname rlsdb2 rengine=db2 server=remos390
        roptions="ssid=db2a authid=testid";

proc print data=rlsdb2.employees;
run;
```

See Also

“Overview of the LIBNAME Statement for Relational Databases” on page 73

LIBNAME Options for Relational Databases

11OCT2007: Added “Alias” label for all options having aliases. —Kat Turk

When you specify an option in the LIBNAME statement, it applies to all objects (such as tables and views) in the database that the libref represents. For information

about options that you specify on individual SAS data sets, see the chapter about data set options.

Many LIBNAME options are also available for use with the Pass-Through Facility. See the section on the Pass-Through Facility in the documentation for your SAS/ACCESS interface to determine which LIBNAME options are available in the Pass-Through Facility for your DBMS.

For a list of the LIBNAME options available in your SAS/ACCESS interface, see the documentation for your SAS/ACCESS interface.

Note: When a like-named option is specified in both the LIBNAME statement and after a data set name, SAS uses the value that is specified on the data set name. △

ACCESS= LIBNAME Option

Determines the access level with which a libref connection is opened

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

ACCESS=READONLY

Syntax Description

READONLY

specifies that tables and views can be read but not updated.

Details

Using this option prevents writing to the DBMS. If this option is omitted, tables and views can be read and updated if you have the necessary DBMS privileges.

AUTHID= LIBNAME Option

Lets you qualify your table names with an authorization ID, user ID, or group ID

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 z/OS

Alias: SCHEMA=

Default value: none

Syntax

AUTHID=*authorization-ID*

Syntax Description

authorization-ID

cannot exceed eight characters.

Details

When you specify the AUTHID= option, every table that is referenced by the libref is qualified as *authid.tablename* before any SQL code is passed to the DBMS. If you do not specify a value for AUTHID=, the table name is not qualified before it is passed to the DBMS. After the DBMS receives the table name, it automatically qualifies it with your user ID. You can override the LIBNAME AUTHID= option by using the AUTHID= data set option. This option is not validated until you access a table.

See Also

To apply this option to an individual data set, see the data set option “AUTHID= Data Set Option” on page 158.

AUTOCOMMIT= LIBNAME Option

Indicates whether updates are committed immediately after they are submitted

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Sybase

Default value: DBMS-specific, see details below

Syntax

AUTOCOMMIT= YES | NO

Syntax Description

YES

specifies that all updates, deletes, and inserts are committed (that is, saved to a table) immediately after they are submitted, and no rollback is possible.

NO

specifies that the SAS/ACCESS engine automatically performs the commit when it reaches the DBCOMMIT= value, or the default number of rows if DBCOMMIT is not set.

Details

If you are using the SAS/ACCESS LIBNAME statement, the default is NO if the data source provider supports transactions and the connection is used for updating data. For read-only connections and the Pass-Through Facility, the default is YES.

Informix and MySQL Details: The default is YES.

Netezza Details: The default is YES for PROC PRINT but NO for updates and for the main LIBNAME connection.

See Also

To apply this option to an individual data set, see the data set option “AUTOCOMMIT= Data Set Option” on page 158.

BL_KEEPIDENTITY= LIBNAME Option

Determines whether the identity column that is created during bulk loading is populated with values generated by Microsoft SQL Server or with values provided by the user

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: OLE DB

Default value: NO

Syntax

BL_KEEPIDENTITY= YES | NO

Syntax Description

YES

specifies that the user must provide values for the identity column.

NO

specifies that Microsoft SQL Server generates values for an identity column in the table.

Details

This option is valid only when you use the Microsoft SQL Server provider.

See Also

To apply this option to an individual data set, see the data set option “BL_KEEPIDENTITY= Data Set Option” on page 179.

BL_KEEPNULLS= LIBNAME Option

Indicates how NULL values in Microsoft SQL Server columns that accept NULL are handled during bulk loading

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: OLE DB

Default value: YES

Syntax

BL_KEEPNULLS= YES | NO

Syntax Description

YES

specifies that Microsoft SQL Server preserves NULL values inserted by the OLE DB interface.

NO

specifies that Microsoft SQL Server replaces NULL values that are inserted by the OLE DB interface with a default value (as specified in the DEFAULT constraint).

Details

This option only affects values in Microsoft SQL Server columns that accept NULL and have a DEFAULT constraint.

See Also

To apply this option to an individual data set, see the data set option “BL_KEEPNULLS= Data Set Option” on page 180.

BL_LOG= LIBNAME Option

Specifies the name of the error file to which all errors are written when BULKLOAD=YES

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Microsoft SQL Server, ODBC

Default value: none

Syntax

BL_LOG=*filename*

Details

This option is valid only for connections to Microsoft SQL Server. If BL_LOG= is not specified, errors are not recorded during bulk loading.

See Also

To apply this option to an individual data set, see the data set option “BL_LOG= Data Set Option” on page 182.

BL_OPTIONS= LIBNAME Option

Passes options to the DBMS bulk load facility, affecting how it loads and processes data

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: ODBC, OLE DB

Default value: not specified

Syntax

BL_OPTIONS='option <..., option>'

Details

You can use BL_OPTIONS= to pass options to the DBMS bulk load facility when it is invoked, thereby affecting how data is loaded and processed. You must separate multiple options with commas and enclose the entire string of options in quotation marks.

By default, no options are specified. This option takes the same values as the *-h* HINT option of the Microsoft BCP utility. Refer to the Microsoft SQL Server documentation for more information about bulk copy options.

This option is valid only when you are using the Microsoft SQL Server driver or the Microsoft SQL Server provider on Windows platforms.

ODBC Details: The supported hints are ORDER, ROWS_PER_BATCH, KILOBYTES_PER_BATCH, TABLOCK, and CHECK_CONSTRAINTS. If you specify UPDATE_LOCK_TYPE=TABLE, the TABLOCK hint is automatically added.

See Also

To apply this option to an individual data set, see the data set option “BL_OPTIONS= Data Set Option” on page 183.

BULKLOAD= LIBNAME Option

Determines whether SAS uses a DBMS facility to insert data into a DBMS table

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: ODBC, OLE DB, Teradata

Default value: NO

Syntax

BULKLOAD=YES | NO

Syntax Description

YES

calls a DBMS-specific bulk load facility in order to insert or append rows to a DBMS table.

NO

does not call the DBMS bulk load facility.

Details

See the SAS/ACCESS documentation for your DBMS for additional DBMS-specific details.

See Also

To apply this option to an individual data set, see the data set option “BULKLOAD= Data Set Option” on page 194.

CAST= LIBNAME Option

Specifies whether data conversions should be performed on the Teradata DBMS server or by SAS

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Teradata

Default value: none

Syntax

CAST=YES | NO

Syntax Description

YES

forces data conversions (casting) to be done on the Teradata DBMS server and overrides any data overhead percentage limit.

NO

forces data conversions to be done by SAS and overrides any data overhead percentage limit.

Details

Internally, SAS numbers and dates are floating point values. Teradata has varying formats for numbers, including integers, floating point values, and decimal values. Number conversion must occur when you are reading Teradata numbers that are not floating point (Teradata FLOAT). SAS/ACCESS can use the Teradata CAST= function to cause Teradata to perform numeric conversions. The parallelism of Teradata makes it well suited to perform this work. This is especially true when running SAS on z/OS (MVS) where CPU activity can be costly.

CAST= can cause more data to be transferred from Teradata to SAS, as a result of the option forcing the Teradata type into a larger SAS type. For example, the CAST= transfer of a Teradata BYTEINT to SAS floating point adds seven overhead bytes to each row transferred.

The following Teradata types are candidates for casting:

- INTEGER
- BYTEINT
- SMALLINT
- DECIMAL
- DATE.

SAS/ACCESS limits data expansion for CAST= to 20 percent in order to trade rapid data conversion by Teradata for extra data transmission. If casting does not exceed a 20 percent data increase, all candidate columns are cast. If the increase exceeds this limit, then SAS attempts to cast Teradata DECIMAL types only. If casting only DECIMAL types still exceeds the increase limit, data conversions are done by SAS.

You can alter the casting rules by using either the CAST= or “CAST_OVERHEAD_MAXPERCENT= LIBNAME Option” on page 86 option. With CAST_OVERHEAD_MAXPERCENT=, you can change the 20 percent overhead limit. With CAST=, you can override the percentage rules:

- CAST=YES forces Teradata to cast all candidate columns
- CAST=NO cancels all Teradata casting

CAST= only applies when you are reading Teradata tables into SAS. It does not apply when you are writing Teradata tables from SAS.

Also, CAST= only applies to SQL that SAS generates for you. If you supply your own SQL with the explicit SQL feature of PROC SQL, you must code your own casting clauses to force data conversions to occur in Teradata instead of SAS.

Examples

The following example demonstrates the use of the CAST= option in a LIBNAME statement to force casting for all tables referenced:

```
libname mydblib teradata user=testuser pw=testpass cast=yes;
proc print data=mydblib.emp;
where empno<1000;
run;

proc print data=mydblib.sal;
where salary>50000;
run;
```

The following example demonstrates the use of the CAST= option in a table reference in order to turn off casting for that table:

```
proc print data=mydblib.emp (cast=no);
where empno<1000;
run;
```

See Also

“CAST= Data Set Option” on page 195

CAST_OVERHEAD_MAXPERCENT= LIBNAME Option

Specifies the overhead limit for data conversions to be performed in Teradata instead of SAS

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Teradata

Default value: 20 percent

Syntax

CAST_OVERHEAD_MAXPERCENT=<n>

Syntax Description

<n>

Any positive numeric value. The engine default is 20.

Details

Teradata INTEGER, BYTEINT, SMALLINT, and DATE columns require conversion when read in to SAS. Conversions can be performed either by Teradata or by SAS. When performed in Teradata, using Teradata’s CAST operator, the row size transmitted to SAS can increase. CAST_OVERHEAD_MAXPERCENT= limits the allowable increase, also called conversion overhead.

Examples

The following example demonstrates the use of `CAST_OVERHEAD_MAXPERCENT=` to increase the allowable overhead to 40 percent:

```
proc print data=mydblib.emp (cast_overhead_maxpercent=40);
  where empno<1000;
run;
```

See Also

“`CAST= LIBNAME` Option” on page 84 for more information about conversions, conversion overhead, and casting.

CELLPROP= LIBNAME Option

Modifies the metadata and content of a result data set that is defined through an MDX command

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: OLE DB

Default value: VALUE

Syntax

`CELLPROP=VALUE | FORMATTED_VALUE`

Syntax Description

VALUE

specifies that the SAS/ACCESS engine tries to return actual data values. If all of the values in a column are numeric, then that column is defined as `NUMERIC`.

FORMATTED_VALUE

specifies that the SAS/ACCESS engine returns formatted data values. All of the columns are defined as `CHARACTER`.

Details

When an MDX command is issued, the resulting data set might have columns that contain one or more types of data values — the actual value of the cell or the formatted value of the cell.

For example, if you issue an MDX command and the resulting data set contains a column named `SALARY`, the column could contain data values of two types. It could contain numeric values, such as **50000**, or it could contain formatted values, such as **\$50,000**. Setting the `CELLPROP=` option determines how the values are defined and the value of the column.

It is possible for a column in a result set to contain both `NUMERIC` and `CHARACTER` data values. For example, a data set might return the data values of **50000**, **60000**, and **UNKNOWN**. SAS data sets cannot contain both types of data. In this

situation, even if you specify CELLPROP=VALUE, the SAS/ACCESS engine defines the column as CHARACTER and returns formatted values for that column.

For more information about MDX commands, see the SAS/ACCESS documentation for OLE DB.

COMMAND_TIMEOUT= LIBNAME Option

Specifies the number of seconds to wait before a data source command times out

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: OLE DB

Default value: 0

Syntax

COMMAND_TIMEOUT=*number-of-seconds*

Syntax Description

number-of-seconds

is an integer greater than or equal to 0.

Details

The default value is 0, which means there is no time-out.

See Also

To apply this option to an individual data set, see the data set option “COMMAND_TIMEOUT= Data Set Option” on page 198.

CONNECTION= LIBNAME Option

Specifies whether operations on a single libref share a connection to the DBMS, and whether operations on multiple librefs share a connection to the DBMS

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, MySQL, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

CONNECTION= SHAREDREAD | UNIQUE | SHARED | GLOBALREAD | GLOBAL

Syntax Description

Not all values are valid for all SAS/ACCESS interfaces. See details below.

SHAREDREAD

specifies that all READ operations that access DBMS tables *in a single libref* share a single connection. A separate connection is established for every table that is opened for update or output operations.

Where available, this is usually the default value because it offers the best performance and it guarantees data integrity.

UNIQUE

specifies that a separate connection is established every time a DBMS table is accessed by your SAS application.

Use UNIQUE if you want each use of a table to have its own connection.

SHARED (valid for DB2 UNIX/PC, DB2 z/OS, HP Neoview, Microsoft SQL Server, Netezza, ODBC, Oracle)

specifies that *all* operations that access DBMS tables *in a single libref* share a single connection.

Use this option with caution. When a single SHARED connection is used for multiple table opens, a commit or rollback performed on one table being updated also applies to all other tables opened for update. Even if a table is just opened for READ, its READ cursor might get resynchronized as a result of this commit or rollback. If the cursor is resynchronized, there is no guarantee that the new solution table will match the original solution table that was being read.

Use SHARED to eliminate the deadlock that can occur when you create and load a DBMS table from an existing table that resides in the same database or tablespace. This only happens in certain output processing situations and is the only recommended use for CONNECTION=SHARED.

Note: The CONNECTION= option only influences the connections used for opening tables with a libref. Setting CONNECTION=SHARED has no influence on utility connections or explicit pass-through connections. △

GLOBALREAD

specifies that all READ operations that access DBMS tables *in multiple librefs* share a single connection if the following is true:

- the participating librefs are created by LIBNAME statements that specify identical values for the CONNECTION=, CONNECTION_GROUP=, DBCONINIT=, DBCONTERM=, DBLIBINIT=, and DBLIBTERM= options
- the participating librefs are created by LIBNAME statements that specify identical values for any DBMS connection options.

A separate connection is established for each table that is opened for update or output operations.

GLOBAL (valid for DB2 UNIX/PC, DB2 z/OS, HP Neoview, Microsoft SQL Server, Netezza, ODBC, Oracle)

specifies that *all* operations that access DBMS tables *in multiple librefs* share a single connection if the following is true:

- all of the participating librefs are created by LIBNAME statements that specify identical values for the CONNECTION=, CONNECTION_GROUP=, DBCONINIT=, DBCONTERM=, DBLIBINIT=, and DBLIBTERM= options
- all of the participating librefs are created by LIBNAME statements that specify identical values for any DBMS connection options.

One connection is shared for all tables that are referenced by any of the librefs for which CONNECTION=GLOBAL is specified.

Use this option with caution. When a GLOBAL connection is used for multiple table opens, a commit/rollback performed on one table being updated also applies to all other tables opened for update. Even if a table is just opened for READ, its READ cursor might get resynchronized as a result of this commit/rollback. If the cursor is resynchronized, there is no guarantee that the new solution table will match the original solution table that was being read.

Details

For most SAS/ACCESS interfaces, there must be a connection, also known as an *attach*, to the DBMS server before any data can be accessed. Typically, each DBMS connection has one transaction, or work unit, that is active in the connection. This transaction is affected by any SQL commits or rollbacks that the engine performs within the connection while executing the SAS application.

You can use the CONNECTION= option to control the number of connections, and therefore the transactions, that your SAS/ACCESS interface executes and supports for each LIBNAME statement.

GLOBALREAD is the default value for CONNECTION= when you specify CONNECTION_GROUP=.

This option is supported by the SAS/ACCESS interfaces that support single connections or multiple, simultaneous connections to the DBMS.

Microsoft SQL Server and ODBC Details: If the data source supports only one active open cursor per connection, the default value is CONNECTION=UNIQUE; otherwise, the default value is CONNECTION=SHAREDREAD.

MySQL & Netezza Details: The default value is UNIQUE.

Teradata Details: for channel-attached systems (MVS), the default is SHAREDREAD; for network attached systems (UNIX and PC platforms), the default is UNIQUE

Examples

In the following SHAREDREAD example, MYDBLIB makes the first connection to the DBMS. This connection is used to print the data from MYDBLIB.TAB. MYDBLIB2 makes the second connection to the DBMS. A third connection is used to update MYDBLIB.TAB. The third connection is closed at the end of the PROC SQL UPDATE statement. The first and second connections are closed with the CLEAR option.

```
libname mydblib oracle user=testuser /* connection 1 */
      pw=testpass path='myorapath'
      connection=sharedread;

libname mydblib2 oracle user=testuser /* connection 2 */
      pw=testpass path='myorapath'
      connection=sharedread;

proc print data=mydblib.tab ...
proc sql;
      update mydblib.tab ... /* connection 3 */

libname mydblib clear;
libname mydblib2 clear;
```

In the following GLOBALREAD example, the two librefs, MYDBLIB and MYDBLIB2, share the same connection for read access because CONNECTION=GLOBALREAD and the connection options are identical. The first

connection is used to print the data from MYDBLIB.TAB while a second connection is made for updating MYDBLIB.TAB. The second connection is closed at the end of the step. Note that the first connection is closed with the final LIBNAME statement.

```
libname mydblib oracle user=testuser /* connection 1 */
      pw=testpass path='myorapath'
      connection=globalread;

libname mydblib2 oracle user=testuser
      pw=testpass path='myorapath'
      connection=globalread;

proc print data=mydblib.tab ...
proc sql;                                /* connection 2 */
      update mydblib.tab ...

libname mydblib clear;                   /* does not close connection 1 */
libname mydblib2 clear;                  /* closes connection 1 */
```

In the following UNIQUE example, the libref, MYDBLIB, does not establish a connection. A connection is established in order to print the data from MYDBLIB.TAB. That connection is closed at the end of the print procedure. Another connection is established to updated MYDBLIB.TAB. That connection is closed at the end of the PROC SQL. The CLEAR option in the LIBNAME statement at the end of this example does not close any connections.

```
libname mydblib oracle user=testuser
      pw=testpass path='myorapath'
      connection=unique;

proc print data=mydblib.tab ...
proc sql;
      update mydblib.tab ...

libname mydblib clear;
```

The following example is a GLOBAL example for DB2 under z/OS. Both PROC DATASETS invocations appropriately report “no members in directory” because SESSION.B, as a temporary table, has no entry in the system catalog SYSIBM.SYSTABLES. However, the DATA_NULL_ step and SELECT * from PROC SQL step both return the expected rows. For DB2 under z/OS, when SCHEMA=SESSION the database first looks for a temporary table before attempting to access any physical schema named SESSION.

```
libname x db2 connection=global schema=SESSION;
proc datasets lib=x;
quit;

/*
 * DBMS-specific code to create a temporary table impervious
 * to commits, and then populate the table directly in the
 * DBMS from another table.
 */
proc sql;
connect to db2(connection=global schema=SESSION);
execute ( DECLARE GLOBAL TEMPORARY TABLE SESSION.B LIKE SASDXS.A
```

```

        ON COMMIT PRESERVE ROWS
        ) by db2;
execute ( insert into SESSION.B select * from SASDXS.A
        ) by db2;
quit;

/*
 * Get at the temp table through the global libref.
 */
data _null_;
set x.b;
put _all_;
run;

/*
 * Get at the temp table through the global connection.
 */
proc sql;
connect to db2 (connection=global schema=SESSION);
select * from connection to db2
( select * from SESSION.B );
quit;

proc datasets lib=x;
quit;

```

In the following SHARED example, DB2DATA.NEW is created in the database TEST. Because the table DB2DATA.OLD exists in the same database. The option CONNECTION=SHARED enables the DB2 engine to share the connection both for reading the old table and for creating and loading the new table.

```

libname db2data db2 connection=shared;
data db2data.new (in = 'database test');
    set db2data.old;
run;

```

In the following GLOBAL example, two different librefs share one connection.

```

libname db2lib db2 connection=global;
libname db2data db2 connection=global;
data db2lib.new(in='database test');
    set db2data.old;
run;

```

If you did not use the CONNECTION= option in the above two examples, you would deadlock in DB2 and get the following error:

```

ERROR: Error attempting to CREATE a DBMS table.
ERROR: DB2 execute error DSNT408I SQLCODE = --911,
ERROR: THE CURRENT UNIT OF WORK HAS BEEN ROLLED
BACK DUE TO DEADLOCK.

```


See Also

“DEFER= LIBNAME Option” on page 111

“ACCESS= LIBNAME Option” on page 79

“CONNECTION_GROUP= LIBNAME Option” on page 93

CONNECTION_GROUP= LIBNAME Option

Causes operations on multiple librefs to share a connection to the DBMS. Also causes operations on multiple Pass-Through Facility CONNECT statements to share a connection to the DBMS

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

CONNECTION_GROUP= *connection-group-name*

Syntax Description

connection-group-name

is the name of a connection group.

Details

This option causes a DBMS connection to be shared by all READ operations on multiple librefs if the following is true:

- all of the participating librefs are created by LIBNAME statements that specify the same value for CONNECTION_GROUP=
- all of the participating librefs are created by LIBNAME statements that specify identical DBMS connection options.

To share a connection for *all* operations against multiple librefs, specify CONNECTION= GLOBAL on all participating LIBNAME statements. Not all SAS/ACCESS interfaces support CONNECTION=GLOBAL.

Note: If you specify CONNECTION=GLOBAL or CONNECTION=GLOBALREAD, operations on multiple librefs can share a connection even if you omit CONNECTION_GROUP=. △

Informix Details: This option lets multiple librefs or multiple Pass-Through Facility CONNECT statements share a connection to the DBMS. It overcomes the Release 8.2 limitation that prevents users from accessing scratch tables across step boundaries because new connections had to be established with every procedure.

Example

In the following example, the MYDBLIB libref shares a connection with MYDBLIB2 by specifying CONNECTION_GROUP=MYGROUP and by specifying identical connection options. The libref MYDBLIB3 makes a second connection to another connection group called ABC. The first connection is used to print the data from MYDBLIB.TAB, and is also used for updating MYDBLIB.TAB. The third connection is closed at the end of the step. Note that the first connection is closed by the final LIBNAME statement for that connection. Similarly, the second connection is closed by the final LIBNAME statement for that connection.

```
libname mydblib oracle user=testuser      /* connection 1 */
      pw=testpass
      connection_group=mygroup;

libname mydblib2 oracle user=testuser
      pw=testpass
      connection_group=mygroup;

libname mydblib3 oracle user=testuser     /* connection 2 */
      pw=testpass
      connection_group=abc;

proc print data=mydblib.tab ...
proc sql;                                /* connection 1 */
      update mydblib.tab ...

libname mydblib clear;                   /* does not close connection 1*/
libname mydblib2 clear;                  /* closes connection 1 */
libname mydblib3 clear;                  /* closes connection 2 */
```

CONNECTION_TIMEOUT= LIBNAME Option

Specifies the number of seconds to wait before a connection times out

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: HP Neoview

Alias: CON_TIMEOUT=

Default value: 0

Syntax

CONNECTION_TIMEOUT=*number-of-seconds*

Syntax Description

number-of-seconds

A number greater than or equal to 0. It represents the number of seconds that SAS/ACCESS Interface to HP Neoview waits for any operation on the connection to complete before returning to SAS. If the value is 0, which is the default, no timeout occurs.

CURSOR_TYPE= LIBNAME Option

Specifies the cursor type for read-only and updatable cursors

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: DB2 UNIX/PC, Microsoft SQL Server, ODBC, OLE DB

Default value: DBMS-specific

Syntax

CURSOR_TYPE= DYNAMIC | FORWARD_ONLY | KEYSET_DRIVEN | STATIC

Syntax Description

DYNAMIC

specifies that the cursor reflects all of the changes that are made to the rows in a result set as you move the cursor. The data values and the membership of rows in the cursor can change dynamically on each fetch. This is the default for the DB2 UNIX/PC, ODBC, and Microsoft SQL Server interfaces.

FORWARD_ONLY

specifies that the cursor functions like a DYNAMIC cursor except that it only supports fetching the rows sequentially. (This value is not valid in OLE DB.)

KEYSET_DRIVEN

specifies that the cursor determines which rows belong to the result set when the cursor is opened. However, changes that are made to these rows are reflected as you scroll around the cursor.

STATIC

specifies that the cursor builds the complete result set when the cursor is opened. No changes that are made to the rows in the result set after the cursor is opened are reflected in the cursor. Static cursors are read-only.

Details

Not all drivers support all cursor types. An error is returned if the specified cursor type is not supported.

The driver is allowed to modify the default without an error.

OLE DB Details: By default, this option is not set and the provider uses a default. See your provider documentation for more information. See OLE DB programmer reference documentation for details about these properties. The OLE DB properties applied to an open row set are as follows:

CURSOR_TYPE=	OLE DB Properties Applied
DYNAMIC	DBPROP_OTHERINSERT=TRUE, DBPROP_OTHERUPDATEDELETE=TRUE
KEYSET_DRIVEN	DBPROP_OTHERINSERT=FALSE, DBPROP_OTHERUPDATEDELETE=TRUE
STATIC	DBPROP_OTHERINSERT=FALSE, DBPROP_OTHERUPDATEDELETE=FALSE

See Also

To apply this option to an individual data set, see the data set option “CURSOR_TYPE= Data Set Option” on page 198.

DBCOMMIT= LIBNAME Option

Causes an automatic COMMIT (a permanent writing of data to the DBMS) after a specified number of rows have been processed

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: 1000 when inserting rows into a DBMS table; 0 when updating a DBMS table

Syntax

DBCOMMIT=*n*

Syntax Description

n
is an integer greater than or equal to 0.

Details

DBCOMMIT= affects update, delete, and insert processing. The number of rows that are processed includes rows that are not processed successfully. If you set DBCOMMIT=0, a commit is issued only once (after the procedure or DATA step completes).

If the DBCOMMIT= option is explicitly set, SAS/ACCESS fails any update that has a WHERE clause.

Note: If you specify both DBCOMMIT= and ERRLIMIT=, and these options collide during processing, then the COMMIT is issued first and the ROLLBACK is issued second. Because the COMMIT (caused by the DBCOMMIT= option) is issued prior to the ROLLBACK (caused by the ERRLIMIT= option), the DBCOMMIT= option is said to override the ERRLIMIT= option in this situation. △

DB2 under UNIX and PC Hosts Details: When BULKLOAD=YES, the default is 10000.

See Also

To apply this option to an individual data set, see the data set option “DBCOMMIT= Data Set Option” on page 199.

DBCONINIT= LIBNAME Option

Specifies a user-defined initialization command to be executed immediately after every connection to the DBMS that is within the scope of the LIBNAME statement or libref

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBCONINIT=<'>DBMS-user-command<'>

Syntax Description

DBMS-user-command

is any valid command that can be executed by the SAS/ACCESS engine and that does not return a result set or output parameters.

Details

The initialization command that you select can be a stored procedure or any DBMS SQL statement that might provide additional control over the interaction between your SAS/ACCESS interface and the DBMS.

The command executes immediately after each DBMS connection is successfully established. If the command fails, then a disconnect occurs and the libref is not assigned. You must specify the command as a single, quoted string.

Note: The initialization command might execute more than once, because one LIBNAME statement might have multiple connections; for example, one for reading and one for updating. △

Examples

In the following example, the DBCONINIT= option causes the DBMS to apply the SET statement to every connection that uses the MYDBLIB libref.

```
libname mydblib db2
      dbconinit="SET CURRENT SQLID='myauthid'";

proc sql;
  select * from mydblib.customers;

  insert into mydblib.customers
    values('33129804', 'VA', '22809', 'USA',
          '540/545-1400', 'BENNETT SUPPLIES', 'M. JONES',
          '2199 LAUREL ST', 'ELKTON', '22APR97'd);

  update mydblib.invoices
    set amtbill = amtbill*1.10
    where country = 'USA';

quit;
```

In the following example, a stored procedure is passed to DBCONINIT=.

```
libname mydblib oracle user=testuser pass=testpass
      dbconinit="begin dept_test(1001,25)";
end;
```

The SAS/ACCESS engine retrieves the stored procedure and executes it.

See Also

“DBCONTERM= LIBNAME Option” on page 98

DBCONTERM= LIBNAME Option

Specifies a user-defined termination command to be executed before every disconnect from the DBMS that is within the scope of the LIBNAME statement or libref

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBCONTERM=<'>*DBMS-user-command*<'>

Syntax Description

DBMS-user-command

is any valid command that can be executed by the SAS/ACCESS engine and that does not return a result set or output parameters.

Details

The termination command that you select can be a stored procedure or any DBMS SQL statement that might provide additional control over the interaction between the SAS/ACCESS engine and the DBMS. The command executes immediately before SAS terminates each connection to the DBMS. If the command fails, then SAS provides a warning message but the library deassignment and disconnect still occur. You must specify the command as a single, quoted string.

Note: The termination command might execute more than once, because one LIBNAME statement might have multiple connections; for example, one for reading and one for updating. △

Examples

In the following example, the DBMS drops the Q1_SALES table before SAS disconnects from the DBMS.

```
libname mydblib db2 user=testuser using=testpass
      db=invoice bconterm='drop table q1_sales';
```

In the following example, the stored procedure, SALESTAB_STORED_PROC, is executed each time SAS connects to the DBMS, and the BONUSSES table is dropped when SAS terminates each connection.

```
libname mydblib db2 user=testuser
      using=testpass db=sales
      dbconinit='exec salestab_stored_proc'
      dbconterm='drop table bonuses';
```

See Also

“DBCONINIT= LIBNAME Option” on page 97

DBCREATE_TABLE_OPTS= LIBNAME Option

Specifies DBMS-specific syntax to be added to the CREATE TABLE statement

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBCREATE_TABLE_OPTS='DBMS-SQL-clauses'

DBMS-SQL-clauses

are one or more DBMS-specific clauses that can be appended to the end of an SQL CREATE TABLE statement.

Details

You can use DBCREATE_TABLE_OPTS= to add DBMS-specific clauses to the end of the SQL CREATE TABLE statement. The SAS/ACCESS engine passes the SQL CREATE TABLE statement and its clauses to the DBMS, which executes the statement and creates the DBMS table. DBCREATE_TABLE_OPTS= applies only when you are creating a DBMS table by specifying a libref associated with DBMS data.

See Also

To apply this option to an individual data set, see the data set option “DBCREATE_TABLE_OPTS= Data Set Option” on page 201.

DBGEN_NAME= LIBNAME Option

Specifies how SAS automatically renames DBMS columns that contain characters that SAS does not allow, such as \$, to valid SAS variable names

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS

Syntax

DBGEN_NAME= DBMS | SAS

Syntax Description**DBMS**

specifies that the DBMS columns are renamed to valid SAS variable names. Disallowed characters are converted to underscores. If a column is converted to a name that already exists, then a sequence number is appended to the end of the new name.

SAS

specifies that DBMS columns are renamed to the format `_COL n` , where n is the column number (starting with zero).

Details

SAS retains column names when it reads data from tables, unless a column name contains characters that SAS does not allow, such as \$. SAS allows alphanumeric characters and the underscore (_).

For example, if you specify DBGEN_NAME=SAS, a DBMS column named **Dept\$Amt** is renamed to **_COLn**. If you specify DBGEN_NAME=DBMS, the **Dept\$Amt** column is renamed to **Dept_Amt**.

This option is intended primarily for National Language Support, notably for the conversion of kanji to English characters. English characters that are converted from kanji are often those that are not allowed in SAS.

Note: The various SAS/ACCESS interfaces each handled name collisions differently in SAS Version 6. Some interfaces appended to the end of the name; other interfaces replaced the last character(s) in the name. Some interfaces used a single sequence number, other interfaces used unique counters. If you specify VALIDVARNAME=V6 ; then name collisions must be handled the same as they were in SAS Version 6. △

See Also

To apply this option to an individual data set, see the data set option “DBGEN_NAME= Data Set Option” on page 204

DBINDEX= LIBNAME Option

Improves performance when processing a join that involves a large DBMS table and a small SAS data set

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

DBINDEX= YES | NO

Syntax Description

YES

specifies that SAS uses columns in the WHERE clause that have defined DBMS indexes.

NO

specifies that SAS does not use indexes that are defined on DBMS columns.

Details

When you are processing a join that involves a large DBMS table and a relatively small SAS data set, you might be able to use DBINDEX= to improve performance.

CAUTION:

Improper use of this option can degrade performance. See “Using the DBINDEX=, DBKEY=, and MULTI_DATASRC_OPT= Options” on page 43 for detailed information about using this option. △

See Also

To apply this option to an individual data set, see the data set option “DBINDEX= Data Set Option” on page 205.

DBLIBINIT= LIBNAME Option

Specifies a user-defined initialization command to be executed once within the scope of the LIBNAME statement or libref that established the first connection to the DBMS

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBLIBINIT= <'>*DBMS-user-command*<'>

Syntax Description

DBMS-user-command

is any DBMS command that can be executed by the SAS/ACCESS engine and that does not return a result set or output parameters.

Details

The initialization command that you select can be a script, stored procedure, or any DBMS SQL statement that might provide additional control over the interaction between your SAS/ACCESS interface and the DBMS.

The command executes immediately after the first DBMS connection is successfully established. If the command fails, then a disconnect occurs and the libref is not assigned. You must specify the command as a single, quoted string, unless it is an environment variable.

DBLIBINIT= fails if either CONNECTION=UNIQUE or DEFER=YES, or if both of these LIBNAME options are specified.

When multiple LIBNAME statements share a connection, the initialization command executes only for the first LIBNAME statement, immediately after the DBMS connection is established. (Multiple LIBNAME statements that use CONNECTION=GLOBALREAD and identical values for CONNECTION_GROUP=, DBCONINIT=, DBCONTERM=, DBLIBINIT=, and DBLIBTERM= options and any DBMS connection options can share the same connection to the DBMS.)

Example

In the following example, CONNECTION=GLOBALREAD is specified in both LIBNAME statements, but the DBLIBINIT commands are different. Therefore, the second LIBNAME statement fails to share the same physical connection.

```
libname mydblib oracle user=testuser pass=testpass
        connection=globalread dblinkinit='Test';

libname mydblib2 oracle user=testuser pass=testpass
        connection=globalread dblinkinit='NoTest';
```

See Also

“DBLIBTERM= LIBNAME Option” on page 103

DBLIBTERM= LIBNAME Option

Specifies a user-defined termination command to be executed once, before the DBMS that is associated with the first connection made by the LIBNAME statement or libref disconnects

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBLIBTERM= <'>*DBMS-user-command*<'>

Syntax Description

DBMS-user-command

is any DBMS command that can be executed by the SAS/ACCESS engine and that does not return a result set or output parameters.

Details

The termination command that you select can be a script, stored procedure, or any DBMS SQL statement that might provide additional control over the interaction between the SAS/ACCESS engine and the DBMS. The command executes immediately before SAS terminates the last connection to the DBMS. If the command fails, then SAS provides a warning message but the library deassignment and disconnect still occurs. You must specify the command as a single, quoted string.

DBLIBTERM= fails if either CONNECTION=UNIQUE or DEFER=YES or both of these LIBNAME options are specified.

When two LIBNAME statements share the same physical connection, the termination command is executed only once. (Multiple LIBNAME statements that use CONNECTION=GLOBALREAD and identical values for CONNECTION_GROUP=, DBCONINIT=, DBCONTERM=, DBLIBINIT=, and DBLIBTERM= options and any DBMS connection options can share the same connection to the DBMS.)

Example

In the following example, CONNECTION=GLOBALREAD is specified on both LIBNAME statements, but the DBLIBTERM commands are different. Therefore, the second LIBNAME statement fails to share the same physical connection.

```
libname mydblib oracle user=testuser pass=testpass
      connection=globalread dblibterm='Test';
```

```
libname mydblib2 oracle user=testuser pass=testpass
      connection=globalread dblibterm='NoTest';
```

See Also

“DBLIBINIT= LIBNAME Option” on page 102

DBLINK= LIBNAME Option

Specifies a link from your local database to database objects on another server (in the Oracle interface); specifies a link from your default database to another database on the server to which you are connected (in the Sybase interface)

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Oracle, Sybase

Default value: none

Syntax

DBLINK=*database-link*

Details

Oracle Details: A link is a database object that is used to identify an object stored in a remote database. A link contains stored path information and may also contain user name and password information for connecting to the remote database. If you specify a link, SAS uses the link to access remote objects. If you omit DBLINK=, SAS accesses objects in the local database.

Sybase Details: You can use this option to link to another database within the same server to which you are connected. If you omit DBLINK=, SAS can access only objects that are in your default database.

See Also

To apply this option to an individual data set, see the data set option “DBMASTER= Data Set Option” on page 210.

DBMAX_TEXT= LIBNAME Option

Determines the length of any very long DBMS character data type that is read into SAS or written from SAS when using a SAS/ACCESS engine

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase

Default value: 1024

Syntax

DBMAX_TEXT=<integer>

Syntax Description

integer

is an integer between 1 and 32,767.

Details

This option applies to reading, appending, and updating rows in an existing table. It does not apply when you are creating a table.

Examples of a DBMS data type are the Sybase TEXT data type or the Oracle LONG RAW data type.

See Also

To apply this option to an individual data set, see the data set option “DBMAX_TEXT= Data Set Option” on page 211.

DBNULLKEYS= LIBNAME Option

Controls the format of the WHERE clause when you use the DBKEY= data set option

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle

Default value: DBMS-specific

Syntax

DBNULLKEYS= YES | NO

Details

If there might be NULL values in the transaction table or the master table for the columns that you specify in the DBKEY= option, use DBNULLKEYS=YES. This is the default for most interfaces. When you specify DBNULLKEYS=YES and specify a column that is not defined as NOT NULL in the DBKEY= data set option, SAS generates a WHERE clause that can find NULL values. For example, if you specify DBKEY=COLUMN and COLUMN is not defined as NOT NULL, SAS generates a WHERE clause with the following syntax:

```
WHERE ((COLUMN = ?) OR ((COLUMN IS NULL) AND (? IS NULL)))
```

This syntax enables SAS to prepare the statement once and use it for any value (NULL or NOT NULL) in the column. Note that this syntax has the potential to be much less efficient than the shorter form of the WHERE clause (presented below). When you specify DBNULLKEYS=NO or specify a column that is defined as NOT NULL in the DBKEY= option, SAS generates a simple WHERE clause.

If you know that there are no NULL values in the transaction table or the master table for the columns that you specify in the DBKEY= option, then you can use DBNULLKEYS=NO. This is the default for the interface to Informix. If you specify DBNULLKEYS=NO and specify DBKEY=COLUMN, SAS generates a shorter form of the WHERE clause (regardless of whether or not the column specified in DBKEY= is defined as NOT NULL):

```
WHERE (COLUMN = ?)
```

See Also

To apply this option to an individual data set, see the data set option “DBNULLKEYS= Data Set Option” on page 213.

DBPROMPT= LIBNAME Option

Specifies whether SAS displays a window that prompts the user to enter DBMS connection information prior to connecting to the DBMS in interactive mode

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, Oracle, Sybase, Teradata

Default value: NO

Syntax

DBPROMPT=YES | NO

Syntax Description

YES

specifies that SAS displays a window that interactively prompts you for the DBMS connection options the first time the libref is used.

NO

specifies that SAS does not display the prompting window.

Details

If you specify DBPROMPT=YES, it is not necessary to provide connection options with the LIBNAME statement. If you specify connection options with the LIBNAME statement and you specify DBPROMPT=YES, then the connection option values are displayed in the window (except for the password value. This value appears as a series of asterisks). All of these values can be overridden interactively.

The DBPROMPT= option interacts with the DEFER= option to determine when the prompt window appears. If DEFER=NO, the DBPROMPT window opens when the LIBNAME statement is executed. If DEFER=YES, the DBPROMPT window opens the first time a table or view is opened. The DEFER= option normally defaults to NO but defaults to YES if DBPROMPT=YES. You can override this default by explicitly setting DEFER=NO.

The DBPROMPT window usually opens only once for each time that the LIBNAME statement is specified. It might open multiple times if DEFER=YES and the connection fails when SAS tries to open a table. In these cases, the DBPROMPT window opens until a successful connection occurs or until you click **Cancel**.

The maximum password length for most of the SAS/ACCESS LIBNAME interfaces is 32 characters.

Oracle Details: You can enter 30 characters for the USERNAME and PASSWORD and up to 70 characters for the PATH, depending on your platform.

Teradata Details: You can enter up to 30 characters for the USERNAME and PASSWORD.

Examples

In the following example, the DBPROMPT window does not open when the LIBNAME statement is submitted because DEFER=YES. The DBPROMPT window opens when the PRINT procedure is processed, a connection is made, and the table is opened.

```
libname mydblib oracle dbprompt=yes
      defer=yes;

proc print data=mydblib.staff;
run;
```

In the following example, the DBPROMPT window opens while the LIBNAME statement is processing. The DBPROMPT window does not open in subsequent statements because the DBPROMPT window opens only once per LIBNAME statement.

```
libname mydblib oracle dbprompt=yes
      defer=no;
```

In the following example, values provided in the LIBNAME statement are pulled into the DBPROMPT window. The values `testuser` and `ABC_server` appear in the DBPROMPT window and can be edited and confirmed by the user. The password value appears in the DBPROMPT window as a series of asterisks; it can also be edited by the user.

```
libname mydblib oracle
      user=testuser pw=testpass
      path='ABC_server' dbprompt=yes defer=no;
```

See Also

To apply this option to a view descriptor, see the data set option “DBPROMPT= Data Set Option” on page 214.

DBSASLABEL= LIBNAME Option

Specifies the column labels an engine uses

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, Informix, Microsoft SQL Server, MySQL, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: COMPAT

Syntax

DBSASLABEL=COMPAT | NONE

Syntax Description

COMPAT

specifies that the labels returned should be compatible with what the application normally receives—meaning that engines exhibit their normal behavior.

NONE

specifies that the engine does not return a column label. The engine returns blanks for the column labels.

Details

By default, the SAS/ACCESS interface for your DBMS generates column labels from the column names, rather than from the real column labels.

This option enables the user to override this default behavior. It is useful in the PROC SQL context where column labels instead of column aliases are used as headers.

Examples

The following example demonstrates how DBSASLABEL= is used as a LIBNAME option to return blank column labels so that PROC SQL can use the column aliases as the column headers.


```
libname x oracle user=scott pw=tiger;
proc sql;
    select deptno as Department ID, loc as Location from mylib.dept(dbsaslabel=none);
```

Without the DBSASLABEL= option set to NONE, the aliases would be ignored, and DEPTNO and LOC would be used as column headers in the result set.

DBSLICEPARM= LIBNAME Option

Controls the scope of DBMS threaded reads and the number of threads

Valid in: the SAS/ACCESS LIBNAME statement (Also available as a SAS configuration option, SAS invocation option, global SAS option, or data set option)

DBMS support: DB2 UNIX/PC, DB2 z/OS, Informix, Microsoft SQL Server, ODBC, Oracle, Sybase, Teradata

Default value: THREADED_APPS,2 (DB2 z/OS, Oracle, Teradata) THREADED_APPS,2 or 3 (DB2 UNIX/PC, Informix, Microsoft SQL Server, ODBC, Sybase)

Syntax

DBSLICEPARM=NONE | THREADED_APPS | ALL

DBSLICEPARM=(NONE | THREADED_APPS | ALL <, *max-threads*>)

Syntax Description

Two syntax diagrams are shown here in order to highlight the simpler version. In most cases, the simpler version suffices.

NONE

disables DBMS threaded read. SAS reads tables on a single DBMS connection, as it did with SAS Version 8 and earlier.

THREADED_APPS

Makes fully threaded SAS procedures (threaded applications) eligible for threaded reads.

ALL

Makes all read-only librefs eligible for threaded reads. This includes SAS threaded applications, as well as the SAS DATA step and numerous SAS procedures.

max-threads

positive integer value that specifies the maximum number of connections per table read. A partition or portion of the data is read on each connection. The combined rows across all partitions are the same irrespective of the number of connections. Changes to the number of connections do not change the result set. Increasing the number of connections instead redistributes the same result set across more connections.

There are diminishing returns when increasing the number of connections. With each additional connection, more burden is placed on the DBMS, and a smaller percentage of time saved on the SAS step. Therefore, you should consult your DBMS-specific documentation concerning partitioned reads before using this parameter.

Details

DBSLICEPARAM= can be used in numerous locations, and the usual rules of option precedence apply. Table option has the highest precedence, then LIBNAME option, and so on. SAS configuration file option has the lowest precedence as DBSLICEPARAM= in any of the other locations overrides that configuration setting.

DBSLICEPARAM=ALL and DBSLICEPARAM=THREADED_APPS make SAS programs eligible for threaded reads. To determine whether threaded reads are actually generated, turn on SAS tracing and run a program, as shown in the following example:

```
options sastrace='',,t' sastraceloc=saslog nostsuffix;
proc print data=lib.dbtable(dbsliceparam=(ALL));
  where dbcol>1000;
run;
```

If you want to directly control the threading behavior, use the DBSLICE= data set option.

For DB2 UNIX/PC, Informix, Microsoft SQL Server, ODBC, and Sybase, the default thread number is dependent on whether an application passes in the number of threads (CPUCOUNT=) and whether the data type of the column selected for the data partitioning purpose is binary.

Examples

The following code demonstrates how to use DBSLICEPARAM= in a PC SAS configuration file entry to turn off threaded reads for all SAS users:

```
--dbsliceparam NONE
```

The following code demonstrates how to use DBSLICEPARAM= as an z/OS invocation option to turn on threaded reads for read-only references to DBMS tables throughout a SAS job:

```
sas o(dbsliceparam=ALL)
```

The following code demonstrates how to use DBSLICEPARAM= as a SAS global option, most likely as one of the first statements in your SAS code, to increase maximum threads to three for SAS threaded apps:

```
option dbsliceparam=(threaded_apps,3);
```

The following code demonstrates how to use DBSLICEPARAM= as a LIBNAME option to turn on threaded reads for read-only table references that use this particular libref:

```
libname dblib oracle user=scott password=tiger dbsliceparam=ALL;
```

The following code demonstrates how to use DBSLICEPARAM= as a table level option to turn on threaded reads for this particular table, requesting up to four connections:

```
proc reg SIMPLE;
  data=dblib.customers (dbsliceparam=(all,4));
  var age weight;
  where years_active>1;
run;
```

See Also

“DBSLICEPARAM= Data Set Option” on page 219

DEFER= LIBNAME Option

Specifies when the connection to the DBMS occurs

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: NO

Syntax

DEFER= NO | YES

Syntax Description

NO

specifies that the connection to the DBMS occurs when the libref is assigned by a LIBNAME statement.

YES

specifies that the connection to the DBMS occurs when a table in the DBMS is opened.

Details

The default value of NO is overridden if DBPROMPT=YES.

The DEFER= option is ignored when CONNECTION=UNIQUE, because a connection is performed every time a table is opened.

Microsoft SQL Server, Netezza, and ODBC Details: When you set DEFER=YES, you must also set the PRESERVE_TAB_NAMES= and PRESERVE_COL_NAMES= options to their desired values. Normally, SAS queries the data source to determine the correct defaults for these options during LIBNAME assignment, but setting DEFER=YES postpones the connection. Because these values must be set at the time of LIBNAME assignment, you must assign them explicitly when you set DEFER=YES.

See Also

“CONNECTION= LIBNAME Option” on page 88

DEGREE= LIBNAME Option

Determines whether DB2 uses parallelism

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 z/OS

Default value: ANY

Syntax

DEGREE= ANY | 1

Syntax Description

ANY

enables DB2 to use parallelism, and issues the SET CURRENT DEGREE ='xxx' for all DB2 threads that use that libref.

1

explicitly disables the use of parallelism.

Details

When DEGREE=ANY, DB2 has the option of using parallelism, when it is appropriate.

Setting DEGREE=1 prevents DB2 from performing parallel operations. Instead, DB2 is restricted to performing one task that, while perhaps slower, uses less system resources.

DELETE_MULT_ROWS= LIBNAME Option

Indicates whether to allow SAS to delete multiple rows from a data source, such as a DBMS table

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB

Default value: NO

Syntax

DELETE_MULT_ROWS=YES | NO

Syntax Description

YES

specifies that SAS/ACCESS processing continues if multiple rows are deleted. This might produce unexpected results.

NO

specifies that SAS/ACCESS processing does not continue if multiple rows are deleted.

Details

Some providers do not handle the following DBMS SQL statement well and, therefore, delete more than the current row:

```
DELETE ... WHERE CURRENT OF CURSOR
```

DIRECT_EXE= LIBNAME Option

Lets you pass an SQL delete statement directly to a DBMS via implicit pass-through.

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DIRECT_EXE=DELETE

Syntax Description

DELETE

specifies that an SQL delete statement is passed directly to the DBMS for processing.

Details

Performance is greatly increased by using DIRECT_EXE=. The SQL delete statement is passed directly to the DBMS, instead of SAS reading the entire result set and deleting one row at a time.

Examples

The following example demonstrates the use of DIRECT_EXE= to empty a table from a database.

```
libname x oracle user=scott password=tiger
    path=oracllev8 schema=dbitest
direct_exe=delete; /* create an oracle table of 5 rows */data x.dbi_dft;
do coll=1 to 5;
output;
end;
run;

options sastrace=",,,d" sastraceloc=saslog nostsuffix;
proc sql;
delete * from x.dbi_dft;
quit;
```

By turning trace on, you should see something similar to the following:

Output 9.1 SAS Log Output

```
ORACLE_9: Executed:
delete from dbi_dft
```

DIRECT_SQL= LIBNAME Option

Lets you specify whether generated SQL is passed to the DBMS for processing

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: YES

Syntax

DIRECT_SQL= YES | NO | NONE | NOGENSQL | NOWHERE | NOFUNCTIONS
| NOMULTOUTJOINS

Syntax Description

YES

specifies that generated SQL from PROC SQL is passed directly to the DBMS for processing.

NO

specifies that generated SQL from PROC SQL is not passed to the DBMS for processing. This is the same as specifying the value NOGENSQL.

NONE

specifies that generated SQL is not passed to the DBMS for processing. This includes SQL that is generated from PROC SQL, SAS functions that can be converted into DBMS functions, joins, and WHERE clauses.

NOGENSQL

prevents PROC SQL from generating SQL to be passed to the DBMS for processing.

NOWHERE

prevents WHERE clauses from being passed to the DBMS for processing. This includes SAS WHERE clauses and PROC SQL generated or PROC SQL specified WHERE clauses.

NOFUNCTIONS

prevents SQL statements from being passed to the DBMS for processing when they contain functions.

NOMULTOUTJOINS

specifies that PROC SQL does not attempt to pass any multiple outer joins to the DBMS for processing. Other join statements may be passed down however, including portions of a multiple outer join.

Details

By default, processing is passed to the DBMS whenever possible, because the database might be able to process the functionality more efficiently than SAS does. In some instances, however, you might not want the DBMS to process the SQL. For example, the presence of null values in the DBMS data might cause different results depending on whether the processing takes place in SAS or in the DBMS. If you do not want the

DBMS to handle the SQL, use `DIRECT_SQL=` to force SAS to handle some or all of the SQL processing.

If you specify `DIRECT_SQL=NOGENSQL`, then PROC SQL does not generate DBMS SQL. This means that SAS functions, joins, and DISTINCT processing that occur *within* PROC SQL are not passed to the DBMS for processing. (SAS functions *outside* PROC SQL can still be passed to the DBMS.) However, if PROC SQL contains a WHERE clause, the WHERE clause *is* passed to the DBMS, if possible. Unless you specify `DIRECT_SQL=NOWHERE`, SAS attempts to pass all WHERE clauses to the DBMS.

If you specify more than one value for this option, separate the values with spaces and enclose the list of values in parentheses. For example, you could specify `DIRECT_SQL=(NOFUNCTIONS, NOWHERE)`.

`DIRECT_SQL=` overrides the LIBNAME option “SQL_FUNCTIONS= LIBNAME Option” on page 146. If you specify `SQL_FUNCTIONS=ALL` and `DIRECT_SQL=NONE`, no functions are passed.

Examples

The following example prevents a join between two tables from being processed by the DBMS, by setting `DIRECT_SQL=NOGENSQL`. Instead, SAS processes the join.

```
proc sql;
create view work.v as
  select tabl.deptno, dname from
         mydblib.table1 tabl,
         mydblib.table2 tab2
  where tabl.deptno=tab2.deptno
  using libname mydblib oracle user=testuser
         password=testpass path=myserver direct_sql=nogensql;
```

The following example prevents a SAS function from being processed by the DBMS.

```
libname mydblib oracle user=testuser password=testpass direct_sql=nofunctions;
proc print data=mydblib.tab1;
  where lastname=soundex ('Paul');
```

ENABLE_BULK= LIBNAME Option

Enables the connection to process bulk copy when you load data into a Sybase table

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Sybase

Default value: YES

Syntax

`ENABLE_BULK=YES | NO`

Syntax Description

NO

disables the bulk copy ability for the libref.

YES

enables the connection to perform a bulk copy of SAS data into Sybase.

Details

Bulk copy groups rows so that they are inserted as a unit into the Sybase table. Using bulk copy can improve performance.

If you use both the LIBNAME option, ENABLE_BULK=, and the data set option, BULKLOAD=, the values of the two options must be the same or an error is returned. However, since the default value of ENABLE_BULK= is YES, you do not have to specify ENABLE_BULK= in order to use the BULKLOAD= data set option.

Note: In Version 7 and previous releases, this option was called BULKCOPY=. In Version 8 and later, an error is returned if you specify BULKCOPY=. Δ

ERRLIMIT= LIBNAME Option

Specifies the number of errors that are allowed while using the Fastload utility before SAS stops loading data to Teradata.

Valid in: DATA and PROC steps (wherever Fastload is used)

DBMS support: Teradata

Default value: 1 million

Syntax

ERRLIMIT=*integer*

Syntax Description

integer

positive integer that represents the number of errors after which SAS stops loading data.

Details

SAS stops loading data when the specified number of errors is reached and Fastload is paused. When Fastload is paused, the table being loaded cannot be used. Since restart capability for Fastload is not yet supported, the error tables must be manually deleted before the table can be loaded again.

Example

In the following example, SAS stops processing and pauses Fastload at the occurrence of the tenth error.


```

libname mydblib teradata user=terauser pw=XXXXXX ERRLIMIT=10;

data mydblib.trfload(bulkload=yes dbtype=(i='int check (i > 11)') );
  do
    i=1 to 50000;output;
  end;
run;

```

IGNORE_READ_ONLY_COLUMNS= LIBNAME Option

Specifies whether to ignore or include columns whose data types are read-only when generating an SQL statement for inserts or updates

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB

Default value: NO

Syntax

IGNORE_READ_ONLY_COLUMNS=YES | NO

Syntax Description

YES

specifies that the SAS/ACCESS engine ignores columns whose data types are read-only when you are generating insert and update SQL statements

NO

specifies that the SAS/ACCESS engine does not ignore columns whose data types are read-only when you are generating insert and update SQL statements

Details

Several databases include data types that can be read-only, such as Microsoft SQL Server's timestamp data type. Several databases also have properties that allow certain data types to be read-only, such as Microsoft SQL Server's identity property.

When the IGNORE_READ_ONLY_COLUMNS option is set to NO (default) and a DBMS table contains a read-only column, an error is returned indicating that the data could not be modified for that column.

Example

For the following example, a database that contains the table Products is created with two columns: ID and PRODUCT_NAME. The ID column is defined by a read-only data type and PRODUCT_NAME is a character column.

```
CREATE TABLE products (id int IDENTITY PRIMARY KEY, product_name varchar(40))
```

Assume you have a SAS data set that contains the name of your products, and you would like to insert the data into the Products table:

```
data work.products;
  id=1;
  product_name='screwdriver';
  output;
  id=2;
  product_name='hammer';
  output;
  id=3;
  product_name='saw';
  output;
  id=4;
  product_name='shovel';
  output;
run;
```

With IGNORE_READ_ONLY_COLUMNS=NO (the default), an error is returned by the database because in this example, the ID column cannot be updated. However, if you set the option to YES and execute a PROC APPEND, the append succeeds, and the SQL statement that is generated does not contain the ID column.

```
libname x odbc uid=dbitest pwd=dbigrpl dsn=lupinss
          ignore_read_only_columns=yes;
options sastrace=',,,d' sastraceloc=saslog nostsuffix;
proc append base=x.PRODUCTS data=work.products;
run;
```

See Also

To apply this option to an individual data set, see the data set option “IGNORE_READ_ONLY_COLUMNS= Data Set Option” on page 224.

IN= LIBNAME Option

Lets you specify the database and tablespace in which you want to create a new table

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS

Default value: none

Syntax

IN='database-name.tablespace-name' | 'DATABASE database-name'

Syntax Description

database-name.tablespace-name

specifies the names of the database and tablespace, which are separated by a period. Enclose the entire specification in single quotation marks.

DATABASE *database-name*

specifies only the database name. Specify the word DATABASE, then a space, then the database name. Enclose the entire specification in single quotation marks.

Details

The IN= option is relevant only when you are creating a new table. If you omit this option, the default is to create the table in the default database, implicitly creating a simple tablespace.

See Also

To apply this option to an individual data set, see the data set option “IN= Data Set Option” on page 225.

INSERT_SQL= LIBNAME Option

Determines the method that is used to insert rows into a data source

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Microsoft SQL Server, ODBC, OLE DB

Default value: DBMS-specific, see details below

Syntax

INSERT_SQL=YES | NO

Syntax Description

YES

specifies that SAS/ACCESS uses the data source’s SQL insert method to insert new rows into a table.

NO

specifies that SAS/ACCESS uses an alternate (DBMS-specific) method to insert new rows into a table.

Details

Flat-file databases (such as dBASE, FoxPro, and text files) generally have improved insert performance when INSERT_SQL=NO. Other databases might have inferior insert performance (or might fail) with this setting, so you should experiment to determine the optimal setting for your situation.

Microsoft SQL Server Details: The Microsoft SQL Server default is YES. When INSERT_SQL=NO, the SQLSetPos (SQL_ADD) function inserts rows in groups that are the size of the INSERTBUFF= option value. The SQLSetPos (SQL_ADD) function does not work unless it is supported by your driver.

ODBC Details: The ODBC default is YES, except for Microsoft Access, which has a default of NO. When INSERT_SQL=NO, the SQLSetPos (SQL_ADD) function inserts rows in groups that are the size of the INSERTBUFF= option value. The SQLSetPos (SQL_ADD) function does not work unless it is supported by your driver.

OLE DB Details: By default, the OLE DB interface attempts to use the most efficient row insertion method for each data source. You can use the INSERT_SQL option to override the default in the event that it is not optimal for your situation. The OLE DB alternate method (used when this option is set to NO) uses the OLE DB IRowsetChange interface.

See Also

To apply this option to an individual data set, see the data set option “INSERT_SQL= Data Set Option” on page 226.

INSERTBUFF= LIBNAME Option

Specifies the number of rows in a single insert operation

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle

Default value: DBMS-specific

Syntax

INSERTBUFF=*positive-integer*

Syntax Description

positive-integer

specifies the number of rows to insert. SAS allows the maximum that is allowed by the DBMS.

Details

The optimal value for this option varies with factors such as network type and available memory. You might need to experiment with different values in order to determine the best value for your site.

The SAS application messages that indicate the success or failure of an insert operation represent information for only a single insert, even when multiple inserts are performed. Therefore, when you assign a value that is greater than INSERTBUFF=1, these messages might be incorrect.

If you specify the DBCOMMIT= option with a value that is less than the value of INSERTBUFF=, then DBCOMMIT= overrides INSERTBUFF=.

Note: When you insert rows by using the VIEWTABLE window or the FSVIEW or FSEEDIT procedure, use INSERTBUFF=1 to prevent the DBMS interface from trying to insert multiple rows. These features do not support inserting more than one row at a time. △

Note: Additional driver-specific restrictions might apply. △

DB2 UNIX/PC Details: You must specify INSERT_SQL=YES in order to use this option. If one row in the insert buffer fails, all rows in the insert buffer fail. The default is calculated based upon the row length of your data.

Microsoft SQL Server Details: You must specify INSERT_SQL=YES in order to use this option. The default is 1.

MySQL Details: The default is 0. Any value greater than 0 turns on the INSERTBUFF= option. The engine then calculates how many rows it can insert at one time, based on the row size. If one row in the insert buffer fails, all rows in the insert buffer might fail, depending on your storage type.

HP Neoview and Netezza Details: The default is automatically calculated based on row length.

ODBC Details: The default is 1.

OLE DB Details: The default is 1.

Oracle Details: The default is 10.

See Also

To apply this option to an individual data set, see the data set option “INSERTBUFF= Data Set Option” on page 227.

INTERFACE= LIBNAME Option

Specifies the name and location of the interfaces file that is searched when you connect to the Sybase server

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Sybase

Default value: none

Syntax

INTERFACE=<'>filename<'>

Details

The interfaces file contains names and access information for the available servers on the network. If you omit a filename, the default action for your operating system occurs. INTERFACE= is not used in some operating environments. Contact your database administrator to see whether this statement applies to your computing environment.

KEYSET_SIZE= LIBNAME Option

Specifies the number of rows that are keyset driven

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: Microsoft SQL Server, ODBC

Default value: 0

Syntax

KEYSET_SIZE=*number-of-rows*

Syntax Description

number-of-rows

is an integer with a value between 0 and the number of rows in the cursor.

Details

This option is valid only when `CURSOR_TYPE=KEYSET_DRIVEN`. See “`CURSOR_TYPE= LIBNAME Option`” on page 95 for more information about keyset-driven cursors.

If `KEYSET_SIZE=0`, then the entire cursor is keyset driven. If a value greater than 0 is specified for `KEYSET_SIZE=`, then the value chosen indicates the number of rows within the cursor that will function as a keyset-driven cursor. When you scroll beyond the bounds that are specified by `KEYSET_SIZE=`, then the cursor becomes dynamic and new rows might be included in the cursor. This becomes the new keyset and the cursor functions as a keyset-driven cursor again. Whenever the specified value is between 1 and the number of rows in the cursor, the cursor is considered to be a mixed cursor. The reason is part of the cursor functions as a keyset-driven cursor, and part of the cursor functions as a dynamic cursor.

See Also

To apply this option to an individual data set, see the data set option “`KEYSET_SIZE= Data Set Option`” on page 228.

LOCATION= LIBNAME Option

Lets you further qualify exactly where a table resides

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 z/OS

Default value: none

Syntax

LOCATION=*location*

Details

The location name maps to the location in the SYSIBM.SYSLOCATIONS catalog in the communication database.

In the DB2 interface, the location is converted to the first level of a three-level table name: *location.authid.table*. The connection to the remote DB2 subsystem is done implicitly by DB2's Distributed Data Facility (DDF) when DB2 receives a three-level table name in an SQL statement.

If you omit this option, SAS accesses the data from the local DB2 database unless you have specified a value for the SERVER= option. This option is not validated until you access a DB2 table. If you specify LOCATION=, you must also specify the AUTHID= option.

See Also

- To apply this option to an individual data set, see the data set option “LOCATION= Data Set Option” on page 229
- For information about accessing a database server on Linux, UNIX, or Windows using a libref, see the REMOTE_DBTYPE= LIBNAME option.

LOCKTABLE= LIBNAME Option

Places exclusive or shared locks on tables

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Informix

Default value: no locking

Syntax

LOCKTABLE= EXCLUSIVE | SHARE

Syntax Description

EXCLUSIVE

specifies that other users are prevented from accessing each table that you open in the libref.

SHARE

specifies that other users or processes can read data from the tables, but they cannot update the data.

Details

You can lock tables only if you are the owner or have been granted the necessary privilege.

See Also

To apply this option to an individual data set, see the data set option “LOCKTABLE= Data Set Option” on page 230.

LOCKTIME= LIBNAME Option

Specifies the number of seconds to wait until rows are available for locking

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Informix

Default value: none

Syntax

LOCKTIME=*positive-integer*

Details

You must specify LOCKWAIT=YES for LOCKTIME= to have an effect. If you omit the LOCKTIME= option and use LOCKWAIT=YES, SAS suspends your process indefinitely until a lock can be obtained.

See Also

“LOCKWAIT= LIBNAME Option” on page 124

LOCKWAIT= LIBNAME Option

Specifies whether to wait indefinitely until rows are available for locking

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Informix, Oracle

Default value: DBMS-specific

Syntax

LOCKWAIT= YES | NO

Syntax Description

YES

specifies that SAS waits until rows are available for locking.

NO

specifies that SAS does not wait and returns an error to indicate that the lock is not available.

LOGDB= LIBNAME Option

Redirects the log tables created by Teradata's FastExport utility to an alternate database

Valid in: DATA and PROC steps (wherever Fastload is used)

DBMS support: Teradata

Default value: Default Teradata database for the libref

Syntax

LOGDB=<*database-name*>

Syntax Description

database-name

the name of the Teradata database.

Details

This option enables the restart log tables that are used by Teradata's FastExport utility to be created in an alternate database. You must have necessary permissions to create tables in the specified database. Only restart tables are created in the specified database. Restart capability of FastExport is not yet supported.

Example

In the following example, PROC PRINT invokes Teradata's FastExport utility if it is installed. The LOGDB= option in the LIBNAME statement causes the restart log tables that are created by Teradata's FastExport to be created in the database ALTDB.

```
libname mydblib teradata user=testuser pw=testpass logdb=altdb;
proc print data=mydblib.mytable(dbsliceparm=all);
run;
```

MAX_CONNECTS= LIBNAME Option

Specifies the maximum number of simultaneous connections that Sybase allows

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Sybase

Default value: 25

Syntax

MAX_CONNECTS=*numeric-value*

Details

If you omit MAX_CONNECTS=, the default for the maximum number of connections is 25. Note that increasing the number of connections has a direct impact on memory.

MULTI_DATASRC_OPT= LIBNAME Option

Used in place of DBKEY to improve performance when processing a join between two data sources

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: NONE

Syntax

MULTI_DATASRC_OPT=NONE | IN_CLAUSE

Syntax Description

NONE

turns off the functionality of the option.

IN_CLAUSE

specifies that an IN clause containing the values read from a smaller table are used to retrieve the matching values in a larger table based on a key column designated in an equi-join.

Details

When processing a join between a SAS data set and a DBMS table, the SAS data set should be smaller than the DBMS table for optimal performance. However, in the event that the SAS data set is *larger* than that DBMS table, the SAS data set is still used in the IN clause.

When SAS is processing a join between two DBMS tables, SELECT COUNT (*) is issued to determine which table is smaller and if it qualifies for an IN clause. You can use “DBMASTER= Data Set Option” on page 210 to prevent the SELECT COUNT (*) from being issued.

Currently, the IN clause has a limit of 4,500 unique values.

Oracle Details: Oracle can handle an IN clause of only 1,000 values. Therefore, it divides larger IN clauses into multiple, smaller IN clauses. The results are combined into a single result set. For example if an IN clause contained 4,000 values, Oracle produces four IN clauses that each contain 1,000 values. A single result is produced, as if all 4,000 values were processed as a whole.

OLE DB Details: OLE DB restricts the number of values allowed in an IN clause to 255.

Setting DBKEY= automatically overrides MULTI_DATASRC_OPT=.

DIRECT_SQL= can impact this option as well. If DIRECT_SQL=NONE or NOWHERE, the IN clause cannot be built and passed to the DBMS, regardless of the value of MULTI_DATASRC_OPT=. These setting for DIRECT_SQL= prevent a WHERE clause from being passed.

Examples

The following example builds and passes an IN clause from the SAS table to the DBMS table, retrieving only the necessary data to process the join:

```
proc sql;
create view work.v as
select tab2.deptno, tab2.dname from
work.sastable tab1, dblink.table2 tab2
where tab12.deptno = tab2.deptno
using libname dblink oracle user=testuser password=testpass
multi_datasrc_opt=in_clause;
quit;
```

The following example prevents the building and passing of the IN clause to the DBMS. This action requires all rows from the DBMS table to be brought into SAS for processing the join:

```
libname dblink oracle user=testuser password=testpass multi_datasrc_opt=none;
proc sql;
select tab2.deptno, tab2.dname from
work.table1 tab1,
dblib.table2 tab2
where tab1.deptno=tab2.deptno;
quit;
```

See Also

“DBMASTER= Data Set Option” on page 210

OR_UPD_NOWHERE= LIBNAME Option

Specifies whether SAS uses an extra WHERE clause when updating rows with no locking

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Oracle

Alias: ORACLE_73_OR_ABOVE=

Default value: YES

Syntax

OR_UPD_NOWHERE= YES | NO

Syntax Description

YES

specifies that SAS does not use an additional WHERE clause to determine whether each row has changed since it was read. Instead, SAS uses the **SERIALIZABLE** isolation level (available with Oracle 7.3 and above) for update locking. If a row changes after the serializable transaction starts, the update on that row fails.

NO

specifies that SAS uses an additional WHERE clause to determine whether each row has changed since it was read. If a row has changed since being read, the update fails.

Details

Use this option when you are updating rows without locking (UPDATE_LOCK_TYPE=NOLOCK).

By default (OR_UPD_NOWHERE=YES), updates are performed in serializable transactions so you can avoid extra WHERE clause processing and potential WHERE clause floating point precision problems.

Note: Due to the published Oracle bug 440366, sometimes an update on a row fails even if the row has not changed. Oracle offers the following solution: When creating a table, increase the number of INITRANS to at least 3 for the table. Δ

See Also

To apply this option to an individual data set or a view descriptor, see the data set option “OR_UPD_NOWHERE= Data Set Option” on page 249.

PACKETSIZE= LIBNAME Option

Lets you specify the packet size for Sybase to use

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Sybase

Default value: current server setting

Syntax

PACKETSIZE=*numeric-value*

Syntax Description

numeric-value

is any multiple of 512, up to the limit of the maximum network packet size setting on your server.

Details

If you omit **PACKETSIZE=**, the default is the current server setting. You can query the default network packet value in ISQL by using the Sybase **sp_configure** command.

PREFETCH= LIBNAME Option

Enables the PreFetch facility on tables that are accessed by the libref defined with the LIBNAME statement

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Teradata

Default value: not enabled

Syntax

PREFETCH=*'unique_storename, [#sessions,algorithm]'*

Syntax Description

unique_storename

is a unique name that you specify. This value names the Teradata macro that PreFetch creates to store selected SQL statements in the first run of a job. During subsequent runs of the job, SAS/ACCESS presubmits the stored SQL statements in parallel to the Teradata DBMS.

#sessions

controls the number of statements that PreFetch submits in parallel to Teradata. A valid value is 1 through 9. If you do not specify a *#sessions* value, the default is 3.

algorithm

specifies the algorithm that PreFetch uses to order the selected SQL statements. Currently, the only valid value is SEQUENTIAL.

Details

Before using PreFetch, see the PreFetch description in the SAS/ACCESS documentation for Teradata for a complete discussion. Review when and how the option enhances read performance of a job that is run more than once.

PRESERVE_COL_NAMES= LIBNAME Option

Preserves spaces, special characters, and case-sensitivity in DBMS column names when you create DBMS tables

Valid in: the SAS/ACCESS LIBNAME statement (when you create DBMS tables)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Teradata

Default value: DBMS-specific

Syntax

PRESERVE_COL_NAMES= NO | YES

Syntax Description**NO**

specifies that column names used in DBMS table creation are derived from SAS variable names by using the SAS variable name normalization rules. See the VALIDVARNAME system option for more information. However, the database applies its DBMS-specific normalization rules to the SAS variable names when creating the DBMS column names.

The use of N-Literals to create column names that use database keywords or special symbols other than the underscore character might be illegal when DBMS normalization rules are applied. To include nonstandard SAS symbols or database keywords, specify PRESERVE_COL_NAMES=YES.

NO is the default for most DBMS interfaces.

YES

specifies that column names that are used in table creation are passed to the DBMS with special characters and the exact, case-sensitive spelling of the name preserved.

Details

This option applies only when you use SAS/ACCESS to create a new DBMS table. When you create a table, you assign the column names by using one of the following methods:

- To control the case of the DBMS column names, specify variables using the desired case and set PRESERVE_COL_NAMES=YES. If you use special symbols or blanks, you must set VALIDVARNAME= to ANY and use N-Literals. See the section about names in *SAS/ACCESS for Relational Databases: Reference* and the section about system options in *SAS Language Reference: Dictionary* for further information.
- To enable the DBMS to normalize the column names according to its naming conventions, specify variables using any case and set PRESERVE_COLUMN_NAMES= NO.

Note: When you use SAS/ACCESS to read from, insert rows into, or modify data in an existing DBMS table, SAS identifies the database column names by their spelling. Therefore, when the database column exists, the case of the variable does not matter. Δ

See the topic about naming in the documentation for your SAS/ACCESS interface for additional details.

Specify the alias PRESERVE_NAMES= if you plan to specify both the PRESERVE_COL_NAMES= and PRESERVE_TAB_NAMES= options in your LIBNAME statement. Using this alias saves you some time when coding.

To use column names in your SAS program that are not valid SAS names, you must use one of the following techniques:

- Use the DQUOTE= option in PROC SQL and then reference your columns using double quotation marks. For example:

```
proc sql dquote=ansi;
  select "Total$Cost" from mydblib.mytable;
```

- Specify the global system option VALIDVARNAME=ANY and use name literals in the SAS language. For example:

```
proc print data=mydblib.mytable;
  format 'Total$Cost'n 22.2;
```

Note that if you are *creating* a table in PROC SQL, you must also include the PRESERVE_COL_NAMES=YES option in your LIBNAME statement. For example:

```
libname mydblib oracle user=testuser password=testpass
  preserve_col_names=yes;
proc sql dquote=ansi;
  create table mydblib.mytable ("my$column" int);
```

PRESERVE_COL_NAMES= does not apply to the Pass-Through Facility.

See Also

To apply this option to an individual data set, see the PRESERVE_COL_NAMES= data set option “PRESERVE_COL_NAMES= Data Set Option” on page 251 naming in your DBMS interface.

PRESERVE_TAB_NAMES= LIBNAME Option

Preserves spaces, special characters, and case-sensitivity in DBMS table names

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza ODBC, OLE DB, Oracle, Teradata

Default value: DBMS-specific

Syntax

PRESERVE_TAB_NAMES= NO | YES

Syntax Description

NO

specifies that when you create DBMS tables or refer to an existing table, the table names are derived from SAS member names by using SAS member name normalization. However, the database applies DBMS-specific normalization rules to the SAS member names. Therefore, the table names are created or referenced in the database following the DBMS-specific normalization rules.

When you use SAS to read a list of table names (for example, in the SAS Explorer window), the tables, whose names do not conform to the SAS member name normalization rules, do not appear in the output. In SAS line mode, the number of tables that do not display from PROC DATASET due to this restriction is noted as a note:

Note: "Due to the PRESERVE_TAB_NAMES=NO LIBNAME option setting, 12 table(s) have not been displayed." Δ

You do not get this warning when using SAS Explorer.

SAS Explorer displays DBMS table names in capitalized form when PRESERVE_TAB_NAMES=NO. This is now how the tables are represented in the DBMS.

NO is the default for most DBMS interfaces.

YES

specifies that table names are read from and passed to the DBMS with special characters, and the exact, case-sensitive spelling of the name is preserved.

Details

See the topic on naming in the documentation for your SAS/ACCESS interface for additional details.

To use table names in your SAS program that are not valid SAS names, use one of the following techniques:

- Use the PROC SQL option DQUOTE= and place double quotation marks around the table name. The libref must specify PRESERVE_TAB_NAMES=YES. For example:

```
libname mydblib oracle user=testuser password=testpass
preserve_tab_names=yes;
```



```
proc sql dquote=ansi;
  select * from mydblib."my table";
```

- Use name literals in the SAS language. The libref must specify PRESERVE_TAB_NAMES=YES. For example:

```
libname mydblib oracle user=testuser password=testpass preserve_tab_names=yes;
proc print data=mydblib.'my table'n;
run;
```

Specify the alias PRESERVE_NAMES= to save time if you are specifying both PRESERVE_COL_NAMES= and PRESERVE_TAB_NAMES= in your LIBNAME statement.

Oracle Details: Unless you specify PRESERVE_TAB_NAMES=YES, the table name that you enter for SCHEMA= or for the DBINDEX= data set option is converted to uppercase.

Example

If you use PROC DATASETS to read the table names in an Oracle database that contains three tables, My_Table, MY_TABLE, and MY TABLE. The results differ depending on the setting of PRESERVE_TAB_NAMES.

If the libref specifies PRESERVE_TAB_NAMES=NO, then the PROC DATASETS output is one table name, MY_TABLE. This is the only table name that is in Oracle normalized form (uppercase letters and a valid symbol, the underscore). My_Table is not displayed because it is not in Oracle - MY TABLE is not displayed because it is not in SAS member normalized form (the embedded space is a nonstandard SAS character).

If the libref specifies PRESERVE_TAB_NAMES=YES, then the PROC DATASETS output includes all three table names, My_Table, MY_TABLE, and MY TABLE.

See Also

To apply this option to an individual data set, see the PRESERVE_TAB_NAMES= data set option “PRESERVE_COL_NAMES= LIBNAME Option” on page 130 naming in your DBMS interface.

QUALIFIER= LIBNAME Option

Lets you identify database objects, such as tables and views, using the specified qualifier

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB

Default value: none

Syntax

QUALIFIER=<qualifier-name>

Details

If this option is omitted, the default is the default DBMS qualifier name, if any. QUALIFIER= can be used for any DBMS that allows three-part identifier names, such as *qualifier.schema.object*.

MySQL Details: The MySQL interface does not support three-part identifier names, so a two-part name is used (such as *qualifier.object*).

Examples

In the following LIBNAME statement, the QUALIFIER= option causes any reference in SAS to mydblib.employee to be interpreted by ODBC as mydept.scott.employee.

```
libname mydblib odbc dsn=myoracle
      password=testpass schema=scott
      qualifier=mydept;
```

In the following example, the QUALIFIER= option causes any reference in SAS to mydblib.employee to be interpreted by OLE DB as pcdivision.raoul.employee.

```
libname mydblib oledb provider=SQLOLEDB
      properties=("user id=dbajorge "data source"=SQLSERVER)
      schema=raoul qualifier=pcdivision;
proc print data=mydblib.employee;
run;
```

See Also

To apply this option to an individual data set, see the data set option “QUALIFIER= Data Set Option” on page 253.

QUALIFY_ROWS= LIBNAME Option

Uniquely qualifies all member values in a result set

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: OLE DB

Default value: NO

Syntax

QUALIFY_ROWS= YES | NO

Syntax Description

YES

specifies that when the OLE DB interface flattens the result set of an MDX command, the values in each column are uniquely identified using a hierarchical naming scheme.

NO

specifies that when the OLE DB interface flattens the result set of an MDX command, the values in each column are not qualified, which means they might not be unique.

Details

For example, when this option is set to NO, a GEOGRAPHY column might have a value of **PORTLAND** for **Portland, Oregon**, and the same value of **PORTLAND** for **Portland, Maine**. When you set this option to YES, the two values might become **[USA].[Oregon].[Portland]** and **[USA].[Maine].[Portland]**, respectively.

Note: Depending on the size of the result set, QUALIFY_ROWS=YES can have a significant, negative impact on performance. The reason is that because it forces the OLE DB interface to search through various schemas to gather the information needed to create unique qualified names. △

For more information about MDX commands, see the SAS/ACCESS documentation for OLE DB.

QUERY_TIMEOUT= LIBNAME Option

Specifies the number of seconds of inactivity to wait before canceling a query

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, Netezza, ODBC

Default value: 0

Syntax

QUERY_TIMEOUT=*number-of-seconds*

Details

The default value of 0 indicates that there is no time limit for a query. This option is useful when you are testing a query or if you suspect that a query might contain an endless loop.

See Also

To apply this option to an individual data set, see the data set option “QUERY_TIMEOUT= Data Set Option” on page 253.

QUOTE_CHAR= LIBNAME Option

Specifies which quotation mark character to use when delimiting identifiers

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB

Default value: none

Syntax

QUOTE_CHAR=*character*

Syntax Description

character

is the quotation mark character to use when delimiting identifiers, such as the double quotation mark (").

Details

The provider usually specifies the delimiting character. However, when there is a difference between what the provider allows for this character and what the DBMS allows, the QUOTE_CHAR= option overrides the character returned by the provider.

Microsoft SQL Server Details: QUOTE_CHAR= overrides the Microsoft SQL Server default.

ODBC Details: This option is mainly for the ODBC interface to Sybase and should be used in conjunction with the DBCONINIT and DBLIBINIT LIBNAME options.

QUOTE_CHAR= overrides the ODBC default because some drivers return a blank for the identifier delimiter even though the DBMS uses a quotation mark (for example, ODBC to Sybase).

Examples

If you would like your quotation character to be a single quotation mark, then specify the following:

```
libname x odbc dsn=mydsn pwd=mypassword quote_char=''';
```

If you would like your quotation character to be a double quotation mark, then specify the following:

```
libname x odbc dsn=mydsn pwd=mypassword quote_char='\"';
```

QUOTED_IDENTIFIER= LIBNAME Option

Lets you specify table and column names with embedded spaces and special characters

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Sybase

Default value: NO

Syntax

QUOTED_IDENTIFIER=YES|NO

Details

This option is used in place of the PRESERVE_TAB_NAMES and PRESERVE_COL_NAMES options, which have no effect on the Sybase interface, due to Sybase's default to case-sensitivity.

READ_ISOLATION_LEVEL= LIBNAME Option

Defines the degree of isolation of the current application process from other concurrently running application processes

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: DB2 UNIX/PC, DB2 z/OS, Informix, Microsoft SQL Server, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

READ_ISOLATION_LEVEL=*DBMS-specific value*

Syntax Description

See the documentation for your SAS/ACCESS interface for the values for your DBMS.

Details

The degree of isolation defines

- the degree to which rows that are read and updated by the current application are available to other concurrently executing applications
- the degree to which update activity of other concurrently executing application processes can affect the current application.

This option is ignored in the interfaces to DB2 under UNIX and PC Hosts and ODBC if READ_LOCK_TYPE= is not set to ROW.

For DBMS-specific details, see locking in the documentation for your SAS/ACCESS interface.

See Also

To apply this option to an individual data set, see the data set option “READ_ISOLATION_LEVEL= Data Set Option” on page 254.

READ_LOCK_TYPE= LIBNAME Option

Specifies how data in a DBMS table is locked during a READ transaction

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, Microsoft SQL Server, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

READ_LOCK_TYPE= ROW | PAGE | TABLE | NOLOCK | VIEW

Syntax Description

ROW

locks a row if any of its columns are accessed. If you are using the interface to ODBC or DB2 under UNIX and PC hosts, READ_LOCK_TYPE=ROW indicates that locking is based on the READ_ISOLATION_LEVEL= option. (This value is valid in the DB2 under UNIX and PC hosts, Microsoft SQL Server, ODBC, and Oracle interfaces.)

PAGE

locks a page of data, which is a DBMS-specific number of bytes. (This value is valid in the Sybase interface.)

TABLE

locks the entire DBMS table. If you specify READ_LOCK_TYPE=TABLE, you must also specify CONNECTION= UNIQUE, or you receive an error message. Setting CONNECTION=UNIQUE ensures that your table lock is not lost, for example, due to another table closing and committing rows in the same connection. (This value is valid in the DB2 under z/OS, DB2 under UNIX and PC hosts, ODBC, Oracle, Microsoft SQL Server, and Teradata interfaces.)

NOLOCK

does not lock the DBMS table, pages, or rows during a read transaction. (This value is valid in the Oracle and Sybase interfaces, and in the ODBC and Microsoft SQL Server interfaces when using the Microsoft SQL Server driver.)

VIEW

locks the entire DBMS view. (This value is valid in the Teradata interface.)

Details

If you omit READ_LOCK_TYPE=, the default is the DBMS' default action. You can set a lock for one DBMS table by using the data set option or for a group of DBMS tables by using the LIBNAME option.

For DBMS-specific details, see locking in the documentation for your SAS/ACCESS interface.

See the documentation for your SAS/ACCESS interface for additional details.

Example

In the following example, the libref MYDBLIB uses the SAS/ACCESS interface to Oracle to connect to an Oracle database. USER=, PASSWORD=, and PATH= are SAS/ACCESS connection options. The LIBNAME options specify that row-level locking is used when data is read or updated:

```
libname mydblib oracle user=testuser password=testpass
      path=myoraph read_lock_type=row update_lock_type=row;
```

See Also

To apply this option to an individual data set, see the data set option “READ_LOCK_TYPE= Data Set Option” on page 254.

READ_MODE_WAIT= LIBNAME Option

Specifies during SAS/ACCESS read operations whether Teradata should wait to acquire a lock or should fail the request when the DBMS resource is already locked by a different user

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Teradata

Default value: none

Syntax

READ_MODE_WAIT= YES|NO

Syntax Description

YES

specifies for Teradata to wait to acquire the lock, so SAS/ACCESS waits indefinitely until it can acquire the lock.

NO

specifies for Teradata to fail the lock request if the specified DBMS resource is locked.

Details

If you specify READ_MODE_WAIT=NO and if a different user holds a restrictive lock, then the executing SAS step fails. SAS/ACCESS continues processing the job by

executing the next step. For more information, see the SAS/ACCESS documentation for Teradata.

If you specify READ_MODE_WAIT=YES, SAS/ACCESS waits indefinitely until it can acquire the lock.

A *restrictive* lock means that another user is holding a lock that prevents you from obtaining your desired lock. Until the other user releases the restrictive lock, you cannot obtain your lock. For example, another user's table level WRITE lock prevents you from obtaining a READ lock on the table.

See Also

To apply this option to an individual data set, see the data set option "READ_MODE_WAIT= Data Set Option" on page 255.

READBUFF= LIBNAME Option

Specifies the number of rows of DBMS data to read into the buffer

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle, Sybase

Default value: DBMS-specific

Syntax

READBUFF=*integer*

Syntax Description

integer

is the positive number of rows to hold in memory. SAS allows the maximum number that is allowed by the DBMS.

Details

This option improves performance by specifying a number of rows that can be held in memory for input into SAS. Buffering data reads can decrease network activities and increase performance. However, because SAS stores the rows in memory, higher values for READBUFF= use more memory. In addition, if too many rows are selected at once, then the rows that are returned to the SAS application might be out of date. For example, if someone else modifies the rows, you do not see the changes.

When READBUFF=1, only one row is retrieved at a time. The higher the value for READBUFF=, the more rows the DBMS engine retrieves in one fetch operation.

DB2 UNIX/PC Details: If you do not specify this option, the buffer size is automatically calculated based upon the row length of your data and the SQLExtendedFetch API call is used (this is the default). ROWSET_SIZE= is an alias for this option.

Microsoft SQL Server and ODBC Details: If you do not specify this option, the SQLFetch API call is used and no internal SAS buffering is performed (this is the default). Setting READBUFF=1 or greater causes the SQLExtendedFetch API call to be used. ROWSET_SIZE= is an alias for this option.

HP Neoview and Netezza Details: The default is automatically calculated based on row length.

OLE DB Details: The default is 1. ROWSET_SIZE= is an alias for this option.

Oracle Details: The default is 250.

Sybase Details: The default is 100.

See Also

To apply this option to an individual data set, see the data set option “READBUFF= Data Set Option” on page 256.

REMOTE_DBTYPE= LIBNAME Option

Specifies whether the libref points to a database server on z/OS or to one on Linux, UNIX, or Windows (LUW)

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 z/OS

Default value: ZOS

Syntax

REMOTE_DBTYPE= LUW | ZOS

Syntax Description

LUW

specifies that the database server that is accessed through the libref resides on Linux, UNIX, or Windows.

ZOS

specifies that the database server that is accessed through the libref runs on z/OS (default).

Details

Specifying REMOTE_DBTYPE= in the LIBNAME statement ensures that the SQL that is used by some SAS procedures to access the DB2 catalog tables is generated properly, and that it is based upon the database server type.

This option also enables special catalog calls (such as DBMS::Indexes) to function properly when the target database does not reside on a mainframe computer.

REMOTE_DBTYPE= is used in conjunction with the SERVER= CONNECT statement option or the LOCATION= LIBNAME option. If neither option is used, then REMOTE_DBTYPE= is ignored.

Example

The following is an example of using REMOTE_DBTYPE= with the SERVER= option.

```
libname mylib db2 ssid=db2a server=db2_udb remote_dbtype=lw;
proc datasets lib=mylib;

quit;
```

This SAS code, by specifying REMOTE_DBTYPE=LUW, enables the catalog call to work properly for this remote connection.

```
proc sql;
  connect to db2 (ssid=db2a server=db2_udb remote_dbtype=lw);

  select * from connection to db2
  select * from connection to db2
    (DBMS::PrimaryKeys ("", "JOSMITH", ""));
quit;
```

See Also

For more information about the other options that work in conjunction with REMOTE_DBTYPE=, see the LOCATION= LIBNAME option and the SERVER= CONNECT statement option.

REREAD_EXPOSURE= LIBNAME Option

Specifies whether the SAS/ACCESS engine functions like a random access engine for the scope of the LIBNAME statement

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: NO

Syntax

REREAD_EXPOSURE= NO | YES

Syntax Description

NO

specifies that the SAS/ACCESS engine functions as an RMOD engine, which means that your data is protected by the normal data protection that SAS provides.

YES

specifies that the SAS/ACCESS engine functions like a random access engine when rereading a row so that you cannot guarantee that the same row is returned. For example, if you read row 5 and someone else deletes it, then the next time you read

row 5, you will read a different row. You have the potential for data integrity exposures within the scope of your SAS session.

Details

CAUTION:

Using REREAD_EXPOSURE= could cause data integrity exposures △

Netezza, ODBC, and OLE DB Details: To avoid data integrity problems, it is advisable to set UPDATE_ISOLATION_LEVEL=S (serializable) if you set REREAD_EXPOSURE=YES.

Oracle Details: To avoid data integrity problems, it is advisable to set UPDATE_LOCK_TYPE= TABLE if you set REREAD_EXPOSURE=YES.

SCHEMA= LIBNAME Option

Lets you read database objects, such as tables and views, in the specified schema

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

SCHEMA=*schema-name*

Details

If this option is omitted, you connect to the default schema for your DBMS.

The values for SCHEMA= are usually case-sensitive, so use care when you specify this option.

Oracle Details: Specify a schema name to be used when referring to database objects. SAS can access another user's database objects by using a specified schema name. If PRESERVE_TAB_NAMES=NO, SAS converts the SCHEMA= value to uppercase because all values in the Oracle data dictionary are uppercase unless quoted.

Teradata Details: If you omit this option, a libref points to your default Teradata database, which often has the same name as your user name. You can use this option to point to a different database. You can use this option to view or modify a different user's DBMS tables or views if you have the required Teradata privileges. (For example, to read another user's tables, you must have the Teradata privilege SELECT for that user's tables.) The Teradata alias for SCHEMA= is DATABASE=. For more information about changing the default database, see the DATABASE statement in your Teradata documentation.

Examples

In the following LIBNAME statement example, the SCHEMA= option causes any reference in SAS to mydb.employee to be interpreted by DB2 as scott.employee.

```
libname mydb db2 SCHEMA=SCOTT;
```

To access an Oracle object in another schema, use the SCHEMA= option as in the following example. The schema name is typically a person's user name or ID.

```
libname mydblib oracle user=testuser
password=testpass path='hrdept_002' schema=john;
```

In the following example, the Oracle SCHEDULE table resides in the AIRPORTS schema, and is specified as AIRPORTS.SCHEDULE. To access this table in PROC PRINT and still use the libref (CARGO) in the SAS/ACCESS LIBNAME statement, you specify the schema in the SCHEMA= option. Then you put the *libref.table* in the procedure's DATA statement.

```
libname cargo oracle schema=airports user=testuser password=testpass
path="myorapath";
proc print data=cargo.schedule;
run;
```

In the Teradata interface example that follows, the user testuser prints the emp table, which is located in the otheruser database.

```
libname mydblib teradata user=testuser pw=testpass schema=otheruser;
proc print data=mydblib.emp;
run;
```

See Also

To apply this option to an individual data set, see the data set option "SCHEMA= Data Set Option" on page 259.

SHOW_SYNONYMS= LIBNAME Option

Specifies whether PROC DATASETS shows only tables and views for the current user (or schema if the SCHEMA= option is specified)

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Oracle

Default value: NO

Syntax

SHOW_SYNONYMS= <YES | NO>

Syntax Description

YES

specifies that PROC DATASETS shows only the synonyms that represent the tables and views for the current user (or schema if the SCHEMA= option is specified).

NO

specifies that PROC DATASETS shows only tables and views for the current user (or schema if the SCHEMA= option is specified).

Details

Instead of submitting PROC DATASETS, you can click the libref for the SAS Explorer window to get this same information.

SPOOL= LIBNAME Option

Specifies whether SAS creates a utility spool file during read transactions that read data more than once

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: YES

Syntax

SPOOL= YES | NO | DBMS

Syntax Description

YES

specifies that SAS creates a utility spool file into which it writes the rows that are read the first time. For subsequent passes through the data, the rows are read from the utility spool file rather than being re-read from the DBMS table. This guarantees that the row set is the same for every pass through the data.

NO

specifies that the required rows for all passes of the data are read from the DBMS table. No spool file is written. There is no guarantee that the row set is the same for each pass through the data.

DBMS

is valid for Oracle only. The required rows for all passes of the data are read from the DBMS table. Additional enforcements are made on the DBMS server side to ensure that the row set is the same for every pass through the data. This setting causes the SAS/ACCESS interface to Oracle to satisfy the two-pass requirement by starting a read-only transaction. SPOOL=YES and SPOOL=DBMS have comparable performance results for Oracle; however, SPOOL=DBMS does not use any disk space.

When SPOOL is set to DBMS, the CONNECTION option must be set to UNIQUE. If not, an error occurs.

Details

In some cases, SAS processes data in more than one pass through the same set of rows. Spooling is the process of writing rows that have been retrieved during the first pass of a data read to a spool file. In the second pass, rows can be reread without performing I/O to the DBMS a second time. When data must be read more than once, spooling improves performance. Spooling also guarantees that the data remains the same between passes, as most SAS/ACCESS interfaces do not support member-level locking.

MySQL Details: Do not use SPOOL=NO with the MySQL interface.

Teradata Details: SPOOL=NO requires SAS/ACCESS to issue identical SELECT statements to Teradata twice. Additionally, because the Teradata table can be modified between passes, SPOOL=NO can cause data integrity problems. Use SPOOL=NO with discretion.

SQL_FUNCTIONS= LIBNAME Option

Specifies that the functions that match those supported by SAS should be passed to the DBMS

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, HP Neoview, MySQL, Netezza, ODBC, OLE DB

Default value: NONE

Syntax

SQL_FUNCTIONS=ALL

Syntax Description

ALL

specifies that functions that match those that are supported by SAS should be passed to the DBMS.

Details

DB2 UNIX/PC, HP Neoview, Netezza, and ODBC Details: When SQL_FUNCTIONS= is set to ALL, only the functions that the DBMS drivers support are passed. Only a fraction of the functions might be available, based on your provider:

DATE

DATEPART

DATETIME

TIME

TIMEPART

TODAY

DAY
HOUR
MINUTE
MONTH
QRT
SECOND
WEEKDAY
YEAR
BYTE
COMPRESS
INDEX
LENGTH
REPEAT
SOUNDEX
SUBSTR
TRANWRD
TRIMN
MOD

OLE DB Details: When SQL_FUNCTIONS= is set to ALL, the following functions are passed to the DBMS, regardless of whether they are supported by the driver:

LOWCASE
UPCASE
ABS
ARCOS
ARSIN
ATAN
CEIL
COS
EXP
FLOOR
LOG
LOG10
SIGN
SIN
SQRT
TAN
DATE

DATEPART
DATETIME
TIME
TIMEPART
TODAY
DAY
HOUR
MINUTE
MONTH
QRT
SECOND
WEEKDAY
YEAR
BYTE
COMPRESS
INDEX
LENGTH
REPEAT
SOUNDEX
SUBSTR
TRANWRD
TRIMN
MOD

Use of this option can cause unexpected results, especially if used for NULL processing and date/time/timestamp handling. For example, the following SAS code executed without SQL_FUNCTIONS= enabled returns the SAS date 15308:

```
proc sql;  
  select distinct DATE () from x.test;  
quit;
```

However, the same code with SQL_FUNCTIONS=ALL, returns 2001-1-29, which is an ODBC date format. Care should be exercised when using this option.

STRINGDATES= LIBNAME Option

Specifies whether to read date and time values from the database as character strings or as numeric date values

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB

Default value: NO

Syntax

STRINGDATES= YES | NO

Syntax Description

YES

specifies that SAS reads date and time values as character strings.

NO

specifies that SAS reads date and time values as numeric date values.

Details

STRINGDATES=NO is used for Version 6 compatibility.

TRACE= LIBNAME Option

Specifies whether to turn on tracing information that is used in debugging

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: HP Neoview, Microsoft SQL Server, Netezza, ODBC

Default value: NO

Syntax

TRACE= YES | NO

Syntax Description

YES

specifies that tracing is turned on, and the DBMS driver manager writes each function call to the trace file that is specified by TRACEFILE=.

NO

specifies that tracing is not turned on.

Details

This option is not supported on UNIX platforms.

See Also

“TRACEFILE= LIBNAME Option” on page 150

TRACEFILE= LIBNAME Option

Specifies the filename to which the DBMS driver manager writes trace information

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: HP Neoview, Microsoft SQL Server, Netezza, ODBC

Default value: none

Syntax

TRACEFILE= *filename* | *<'>path-and-filename<'>*

Details

TRACEFILE= is used only when TRACE=YES. If you specify a filename without a path, the SAS trace file is stored with your data files. If you specify a directory, enclose the fully qualified filename in single quotation marks.

If you do not specify the TRACEFILE= option, output is directed to a default file.

This option is not supported on UNIX platforms.

See Also

“TRACE= LIBNAME Option” on page 149

UPDATE_ISOLATION_LEVEL= LIBNAME Option

Defines the degree of isolation of the current application process from other concurrently running application processes

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, Microsoft SQL Server, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

UPDATE_ISOLATION_LEVEL= *DBMS-specific-value*

Syntax Description

The values for this option are DBMS-specific. See the documentation for your SAS/ACCESS interface.

Details

The degree of isolation defines the following

- the degree to which rows that are read and updated by the current application are available to other concurrently executing applications
- the degree to which update activity of other concurrently executing application processes can affect the current application.

This option is ignored in the interfaces to DB2 under UNIX and PC Hosts and ODBC if UPDATE_LOCK_TYPE= is not set to ROW.

For DBMS-specific details, see locking in the documentation for your SAS/ACCESS interface.

See Also

To apply this option to an individual data set, see the data set option “UPDATE_ISOLATION_LEVEL= Data Set Option” on page 264.

UPDATE_LOCK_TYPE= LIBNAME Option

Specifies how data in a DBMS table is locked during an update transaction

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: DB2 UNIX/PC, DB2 z/OS, Microsoft SQL Server, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

UPDATE_LOCK_TYPE= ROW | PAGE | TABLE | NOLOCK | VIEW

Syntax Description

ROW

locks a row if any of its columns are going to be updated. (This value is valid in the DB2 under UNIX and PC hosts, Microsoft SQL Server, ODBC, OLE DB, and Oracle interfaces.)

PAGE

locks a page of data, which is a DBMS-specific number of bytes. (This value is valid in the Sybase interface.)

TABLE

locks the entire DBMS table. (This value is valid in the DB2 under z/OS, DB2 under UNIX and PC hosts, ODBC, Oracle, Microsoft SQL Server, and Teradata interfaces.)

NOLOCK

does not lock the DBMS table, page, or any rows when reading them for update. (This value is valid in the Microsoft SQL Server, ODBC, Oracle, and Sybase interfaces.)

VIEW

locks the entire DBMS view. (This value is valid in the Teradata interface.)

Details

You can set a lock for one DBMS table by using the data set option or for a group of DBMS tables by using the LIBNAME option.

For DBMS-specific details, see locking in the documentation for your SAS/ACCESS interface.

See the documentation for your SAS/ACCESS interface for additional details.

See Also

To apply this option to an individual data set, see the data set option “UPDATE_LOCK_TYPE= Data Set Option” on page 265.

UPDATE_MODE_WAIT= LIBNAME Option

Specifies during SAS/ACCESS update operations whether Teradata should wait to acquire a lock or fail the request when the DBMS resource is locked by a different user

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Teradata

Default value: none

Syntax

UPDATE_MODE_WAIT= YES|NO

Syntax Description**YES**

specifies for Teradata to wait to acquire the lock, so SAS/ACCESS waits indefinitely until it can acquire the lock.

NO

specifies for Teradata to fail the lock request if the specified DBMS resource is locked.

Details

If you specify UPDATE_MODE_WAIT=NO and if a different user holds a restrictive lock, then the executing SAS step fails. SAS/ACCESS continues processing the job by executing the next step.

A *restrictive* lock means that a different user is holding a lock that prevents you from obtaining your desired lock. Until the other user releases the restrictive lock, you cannot obtain your lock. For example, another user's table level WRITE lock prevents you from obtaining a READ lock on the table.

Use SAS/ACCESS locking options only when Teradata's standard locking is undesirable.

For more information, see the documentation for the Teradata interface.

See Also

To apply this option to an individual data set, see the data set option "UPDATE_MODE_WAIT= Data Set Option" on page 266.

UPDATE_MULT_ROWS= LIBNAME Option

Indicates whether to allow SAS to update multiple rows from a data source, such as a DBMS table

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB

Default value: NO

Syntax

UPDATE_MULT_ROWS= YES | NO

Syntax Description

YES

specifies that SAS/ACCESS processing continues if multiple rows are updated. This might produce unexpected results.

NO

specifies that SAS/ACCESS processing does not continue if multiple rows are updated.

Details

Some providers do not handle the following DBMS SQL statement well, and therefore update more than the current row with this statement:

```
UPDATE ... WHERE CURRENT OF CURSOR
```

UPDATE_MULT_ROWS= enables SAS/ACCESS to continue if multiple rows were updated.

UPDATE_SQL= LIBNAME Option

Determines the method that is used to update and delete rows in a data source

Valid in: SAS/ACCESS LIBNAME statement

DBMS support: Microsoft SQL Server, ODBC

Default value: YES (except for the Oracle drivers from Microsoft and Oracle)

Syntax

UPDATE_SQL=YES | NO

Syntax Description

YES

specifies that SAS/ACCESS uses Current-of-Cursor SQL to update or delete rows in a table.

NO

specifies that SAS/ACCESS uses the SQLSetPos() application programming interface (API) to update or delete rows in a table.

Details

This is the update/delete equivalent of the INSERT_SQL= option. The default for the Oracle drivers from Microsoft and Oracle is NO because these drivers do not support Current-Of-Cursor operations.

See Also

To apply this option to an individual data set, see the data set option “UPDATE_SQL= Data Set Option” on page 267.

UPDATEBUFF= LIBNAME Option

Specifies the number of rows that are processed in a single DBMS update or delete operation

Valid in: the SAS/ACCESS LIBNAME statement

DBMS support: Oracle

Default value: 1

Syntax

UPDATEBUFF=*positive-integer*

Syntax Description

positive-integer

is the number of rows in an operation. SAS allows the maximum that the DBMS allows.

Details

When updating with the VIEWTABLE window or the FSVIEW procedure, use UPDATEBUFF=1 to prevent the DBMS interface from trying to update multiple rows. By default, these features update only observation at a time (since by default they use record-level locking, they lock only the observation that is currently being edited).

See Also

To apply this option to an individual data set, see the data set option “UPDATEBUFF= Data Set Option” on page 268.

USE_ODBC_CL= LIBNAME Option

Indicates whether the Driver Manager uses the ODBC Cursor Library

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: Microsoft SQL Server, Netezza, ODBC

Default value: NO

Syntax

USE_ODBC_CL= YES | NO

Syntax Description

YES

specifies that the Driver Manager uses the ODBC Cursor Library. The ODBC Cursor Library supports block scrollable cursors and positioned update and delete statements.

NO

specifies that the Driver Manager uses the scrolling capabilities of the driver.

Details

For more information about the ODBC Cursor Library, see your vendor-specific documentation.

UTILCONN_TRANSIENT= LIBNAME Option

Enables utility connections to be maintained or dropped, as needed

Valid in: the SAS/ACCESS LIBNAME statement and some DBMS-specific connection options. Please refer to your DBMS for details.

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific (see the details in this section)

Syntax

UTILCONN_TRANSIENT= NO | YES

Syntax Description

NO

specifies that a utility connection is maintained for the lifetime of the libref.

YES

specifies that a utility connection is automatically dropped as soon as it is no longer in use.

Details


For engines that can lock system resources as a result of operations such DELETE or RENAME, or as a result of queries on system tables or table indices, a utility connection is used. The utility connection prevents the COMMIT statements that are issued to unlock system resources from being submitted on the same connection that is being used for table processing. Keeping the COMMIT statements off of the table processing connection alleviates the problems they can cause such as invalidating cursors and committing pending updates on the tables being processed.

Because a utility connection exists for each LIBNAME statement, the number of connection to a DBMS can get large as multiple librefs are assigned across multiple SAS sessions. Setting UTILCONN_TRANSIENT= to YES keeps these connections from existing when they are not being used, thus reducing the number of current connections to the DBMS at any given point in time.

UTILCONN_TRANSIENT= has no effect on engines that do not support utility connections.

DB2 under z/OS Details: The default is YES.

DB2 under UNIX and PC Hosts, Informix, HP Neoview, Microsoft SQL Server, Netezza, ODBC, MySQL, OLEDB, Oracle, Sybase, and Teradata Details: The default is NO.



CHAPTER

10

Data Set Options for Relational Databases

Overview of Data Set Options for Relational Databases 157

Overview of Data Set Options for Relational Databases

You can specify SAS/ACCESS data set options on a SAS data set when you access DBMS data with the SAS/ACCESS LIBNAME statement . A data set option applies only to the data set on which it is specified, and it remains in effect for the duration of the DATA step or procedure. (See “LIBNAME Options for Relational Databases” on page 78 for options that can be assigned to a *group* of relational DBMS tables or views.)

For example, SAS/ACCESS data set options can be used as follows:

```
libname myoralib oracle;
proc print myoralib.mytable(data-set-option=value)
```

You can also use SAS/ACCESS data set options on a SAS data set when you access DBMS data using access descriptors, see “Using Descriptors with the ACCESS Procedure” on page 325. For example:

```
proc print mylib.myviewd(data-set-option=value)
```

Most data set options *cannot* be used on a PROC SQL DROP (table or view) statement.

You can use the CNTLLEV=, DROP=, FIRSTOBS=, IN=, KEEP=, OBS=, RENAME=, and WHERE= SAS data set options when you access DBMS data. The REPLACE= SAS data set option is not supported by SAS/ACCESS interfaces. For information about using SAS data set options, see the *SAS Language Reference: Dictionary*.

The information in this section explains all the applicable data set option. The information includes DBMS support and the corresponding LIBNAME options, and refers you to the documentation for your SAS/ACCESS interface when appropriate. The documentation for your SAS/ACCESS interface lists the data set options that are available for your DBMS and provides their default values.

Note: Specifying data set options in PROC SQL might reduce performance, because it prevents operations from being passed to the DBMS for processing. For more information, see “Overview of Optimizing Your SQL Usage” on page 37. △

AUTHID= Data Set Option

Lets you qualify the specified table with an authorization ID, user ID, or group ID

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 z/OS

Alias: SCHEMA=

Default value: LIBNAME setting

Syntax

AUTHID=*authorization-ID*

Syntax Description

authorization-ID

is limited to eight characters.

Details

If you specify a value for the AUTHID= option, the table name is qualified as *authid.tablename* before any SQL code is passed to the DBMS. If AUTHID= is not specified, the table name is not qualified before it is passed to the DBMS, and the DBMS uses your user ID as the qualifier. If you specify AUTHID= in a SAS/SHARE LIBNAME statement, the ID of the active server is the default ID.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “AUTHID= LIBNAME Option” on page 79.

AUTOCOMMIT= Data Set Option

Specifies whether to enable the DBMS autocommit capability

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: MySQL, Sybase

Default value: LIBNAME setting

Syntax

AUTOCOMMIT=YES | NO

Syntax Description

YES

specifies that all updates, inserts, and deletes are committed immediately after they are executed and no rollback is possible.

NO

specifies that SAS performs the commit after processing the number of row that are specified by using DBCOMMIT=, or the default number of rows if DBCOMMIT= is not specified.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “AUTOCOMMIT= LIBNAME Option” on page 80.

BL_BADFILE= Data Set Option

Identifies a file that contains records that were rejected during a bulk load

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: Oracle

Default value: creates a data file in the current directory or with the default file specifications

Syntax

BL_BADFILE= *path-and-filename*

Syntax Description

path-and-filename

is an SQL*Loader file to which rejected rows of data are written. On most platforms, the default filename takes the form BL_<table>_<unique-ID>.bad, where

table is the table name

unique-ID is a number that is used to prevent collisions in the event of two or more simultaneous bulk loads of a particular table. The SAS/ACCESS engine generates the number.

Details

If you do not specify this option and a BAD file does not exist, a file is created in the current directory (or with the default file specifications). If you do not specify this option and a BAD file already exists, the Oracle bulk loader reuses the file, replacing the contents with rejected rows from the new load.

Records can be rejected by either the SQL*Loader or by Oracle. For example, the SQL*Loader can reject a record that contains invalid input, and Oracle can reject a

record because it does not contain a unique key. If no records are rejected, the BAD file is not created.

On most operating systems, the BAD file is created in the same format as the DATA file, so the rejected records can be loaded after corrections have been made.

Operating Environment Information: On z/OS operating systems, the BAD file is created with default DCB attributes. For information about how to overcome this, see the section about SQL*Loader file attributes in the SQL*Loader chapter in the Oracle user's guide for z/OS. Δ

To specify this option, you must first specify YES for "BULKLOAD= Data Set Option" on page 194.

BL_CODEPAGE= Data Set Option

Identifies the codepage that the DBMS engine uses to convert SAS character data to the current database codepage during a bulk load

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC

Default value: the current window's codepage ID.

Syntax

BL_CODEPAGE= *numeric-codepage-ID*

Syntax Description

numeric-codepage-ID

is a numeric value that represents a character set that is used to interpret multibyte character data and determine the character values.

Details

The value for this option must never be 0. If you do not wish any codepage conversions to take place, use the BL_OPTIONS= option to specify 'FORCEIN'. Codepage conversions only occur for DB2 character data types.

To specify this option, you must first specify YES for "BULKLOAD= Data Set Option" on page 194.

BL_CONTROL= Data Set Option

Identifies a file containing SQLLDR control statements that describe the data to be included in a bulk load

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: creates a control file in the current directory or with the default file specifications

Syntax

BL_CONTROL= *path-and-control-filename*

Syntax Description

path-and-control-filename

is an SQL*Loader file to which SQLLDR control statements are written. On most platforms, the default filename takes the form BL_<table>_<unique-ID>.ctl, where *table* is the table name

unique-ID is a number used to prevent collisions in the event of two or more simultaneous bulk loads of a particular table. The SAS/ACCESS engine generates the number.

Details

If you do not specify this option and a control file does not exist, a file is created in the current directory (or with the default file specifications). If you do not specify this option and a control file already exists, the interface to Oracle reuses the file, replacing the contents with the new control statements.

The SAS/ACCESS interface for Oracle creates the control file by using information from the input data and SAS/ACCESS options. The file contains Data Definition Language (DDL) definitions that specify the location of the data and how the data corresponds to the database table. It is used to specify exactly how the loader should interpret the data that you are loading from the DATA file (.DAT file).

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_COPY_LOCATION= Data Set Option

Specifies the directory to which DB2 saves a copy of the loaded data. This option is only valid when used in conjunction with BL_RECOVERABLE=YES.

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC

Default value: none

Syntax

BL_COPY_LOCATION=*pathname*

BL_DATAFILE= Data Set Option

Identifies the file that contains the data to be loaded or appended into a DBMS table during a bulk load

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, Netezza, Oracle

Default value: creates a data file in the current directory or with the default file specifications

Syntax

BL_DATAFILE= *path-and-data-filename*

Syntax Description

path-and-data-filename

is a file that contains the rows of data to be loaded. On most platforms, the default filename takes the form BL_<table>_<unique-ID>.ext, where:

table is the table name

unique-ID is a number used to prevent collisions in the event of two or more simultaneous bulk loads of a particular table. The SAS/ACCESS engine generates the number.

ext is the file extension (.DAT or .IXF) for the data file.

Details

To specify this option, you must first specify YES for the BULKLOAD= data set option “BULKLOAD= Data Set Option” on page 194.

The SAS/ACCESS engine creates this data file from the input SAS data set before calling the bulk loader. The data file contains SAS data that is ready to load into the DBMS.

By default, the data file is deleted after the load is completed. To override this behavior, specify `BL_DELETE_DATAFILE=NO`.

Netezza Details: You can use this option only when `BL_USE_PIPE="BL_USE_PIPE=Data Set Option"` on page 191NO.

Oracle Details: If you do not specify this option and a data file does not exist, the file is created in the current directory or with the default file specifications. If you do not specify this option and a data file already exists, SAS/ACCESS reuses the file, replacing the contents with the new data. SAS/ACCESS Interface to Oracle on z/OS is the exception: The data file is never reused because the interface causes bulk load to fail instead of reusing a data file.

BL_DB2CURSOR= Data Set Option

Specifies a string that contains a valid DB2 SELECT statement that points to either local or remote objects (tables or views).

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: none

Syntax

BL_DB2CURSOR= 'SELECT * from *filename*'

Details

This option is used only if you specify it. You can use it to select load DB2 tables directly from other DB2 and non-DB2 objects. However, before you can select data from a remote location, your database administrator must first populate the communication database with the appropriate entries.

BL_DB2DEVT_PERM= Data Set Option

Specifies the unit address or generic device type that is used for permanent data sets that the LOAD utility creates—also SYSIN, SYSREC, and SYSPRINT when SAS allocates them

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: SYSDA

Syntax

BL_DB2DEVT_PERM= *unit-specification*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2DEVT_TEMP= Data Set Option

Specifies the unit address or generic device type that is used for temporary data sets that the LOAD utility creates (PNCH, COPY1, COPY2, RCPY1, RCPY2, WORK1, WORK2)

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: SYSDA

Syntax

BL_DB2DEVT_TEMP= *unit-specification*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2DISC= Data Set Option

Specifies the SYSDISC data set name for the LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: a generated data set name

Syntax

BL_DB2DISC= *data-set-name*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2ERR= Data Set Option

Specifies the SYSERR data set name for the LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: a generated data set name

Syntax

BL_DB2ERR= *data-set-name*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2IN= Data Set Option

Specifies the SYSIN data set name for the LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: a generated data set name

Syntax

BL_DB2IN= *data-set-name*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2LDCT1= Data Set Option

Specifies a string in the LOAD utility control statement, between LOAD DATA and INTO TABLE

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: nothing

Syntax

BL_DB2LDCT1=*'string'*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2LDCT2= Data Set Option

Specifies a string in the LOAD utility control statement, between INTO TABLE *table-name* and (*field-specification*)

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: nothing

Syntax

BL_DB2LDCT2= 'string'

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2LDEXT= Data Set Option

Specifies the mode of execution for the DB2 LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: GENRUN

Syntax

BL_DB2LDEXT= GENRUN | GENONLY | USERUN

Syntax Description

GENRUN

generates the control file (SYSIN) and the data file (SYSREC) and invokes the utility with them.

GENONLY

generates the control file (SYSIN) and the data file (SYSREC) but does not invoke the utility. Use this method when you need to edit the control file or verify the generated control statement or data before you run the utility.

USERUN

uses existing control and data files, and runs the utility with them. The existing files can be from a previous run or from previously run batch utility jobs. Use this execution method when you are restarting an invocation of the utility.

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2MAP= Data Set Option

Specifies the SYSMAP data set name for the LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: a generated data set name

Syntax

BL_DB2MAP= *data-set-name*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2PRINT= Data Set Option

Specifies the SYSPRINT data set name for the LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: a generated data set name

Syntax

BL_DB2PRINT= *data-set-name*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2PRNLOG= Data Set Option

Determines whether the SYSPRINT output is written to the SAS log

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: YES

Syntax

BL_DB2PRNLOG= YES | NO

Syntax Description

YES

specifies that the SYSPRINT output is written to the SAS log.

NO

specifies that the SYSPRINT output is not written to the SAS log.

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2REC= Data Set Option

Specifies the SYSREC data set name for the LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: a generated data set name

Syntax

BL_DB2REC= *data-set-name*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2RECSP= Data Set Option

Determines the number of cylinders to specify as the primary allocation for the SYSREC data set when it is created

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: 10

Syntax

BL_DB2RECSP= *primary-allocation*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2RSTRT= Data Set Option

Tells the LOAD utility whether the current load is a restart and, for a restart, indicates where to begin

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: NO

Syntax

BL_DB2RSTRT= NO | CURRENT | PHASE

Syntax Description

NO

specifies a new invocation of the LOAD utility, not a restart.

CURRENT

specifies to restart at the last commit point.

PHASE

specifies to restart at the beginning of the current phase.

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2SPC_PERM= Data Set Option

Determines the number of cylinders to specify as the primary allocation for the permanent data sets that are created by the LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: 10

Syntax

BL_DB2SPC_PERM= *primary-allocation*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2SPC_TEMP= Data Set Option

Determines the number of cylinders to specify as the primary allocation for the temporary data sets that are created by the LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: 10

Syntax

BL_DB2SPC_TEMP= *primary-allocation*

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2TBLXST= Data Set Option

Indicates whether the LOAD utility runs against an existing table

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: NO

Syntax

BL_DB2TBLXST= YES | NO

Syntax Description

YES

specifies that the LOAD utility runs against an existing table. This is *not* a replacement operation. See details below.

NO

specifies that the LOAD utility does not run against an existing table.

Details

For details about this option, see the bulk load topic in the documentation for the interface to DB2 under z/OS.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DB2UTID= Data Set Option

Specifies a unique identifier for a given run of the DB2 LOAD utility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS Support: DB2 z/OS

Default value: user ID and second level DSN qualifier

Syntax

BL_DB2UTID= *utility-ID*

Details

For details about this option, see

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DELETE_DATAFILE= Data Set Option

Deletes the data file that is created for the DBMS bulk-load facility

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, Netezza, Oracle

Alias: BL_DELETE_FILES= (Oracle)

Default value: YES

Syntax

BL_DELETE_DATAFILE=YES | NO

Syntax Description

YES

deletes the data file that the SAS/ACCESS engine creates for the DBMS bulk-load facility.

NO

saves the data file from deletion.

Details

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

DB2 UNIX/PC Details: Setting BL_DELETE_DATAFILE= to YES deletes the temporary data file that is created after the load is completed. Only the data file is deleted.

Netezza Details: You can use this option only when BL_USE_PIPE=“BL_USE_PIPE= Data Set Option” on page 191NO.

BL_DELIMITER= Data Set Option

Specifies override of the default delimiter character for separating columns of data during data transfer or retrieval during bulk load or bulk unload

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Netezza

Default value: | (the pipe symbol)

Syntax

BL_DELIMITER='<any single character>'

Details

To specify this option, you must first specify YES for the BULKLOAD= data set option “BULKLOAD= Data Set Option” on page 194.

Here is when you might want to use this option:

- to override the default delimiter character that SAS/ACCESS Interface to Netezza uses to separate columns of data that it transfers to or retrieves from Netezza during bulk load or bulk unload
- if your character data contains the pipe symbol (|) to avoid any problems while parsing the data stream

Example

Data in this example contains the pipe symbol:

```
data work.testdel;
col1='my|data';col2=12;
run;
```

This example shows how you can override this default when BULKLOAD=YES:

```
/* Using a comma to delimit data */
proc append base=netlib.mydat(BULKLOAD=YES BL_DELIMITER=',')
data=work.testdel;
run;
```

See Also

BL_DATAFILE= data set option“BL_DATAFILE= Data Set Option” on page 162

BL_DELETE_DATAFILE= data set option“BL_DELETE_DATAFILE= Data Set Option” on page 173

BL_OPTIONS= data set option“BL_OPTIONS= Data Set Option” on page 183

BL_USE_PIPE= data set option“BL_USE_PIPE= Data Set Option” on page 191

BULKUNLOAD= data set option“BULKUNLOAD= Data Set Option” on page 195

BL_DIRECT_PATH= Data Set Option

Sets the Oracle SQL*Loader **DIRECT** option

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: YES

Syntax

BL_DIRECT_PATH= YES | NO

Syntax Description

YES

sets the Oracle SQL*Loader option **DIRECT** to **TRUE**, enabling the SQL*Loader to use Direct Path Load to insert rows into a table.

NO

sets the Oracle SQL*Loader option **DIRECT** to **FALSE**, enabling the SQL*Loader to use Conventional Path Load to insert rows into a table.

Details

The Conventional Path Load reads in multiple data records and places them in a binary array. When the array is full, it is passed to Oracle for insertion, and Oracle uses the SQL interface with the array option.

The Direct Path Load creates data blocks that are already in the Oracle database block format. The blocks are then written directly into the database. This method is significantly faster, but there are restrictions. For more information about the SQL*Loader Direct and Conventional Path loads, see your Oracle utilities documentation for SQL*Loader.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_DISCARDFILE= Data Set Option

Identifies the file that contains the records that were filtered out of a bulk load because they did not match the criteria specified in the CONTROL file

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: creates a file in the current directory or with the default file specifications

Syntax

BL_DISCARDFILE= *path-and-discard-filename*

Syntax Description

path-and-discard-filename

is an SQL*Loader discard file containing rows that did not meet the specified criteria. On most platforms, the default filename takes the form

BL_<table>_<unique-ID>.dsc, where

table is the table name

unique-ID is a number used to prevent collisions in the event of two or more simultaneous bulk loads of a particular table. The SAS/ACCESS engine generates the number.

Details

SQL*Loader creates the file of discarded rows only if there are discarded rows and if a discard file is requested. If you do not specify this option and a discard file does not exist, a discard file is created in the current directory or with the default file specifications. If you do not specify this option and a discard file already exists, the Oracle bulk loader reuses the existing file and replaces the contents with discarded rows from the new load.

On most operating systems, the discard file has the same format as the data file, so the discarded records can be loaded after corrections are made.

Operating Environment Information: On z/OS operating systems, the discard file is created with default DCB attributes. For information about how to overcome this, see the section about SQL*Loader file attributes in the SQL*Loader chapter in the Oracle user's guide for z/OS. △

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

See Also

See “BL_BADFILE= Data Set Option” on page 159 to set the name and location of the file that contains *rejected* rows.

BL_INDEX_OPTIONS= Data Set Option

Lets you specify SQL*Loader Index options with bulk loading

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Alias: SQLLDR_INDEX_OPTION=

Default value: none

Syntax

BL_INDEX_OPTIONS=*any valid SQL*Loader Index option***segment-name**

Syntax Description

*any valid SQL*Loader Index option*

The value specified for this option must be a valid SQL*Loader index option, such as one of the following. Otherwise, an error occurs.

SINGLEROW Use this option when loading either a direct path with APPEND on systems with limited memory or a small number of records into a large table. It inserts each index entry directly into the index, one record at a time.

By default, SQL*Loader does not use this option to append records to a table.

SORTED INDEXES This clause applies when you are loading a direct path. It tells the SQL*Loader that the incoming data has already been sorted on the specified indexes, allowing SQL*Loader to optimize performance. It allows the SQL*Loader to optimize index creation by eliminating the sort phase for this data when using the direct-path load method.

Details

You can now pass in SQL*Loader index options when bulk loading. For details about these options, see the Oracle utilities documentation.

Example

The following example shows how you can use this option.

```
proc sql;
connect to oracle ( user=scott pw=tiger path=alien);
execute ( drop table blidxopts) by oracle;
execute ( create table blidxopts ( empno number, empname varchar2(20))) by
oracle;
execute ( drop index blidxopts_idx) by oracle;
execute ( create index blidxopts_idx on blidxopts ( empno ) ) by oracle;

quit;

libname x oracle user=scott pw=tiger path=alien;

data new;
empno=1; empname='one';
output;
empno=2; empname='two';
output;
run;

proc append base= x.blidxopts( bulkload=yes bl_index_options='sorted indexes
( blidxopts_idx)' ) data= new;
run;
```

BL_INDEXING_MODE= Data Set Option

Used to indicate which scheme the DB2 load utility should use with respect to index maintenance

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC

Default value: AUTOSELECT

Syntax

BL_INDEXING_MODE= AUTOSELECT | REBUILD | INCREMENTAL |
DEFERRED

Syntax Description

AUTOSELECT

The load utility automatically decides between REBUILD or INCREMENTAL mode. This is the default.

REBUILD

All indexes are rebuilt.

INCREMENTAL

Indexes are extended with new data

DEFERRED

The load utility does not attempt index creation if this mode is specified. Indexes are marked as needing a refresh.

See Also

For more information about the usage of these values, see the DB2 *Data Movement Utilities Guide and Reference*.

BL_KEEPIENTITY= Data Set Option

Determines whether the identity column that is created during a bulk load is populated with values generated by Microsoft SQL Server or with values provided by the user

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: OLE DB

Default value: LIBNAME setting

Syntax

BL_KEEPIENTITY=YES | NO

Syntax Description**YES**

specifies that the user must provide values for the identity column.

NO

specifies that the Microsoft SQL Server generates values for an identity column in the table.

Details

This option is only valid when you use the Microsoft SQL Server provider.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “BL_KEEPIENTITY= LIBNAME Option” on page 81.

BL_KEEPNULLS= Data Set Option

Indicates how NULL values in Microsoft SQL Server columns that accept NULL are handled during a bulk load

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: OLE DB

Default value: LIBNAME setting

Syntax

BL_KEEPNULLS= YES | NO

Syntax Description

YES

preserves NULL values inserted by the OLE DB interface.

NO

replaces NULL values that are inserted by the OLE DB interface with a default value (as specified in the DEFAULT constraint).

Details

This option affects only values in Microsoft SQL Server columns that accept NULL and that have a DEFAULT constraint.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “BL_KEEPNULLS= LIBNAME Option” on page 82.

BL_LOAD_METHOD= Data Set Option

Specifies the method by which data is loaded into an Oracle table during bulk loading

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: INSERT when loading an empty table; APPEND when loading a table that contains data

Syntax

BL_LOAD_METHOD= INSERT | APPEND | REPLACE | TRUNCATE

Syntax Description

INSERT

requires the DBMS table to be empty before loading.

APPEND

appends rows to an existing DBMS table.

REPLACE

deletes all rows in the existing DBMS table and then loads new rows from the data file.

TRUNCATE

uses the *SQL truncate* command to achieve the best possible performance. The DBMS table's referential integrity constraints must first be disabled.

Details

The REPLACE and TRUNCATE values apply only when you are loading data into a table that already contains data. In this case, you can use REPLACE and TRUNCATE to override the default value of APPEND. See your Oracle utilities documentation for information about using the TRUNCATE and REPLACE load methods.

To specify this option, you must first specify YES for "BULKLOAD= Data Set Option" on page 194.

BL_LOAD_REPLACE= Data Set Option

Specifies whether DB2 appends or replaces rows during bulk loading

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC

Default value: NO

Syntax

BL_LOAD_REPLACE= NO | YES

Syntax Description

NO

the CLI LOAD interface appends new rows of data to the DB2 table.

YES

the CLI LOAD interface replaces the existing data in the table.

BL_LOG= Data Set Option

Identifies a log file that contains information such as statistics and error information for a bulk load

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, Oracle, Teradata

Default value: DBMS-specific

Syntax

BL_LOG= *path-and-log-filename*

Syntax Description

path-and-log-filename

is a file to which information about the loading process is written.

Details

When the DBMS bulk-load facility is invoked, it creates a log file. The contents of the log file are DBMS-specific. The BL_ prefix distinguishes this log file from the one created by the SAS log. If BL_LOG= is specified with the same path and filename as an existing log, the new log replaces the existing log.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194. See the documentation for your SAS/ACCESS interface for additional details.

Oracle Details: When the SQL*Loader is invoked, it creates a log file. This file contains a detailed summary of the load, including a description of any errors. If SQL*Loader cannot create a log file, execution of the bulk load terminates. If a log file does not exist, the default action is to create a log file in the current directory or with the default file specifications. If a log file already exists, the Oracle bulk loader reuses the file, replacing the contents with information from the new load. On most platforms, the default filename takes the form BL_<table>_<unique-ID>.log, where

table is the table name

unique-ID is a number used to prevent collisions in the event of two or more simultaneous bulk loads of a particular table. The SAS/ACCESS engine generates the number.

DB2 UNIX/PC Details: If BL_LOG= is not specified, the log file is deleted automatically after a successful operation. See for more information.

Teradata Details: See the bulk load topic in the documentation for the Teradata interface for more information.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “BL_LOG= LIBNAME Option” on page 82.

BL_METHOD= Data Set Option

Specifies which bulk loading method to use for DB2

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC

Default value: None

Syntax

BL_METHOD=CLILoad

Syntax Description

CLILoad

activates the CLI LOAD interface to the LOAD utility. BULKLOAD=YES must also be specified in order to use the CLI LOAD interface.

BL_OPTIONS= Data Set Option

Passes options to the DBMS bulk-load facility, affecting how it loads and processes data

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, Netezza, OLE DB, Oracle

Default value: DBMS-specific

Syntax

BL_OPTIONS='<option...,option>'

Syntax Description

option

specifies an option from the available options that are specific to each SAS/ACCESS interface. See details below.

Details

To specify this option, you must first specify YES for the BULKLOAD= data set option “BULKLOAD= Data Set Option” on page 194.

You can use BL_OPTIONS= to pass options to the DBMS bulk-load facility when it is invoked, thereby affecting how data is loaded and processed. You must separate

multiple options with commas and enclose the entire string of options in single quotation marks.

DB2 under UNIX and PC Hosts Details: This option passes DB2 file-type modifiers to DB2 LOAD or IMPORT commands to affect how data is loaded and processed. Not all DB2 file type modifiers are appropriate for all situations. You can specify one or more DB2 file type modifiers with .IXF files. For a list of file type modifiers, see the description of the LOAD and IMPORT utilities in the *IBM® DB2 Universal Database Data Movement Utilities Guide and Reference*.

Netezza Details: Any text that you enter for this option is appended to the USING clause of the CREATE EXTERNAL TABLE statement—namely, any external_table_options in the *Netezza Database Users's Guide*.

OLE DB Details: By default, no options are specified. This option is valid only when you are using the Microsoft SQL Server provider. This option takes the same values as the *-h* HINT option of the Microsoft BCP utility. For example, the ORDER= option sets the sort order of data in the data file; you can use it to improve performance if the file is sorted according to the clustered index on the table. See the Microsoft SQL Server documentation for a complete list of supported bulk copy options.

Oracle Details: You can use this option to specify the SQL*Loader options ERRORS= and LOAD=. The ERRORS= option specifies the number of insert errors that terminates the load. The default value of ERRORS=1000000 overrides the default value for the Oracle SQL*Loader ERRORS= option, which is 50. LOAD= specifies the maximum number of logical records to load. If the LOAD= option is not specified, all rows are loaded. See your Oracle utilities documentation for a complete list of SQL*Loader options that you can specify in BL_OPTIONS=.

Example

In the following Oracle example, BL_OPTIONS= specifies the number of errors that are permitted during a load of 2,000 rows of data. Notice that the entire listing of options is enclosed in quotation marks.

```
bl_options='ERRORS=999,LOAD=2000'
```

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “BL_OPTIONS= LIBNAME Option” on page 83.

BL_PARFILE= Data Set Option

Creates a file that contains the SQL*Loader command line options

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: none

Syntax

BL_PARFILE=<parse-file>

Syntax Description

parse-file

the name you give the file that contains the SQL*Loader command line options. The name can also specify the path. If no path is specified, the file is created in the current directory.

Details

This option prompts the SQL*Loader to use the PARFILE= option. You can use this SQL*Loader option to specify SQL*Loader command line options in a file instead of as command line options. For example the SQL*Loader can be invoked by specifying user ID and control options as follows:

```
sqlldr userid=scott/tiger control=example.ctl
```

However, you can invoke it using the PARFILE = option as follows:

```
sqlldr parfile=example.par
```

Example.par now contains the USERID= and CONTROL= options. One of the biggest advantages of using the BL_PARFILE= option is security, since the user ID and password are stored in a separate file.

The permissions on the file default to the operating system defaults. It is advisable to create this file in a protected directory so that the contents of the file are not accessible to unauthorized users.

You can display the contents of the parse file in the SAS log by using the option **SASTRACE=",,,d"**. However, the password is blocked out and replaced with **xxxx**.

Note: The parse file is deleted at the end of SQL*Loader processing. Δ

Example

The following example demonstrates how SQL*Loader invocation is different when the BL_PARFILE= option is specified.

```
libname x oracle user=scott pw=tiger;
/* SQL*Loader is invoked as follows without BL_PARFILE= */
sqlldr userid=scott/tiger@oraclev9 control=bl_bltst_0.ctl log=bl_bltst_0.log
bad=bl_bltst_0.bad discard=bl_bltst_0.dsc */

data x.bltst ( bulkload=yes);
c1=1;
run;
/* Note in the DATA step below, which uses BL_PARFILE=, how SQL*Loader is invoked */
sqlldr parfile=test.par
/* In this case all the options are written to the test.par file. */
data x.bltst2 ( bulkload=yes bl_parfile='test.par');
c1=1;
run;
```

BL_PRESERVE_BLANKS= Data Set Option

Determines how the SQL*Loader handles requests to insert blank spaces into CHAR/VARCHAR2 columns with the NOT NULL constraint

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: NO

Syntax

BL_PRESERVE_BLANKS=YES | NO

Syntax Description

YES

specifies that blank values are inserted as blank spaces.

CAUTION:

When this option is set to YES, *any trailing blank spaces are also inserted*. For this reason, use this option with caution. It is recommended that you only set this option to YES for CHAR columns. It is not recommended that you set this option to YES for VARCHAR2 columns, because trailing blank spaces are significant in VARCHAR2 columns. Δ

NO

specifies that blank values are inserted as NULL values.

Details

Operating Environment Information: This option is not supported on z/OS. Δ

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_RECOVERABLE= Data Set Option

Determines whether the LOAD process is recoverable

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, Oracle

Default value: DBMS-specific (see details in this section)

Syntax

BL_RECOVERABLE=YES | NO

Syntax Description

YES

specifies that the LOAD process is recoverable. For DB2, YES also specifies that the copy location for the data should be specified by **BL_COPY_LOCATION=**.

NO

specifies that the LOAD process is not recoverable. For Oracle, NO adds the **UNRECOVERABLE** keyword before the LOAD keyword in the control file.

Details for Oracle

DB2 for UNIX and PC Hosts Details: The default value is NO.

Oracle Details: The default value is YES. Specify NO to improve direct load performance.

CAUTION:

Be aware that an unrecoverable load does not log loaded data into the redo log file and so Media recovery is disabled for the loaded table. For more information about the implications of using the UNRECOVERABLE parameter in Oracle, see your Oracle utilities manual. △

Example

This Oracle example shows how to use **BL_RECOVERABLE=** to specify that the load is unrecoverable.

```
data x.recover_no (bulkload=yes bl_recoverable=no); c1=1; run;
```

BL_REMOTE_FILE= Data Set Option

Specifies the base filename and location of DB2 LOAD temporary files

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC

Default value: none

Syntax

BL_REMOTE_FILE=*pathname-and-base-filename*

Syntax Description

pathname-and-base-filename

is the full pathname and base filename to which DB2 appends extensions (such as .log, .msg, and .dat files) to create temporary files during load operations. By default, the base filename takes the form BL_<table>_<unique-ID>, where

table is the table name

unique-ID is a number used to prevent collisions in the event of two or more simultaneous bulk loads of a particular table. The SAS/ACCESS engine generates the number.

Details

When you specify this option, the DB2 LOAD command is used (instead of the IMPORT command). See the bulk load topic in the SAS/ACCESS documentation for DB2 under UNIX and PC hosts for more information about these commands.

For *pathname*, specify a location on a DB2 server that is accessed exclusively by a single DB2 server instance, and for which the instance owner has read and write permissions. Make sure that each LOAD command is associated with a unique *pathname-and-base-filename* value.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

Note: Do NOT use BL_REMOTE_FILE= unless you have SAS Release 6.1 or later for both the DB2 client and server. Using the LOAD facility with a DB2 client or server before Release 6.1 might cause the tablespace to become unusable in the event of a load error. This might affect tables other than the table being loaded. △

BL_SERVER_DATAFILE= Data Set Option

Specifies the name and location of the data file as seen by the DB2 server instance

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC

Default value: same as BL_DATAFILE

Syntax

BL_SERVER_DATAFILE=*path-and-data-filename*

Syntax Description

pathname-and-data-filename

is the fully-qualified pathname and filename of the data file to be loaded, as seen by the DB2 server instance. By default, the base filename takes the form

BL_<table>_<unique-ID>, where

table is the table name

unique-ID is a number used to prevent collisions in the event of two or more simultaneous bulk loads of a particular table. The SAS/ACCESS engine generates the number.

Details

If the path to the data file from the DB2 server instance is different from the path to the data file from the client, you must use BL_SERVER_DATAFILE= to specify the path from the DB2 server.

By enabling the DB2 server instance to directly access the data file specified by BL_DATAFILE=, this option facilitates use of the DB2 LOAD command. See the bulk load topic in the SAS/ACCESS documentation for DB2 under UNIX and PC hosts for more information about the LOAD command.

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194 and specify a value for “BL_REMOTE_FILE= Data Set Option” on page 187.

See Also

To specify the path from the client, see “BL_DATAFILE= Data Set Option” on page 162.

BL_SQLLDR_PATH= Data Set Option

Specifies the location of the SQLLDR executable file

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: SQLLDR

Syntax

BL_SQLLDR_PATH=*pathname*

Syntax Description

pathname

is the full pathname to the SQLLDR executable file so that the SAS/ACCESS interface for Oracle can invoke SQL*Loader.

Details

Normally there is no need to specify this option because the environment is set up to find the Oracle SQL*Loader automatically.

Operating Environment Information: This option is ignored on z/OS. Δ

To specify this option, you must first specify YES for “BULKLOAD= Data Set Option” on page 194.

BL_SUPPRESS_NULLIF= Data Set Option

Indicates whether to suppress the NULLIF clause for the specified columns when a table is created in order to increase performance

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: NO

Syntax

BL_SUPPRESS_NULLIF=< ALL =YES | NO > | (<column-name-1=YES | NO > <column-name-n=YES | NO >...)

Syntax Description

YES

column-name-1=YES indicates that the NULLIF clause should be suppressed for the specified column in the table.

NO

column-name-1=NO indicates that the NULLIF clause should not be suppressed for the specified column in the table.

ALL

specifies that the YES or NO applies to all columns in the table.

Details

If you specify more than one column name, the names must be separated with spaces.

The BL_SUPPRESS_NULLIF= option processes values from left to right. If you specify a column name twice or use the ALL value, the last value overrides the first value that you specified for the column.

Example

This example uses the BL_SUPPRESS_NULLIF= option in the DATA step to suppress the NULLIF clause for columns C1 and C5 in the table.

```
data x.suppressnullif2_yes (bulkload=yes BL_SUPPRESS_NULLIF=(c1=yes c5=yes));
run;
```

The next example uses the BL_SUPPRESS_NULLIF= option in the DATA step to suppress the NULLIF clause for all columns in the table.

```
libname x oracle user=dbitest pw=tiger path=lupin_o9010;

%let num=1000000; /* 1 million rows */

data x.testlmm ( bulkload=yes
                 BL_SUPPRESS_NULLIF=( _all_ =yes )
                 rename=(year=yearx) );
  set x.biglmil (obs= &num ) ;
run;
```

BL_USE_PIPE= Data Set Option

Specifies a named pipe for data transfer

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Netezza

Default value: YES

Syntax

BL_USE_PIPE=YES | NO

Syntax Description

YES

specifies that a named pipe is used to transfer data between SAS/ACCESS Interface to Netezza and the Netezza client interface when BULKLOAD=YES or BULKUNLOAD=YES.

NO

specifies that a flat file is used to transfer data when BULKLOAD=YES or BULKUNLOAD=YES. You can specify the file in BL_DATAFILE.

Details

To specify this option, you must first specify YES for the BULKLOAD= data set option “BULKLOAD= Data Set Option” on page 194.

By default, SAS/ACCESS Interface to Netezza uses a named pipe interface to transfer large amounts of data between SAS and Netezza when using bulk load or bulk unload. If you prefer to use a flat data file that you can save for later use or examination, specify BL_USE_PIPE=NO.

BL_WARNING_COUNT= Data Set Option

Specifies the maximum number of row warnings to allow before the load fails

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC

Default value: 2147483646

Syntax

BL_WARNING_COUNT=*warning-count*

Details

Use this option to limit the maximum number of rows that generate warnings. Consult the log file for information about why the rows generated warnings.

To specify this option, you must first specify a value for “BL_REMOTE_FILE= Data Set Option” on page 187 and specify YES for “BULKLOAD= Data Set Option” on page 194.

BUFFERS= Data Set Option

Specifies the number of shared memory buffers to be used for transferring data from SAS to Teradata.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: 2

Syntax

BUFFERS=*number-of-shared-memory-buffers*

Syntax Description

number-of-shared-memory-buffers

a numeric value between 1 and 8 that specifies the number of buffers used for transferring data from SAS to Teradata.

Details

BUFFERS= specifies the number of data buffers to use for transferring data from SAS to Teradata. When using MULTILoad=, data is transferred from SAS to Teradata using shared memory segments. The default shared memory buffer size is 64K. The default number of shared memory buffers used for the transfer is 2.

Use BUFFERS= to vary the number of buffers for data transfer from 1 to 8. Use the MBUFFSIZE=data set option to vary the size of the shared memory buffers from the size of each data row up to 1MB.

See Also

For more information about specifying the size of shared memory buffers, see MBUFFSIZE=.

BULK_BUFFER= Data Set Option

Specifies the number of bulk rows that the SAS/ACCESS engine can buffer for output

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Sybase

Default value: 100

Syntax

BULK_BUFFER=*numeric-value*

Syntax Description

numeric-value

is the maximum number of rows that are allowed. This value depends on the amount of memory that is available to your system.

Details

This option improves performance by specifying the number of rows that can be held in memory for efficient retrieval from the DBMS. A higher number signifies that more rows can be held in memory and accessed quickly during output operations.

BULKLOAD= Data Set Option

Loads rows of data as one unit

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, Netezza, ODBC, Oracle, Sybase, Teradata

Alias: BL_DB2LDUTIL= (DB2 z/OS)

Default value: NO

Syntax

BULKLOAD=YES | NO

Syntax Description

YES

calls a DBMS-specific bulk-load facility in order to insert or append rows to a DBMS table.

NO

uses the dynamic SAS/ACCESS engine to insert or append data to a DBMS table.

Details

Using BULKLOAD=YES is the fastest way to insert rows into a DBMS table.

See SAS/ACCESS documentation for your DBMS interface for details.

DB2 UNIX/PC Details: See DB2 UNIX/PC bulk loading.

DB2 z/OS Details: See DB2 z/OS bulk loading.

Netezza Details: See Netezza bulk loading

ODBC Details: See ODBC bulk loading.

Oracle Details: See Oracle bulk loading.

Sybase Details: When BULKLOAD=NO, insertions are processed and rolled back as expected according to DBCOMMIT= and ERRLIMIT= values. If the ERRLIMIT= value is encountered, all uncommitted rows are rolled back. The commit intervals are determined by the DBCOMMIT= data set option.

When BULKLOAD=YES, the first error encountered causes the remaining rows (including the erroneous row) in the buffer to be rejected. No other errors within the same buffer are detected, even if the ERRLIMIT= value is greater than one. In addition, all rows before the error are committed, even if DBCOMMIT= is set larger than the number of the erroneous row.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option BULKLOAD= LIBNAME option “BULKLOAD= LIBNAME Option” on page 84.

BULKUNLOAD= Data Set Option

Rapidly retrieves (fetches) large number of rows from a data set

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Netezza

Default value: NO

Syntax

BULKUNLOAD=YES | NO

Syntax Description

YES

calls the Netezza Remote External Table interface to retrieve data from the Netezza Performance Server.

NO

uses standard Netezza result sets to retrieve data from the DBMS.

Details

Using BULKUNLOAD=YES is the fastest way to retrieve large numbers of rows from a Netezza table. See Netezza bulk unloading for more information.

See Also

BL_DATAFILE= data set option “BL_DATAFILE= Data Set Option” on page 162

BL_DELETE_DATAFILE= data set option “BL_DELETE_DATAFILE= Data Set Option” on page 173

BL_DELIMITER= data set option “BL_DELIMITER= Data Set Option” on page 174

BL_USE_PIPE= data set option “BL_USE_PIPE= Data Set Option” on page 191

BULKLOAD= data set option “BULKLOAD= Data Set Option” on page 194

CAST= Data Set Option

Specifies whether data conversions should be performed on the Teradata DBMS server or by SAS

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Teradata

Default value: none

Syntax

CAST=YES | NO

Syntax Description

YES

forces data conversions (casting) to be done on the Teradata DBMS server and overrides any data overhead percentage limit.

NO

forces data conversions to be done by SAS and overrides any data overhead percentage limit.

Details

Internally, SAS numbers and dates are floating point values. Teradata has varying formats for numbers, including integers, floating point values, and decimal values. Number conversion must occur when you are reading Teradata numbers that are not floating point (Teradata FLOAT). SAS/ACCESS can use the Teradata CAST= function to cause Teradata to perform numeric conversions. The parallelism of Teradata makes it well suited to perform this work. This is especially true if you are running SAS on z/OS (MVS) where CPU activity can be costly.

CAST= can cause more data to be transferred from Teradata to SAS, as a result of the option forcing the Teradata type into a larger SAS type. For example, the CAST= transfer of a Teradata BYTEINT to SAS floating point adds seven overhead bytes to each row transferred.

The following Teradata types are candidates for casting:

- INTEGER
- BYTEINT
- SMALLINT
- DECIMAL
- DATE.

SAS/ACCESS limits data expansion for CAST= to 20 percent in order to trade rapid data conversion by Teradata for extra data transmission. If casting does not exceed a 20 percent data increase, all candidate columns are cast. If the increase exceeds this limit, then SAS attempts to cast Teradata DECIMAL types only. If casting only DECIMAL types still exceeds the increase limit, data conversions are done by SAS.

You can alter the casting rules by using either CAST= or “CAST_OVERHEAD_MAXPERCENT= LIBNAME Option” on page 86. With CAST_OVERHEAD_MAXPERCENT=, you can change the 20 percent overhead limit. With CAST=, you can override the percentage rules:

- CAST=YES forces Teradata to cast all candidate columns
- CAST=NO cancels all Teradata casting

CAST= only applies when you are reading Teradata tables into SAS. It does not apply when you are writing Teradata tables from SAS.

Also, CAST= only applies to SQL that SAS generates for you. If you supply your own SQL with the explicit SQL feature of PROC SQL, you must code your own casting clauses to force data conversions in Teradata instead of SAS.

See Also

“CAST= LIBNAME Option” on page 84

CAST_OVERHEAD_MAXPERCENT= Data Set Option

Specifies the overhead limit for data conversions to be performed in Teradata instead of SAS

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Teradata

Default value: 20 percent

Syntax

CAST_OVERHEAD_MAXPERCENT=<n>

Syntax Description

<n>

Any positive numeric value. The engine default is 20.

Details

Teradata INTEGER, BYTEINT, SMALLINT, and DATE columns require conversion when read in to SAS. Conversions can be performed either by Teradata or by SAS. When performed in Teradata, using Teradata's CAST operator, the row size transmitted to SAS can increase. CAST_OVERHEAD_MAXPERCENT= limits the allowable increase, also called conversion overhead.

Examples

The following example demonstrates the use of CAST_OVERHEAD_MAXPERCENT= to increase the allowable overhead to 40 percent:

```
proc print data=mydblib.emp (cast_overhead_maxpercent=40);  
  where empno<1000;  
run;
```

See Also

“CAST= LIBNAME Option” on page 84 for more information about conversions, conversion overhead, and casting.

COMMAND_TIMEOUT= Data Set Option

Specifies the number of seconds to wait before a command times out

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: OLE DB

Default value: LIBNAME setting

Syntax

COMMAND_TIMEOUT=*number-of-seconds*

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “COMMAND_TIMEOUT= LIBNAME Option” on page 88.

CURSOR_TYPE= Data Set Option

Specifies the cursor type for read only and updatable cursors

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, Microsoft SQL Server, ODBC, OLE DB

Default value: LIBNAME setting

Syntax

CURSOR_TYPE=DYNAMIC | KEYSSET_DRIVEN | FORWARD_ONLY | STATIC

Syntax Description

DYNAMIC

specifies that the cursor reflects all of the changes that are made to the rows in a result set as you move the cursor. The data values and the membership of rows in the cursor can change dynamically on each fetch.

KEYSET_DRIVEN

specifies that the cursor determines which rows belong to the result set when the cursor is opened. However, changes that are made to these rows are reflected as you move the cursor.

FORWARD_ONLY

specifies that the cursor works like a DYNAMIC cursor except that it only supports fetching the rows sequentially. (This is not valid in OLE DB.)

STATIC

specifies that the cursor builds the complete result set when the cursor is opened. No changes made to the rows in the result set after the cursor is opened are reflected in the cursor. Static cursors are read-only.

Details

The driver is allowed to modify the default without an error. Not all database drivers support all cursor types. An error is returned if the specified cursor type is not supported.

OLE DB Details: By default, this option is not set and the provider uses a default. See your provider documentation for more information. See OLE DB programmer reference documentation for details about these properties. The OLE DB properties applied to an open row set are as follows:

CURSOR_TYPE=	OLE DB Properties Applied
DYNAMIC	DBPROP_OTHERINSERT=TRUE, DBPROP_OTHERUPDATEDELETE=TRUE
KEYSET_DRIVEN	DBPROP_OTHERINSERT=FALSE, DBPROP_OTHERUPDATEDELETE=TRUE
STATIC	DBPROP_OTHERINSERT=FALSE, DBPROP_OTHERUPDATEDELETE=FALSE

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “CURSOR_TYPE= LIBNAME Option” on page 95.

DBCOMMIT= Data Set Option

Causes an automatic COMMIT (a permanent writing of data to the DBMS) after a specified number of rows have been processed

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Alias: CHECKPOINT= (Teradata)

Default value: the current LIBNAME setting

Syntax

DBCOMMIT=*n*

Syntax Description

n

is an integer greater than or equal to 0.

Details

DBCOMMIT= affects update, delete, and insert processing. The number of rows processed includes rows that are not processed successfully. When DBCOMMIT=0, a commit is issued only once (after the procedure or DATA step completes).

If the DBCOMMIT= option is explicitly set, SAS/ACCESS fails any update that has a WHERE clause.

Note: If you specify both DBCOMMIT= and ERRLIMIT= and these options collide during processing, COMMIT is issued first and ROLLBACK is issued second. Because the COMMIT that the DBCOMMIT= option caused is issued before the ROLLBACK that the ERRLIMIT= option caused, the DBCOMMIT= option is said to override the ERRLIMIT= option. Δ

DB2 under UNIX and PC Hosts Details: When BULKLOAD=YES, the default is 10000.

Teradata Details: See the FastLoad capability description in the SAS/ACCESS documentation for Teradata for the default behavior of this option.

DBCOMMIT= and ERRLIMIT= are disabled for MultiLoad in order to prevent any conflict with ML_CHECKPOINT=.

Example

In the following example, a commit is issued after every 10 rows are processed:

```
data oracle.dept(dbcommit=10);
    set myoralib.staff;
run;
```

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “DBCOMMIT= LIBNAME Option” on page 96.

DBCONDITION= Data Set Option

Specifies criteria for subsetting and ordering DBMS data

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBCONDITION="DBMS-SQL-query-clause"

Syntax Description

DBMS-SQL-query-clause

is a DBMS-specific SQL query clause, such as WHERE, GROUP BY, HAVING, or ORDER BY.

Details

You can use this option to specify selection criteria in the form of DBMS-specific SQL query clauses, which the SAS/ACCESS engine passes directly to the DBMS for processing. When selection criteria are passed directly to the DBMS for processing, performance is often enhanced. The DBMS checks the criteria for syntax errors when it receives the SQL query.

The DBKEY= and DBINDEX= options are ignored when you use DBCONDITION=.

Example

In the following example, the function that is passed to the DBMS with the DBCONDITION= option causes the DBMS to return to SAS only the rows that satisfy the condition.

```
proc sql;
  create view smithnames as
    select lastname from myoralib.employees
      (dbcondition="where soundex(lastname) = soundex('SMYTHE')" )
      using libname myoralib oracle user=testuser pw=testpass path=dbmssrv;

  select lastname from smithnames;
```

DBCREATE_TABLE_OPTS= Data Set Option

Specifies DBMS-specific syntax to be added to the CREATE TABLE statement

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBCREATE_TABLE_OPTS='DBMS-SQL-clauses'

Syntax Description

DBMS-SQL-clauses

are one or more DBMS-specific clauses that can be appended to the end of an SQL CREATE TABLE statement.

Details

You can use this option to add DBMS-specific clauses to the end of the SQL CREATE TABLE statement. The SAS/ACCESS engine passes the SQL CREATE TABLE statement and its clauses to the DBMS, which executes the statement and creates the DBMS table. This option applies only when you are creating a DBMS table by specifying a libref associated with DBMS data.

Example

In the following example, the DB2 table TEMP is created with the value of the DBCREATE_TABLE_OPTS= option appended to the CREATE TABLE statement.

```
libname mydblib db2 user=testuser
      pwd=testpass dsn=sample;

data mydblib.temp (DBCREATE_TABLE_OPTS='PARTITIONING
      KEY (X) USING HASHING');
x=1; output;
x=2; output;
run;
```

Given this data set option, the following DB2 SQL statement is passed by the SAS/ACCESS interface to DB2 in order to create the DB2 table:

```
CREATE TABLE TEMP (X DOUBLE) PARTITIONING
      KEY (X) USING HASHING
```

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “DBCREATE_TABLE_OPTS= LIBNAME Option” on page 99.

DBFORCE= Data Set Option

Specifies whether to force the truncation of data during insert processing

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: NO

Syntax

DBFORCE=YES | NO

Syntax Description

YES

specifies that the rows which contain data values that exceed the length of the DBMS column are inserted, and the data values are truncated to fit the DBMS column length.

NO

specifies that the rows which contain data values that exceed the DBMS column length are not inserted.

Details

This option determines how the SAS/ACCESS engine handles rows that contain data values that exceed the length of the DBMS column.

The SAS data set option FORCE= overrides this option when it is used with PROC APPEND or the PROC SQL UPDATE statement. The PROC SQL UPDATE statement does not provide a warning before truncating the data.

Example

In the following example, two librefs are associated with Oracle databases; the default databases and schemas are used and therefore are not specified. In the DATA step, MYDBLIB.DEPT is created from the Oracle data referenced by MYORALIB.STAFF. The LASTNAME variable is a character variable of length 20 in MYORALIB.STAFF. During the creation of MYDBLIB.DEPT, the LASTNAME variable is stored as a column of type character and length 10 by using DBFORCE=YES.

```
libname myoralib oracle user=tester1 password=tst1;
libname mydblib oracle user=lee password=dataman;

data mydblib.dept(dbtype=(lastname='char(10)')
  dbforce=yes);
  set myoralib.staff;
run;
```

See Also

“DBTYPE= Data Set Option” on page 221

DBGEN_NAME= Data Set Option

Specifies how SAS renames columns automatically when they contain characters that SAS does not allow

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS

Syntax

DBGEN_NAME=DBMS | SAS

Syntax Description

DBMS

specifies that disallowed characters are converted to underscores.

SAS

specifies that DBMS columns that contain disallowed characters are converted into valid SAS variable names, using the format `_COL n` , where n is the column number (starting with zero). If a name is converted to a name that already exists, a sequence number is appended to the end of the new name.

Details

SAS retains column names when reading data from DBMS tables, unless a column name contains characters that SAS does not allow, such as \$ or @. SAS allows alphanumeric characters and the underscore (_).

This option is intended primarily for National Language Support, notably converting kanji to English characters. English characters converted from kanji are often those that are not allowed in SAS.

Note: Each of the various SAS/ACCESS interfaces handled name collisions differently in SAS Version 6. Some interfaces appended to the end of the name, some replaced the one or more of the last characters in the name, some used a single sequence number, and others used unique counters. If you specify VALIDVARNAME=V6, name collisions are handled the same as in SAS Version 6. Δ

Examples

If you specify DBGEN_NAME=SAS, a DBMS column named DEPT\$AMT is renamed to `_COL n` where n is the column number.

If you specify DBGEN_NAME=DBMS, a DBMS column named DEPT\$AMT is renamed to DEPT_AMT.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “DBGEN_NAME= LIBNAME Option” on page 100.

DBINDEX= Data Set Option

Detects and verifies that indexes exist on a DBMS table. If they do exist and are of the correct type, a join query that is passed to the DBMS might improve performance

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS specific

Syntax

DBINDEX=YES | NO | *<'>index-name<'>*

Syntax Description

YES

triggers the SAS/ACCESS engine to search for all indexes on a table and return them to SAS for evaluation. If a usable index is found, then the join WHERE clause is passed to the DBMS for processing. A usable index is expected to have at least the same attributes as the join column.

NO

no automated index search is performed.

index-name

verifies the index name that is specified for the index columns on the DBMS table. This requires the same type of call as when DBINDEX=YES is used.

Details

When processing a join that involves a large DBMS table and a relatively small SAS data set, you might be able to use DBINDEX= to improve performance.

CAUTION:

Improper use of this option can impair performance. See Using the DBINDEX=, DBKEY=, and MULTI_DATASRC_OPT= Options“Using the DBINDEX=, DBKEY=, and MULTI_DATASRC_OPT= Options” on page 43 for detailed information about using this option. △

Queries must be issued to the necessary DBMS control or system tables to extract index information about a specific table or validate the index that you specified.

You can enter the DBINDEX= option as a LIBNAME option, SAS data set option, or an option with PROC SQL. Here is the order in which the engine processes it option:

- 1 DATA step or PROC SQL specification.
- 2 LIBNAME statement specification.

Note: If you specify the DBKEY= data set option“DBKEY= Data Set Option” on page 207, takes precedence over DBINDEX=. △

Example

The following SAS data set is used in these examples:

```
data s1;
  a=1; y='aaaaa'; output;
  a=2; y='bbbbb'; output;
  a=5; y='ccccc'; output;
run;
```

The following example demonstrates the use of DBINDEX= in the LIBNAME statement:

```
libname mydblib oracle user=myuser password=userpwd dbindex=yes;

proc sql;
  select * from s1 aa, x.dbtabs bb where aa.a=bb.a;
  select * from s1 aa, mydblib.dbtabs bb where aa.a=bb.a;
```

The DBINDEX= values for table dbtabs are retrieved from the DBMS and compared with the join values. In this case, a match was found so the join is passed down to the DBMS using the index. If the index **a** was not found, the join would take place in SAS.

The following example demonstrates the use of DBINDEX= in the SAS DATA step:

```
data a;
  s1;
  set x.dbtabs(dbindex=yes) key=a;
  set mydblib.dbtabs(dbindex=yes) key=a;
run;
```

The key is validated against the list from the DBMS. If **a** is an index, then a pass down occurs. Otherwise the join takes place in SAS.

The following example demonstrates the use of DBINDEX= in PROC SQL:

```
proc sql;
  select * from s1 aa, x.dbtabs(dbindex=yes) bb where aa.a=bb.a;
  select * from s1 aa, mylib.dbtabs(dbindex=yes) bb where aa.a=bb.a;
  /*or*/
  select * from s1 aa, x.dbtabs(dbindex=a) bb where aa.a=bb.a;
  select * from s1 aa, mylib.dbtabs(dbindex=a) bb where aa.a=bb.a;
```

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “DBINDEX= LIBNAME Option” on page 101.

DBKEY= Data Set Option

Specifies a key column to optimize DBMS retrieval. Can improve performance when you are processing a join that involves a large DBMS table and a small SAS data set or DBMS table

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBKEY=(<'>column-1<'><... <'>column-n<'>>)

Syntax Description

column

used by SAS to build an internal WHERE clause to search for matches in the DBMS table based on the key column. For example:

```
select * from sas.a, dbms.b(dbkey=x) where a.x=b.x;
```

In this example, DBKEY= specifies column **x**, which matches the key column designated in the WHERE clause. However, if the DBKEY= column does NOT match the key column in the WHERE clause, then DBKEY= is not used.

Examples

The following example uses DBKEY= with the MODIFY statement in a DATA step:

```
libname invty db2;
data invty.stock;
  set addinv;
  modify invty.stock(dbkey=partno) key=dbkey;
  INSTOCK=instock+nwstock;
  RECDATE=today();
  if _iorc_=0 then replace;
run;
```

To use more than one value for DBKEY=, you must include the second value as a join on the WHERE clause. In the following example, the PROC SQL brings the entire DBMS table into SAS and then proceeds with processing:

```
options sastrace=',,,d' sastraceloc=saslog nostsuffix;

proc sql;
create table work.barbkey as
select keyvalues.empid, employees.hiredate, employees.jobcode
      from mydblib.employees(dbkey=(empid jobcode))
      inner join work.keyvalues on employees.empid = keyvalues.empid;
quit;
```

Details

When processing a join that involves a large DBMS table and a relatively small SAS data set, you might be able to use DBKEY= to improve performance.

When you specify DBKEY=, it is *strongly* recommended that an index exists for the key column in the underlying DBMS table. Performance can be severely degraded without an index.

CAUTION:

Improper use of this option can decrease performance. See “Using the DBINDEX=, DBKEY=, and MULTI_DATASRC_OPT= Options” on page 43 for detailed information about using this option. Δ

DBLABEL= Data Set Option

Specifies whether to use SAS variable labels or SAS variable names as the DBMS column names during output processing

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: NO

Syntax

DBLABEL=YES | NO

Syntax Description

YES

specifies that SAS variable *labels* are used as DBMS column names during output processing.

NO

specifies that SAS variable *names* are used as DBMS column names.

Details

This option is valid only for creating DBMS tables.

Example

In the following example, a SAS data set, NEW, is created with one variable C1. This variable is assigned a label of DEPTNUM. In the second DATA step, the MYDBLIB.MYDEPT table is created by using DEPTNUM as the DBMS column name. Setting DBLABEL=YES enables the label to be used as the column name.

```

data new;
  label c1='deptnum';
  c1=001;
run;

data mydblib.mydept(dblabel=yes);
  set new;
run;

proc print data=mydblib.mydept;
run;

```

DBLINK= Data Set Option

Specifies a link from your default database to another database on the server to which you are connected in the Sybase interface; and specifies a link from your local database to database objects on another server in the Oracle interface

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle, Sybase

Default value: LIBNAME setting

Syntax

DBLINK=*database-link*

Details

This option operates differently in each DBMS.

Oracle Details: A link is a database object that identifies an object that is stored in a remote database. A link contains stored path information and can also contain user name and password information for connecting to the remote database. If you specify a link, SAS uses the link to access remote objects. If you omit DBLINK=, SAS accesses objects in the local database.

Sybase Details: You can use this option to link to another database within the same server to which you are connected. If you omit DBLINK=, SAS can access objects only in your default database.

Example

In this example, SAS sends MYORADB.EMPLOYEES to Oracle as EMPLOYEES@SALES.HQ.ACME.COM.

```
proc print data=myoradb.employees(dblink='sales.hq.acme.com');
run;
```

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “DBLINK= LIBNAME Option” on page 104.

DBMASTER= Data Set Option

Designates which table is the larger table when you are processing a join that involves tables from two different types of databases

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: none

Syntax

DBMASTER=*YES*

Syntax Description

YES

designates which of two tables references in a join operation is the larger table.

Example

In the following example, a table from an Oracle database and a table from a DB2 database are joined. DBMASTER= is set to YES to indicate that the Oracle table is the larger table. The DB2 table is the smaller table.

```
libname mydblib oracle user=testuser /*database 1 */
pw=testpass path='myorapath'

libname mydblib2 db2 user=testuser /*database 2 */
pw=testpass path='mydb2path';

proc sql;
```

```
select * from mydblib.bigtab(dbmaster=yes), mydblib2.smalltab
bigtab.x=smalltab.x;
```

Details

You can use this option with the MULTI_DATASRC_OPT= option to specify which table reference in a join is the larger table. This can improve performance by eliminating the processing normally performed to determine this information. However, this option is ignored when outer joins are processed.

See Also

“MULTI_DATASRC_OPT= LIBNAME Option” on page 126

DBMAX_TEXT= Data Set Option

Determines the length of any very long DBMS character data type that is read into SAS or written from SAS when you are using a SAS/ACCESS engine

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase

Default value: 1024

Syntax

DBMAX_TEXT= *integer*

Syntax Description

integer

is a number between 1 and 32,767.

Details

This option applies to appending and updating rows in an existing table. It does not apply when creating a table.

DBMAX_TEXT= is usually used with a very long DBMS character data type, such as the Sybase TEXT data type or the Oracle LONG RAW data type.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “DBMAX_TEXT= LIBNAME Option” on page 105.

DBNULL= Data Set Option

Indicates whether NULL is a valid value for the specified columns when a table is created

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

```
DBNULL=<_ALL=YES | NO > | ( <column-name-1=YES | NO >
    <...<column-name-n=YES | NO >>)
```

Syntax Description

ALL

specifies that the YES or NO applies to all columns in the table. (This is valid in the interfaces to Informix, Oracle, Sybase, and Teradata only.)

YES

specifies that the NULL value is valid for the specified columns in the DBMS table.

NO

specifies that the NULL value is not valid for the specified columns in the DBMS table.

Details

This option is valid only for creating DBMS tables. If you specify more than one column name, the names must be separated with spaces.

The DBNULL= option processes values from left to right. If you specify a column name twice or if you use the ALL value, the last value overrides the first value that you specified for the column.

Examples

In the following example, by using the DBNULL= option, the EMPID and JOBCODE columns in the new MYDBLIB.MYDEPT2 table are prevented from accepting NULL values. If the EMPLOYEES table contains NULL values in the EMPID or JOBCODE columns, the DATA step fails.

```
data mydblib.mydept2(dbnull=(empid=no jobcode=no));
    set mydblib.employees;
run;
```

In the following example, all columns in the new MYDBLIB.MYDEPT3 table except for the JOBCODE column are prevented from accepting NULL values. If the EMPLOYEES table contains NULL values in any column other than the JOBCODE column, the DATA step fails.


```
data mydblib.mydept3(dbnull=( _ALL_ =no jobcode=YES));
    set mydblib.employees;
run;
```

See Also

“NULLCHAR= Data Set Option” on page 244

“NULLCHARVAL= Data Set Option” on page 245

DBNULLKEYS= Data Set Option

Controls the format of the WHERE clause with regard to NULL values when you use the DBKEY= data set option

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle

Default value: LIBNAME setting

Syntax

DBNULLKEYS= YES | NO

Details

If there might be NULL values in the transaction table or the master table for the columns that you specify in the DBKEY= option, then use DBNULLKEYS=YES. When you specify DBNULLKEYS=YES and also a column that is not defined as NOT NULL in the DBKEY= data set option, SAS generates a WHERE clause to find NULL values. For example, if you specify DBKEY=COLUMN and COLUMN is not defined as NOT NULL, SAS generates a WHERE clause with the following syntax:

```
WHERE ((COLUMN = ?) OR ((COLUMN IS NULL) AND (? IS NULL)))
```

With this syntax, SAS can prepare the statement once and use it for any value (NULL or NOT NULL) in the column. Note that this syntax has the potential to be much less efficient than the shorter form of the WHERE clause that is shown below. When you specify DBNULLKEYS=NO or specify a column that is defined as NOT NULL in the DBKEY= option, SAS generates a simple WHERE clause.

If you know that there are no NULL values in the transaction table or the master table for the columns you specify in the DBKEY= option, you can use DBNULLKEYS=NO. If you specify DBNULLKEYS=NO and specify DBKEY=COLUMN, SAS generates a shorter form of the WHERE clause. It generates this clause regardless of whether or not the column specified in DBKEY= is defined as NOT NULL.

```
WHERE (COLUMN = ?)
```

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “DBNULLKEYS= LIBNAME Option” on page 105.

DBPROMPT= Data Set Option

Specifies whether SAS displays a window that prompts you to enter DBMS connection information

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: HP Neoview, MySQL, Netezza, Oracle, Sybase

Default value: NO

Syntax

DBPROMPT=YES | NO

Syntax Description

YES

displays the prompting window.

NO

does not display the prompting window.

Details

This data set option is supported only for view descriptors.

Oracle Details: In the Oracle interface, you can enter 30 characters each for USERNAME and PASSWORD and up to 70 characters for PATH, depending on your platform and terminal type.

Examples

In the following example, connection information is specified in the ACCESS procedure. The DBPROMPT= data set option defaults to NO during the PRINT procedure because it is not specified.

```
proc access dbms=oracle;
  create alib.mydesc.access;
  user=testuser;
  password=testpass;
  table=dept;
  create vlib.myview.view;
  select all;
run;
```

```
proc print data=vlib.myview;
run;
```

In the following example, the DBPROMPT window opens during connection to the DBMS. Values that were previously specified during the creation of MYVIEW are pulled into the DBPROMPT window fields. You must edit or accept the connection information in the DBPROMPT window to proceed. The password value appears as a series of asterisks; you can edit it.

```
proc print data=vlib.myview(dbprompt=yes);
run;
```

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “DBPROMPT= LIBNAME Option” on page 106.

DBSASLABEL= Data Set Option

Specifies how the engine returns column labels

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: COMPAT

Syntax

DBSASLABEL=COMPAT | NONE

Syntax Description

COMPAT

specifies that the labels returned should be compatible with what the application normally receives. In other words, engines exhibit their normal behavior.

NONE

specifies that the engine does not return a column label. The engine returns blanks for the column labels.

Details

By default, the SAS/ACCESS interface for your DBMS generates column labels from the column names, rather than from the real column labels.

This option enables the user to override this default behavior. It is useful in the PROC SQL context where column labels instead of column aliases are used as headers.

Examples

The following example demonstrates how DBSASLABEL= is used to return blank column labels so PROC SQL can use the column aliases as the column headings.

```
proc sql;
  select deptno as Department ID, loc as Location
  from mylib.dept(dbsaslabel=none);
```

Without the DBSASLABEL= option set to NONE, the aliases would be ignored and DEPTNO and LOC would be used as column headings in the result set.

DBSASTYPE= Data Set Option

Specifies data types to override the default SAS data types during input processing

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Informix, Microsoft SQL Server, SQL, Netezza, ODBC, OLE DB, Oracle

Default value: DBMS-specific

Syntax

DBSASTYPE=(*column-name-1*=<'>*SAS-data-type*<'><...*column-name-n*=<'>*SAS-data-type*<'>>)

Syntax Description

column-name

specifies a DBMS column name.

SAS-data-type

specifies a SAS data type. SAS data types include the following: CHAR(*n*), NUMERIC, DATETIME, DATE, TIME. See the documentation for your SAS/ACCESS interface for details.

Details

By default, the SAS/ACCESS interface for your DBMS converts each DBMS data type to a SAS data type during input processing. When you need a different data type, you can use this option to override the default and assign a SAS data type to each specified DBMS column. Some conversions might not be supported; if a conversion is not supported, SAS prints an error to the log.

Examples

In the following example, DBSASTYPE= specifies a data type to use for the column MYCOLUMN when SAS is printing ODBC data. If the data in this DBMS column is

stored in a format that SAS does not support, such as SQL_DOUBLE(20), this enables SAS to print the values.

```
proc print data=mylib.mytable
  (dbsastype=(mycolumn='CHAR(20)'));
run;
```

In the following example, the data stored in the DBMS FIBERSIZE column has a data type that provides more precision than what SAS could accurately support, such as DECIMAL(20). If you used only PROC PRINT on the DBMS table, the data might be rounded or displayed as a missing value. Instead, you could use DBSASTYPE= to convert the column to a character field of the length 21. Because the DBMS performs the conversion before the data is brought into SAS, there is no loss of precision.

```
proc print data=mylib.specprod
  (dbsastype=(fibersize='CHAR(21)'));
run;
```

The following example, uses DBSASTYPE= to append one table to another when the data types are not comparable. If the EMPID variable in the SAS data set is defined as CHAR(20) and the EMPID column in the DBMS table is defined as DECIMAL(20), you can use DBSASTYPE= to make them match:

```
proc append base=dblib.hrdata (dbsastype=(empid='CHAR(20)'))
  data=saslib.personnel;
run;
```

DBSASTYPE= specifies to SAS that the EMPID is defined as a character field of length 20. When a row is inserted from the SAS data set into a DBMS table, the DBMS performs a conversion of the character field to the DBMS data type DECIMAL(20).

DBSLICE= Data Set Option

Specifies user-supplied WHERE clauses to partition a DBMS query for threaded reads

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, ODBC, Oracle, Sybase, Teradata

Default value: none

Syntax

DBSLICE=(*WHERE-clause-1* " *WHERE-clause-2*" < ... "*WHERE-clause-n*">)

DBSLICE=(<*server=*>"*WHERE-clause-1*" <*server=*>" *WHERE-clause-2*" < ... <*server=*>"*WHERE-clause-n*">)

Syntax Description

Two syntax diagrams are shown here to highlight the simpler version. In many cases, the first, simpler syntax is sufficient. The optional *server=* form is valid only for ODBC and DB2 under UNIX and PC hosts.

WHERE-clause

The WHERE clauses in the syntax signifies DBMS-valid WHERE clauses that partition the data. The clauses should not cause any omissions or duplications of rows in the results set. For example, if EMPNUM can be null, the following DBSLICE= specification omits rows, creating an *incorrect* result set:

```
DBSLICE=( "EMPNUM<1000" "EMPNUM>=1000" )
```

A correct form is:

```
DBSLICE=( "EMPNUM<1000" "EMPNUM>=1000" "EMPNUM IS NULL" )
```

In the following example, DBSLICE= creates an *incorrect* set by duplicating SALES with value zero:

```
DBSLICE=( 'SALES<=0 or SALES=NULL' 'SALES>=0' )
```

server

identifies a particular server node in a DB2 partitioned database or in a Microsoft SQL Server partitioned view. Used for the best possible read performance, this enables your SAS thread to directly connect to the node containing the data partition corresponding to your WHERE clause.

Details

If your table reference is eligible for threaded reads (that is, if it is a read-only LIBNAME table reference), DBSLICE= forces a threaded read to occur, partitioning the table with the WHERE clauses you supply. Use DBSLICE= when SAS is unable to generate threaded reads automatically, or if you can provide better partitioning.

DBSLICE= is appropriate for experienced programmers familiar with the layout of their DBMS tables. A well-tuned DBSLICE= specification usually outperforms SAS automatic partitioning. For example, a well-tuned DBSLICE= specification might better distribute data across threads by taking advantage of a column that SAS/ACCESS cannot use when it automatically generates partitioning WHERE clauses.

DBSLICE= delivers optimal performance for DB2 under UNIX and for Microsoft SQL Server. Conversely, DBSLICE= can degrade performance compared to automatic partitioning. For example, Teradata invokes the FastExport Utility for automatic partitioning. If this is overridden with DBSLICE=, WHERE clauses are generated instead. Even with well planned WHERE clauses, performance is degraded because FastExport is innately faster.

CAUTION:

When using DBSLICE=, you are responsible for data integrity. If your WHERE clauses omit rows from the result set or retrieves the same row on more than one thread, your input DBMS result set is incorrect and your SAS program generates incorrect results. Δ

Examples

In the following example, DBSLICE= partitions on the GENDER column can have only the values **m**, **M**, **f**, and **F**. This DBSLICE= clause does not work for all DBMSs due to the use of UPPER and single quotation marks (some DBMSs require double quotation marks around character literals). Two threads are created.

```
proc reg SIMPLE
data=lib.customers(DBSLICE="UPPER(GENDER)='M' "UPPER(GENDER)='F'");
var age weight;
where years_active>1;
run;
```

The following example partitions on the non-NULL column CHILDREN, the number of children in a family. Three threads are created.

```
data local;
set lib.families(DBSLICE=("CHILDREN<2" "CHILDREN>2" "CHILDREN=2"));
where religion="P";
run;
```

DBSLICEPARM= Data Set Option

Controls the scope of DBMS threaded reads and the number of DBMS connections

Valid in: DATA and PROC Steps (when accessing DBMS data using SAS/ACCESS software) (also available as a SAS configuration file option, SAS invocation option, global SAS option, and LIBNAME option)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, ODBC, Oracle, Sybase, Teradata

Default value: THREADED_APPS,2 (DB2 z/OS, Oracle, and Teradata)

THREADED_APPS,2 or 3 (DB2 UNIX/PC, Informix, Microsoft SQL Server, ODBC, and Sybase)

Syntax

DBSLICEPARM=(NONE | THREADED_APPS | ALL <, *max-threads*>)

Syntax Description

Two syntax diagrams are shown here in order to highlight the simpler version. In most cases, the simpler version suffices.

NONE

disables DBMS threaded reads. SAS reads tables on a single DBMS connection, as it did with SAS Version 8 and earlier.

THREADED_APPS

makes fully threaded SAS procedures (threaded applications) eligible for threaded reads.

ALL

makes all read-only librefs eligible for threaded reads. This includes SAS threaded applications, as well as the SAS DATA step and numerous SAS procedures.

max-threads

specifies with a positive integer value the maximum number of connections per table read. A partition or portion of the data is read on each connection. The combined rows across all partitions are the same irrespective of the number of connections. That is, changes to the number of connections do not change the result set. Increasing the number of connections instead redistributes the same result set across more connections.

There are diminishing returns when increasing the number of connections. With each additional connection, more burden is placed on the DBMS, and a smaller percentage of time is saved in SAS. Therefore, you should consult your DBMS-specific documentation for threaded reads before using this parameter.

Details

DBSLICEPARM= can be used in numerous locations, and the usual rules of option precedence apply. A table option has the highest precedence, then a LIBNAME option, and so on. A SAS configuration file option has the lowest precedence because DBSLICEPARM= in any of the other locations overrides that configuration setting.

DBSLICEPARM=ALL and DBSLICEPARM=THREADED_APPS make SAS programs eligible for threaded reads. To determine whether threaded reads are actually generated, turn on SAS tracing and run a program, as shown in the following example:

```
options sastrace=',,,d' sastraceloc=saslog nostsuffix;
proc print data=lib.dbtable(dbsliceparm=(ALL));
    where dbcol>1000;
run;
```

If you want to directly control the threading behavior, use the DBSLICE= data set option.

For DB2 UNIX/PC, Informix, Microsoft SQL Server, ODBC, and Sybase, the default thread number depends on whether an application passes in the number of threads (CPUCOUNT=) and whether the data type of the column that is selected for the data partitioning purpose is binary.

Examples

This code demonstrates how you can use DBSLICEPARM= in a PC SAS configuration file entry to turn off threaded reads for all SAS users:

```
--dbsliceparm NONE
```

Here is how you can use DBSLICEPARM= as a z/OS invocation option to turn on threaded reads for read-only references to DBMS tables throughout a SAS job:

```
sas o(dbsliceparm=ALL)
```

You can use this code to set DBSLICEPARM= as a SAS global option to increase maximum threads to three for SAS threaded applications. It would most likely be one of the first statements in your SAS code:

```
option dbsliceparm=(threaded_apps,3);
```

The following code demonstrates how to use DBSLICEPARM= as a LIBNAME option to turn on threaded reads for read-only table references that use this particular libref:

```
libname dblib oracle user=scott password=tiger dbsliceparm=ALL;
```

The following code demonstrates how to use DBSLICEPARM= as a table level option to turn on threaded reads for this particular table, requesting up to four connections:

```
proc reg SIMPLE;
    data=dblib.customers (dbsliceparm=(all,4));
    var age weight;
    where years_active>1;
run;
```

See Also

“DBSLICEPARM= LIBNAME Option” on page 109

DBTYPE= Data Set Option

Specifies a data type to use instead of the default DBMS data type when SAS creates a DBMS table

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

```
DBTYPE=(column-name-1=<'>DBMS-type<'>
        <...column-name-n=<'>DBMS-type<'>>)
```

Syntax Description

column-name

specifies a DBMS column name.

DBMS-type

specifies a DBMS data type. See the documentation for your SAS/ACCESS interface for the default data types for your DBMS.

Details

By default, the SAS/ACCESS interface for your DBMS converts each SAS data type to a predetermined DBMS data type when outputting data to your DBMS. When you need a different data type, use DBTYPE= to override the default data type chosen by the SAS/ACCESS engine.

Teradata Details: In Teradata, you can use DBTYPE= to specify data attributes for a column. See your Teradata CREATE TABLE documentation for information about the data type attributes that you can specify. If you specify DBNULL=NO for a column, do not also use DBTYPE= to specify NOT NULL for that column. If you do, 'NOT NULL' is inserted twice in the column definition. This causes Teradata to generate an error message.

Examples

In the following example, DBTYPE= specifies the data types that are used when you create columns in the DBMS table.

```
data mydblib.newdept(dbtype=(deptno='number(10,2)' city='char(25)'));
    set mydblib.dept;
run;
```

The following example creates a new Teradata table, NEWDEPT, specifying the Teradata data types for the DEPTNO and CITY columns.

```
data mydblib.newdept(dbtype=(deptno='byteint' city='char(25)'));
    set dept;
```

```
run;
```

The following example creates a new Teradata table, NEWEMPLOYEES, and specifies a data type and attributes for the EMPNO column. The example encloses the Teradata type and attribute information in double quotation marks. Single quotation marks conflict with those that are required by the Teradata FORMAT attribute. If you use single quotation marks, SAS returns syntax error messages.

```
data mydblib.newemployees(dbtype= (empno="SMALLINT FORMAT '9(5)'  
    CHECK (empno >= 100 AND empno <= 2000)));  
set mydblib.employees;  
run;
```

See Also

“DBFORCE= Data Set Option” on page 202

DISTRIBUTE_ON= Data Set Option

Specifies a column name to use in the DISTRIBUTE ON clause of the CREATE TABLE statement

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Netezza

Alias: DISTRIBUTE=

Default value: none

Syntax

DISTRIBUTE_ON='column-1 <... ,column-n>' | RANDOM

Syntax Description

column

specifies a DBMS column name.

RANDOM

specifies that the Netezza Performance Server distributes data evenly across all SPU's (round-robin distribution).

Details

You can use this option to specify a column name to use in the DISTRIBUTE ON= clause of the CREATE TABLE statement. Each table in the database must have a distribution key that consists of one to four columns. If you do not specify this option, Netezza selects a distribution key.

Examples

This example uses DISTRIBUTE_ON= to create a distribution key on a single column.

```
proc sql;
create table netlib.customtab(DISTRIBUTE_ON='partno')
  as select partno, customer, orderdat from saslib.orders;
quit;
```

To create a distribution key on more than one column, separate the columns with commas.

```
data netlib.mytab(DISTRIBUTE_ON='col1,col2');
col1=1;col2=12345;col4='mytest';col5=98.45;
run;
```

This next example shows how to use the RANDOM keyword.

```
data netlib.foo(distribute_on=random);
mycol1=1;mycol2='test';
run;
```

ERRLIMIT= Data Set Option

Specifies the number of errors that are allowed before SAS stops processing and issues a rollback

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: 1

Syntax

ERRLIMIT=*integer*

Syntax Description

integer

is a positive integer that represents the number of errors after which SAS stops processing and issues a rollback.

Details

SAS calls the DBMS to issue a rollback after a specified number of errors occurs during the processing of inserts, deletes, updates, and appends. If ERRLIMIT= is set to 0, SAS processes all rows, regardless of the number of errors that occur. The SAS log displays the total number of rows processed and the number of failed rows, if applicable.

The DBCOMMIT= option overrides the ERRLIMIT= option. If you specify a value for DBCOMMIT= other than zero, then rollbacks that the ERRLIMIT= option affects might not include records that are processed unsuccessfully because DBCOMMIT= already committed them.

Note: A significant performance impact can result if you use this option from a SAS client session in SAS/SHARE or SAS/CONNECT environment to create or populate a newly created table. To prevent this, use the default setting, ERRLIMIT=1. Δ

Teradata Details: DBCOMMIT= and ERRLIMIT= are disabled for MultiLoad in order to prevent any conflict with ML_CHECKPOINT=.

Example

In the following example, SAS stops processing and issues a rollback to the DBMS at the occurrence of the tenth error. The MYDBLIB libref was assigned in a prior LIBNAME statement.

```
data mydblib.employee3 (errlimit=10);
  set mydblib.employees;
  where salary > 40000;
run;
```

IGNORE_READ_ONLY_COLUMNS= Data Set Option

Specifies whether to ignore or include columns whose data types are read-only when generating an SQL statement for inserts or updates

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB

Default value: NO

Syntax

IGNORE_READ_ONLY_COLUMNS=YES | NO

Syntax Description

YES

specifies that the SAS/ACCESS engine ignores columns whose data types are read-only when you are generating insert and update SQL statements

NO

specifies that the SAS/ACCESS engine does not ignore columns whose data types are read-only when you are generating insert and update SQL statements

Details

Several databases include data types that can be read-only, such as Microsoft SQL Server's timestamp data type. Several databases also have properties that allow certain data types to be read-only, such as Microsoft SQL Server's identity property.

When IGNORE_READ_ONLY_COLUMN=NO (the default) and a DBMS table contains a column that is read-only, an error is returned that the data could not be modified for that column.

Examples

For the following example, a database that contains the table Products is created with two columns: ID and PRODUCT_NAME. The ID column is defined by a read-only data type and PRODUCT_NAME is a character column.

```
CREATE TABLE products (id int IDENTITY PRIMARY KEY, product_name varchar(40))
```

If you have a SAS data set that contains the name of your products, you can insert the data from the SAS data set into the Products table:

```
data work.products;
  id=1;
  product_name='screwdriver';
  output;
  id=2;
  product_name='hammer';
  output;
  id=3;
  product_name='saw';
  output;
  id=4;
  product_name='shovel';
  output;
run;
```

When IGNORE_READ_ONLY_COLUMNS=NO (the default), the database returns an error because the ID column cannot be updated. However, if you set the option to YES and execute a PROC APPEND, the append succeeds and the generated SQL statement does not contain the ID column.

```
libname x odbc uid=dbitest pwd=dbigrpl dsn=lupinss
          ignore_read_only_columns=yes;
options sastrace=',,,' sastraceloc=saslog nostsuffix;
proc append base=x.PRODUCTS data=work.products;
run;
```

See Also

To apply this option to an individual data set, see the LIBNAME option “IGNORE_READ_ONLY_COLUMNS= LIBNAME Option” on page 117.

IN= Data Set Option

Lets you specify the database or tablespace in which you want to create a new table

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS

Default value: LIBNAME setting

Syntax

`IN='database-name.tablespace-name'|'DATABASE database-name'`

Syntax Description

database-name.tablespace-name

specifies the names of the database and tablespace, which are separated by a period.

DATABASE *database-name*

specifies only the database name. In this case, you specify the word DATABASE, then a space and the database name. Enclose the entire specification in single quotation marks.

Details

The IN= option is relevant only when you are creating a new table. If you omit this option, the default is to create the table in the default database or tablespace.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “IN= LIBNAME Option” on page 118.

INSERT_SQL= Data Set Option

Determines the method that is used to insert rows into a data source

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Microsoft SQL Server, ODBC, OLE DB

Default value: LIBNAME setting

Syntax

`INSERT_SQL=YES | NO`

Syntax Description

YES

specifies that the SAS/ACCESS engine uses the data source’s SQL insert method to insert new rows into a table.

NO

specifies that the SAS/ACCESS engine uses an alternate (DBMS-specific) method to add new rows to a table.

Details

Flat-file databases such as dBase, FoxPro, and text files have generally improved insert performance when `INSERT_SQL=NO`. Other databases might have inferior insert performance or might fail with this setting. Therefore, you should experiment to determine the optimal setting for your situation.

ODBC Details: The ODBC default is YES, except for Microsoft Access which has a default of NO. When `INSERT_SQL=NO`, the `SQLSetPos (SQL_ADD)` function inserts rows in groups that are the size of the `INSERTBUFF=` option value. The `SQLSetPos (SQL_ADD)` function does not work unless it is supported by your ODBC driver.

OLE DB Details: By default, the OLE DB interface attempts to use the most efficient row insertion method for each data source. You can use the `INSERT_SQL` option to override the default in the event that it is not optimal for your situation. The OLE DB alternate method (used when this option is set to NO) uses the OLE DB `IRowsetChange` interface.

SQL Server Details: The SQL Server default is YES. When `INSERT_SQL=NO`, the `SQLSetPos (SQL_ADD)` function inserts rows in groups that are the size of the `INSERTBUFF=` option value. The `SQLSetPos (SQL_ADD)` function does not work unless it is supported by your ODBC driver.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the `LIBNAME` option “`INSERT_SQL= LIBNAME` Option” on page 119.

INSERTBUFF= Data Set Option

Specifies the number of rows in a single DBMS insert

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle

Default value: LIBNAME setting

Syntax

`INSERTBUFF=positive-integer`

Syntax Description

positive-integer

specifies the number of rows to insert.

Details

SAS allows the maximum number of rows that is allowed by the DBMS. The optimal value for this option varies with factors such as network type and available memory. You might need to experiment with different values to determine the best value for your site.

When you assign a value that is greater than INSERTBUFF=1, the SAS application notes that indicate the success or failure of the insert operation might be incorrect. This is because these notes represent information for only a single insert, even when multiple inserts are performed.

If the DBCOMMIT= option is specified with a value that is less than the value of INSERTBUFF=, then DBCOMMIT= overrides INSERTBUFF=.

Note: When you insert rows with the VIEWTABLE window or the FSEDIT or FSVIEW procedure, use INSERTBUFF=1 to prevent the DBMS interface from trying to insert multiple rows. These features do not support inserting more than one row at a time. △

Note: Additional driver-specific restrictions might apply. △

DB2 under UNIX and PC Hosts Details: You must specify INSERT_SQL=YES in order to use this option. If one row in the insert buffer fails, all rows in the insert buffer fail.

Microsoft SQL Server Details: You must specify INSERT_SQL=YES in order to use this option.

MySQL Details: The default is 0. Any value greater than 0 turns on the INSERTBUFF= option. The engine then calculates how many rows it can insert at one time, based on the row size. If one row in the insert buffer fails, all rows in the insert buffer might fail, depending on your storage type.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “INSERTBUFF= LIBNAME Option” on page 120.

KEYSET_SIZE= Data Set Option

Specifies the number of rows in the cursor that are keyset driven

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Microsoft SQL Server, ODBC

Default value: LIBNAME setting

Syntax

KEYSET_SIZE=*number-of-rows*

Syntax Description

number-of-rows

is a positive integer from 0 through the number of rows in the cursor.

Details

This option is valid only when CURSOR_TYPE=KEYSET_DRIVEN.

If KEYSET_SIZE=0, then the entire cursor is keyset driven. If a value greater than 0 is specified for KEYSET_SIZE=, then the value chosen indicates the number of rows, within the cursor, that function as a keyset-driven cursor. When you scroll beyond the bounds that are specified by KEYSET_SIZE=, then the cursor becomes dynamic and new rows might be included in the cursor; this becomes the new keyset and the cursor functions as a keyset-driven cursor again. Whenever the value specified is between 1 and the number of rows in the cursor, the cursor is considered to be a mixed cursor. Part of the cursor functions as a keyset-driven cursor and part of the cursor functions as a dynamic cursor.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “KEYSET_SIZE= LIBNAME Option” on page 122.

LOCATION= Data Set Option

Lets you further specify exactly where a table resides

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 z/OS

Default value: LIBNAME setting

Syntax

LOCATION=*location-name*

Details

The location name maps to the location in the SYSIBM.SYSLOCATIONS catalog in the communication database.

In the DBMS engine, the location is converted to the first level of a three-level table name: *location-name*.AUTHID.TABLE. The connection to the remote DBMS subsystem is done implicitly by the DBMS when it receives a three-level name in an SQL statement.

If you specify LOCATION=, you must also specify the AUTHID= option.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “LOCATION= LIBNAME Option” on page 122.

LOCKTABLE= Data Set Option

Places exclusive or shared locks on tables

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Informix

Default value: LIBNAME setting

Syntax

LOCKTABLE=EXCLUSIVE | SHARE

Syntax Description

EXCLUSIVE

locks a table exclusively, preventing other users from accessing any table that you open in the libref.

SHARE

locks a table in shared mode, allowing other users or processes to read data from the tables, but preventing users from updating data.

Details

You can lock tables only if you are the owner or have been granted the necessary privilege. If you omit LOCKTABLE=, no locking occurs.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “LOCKTABLE= LIBNAME Option” on page 123.

MBUFFSIZE= Data Set Option

Specifies the size of the shared memory buffers to be used for transferring data from SAS to Teradata.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: 64K

Syntax

MBUFFSIZE=*size-of-shared-memory-buffers*

Syntax Description

size-of-shared-memory-buffers

a numeric value (between the size of a row being loaded and 1MB) that specifies the buffer size.

Details

MBUFFSIZE= specifies the size of data buffers used for transferring data from SAS to Teradata. Two data set options are available for tuning the number and size of data buffers used for transferring data from SAS to Teradata.

When using MULTILOAD=, data is transferred from SAS to Teradata using shared memory segments. The default shared memory buffer size is 64K. The default number of shared memory buffers used for the transfer is 2.

Use the MBUFFSIZE= data set option to vary the size of the shared memory buffers from the size of each data row up to 1MB.

Use BUFFERS= to vary the number of buffers for data transfer from 1 to 8.

See Also

See BUFFERS= for information about changing the number of shared memory buffers.

ML_CHECKPOINT= Data Set Option

Specifies the interval between checkpoint operations, in minutes.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: 0

Syntax

ML_CHECKPOINT=*checkpoint-rate*

Syntax Description

checkpoint-rate

a numeric value that specifies the interval between checkpoint operations, in minutes.

Details

ML_CHECKPOINT=0 is the default; no checkpoints are taken if the default is used. If ML_CHECKPOINT= is between 1 and 59 inclusive, checkpoints are taken at the

specified intervals, in minutes. If the value of ML_CHECKPOINT= is greater than or equal to 60, then a checkpoint operation occurs after a multiple of the specified rows are loaded.

ML_CHECKPOINT= functions very much like CHECKPOINT in the native Teradata MultiLoad utility, but it differs from the DBCOMMIT= data set option. Note that DBCOMMIT= is disabled for MultiLoad to prevent any conflict.

For more information about using checkpoints for MultiLoad, see Teradata MultiLoad utility documentation.

See Also

For more information about using checkpoints and restarting MultiLoad jobs, see MULTILOAD=.

ML_ERROR1= Data Set Option

Specifies the name of a temporary table that MultiLoad uses to track errors that were generated during the acquisition phase of a bulk-load operation.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: none

Syntax

ML_ERROR1=*temporary-table-name*

Syntax Description

temporary-table-name

specifies the name of a temporary table that MultiLoad uses to track errors that were generated during the acquisition phase of a bulk-load operation.

Details

The Teradata MultiLoad utility uses four different temporary tables when performing the bulk-load operation. MultiLoad uses a log table to track restart information, two error tables to track errors, and a work table to hold data before the insert operation is made.

By default, the SAS/ACCESS MultiLoad facility generates names for these temporary tables:

Restart table	SAS_ML_RS_ <i>randnum</i>
Acquisition error table	SAS_ML_ET_ <i>randnum</i>
Application error table	SAS_ML_UT_ <i>randnum</i>
Work table	SAS_ML_WT_ <i>randnum</i>

In the table names, *randnum* represents a random number.

Use this option to specify the name of a table to use for storing the errors that were generated during the acquisition phase of the MultiLoad bulk-load operation. Upon restarting the job, ML_ERROR1 is used to specify the name of the table that is used for storing errors generated during a previously failed MultiLoad job.

For more information about the temporary table names that MultiLoad uses, see Teradata documentation on the MultiLoad utility.

Note: Do not use ML_ERROR1 with the ML_LOG= data set option. ML_LOG= provides a common prefix for all temporary tables that the Teradata MultiLoad utility uses. △

See Also

To specify a common prefix for all of the temporary tables that are used by the Teradata MultiLoad utility, see ML_LOG=.

ML_ERROR2= Data Set Option

Specifies the name of a temporary table that MultiLoad uses to track errors that were generated during the application phase of a bulk-load operation.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: none

Syntax

ML_ERROR2=*temporary-table-name*

Syntax Description

temporary-table-name

specifies the name of a temporary table that MultiLoad uses to track errors that were generated during the application phase of a bulk-load operation.

Details

The Teradata MultiLoad utility uses four different temporary tables when performing the bulk-load operation. MultiLoad uses a log table to track restart information, two error tables to track errors, and a work table to hold data before the insert operation is made.

By default, the SAS/ACCESS MultiLoad facility generates names for these temporary tables:

Restart table	SAS_ML_RS_ <i>randnum</i>
Acquisition error table	SAS_ML_ET_ <i>randnum</i>
Application error table	SAS_ML_UT_ <i>randnum</i>
Work table	SAS_ML_WT_ <i>randnum</i>

In these table names, *randnum* represents a random number.

Use this option to specify the name of a table to use for storing the errors that were generated during the application phase of the MultiLoad bulk-load operation. When you restart the job, use ML_ERROR2 to specify the name of the table for storing errors that were generated during a previously failed MultiLoad job.

For more information about the temporary table names that MultiLoad uses, see Teradata documentation on the MultiLoad utility.

Note: Do not use ML_ERROR2 with the ML_LOG= data set option. ML_LOG= provides a common prefix for all of the temporary tables that are used by the Teradata MultiLoad utility. Δ

See Also

To specify a common prefix for all temporary tables that the Teradata MultiLoad utility uses, see ML_LOG=

ML_LOG= Data Set Option

Specifies a prefix for the names of the temporary tables that MultiLoad uses during a bulk-load operation.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: none

Syntax

ML_LOG=*prefix-for-MultiLoad-temporary-tables*

Syntax Description

prefix-for-MultiLoad-temporary-tables

specifies the prefix to use when naming Teradata tables that the Teradata MultiLoad utility uses during a bulk-load operation.

Details

The Teradata MultiLoad utility uses four different temporary tables when performing the bulk-load operation. MultiLoad uses a log table to track restart information, two error tables to track errors, and a work table to hold data before the insert operation is made. By default, the SAS/ACCESS MultiLoad facility generates names for these temporary tables:

Restart table	SAS_ML_RS_ <i>randnum</i>
Acquisition error table	SAS_ML_ET_ <i>randnum</i>
Application error table	SAS_ML_UT_ <i>randnum</i>
Work table	SAS_ML_WT_ <i>randnum</i>

In these table names, *randnum* represents a random number.

To specify a different name for the tables listed above, use the ML_RESTART=, ML_ERROR1=, ML_ERROR2=, and ML_WORK= data set options, respectively.

You can use the ML_LOG= data set option to specify a prefix for the temporary table names that MultiLoad uses. For example, if you use **ML_LOG=MY_ERRORS**, these table names are generated:

Restart table	MY_ERRORS_RS
Acquisition error table	MY_ERRORS_ET
Application error table	MY_ERRORS_UT
Work table	MY_ERRORS_WT

ML_LOG= can also be used in a restart step to specify the same prefix that was used by a previous MultiLoad job that used ML_LOG= and failed.

For more information about the temporary table names that MultiLoad uses, see Teradata documentation on the MultiLoad utility.

Note: ML_LOG= is not to be used with data set options ML_RESTART=, ML_ERROR1=, ML_ERROR2=, and ML_WORK=, which provide specific names to the temporary files. △

ML_RESTART= Data Set Option

Specifies the name of a temporary table that is used by MultiLoad to track checkpoint information.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: none

Syntax

ML_RESTART=*temporary-table-name*

Syntax Description

temporary-table-name

specifies the name of the temporary table that is used by the Teradata MultiLoad utility to track checkpoint information.

Details

The Teradata MultiLoad utility uses four different temporary tables when performing the bulk-load operation. MultiLoad uses a log table to track restart information, two error tables to track errors, and a work table to hold data before the insert operation is made.

By default, the SAS/ACCESS MultiLoad facility generates names for these temporary tables:

Restart table	SAS_ML_RS_ <i>randnum</i>
Acquisition error table	SAS_ML_ET_ <i>randnum</i>
Application error table	SAS_ML_UT_ <i>randnum</i>
Work table	SAS_ML_WT_ <i>randnum</i>

In these table names, *randnum* represents a random number.

Use **ML_RESTART=** to specify the name of a table to store checkpoint information. Upon restart, **ML_RESTART=** is used to specify the name of the log table that was used for storing checkpoint information in the earlier failed run.

For more information about the temporary table names that the Teradata MultiLoad utility uses, see Teradata documentation on the MultiLoad utility.

Note: Do not use **ML_RESTART=** with the **ML_LOG=** data set option. **ML_LOG=** provides a common prefix for all temporary tables that the Teradata MultiLoad utility uses. Δ

See Also

To specify a common prefix for all temporary tables that the Teradata MultiLoad utility uses, see **ML_LOG=**.

ML_WORK= Data Set Option

Specifies the name of a temporary table that MultiLoad uses to store intermediate data.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: none

Syntax

ML_WORK=*temporary-table-name*

Syntax Description

temporary-table-name

specifies the name of a temporary table that MultiLoad uses to store intermediate data that is received by the MultiLoad utility during a bulk-load operation.

Details

The Teradata MultiLoad utility uses four different temporary tables when performing the bulk-load operation. MultiLoad uses a log table to track restart information, two error tables to track errors, and a work table to hold data before the insert operation is made.

By default, the SAS/ACCESS MultiLoad facility generates names for these temporary tables:

Restart table	SAS_ML_RS_ <i>randnum</i>
Acquisition error table	SAS_ML_ET_ <i>randnum</i>
Application error table	SAS_ML_UT_ <i>randnum</i>
Work table	SAS_ML_WT_ <i>randnum</i>

In these table names, *randnum* represents a random number.

Use this option to specify the name of the table to use for storing the intermediate data that the MultiLoad utility received during a bulk-load operation. When you restart the job, use ML_WORK= to specify the name of the table for storing intermediate data during a previously failed MultiLoad job.

For more information about the temporary table names that MultiLoad uses, see Teradata documentation on the MultiLoad utility.

Note: Do not use ML_WORK= with the ML_LOG= data set option. ML_LOG= provides a common prefix for all temporary tables that the Teradata MultiLoad utility uses. △

See Also

To specify a common prefix for all temporary tables that the Teradata MultiLoad utility uses, see `ML_LOG=`.

MULTILOAD= Data Set Option

Specifies whether Teradata insert and append operations should make use of the Teradata MultiLoad utility.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: NO

Syntax

`MULTILOAD=`YES | NO

Syntax Description

YES

uses the Teradata MultiLoad utility, if available, to load Teradata tables.

NO

sends inserts to Teradata tables one row at a time.

Details

Bulk Loading

The SAS/ACCESS MultiLoad facility provides a bulk-loading method of loading both empty and existing Teradata tables. Unlike FastLoad, MultiLoad can append data to existing tables.

Data Buffers

Two data set options are available for tuning the number and the size of data buffers that are used for transferring data from SAS to Teradata. Data is transferred from SAS to Teradata using shared memory. The default shared memory buffer size is 64K. The default number of shared memory buffers used for the transfer is 2. The `BUFFERS=` data set option can be used to vary the number of buffers for data transfer from 1 to 8. The `MBUFFSIZE=` data set option can be used to vary the size of the shared memory buffers from the size of each data row up to 1MB.

Temporary Tables

The Teradata MultiLoad utility uses four different temporary tables when performing the bulk-load operation. MultiLoad uses a log table to track restart information, two error tables to track errors, and a work table to hold data before the insert operation is made.

By default, the SAS/ACCESS MultiLoad facility generates names for these temporary tables:

Restart table	SAS_ML_RS_ <i>randnum</i>
Acquisition error table	SAS_ML_ET_ <i>randnum</i>
Application error table	SAS_ML_UT_ <i>randnum</i>
Work table	SAS_ML_WT_ <i>randnum</i>

In these table names, *randnum* represents a random number.

To specify a different name for the tables listed above, use the ML_RESTART=, ML_ERROR1=, ML_ERROR2=, and ML_WORK= data set options, respectively.

You can use the ML_LOG= data set option to specify a prefix for the temporary table names that MultiLoad uses. For example, if you use **ML_LOG=MY_ERRORS**, these table names are generated:

Restart table	MY_ERRORS_RS
Acquisition error table	MY_ERRORS_ET
Application error table	MY_ERRORS_UT
Work table	MY_ERRORS_WT

Here is the order that is used for naming the temporary tables that MultiLoad uses:

- 1 If ML_LOG= is specified, the specified prefix is used when naming temporary tables for MultiLoad.
- 2 If you do not specify ML_LOG=, the values that you specified for ML_ERROR1, ML_ERROR2, ML_WORK, ML_RESTART are used.
- 3 If you do not specify any table naming options, temporary table names are generated by default.

Note: ML_LOG cannot be used with any of these options: ML_ERROR1, ML_ERROR2, ML_WORK, and ML_RESTART. Δ

Restarting MultiLoad

The MultiLoad bulk-load operation (or MultiLoad job) works in two phases: the first phase is called an acquisition phase, during which data is transferred from SAS to Teradata work tables. The second phase is called the application phase, during which data is applied to the target table.

If the MultiLoad job fails during the acquisition phase, you can restart the job from the last successful checkpoint. The exact observation from which the MultiLoad job must be restarted is displayed in the SAS log. If the MultiLoad job fails in the application phase—when data is loaded onto the target tables from the worktable—restart the MultiLoad job outside of SAS. The MultiLoad restart script displays in the SAS log. You can run the generated MultiLoad script outside of SAS to complete the load.

The ML_CHECKPOINT= data set option allows the user to specify the checkpoint rate. Specify a value for ML_CHECKPOINT= if restart capability is desired. If checkpoint tracking is not used and the MultiLoad fails in the acquisition phase, the

load needs to be restarted from the beginning. ML_CHECKPOINT=0 is the default; no checkpoints are taken if the default is used.

If ML_CHECKPOINT is between 1 and 59 inclusive, checkpoints are taken at the specified interval in minutes. If ML_CHECKPOINT is greater than or equal to 60, then a checkpoint operation occurs after a multiple of the specified rows are loaded.

Note: ML_CHECKPOINT= functions very much like the Teradata MultiLoad utility's checkpoint, differs from the DBCOMMIT= data set option. Δ

The following restrictions apply when restarting a failed MultiLoad job:

- \square The failed MultiLoad job must have specified a checkpoint rate other than zero using the ML_CHECKPOINT= data set option. Otherwise, restarting begins from the first record of the source data.

Note, however, that checkpoints are relevant only to the acquisition phase of MultiLoad. Even if ML_CHECKPOINT=0 is specified, a checkpoint takes place at the end of the acquisition phase. If the job fails after that (in the application phase) you must restart the job outside of SAS using the MultiLoad script written to the SAS log.

For example, the following MultiLoad job takes a checkpoint every 1000 records.

```
libname trlib teradata user=testuser pw=XXXXXX server=dbc;

/* Create data to MultiLoad */
data work.testdata;
  do x=1 to 50000;
    output;
  end;
end;

data trlib.mlfloat(MultiLoad=yes ML_CHECKPOINT=1000);
set work.testdata;
run;
```

- \square You must restart the failed MultiLoad job as an append process because the target table already exists. It is also necessary to identify the work tables, restart table, and the error tables used in the original job.

For example, suppose that the DATA step shown above failed with the following error message in the SAS log:

```
ERROR: MultiLoad failed with DBS error 2644 after a checkpoint was
taken for 13000 records.
Correct error and restart as an append process with dataset options
  ML_RESTART=SAS_ML_RS_1436199780, ML_ERROR1=SAS_ML_ET_1436199780,
  ML_ERROR2=SAS_ML_UT_1436199780, and ML_WORK=SAS_ML_WT_1436199780.
If the first run used FIRSTOBS=n, then use the value (7278+n-1) for FIRSTOBS
in the restart.
```

Otherwise use FIRSTOBS=7278.

Note that sometimes, the FIRSTOBS value used on the restart may be an earlier position than the last checkpoint because restart is block-oriented and not record-oriented.

After fixing the error, the job must be restarted as an append process and must specify the same work, error, and restart tables used in the earlier run. A FIRSTOBS= value is used on the source table to specify the record from which to restart.

```
/* Restart a MultiLoad job that failed in the acquisition phase
after correcting the error */
```

```
proc append data=work.testdata(FIRSTOBS=7278)
  base=trmlib.mlfloat(MultiLoad=YES ML_RESTART=SAS_ML_RS_1436199780
    ML_ERROR1=SAS_ML_ET_1436199780 ML_ERROR2=SAS_ML_UT_1436199780
    ML_WORK=SAS_ML_WT_1436199780 ML_CHECKPOINT=1000);
run;
```

- If you used ML_LOG= in the run that failed, you can specify the same value for ML_LOG= on restart. Therefore, you need not specify four data set options to identify the temporary tables that MultiLoad uses.

For example, assume that this is how the original run used ML_LOG=:

```
data trlib.mlfloat(MultiLoad=yes ML_CHECKPOINT=1000 ML_LOG=MY_ERRORS);
  set work.testdata;
run;
```

Assume that the DATA step shown above fails with the following error. The restart capability only needs data set option ML_LOG= to identify all the necessary tables.

ERROR: MultiLoad failed with DBS error 2644 after a checkpoint was taken for 13000 records. Correct error and restart as an append process with dataset options

```
ML_RESTART=SAS_ML_RS_1436199780, ML_ERROR1=SAS_ML_ET_1436199780,
ML_ERROR2=SAS_ML_UT_1436199780, and ML_WORK=SAS_ML_WT_1436199780.
```

If the first run used FIRSTOBS=n, then use the value (7278+n-1) for FIRSTOBS in the restart.

Otherwise use FIRSTOBS=7278.

Note that sometimes, the FIRSTOBS value used on the restart may be an earlier position than the last checkpoint because restart is block-oriented and not record-oriented.

```
proc append data=work.testdata(FIRSTOBS=7278)
  base=trlib.mlfloat(MultiLoad=YES ML_LOG=MY_ERRORS ML_CHECKPOINT=1000);
run;
```

- If the MultiLoad process fails in the application phase, SAS has already transferred all data to be loaded to Teradata. Restarting a MultiLoad job must be performed outside of SAS using the script written in the SAS log. See Teradata MultiLoad documentation for instructions on how to run MultiLoad scripts. Here is an example of a script that was written in the SAS log:

```
==== MultiLoad restart script starts here ====
.LOGTABLE MY_ERRORS_RS;
.LOGON boom/mloaduser,*****;
.begin import mload tables "mlfloat" CHECKPOINT 0 WORKTABLES
  MY_ERRORS_WT ERRORTABLES
  MY_ERRORS_ET MY_ERRORS_UT
/*TIFY HIGH EXIT SASMLNE.DLL TEXT '2180*/;
.layout saslayout indicators;
.FIELD "x" * FLOAT;
.DML Label SASDML;
insert into "mlfloat".*;
.IMPORT INFILE DUMMY
/*SMOD SASMLAM.DLL '2180 2180 2180 */
FORMAT UNFORMAT LAYOUT SASLAYOUT
APPLY SASDML;
.END MLOAD;
```

```
.LOGOFF;
==== MultiLoad restart script ends here ====
ERROR: MultiLoad failed with DBS error 2644 in the application phase.
      Run the MultiLoad restartscript listed above outside of SAS
      to restart the job.
```

- If the original run used a value for FIRSTOBS= for the source data, use the formula from the SAS log error message to calculate the value for FIRSTOBS= upon restart. These examples show how to do this.

```
/* Create data to MultiLoad */
data work.testdata;
  do x=1 to 50000;
    output;
  end;
run;

libname trlib teradata user=testuser pw=testpass server=boom;

/* Load 40,000 rows to the Teradata table */
data trlib.mlfloat(MultiLoad=yes ML_CHECKPOINT=1000 ML_LOG=MY_ERRORS);
set work.testdata(FIRSTOBS=10001);
run;
```

Suppose that the DATA step shown above failed with this error message:

```
ERROR: MultiLoad failed with DBS error 2644 after a checkpoint
was taken for 13000 records.
Correct the error and restart the load as an append process with
data set option ML_LOG=MY_ERRORS.
If the first run used FIRSTOBS=n, then use the value (7278+n-1)
for FIRSTOBS in the restart.
      Otherwise use FIRSTOBS=7278.
Note that sometimes, the FIRSTOBS value specified on the restart
      may be an earlier position than the last checkpoint
      because MultiLoad restart is block-oriented and not
      record-oriented.
```

The FIRSTOBS for the restart step can be calculated using the formula provided, that is, $\text{FIRSTOBS}=7278+100001-1=17278$. Use $\text{FIRSTOBS}=17278$ on the source data.

```
proc append data=work.testdata(FIRSTOBS=17278)
  base=trlib.mlfloat(MultiLoad=YES ML_LOG=MY_ERRORS ML_CHECKPOINT=1000);
run;
```

Note:

- DBCOMMIT= is disabled for MultiLoad in order to prevent any conflict with ML_CHECKPOINT=.
- ERRLIMIT= is not available for MultiLoad because the number of errors are known only at the end of each load phase.
- For restart to work correctly, the data source must return data in the same order. If the order of data that is read varies from one run to another and the load job fails in the application phase, delete temporary tables and restart the load as a new process. If the job fails in the application phase, restart the job outside of SAS as usual since the data needed to complete the load has already been transferred.

- The restart capability in MultiLoad is block-oriented, and not record-oriented. This means that while a checkpoint has been taken at, for example, 5000 records, it might be necessary to restart from an earlier record; for example, record 4000. This is because the block of data containing record 5001 might have started at record 4000. The exact record where restart should take place is displayed in the SAS log.

Δ

Examples

The following example uses MultiLoad to load SAS data to an alternate database. Note that it specifies **database=mloaduser** in the LIBNAME statement.

```
libname trlib teradata user=testuser pw=testpass server=dbc database=mloaduser;
/*MultiLoad 20000 observations into alternate database mloaduser */

data trlib.trmload14(DBCREATE_TABLE_OPTS="PRIMARY INDEX(IDNUM)" MultiLoad=yes
  ML_LOG=TRMLOAD14 ML_CHECKPOINT=5000);
  set permdata.BIG1MIL(drop=year obs=20000);
run;
```

This example extracts data from one table using FastExport and loads data into another table using MultiLoad.

```
libname trlib teradata user=testuser pw=testpass server=dbc;

/* Create data to load */
data trlib.trodd(DBCREATE_TABLE_OPTS="PRIMARY INDEX(IDNUM)" MultiLoad=yes);
  set permdata.BIG1MIL(drop=year obs=10000);
  where mod(IDNUM,2)=1;
run;

/* FastExport from one table and MultiLoad into another */
proc append data=trlib.treven(dbsliceparm=all) base=trlib.trall(MultiLOAD=YES);
run;
```

See Also

- See **SLEEP=** for information about specifying how long to wait before retrying a logon operation.
- See **TENACITY=** for information about specifying how many hours to continue to retry a logon operation.

NULLCHAR= Data Set Option

Indicates how missing SAS character values are handled during insert, update, DBINDEX=, and DBKEY= processing

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: SAS

Syntax

NULLCHAR= SAS | YES | NO

Syntax Description

SAS

indicates that missing character values in SAS data sets are treated as NULL values if the DBMS allows NULLs. Otherwise, they are treated as the NULLCHARVAL= value.

YES

indicates that missing character values in SAS data sets are treated as NULL values if the DBMS allows NULLs. Otherwise, an error is returned.

NO

indicates that missing character values in SAS data sets are treated as the NULLCHARVAL= value (regardless of whether the DBMS allows NULLs for the column).

Details

This option affects insert and update processing and also applies when you use the DBINDEX= and DBKEY= options.

This option works in conjunction with the NULLCHARVAL= data set option, which determines what is inserted when NULL values are not allowed.

All missing SAS numeric values (represented in SAS as '.') are treated by the DBMS as NULLs.

Oracle Details: See the topic on bulk loading in the documentation for the interface to Oracle for interactions between NULLCHAR and BULKLOAD=.

See Also

“NULLCHARVAL= Data Set Option” on page 245

“DBNULL= Data Set Option” on page 212

NULLCHARVAL= Data Set Option

Defines the character string that replaces missing SAS character values during insert, update, DBINDEX=, and DBKEY= processing

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: a blank character

Syntax

NULLCHARVAL=*'character-string'*

Details

This option affects

insert and update processing and also applies when you use the DBINDEX= and DBKEY= options.

This option works with the NULLCHAR= option, which determines whether a missing SAS character value is treated as a NULL value.

If NULLCHARVAL= is longer than the maximum column width, one of the following occurs:

- The string is truncated if DBFORCE=YES.
- The operation fails if DBFORCE=NO.

See Also

“NULLCHAR= Data Set Option” on page 244

“DBFORCE= Data Set Option” on page 202

“DBNULL= Data Set Option” on page 212

OR_PARTITION= Data Set Option

Allows reading, updating, and deleting from a particular partition in a partitioned table, also inserting and bulk loading into a particular partition in a partitioned table

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: none

Syntax

OR_PARTITION =*name of a partition in a partitioned Oracle table*

Syntax Description

name of a partition in a partitioned Oracle table

The name of the partition must be valid or an error occurs.

Details

Use this option in cases where you are working with only one particular partition at a time in a partitioned table. Specifying this option boost performances because you are limiting your access to only one partition of a table instead of the entire table.

This option is appropriate when reading, updating, and deleting from a partitioned table, also when inserting into a partitioned table or bulk loading to a table. You can use it to boost performance.

Example

The following example shows one way you can use this option.

```
libname x oracle user=scott pw=tiger path=oraclev9;

proc delete data=x.orparttest; run;
data x.ORparttest ( dbtype=(NUM='int')
  DBCREATE_TABLE_OPTS='partition by range (NUM)
    (partition p1 values less than (11),
     partition p2 values less than (21),
     partition p3 values less than (31),
     partition p4 values less than (41),
     partition p5 values less than (51),
     partition p6 values less than (61),
     partition p7 values less than (71),
     partition p8 values less than (81)
    )' );
do i=1 to 80;
  NUM=i;
output;
end;
run;

options sastrace=",,t,d" sastraceloc=saslog nostsuffix;

/* input */
proc print data=x.orparttest ( or_partition=p4 );
run;

/* update */
proc sql;

/* update should fail with 14402, 00000, "updating partition key column would
cause a partition change"
// *Cause: An UPDATE statement attempted to change the value of a partition
//          key column causing migration of the row to another partition
// *Action: Do not attempt to update a partition key column or make sure that
```

```

//          the new partition key is within the range containing the old
//          partition key.
*/
update x.orparttest ( or_partition=p4 ) set num=100;

update x.orparttest ( or_partition=p4 ) set num=35;

select * from x.orparttest ( or_partition=p4 );
select * from x.orparttest ( or_partition=p8 );

/* delete */
delete from x.orparttest ( or_partition=p4 );

select * from x.orparttest;
quit;

/* load to an existing table */
data new; do i=31 to 39; num=i; output;end;
run;
data new2; do i=1 to 9; num=i; output;end;
run;

proc append base= x.orparttest ( or_partition=p4 ) data= new;
run;

/* insert should fail 14401, 00000, "inserted partition key is outside
specified partition"
// *Cause: the concatenated partition key of an inserted record is outside
//          the ranges of the two concatenated partition bound lists that
//          delimit the partition named in the INSERT statement
// *Action: do not insert the key or insert it in another partition
*/
proc append base= x.orparttest ( or_partition=p4 ) data= new2;
run;

/* load to an existing table */
proc append base= x.orparttest ( or_partition=p4 bulkload=yes
bl_load_method=truncate ) data= new;
run;

/* insert should fail 14401 */
proc append base= x.orparttest ( or_partition=p4 bulkload=yes
bl_load_method=truncate ) data= new2;
run;

```

Below are a series of sample scenarios that illustrate how you can use this option. The first shows how to create the ORPARTTEST table, on which all remaining examples depend.

```
libname x oracle user=scott pw=tiger path=oraclev9;
proc delete data=x.orparttest; run;
data x.ORparttest ( dbtype=(NUM='int')
  DBCREATE_TABLE_OPTS='partition by range (NUM)
    (partition p1 values less than (11),
     partition p2 values less than (21),
     partition p3 values less than (31),
     partition p4 values less than (41),
     partition p5 values less than (51),
     partition p6 values less than (61),
     partition p7 values less than (71),
     partition p8 values less than (81)
    )' );
do i=1 to 80;
  NUM=i; output;
end;
run;
```

Only the P4 partition is read in the following example.

```
proc print data=x.orparttest ( or_partition=p4 );
run;
```

Rows that belong to only the single P4 partition are updated in the following example.

```
proc sql;
update x.orparttest ( or_partition=p4 ) set num=35;
quit;
```

The above example also illustrates how a particular partition can be updated. However, updates and even inserts to the partition key column are done in such a way that it must be migrated to a different partition in the table. Therefore, the following example fails because the value 100 does not belong to the P4 partition.

```
proc sql;
update x.orparttest ( or_partition=p4 ) set num=100;
quit;
```

All rows in the P4 partition are deleted in the following example.

```
proc sql;
delete from x.orparttest ( or_partition=p4 );
quit;
```

In this next example, rows are added to the P4 partition in the table.

```
data new;
  do i=31 to 39; num=i; output;end;
run;
proc append base= x.orparttest ( or_partition=p4 );
  data= new;
run;
```

The following example also adds rows to the P4 partition but uses the SQL*Loader instead.

```
proc append base= x.orparttest ( or_partition=p4 bulkload=yes );
    data= new;
run;
```

OR_UPD_NOWHERE= Data Set Option

Specifies whether SAS uses an extra WHERE clause when updating rows with no locking

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Alias: ORACLE_73_OR_ABOVE=

Default value: LIBNAME setting

Syntax

OR_UPD_NOWHERE =YES | NO

Syntax Description

YES

SAS does not use an additional WHERE clause to determine whether each row has changed since it was read. Instead, SAS uses the SERIALIZABLE isolation level (available with Oracle 7.3 and later) for update locking. If a row changes after the serializable transaction starts, the update on that row fails.

NO

SAS uses an additional WHERE clause to determine whether each row has changed since it was read. If a row has changed since being read, the update fails.

Details

Use this option when you are updating rows without locking (UPDATE_LOCK_TYPE=NOLOCK).

By default (OR_UPD_NOWHERE=YES), updates are performed in serializable transactions so you can avoid extra WHERE clause processing and potential WHERE clause floating point precision problems.

Specify OR_UPD_NOWHERE=NO for compatibility when you are updating a SAS Version 6 view descriptor.

Note: Due to the published Oracle bug 440366, sometimes an update on a row fails even if the row has not changed. Oracle offers the following solution: When creating a table, increase the number of INITRANS to at least 3 for the table. △

Example

In the following example, you create a small Oracle table called TEST and then update the TEST table once using the default setting (OR_UPD_NOWHERE=YES) and once specifying OR_UPD_NOWHERE=NO.

```
libname oralib oracle user=testuser pw=testpass update_lock_type=no;

data oralib.test;
  c1=1;
  c2=2;
  c3=3;
run;

options sastrace=",,,d" sastraceloc=saslog;

proc sql;
  update oralib.test set c2=22;
  update oralib.test(or_upd_nowhere=no) set c2=222;
quit;
```

This code uses the SASTRACELOC= and SASTRACE= options to send the output to the SAS log.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “OR_UPD_NOWHERE= LIBNAME Option” on page 128.

ORHINTS= Data Set Option

Specifies Oracle hints to pass to Oracle from a SAS statement or SQL procedure

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: none

Syntax

ORHINTS = '*Oracle-hint*'

Syntax Description

Oracle-hint

specifies an Oracle hint for SAS/ACCESS to pass to the DBMS as part of an SQL query.

Details

If you specify an Oracle hint, SAS passes the hint to Oracle. If you omit ORHINTS=, SAS does not send any hints to Oracle.

Examples

The following example runs a SAS procedure on DBMS data and SAS converts the procedure to one or more SQL queries. You can use this option to specify an Oracle hint for SAS to pass as part of the SQL query.

```
libname mydblib oracle user=testuser password=testpass path='myorapath';

proc print data=mydblib.payroll(orphints='/*+ ALL_ROWS */');
run;
```

In the following example, SAS sends the Oracle hint '/*+ ALL_ROWS */' to Oracle as part of the following statement:

```
SELECT /*+ ALL_ROWS */ * FROM PAYROLL
```

PRESERVE_COL_NAMES= Data Set Option

Preserves spaces, special characters, and case sensitivity in DBMS column names when you create DBMS tables.

Valid in: DATA and PROC steps (when creating DBMS tables using SAS/ACCESS software).

DBMS support: DB2 UNIX/PC, DB2 z/OS, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Teradata

Alias: PRESERVE_NAMES= (see Details)

Default value: LIBNAME setting

Syntax

PRESERVE_COL_NAMES=NO | YES

Syntax Description

NO

specifies that column names that are used in DBMS table creation are derived from SAS variable names by using the SAS variable name normalization rules. (See the VALIDVARNAME= system option for more information.) However, the database applies its DBMS-specific normalization rules to the SAS variable names when it creates the DBMS column names.

The use of name literals to create column names that use database keywords or special symbols other than the underscore character might be illegal when DBMS normalization rules are applied. To include nonstandard SAS symbols or database keywords, specify PRESERVE_COL_NAMES=YES.

YES

specifies that column names that are used in table creation are passed to the DBMS with special characters and the exact, case-sensitive spelling of the name preserved.

Details

This option applies only when you use SAS/ACCESS to create a new DBMS table. When you create a table, you assign the column names by using one of the following methods:

- To control the case of the DBMS column names, specify variables with the desired case and set PRESERVE_COL_NAMES=YES. If you use special symbols or blanks, you must set VALIDVARNAME=ANY and use name literals. For more information, see the section about names in *SAS/ACCESS for Relational Databases: Reference* and the section about system options in *SAS Language Reference: Dictionary*.
- To enable the DBMS to normalize the column names according to its naming conventions, specify variables with any case and set PRESERVE_COLUMN_NAMES=NO.

Note: When you use SAS/ACCESS to read from, insert rows into, or modify data in an existing DBMS table, SAS identifies the database column names by their spelling. Therefore, when the database column exists, the case of the variable does not matter. \triangle

For more information, see the topic about naming in the documentation for your SAS/ACCESS interface.

To save some time when coding, specify the PRESERVE_NAMES= alias if you plan to specify both the PRESERVE_COL_NAMES= and PRESERVE_TAB_NAMES= options in your LIBNAME statement.

To use column names in your SAS program that are not valid SAS names, you must use one of the following techniques:

- Use the DQUOTE= option in PROC SQL and then reference your columns using double quotation marks. For example:

```
proc sql dquote=ansi;
  select "Total$Cost" from mydblib.mytable;
```

- Specify the global system option VALIDVARNAME= ANY and use name literals in the SAS language. For example:

```
proc print data=mydblib.mytable;
  format 'Total$Cost'n 22.2;
```

Note that if you are *creating* a table in PROC SQL, you must also include the PRESERVE_COL_NAMES=YES option. For example:

```
libname mydblib oracle user=testuser password=testpass;
proc sql dquote=ansi;
  create table mydblib.mytable (preserve_col_names=yes) ("my$column" int);
```

PRESERVE_COL_NAMES= does not apply to the Pass-Through Facility.

See Also

To assign this option to a *group* of relational DBMS tables or views, see “PRESERVE_COL_NAMES= LIBNAME Option” on page 130 naming in your interface.

QUALIFIER= Data Set Option

Specifies the qualifier to use when you are reading database objects, such as DBMS tables and views

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: HP Neoview, MySQL, Microsoft SQL Server, Netezza, ODBC, OLE DB

Default value: LIBNAME setting

Syntax

QUALIFIER=<qualifier-name>

Details

If this option is omitted, the default qualifier name, if any, is used for the data source. QUALIFIER= can be used for any data source, such as a DBMS object, that allows three-part identifier names: *qualifier.schema.object*.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “QUALIFIER= LIBNAME Option” on page 133.

QUERY_TIMEOUT= Data Set Option

Specifies the number of seconds of inactivity to wait before canceling a query

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, Netezza, ODBC

Default value: LIBNAME setting

Syntax

QUERY_TIMEOUT=*number-of-seconds*

Details

QUERY_TIMEOUT= 0 indicates that there is no time limit for a query. This option is useful when you are testing a query, you suspect that a query might contain an endless loop, or the data is locked by another user.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “QUERY_TIMEOUT= LIBNAME Option” on page 135.

READ_ISOLATION_LEVEL= Data Set Option

Specifies which level of read isolation locking to use when you are reading data

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, Microsoft SQL Server, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

READ_ISOLATION_LEVEL=*DBMS-specific-value*

Syntax Description

dbms-specific-value

See the documentation for your SAS/ACCESS interface for the values for your DBMS.

Details

In the interfaces to ODBC and DB2 under UNIX and PC hosts, this option is ignored if READ_LOCK_TYPE= is not set to ROW.

The degree of isolation defines the following:

- the degree to which rows that are read and updated by the current application are available to other concurrently executing applications
- the degree to which update activity of other concurrently executing application processes can affect the current application.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “READ_ISOLATION_LEVEL= LIBNAME Option” on page 137.

READ_LOCK_TYPE= Data Set Option

Specifies how data in a DBMS table is locked during a read transaction

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, Microsoft SQL Server, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

READ_LOCK_TYPE=ROW | PAGE | TABLE | NOLOCK | VIEW

Syntax Description

Not all values are valid for every interface. See details below.

ROW

locks a row if any of its columns are accessed. (This value is valid in the DB2 under UNIX and PC hosts, Microsoft SQL Server, ODBC, and Oracle interfaces.)

PAGE

locks a page of data, which is a DBMS-specific number of bytes. (This value is valid in the Sybase interface.)

TABLE

locks the entire DBMS table. If you specify READ_LOCK_TYPE=TABLE, you must also specify the LIBNAME option CONNECTION=UNIQUE, or you receive an error message. Setting CONNECTION=UNIQUE ensures that your table lock is not lost, for example, due to another table closing and committing rows in the same connection. (This value is valid in the DB2 under z/OS, DB2 under UNIX and PC hosts, ODBC, Oracle, Microsoft SQL Server, and Teradata interfaces.)

NOLOCK

does not lock the DBMS table, pages, or any rows during a read transaction. (This value is valid in the Oracle and Sybase interfaces and in the ODBC and Microsoft SQL Server interfaces when you use the Microsoft SQL Server driver.)

VIEW

locks the entire DBMS view. (This value is valid in the Teradata interface.)

Details

If you omit READ_LOCK_TYPE=, you get either the default action for the DBMS that you are using, or a lock for the DBMS that was set with the LIBNAME statement.

See the documentation for your SAS/ACCESS interface for additional details.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “READ_LOCK_TYPE= LIBNAME Option” on page 138.

READ_MODE_WAIT= Data Set Option

Specifies during SAS/ACCESS read operations whether Teradata waits to acquire a lock or fails your request when the DBMS resource is locked by a different user

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Teradata

Default value: LIBNAME setting

Syntax

READ_MODE_WAIT= YES | NO

Syntax Description

YES

specifies that Teradata waits to acquire the lock, and SAS/ACCESS waits indefinitely until it can acquire the lock.

NO

specifies that Teradata fails the lock request if the specified DBMS resource is locked.

Details

If you specify READ_MODE_WAIT=NO, and a different user holds a *restrictive* lock, then the executing SAS step fails. SAS/ACCESS continues to process the job by executing the next step. If you specify READ_MODE_WAIT=YES, SAS/ACCESS waits indefinitely until it can acquire the lock.

A *restrictive* lock means that another user is holding a lock that prevents you from obtaining your desired lock. Until the other user releases the restrictive lock, you cannot obtain your lock. For example, another user's table-level WRITE lock prevents you from obtaining a READ lock on the table.

For more information, see the SAS/ACCESS documentation for Teradata.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option "READ_MODE_WAIT= LIBNAME Option" on page 139.

READBUFF= Data Set Option

Specifies the number of rows of DBMS data to read into the buffer

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle, Sybase

Alias: ROWSET_SIZE= (DB2 UNIX/PC, HP Neoview, Microsoft SQL Server, Netezza, ODBC, OLE DB)

Default value: LIBNAME setting

Syntax

READBUFF=*integer*

Syntax Description

integer

is the maximum value that is allowed by the DBMS.

Details

This option improves performance by specifying a number of rows that can be held in memory for input into SAS. Buffering data reads can decrease network activities and increase performance. However, because SAS stores the rows in memory, higher values for READBUFF= use more memory. In addition, if too many rows are selected at once, then the rows that are returned to the SAS application might be out of date.

When READBUFF=1, only one row is retrieved at a time. The higher the value for READBUFF=, the more rows the SAS/ACCESS engine retrieves in one fetch operation.

DB2 UNIX/PC Details: By default, the SQLFetch API call is used and no internal SAS buffering is performed. Setting READBUFF=1 or greater causes the SQLExtendedFetch API call to be used.

HP Neoview, Microsoft SQL Server, Netezza, and ODBC Details: By default, the SQLFetch API call is used and no internal SAS buffering is performed. Setting READBUFF=1 or greater causes the SQLExtendedFetch API call to be used.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “READBUFF= LIBNAME Option” on page 140.

SASDATEFMT= Data Set Option

Changes the SAS date format of a DBMS column

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Informix, Microsoft SQL Server, MySQL, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: DBMS-specific

Syntax

```
SASDATEFMT=(DBMS-date-col-1='SAS-date-format'
<... DBMS-date-col-n='SAS-date-format'>)
```

Syntax Description

DBMS-date-col

specifies the name of a date column in a DBMS table.

SAS-date-format

specifies a SAS date format that has an equivalent (like-named) informat. For example, DATETIME21.2 is both a SAS format and a SAS informat, so it is a valid value for the *SAS-date-format* argument.

Details

If the SAS column date format does not match the date format of the corresponding DBMS column, convert the SAS date values to the appropriate DBMS date values. Use

the SASDATEFMT= option to convert date values from the default SAS date format to another SAS date format that you specify.

Use the SASDATEFMT= option to prevent date type mismatches in the following circumstances:

- during input operations to convert DBMS date values to the correct SAS DATE, TIME, or DATETIME values
- during output operations to convert SAS DATE, TIME, or DATETIME values to the correct DBMS date values.

The column names specified in this option must be DATE, DATETIME, or TIME columns; columns of any other type are ignored.

The format specified must be a valid date format; output with any other format is unpredictable.

If the SAS date format and the DBMS date format match, this option is not needed.

The default SAS date format is DBMS-specific and is determined by the data type of the DBMS column. See the documentation for your SAS/ACCESS interface.

Note: For non-English date types, SAS automatically converts the data to the SAS type of NUMBER. The SASDATEFMT= option does not currently handle these date types. However, you can use a PROC SQL view to convert the DBMS data to a SAS date format as you retrieve the data or use a format statement in other contexts. Δ

Oracle details: It is recommended that “DBSASTYPE= Data Set Option” on page 216 be used instead of SASDATEFMT=.

Examples

In the following example, the APPEND procedure adds SAS data from the SASLIB.DELAY data set to the Oracle table that is accessed by MYDBLIB.INTERNAT. Using SASDATEFMT=, the default SAS format for the Oracle column DATES is changed to the DATE9. format. Data output from SASLIB.DELAY into the DATES column in MYDBLIB.INTERNAT now converts from the DATE9. format to the Oracle format assigned to that type.

```
libname mydblib oracle user=testuser password=testpass;
libname saslib 'your-SAS-library';

proc append base=mydblib.internat(sasdatefmt=(dates='date9.'))force
  data=saslib.delay;
run;
```

In the following example, SASDATEFMT= converts DATE1, a SAS DATETIME value, to a Teradata date column named DATE1.

```
libname x teradata user=testuser password=testpass;

proc sql noerrorstop;
  create table x.dateinfo ( date1 date );
  insert into x.dateinfo
  ( sasdatefmt=( date1='datetime21.' )
  values ( '31dec2000:01:02:30'dt );
```

In the following example, SASDATEFMT= converts DATE1, a Teradata date column, to a SAS DATETIME type named DATE1.

```
libname x teradata user=testuser password=testpass;

data sas_local;
format date1 datetime21.;
set x.dateinfo( sasdatefmt=( date1='datetime21.' ) );
run;
```

SCHEMA= Data Set Option

Lets you read a data source, such as a DBMS table and view, in the specified schema

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, HP Neoview, Informix, Microsoft SQL Server, Netezza, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: LIBNAME setting

Syntax

SCHEMA=*schema-name*

Syntax Description

schema-name

is the name assigned to a logical classification of objects in a relational database.

Details

For this option to work, you must have appropriate privileges to the schema that is specified.

The values for SCHEMA= are usually case-sensitive, so be careful when you specify this option.

Oracle Details: If PRESERVE_TAB_NAMES=NO, SAS converts the SCHEMA= value to uppercase because all values in the Oracle data dictionary are converted to uppercase unless quoted.

Teradata Details: If you omit this option, a libref points to your default Teradata database, which often has the same name as your user name. You can use this option to point to a different database. You can use this option to view or modify a different user's DBMS tables or views if you have the required Teradata privileges. (For example, to read another user's tables, you must have the Teradata privilege SELECT for that user's tables.) For more information about changing the default database, see the DATABASE statement in your Teradata documentation.

Examples

In the following example, SCHEMA= causes MYDB.TEMP_EMPS to be interpreted by DB2 as SCOTT.TEMP_EMPS.

```
proc print data=mydb.temp_emps
  schema=SCOTT;
run;
```

In the following example, SAS sends any reference to Employees as Scott.Employees.

```
libname mydblib oracle user=testuser password=testpass path="myorapath";

proc print data=employees (schema=scott);
run;
```

In the following example, user TESTUSER prints the contents of the Employees table, which is located in the Donna database.

```
libname mydblib teradata user=testuser pw=testpass;

proc print data=mydblib.employees(schema=donna);
run;
```

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “SCHEMA= LIBNAME Option” on page 143.

SEGMENT_NAME= Data Set Option

Lets you control the segment in which you create a table

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Sybase

Default value: none

Syntax

SEGMENT_NAME=*segment-name*

Syntax Description

segment-name

specifies the name of the segment in which to create a table.

SLEEP= Data Set Option

Specifies the number of minutes that MultiLoad waits before it retries logging in to Teradata.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: 6

Syntax

SLEEP=*number-of-minutes*

Syntax Description

number-of-minutes

the number of minutes that MultiLoad waits before it retries logging on to Teradata.

Details

Use the SLEEP= data set option to indicate to MultiLoad how long to wait before it retries logging on to Teradata when the maximum number of utilities are already running. (The maximum number of Teradata utilities that can run concurrently varies from 5 to 15, depending upon the database server setting.) The default value for SLEEP= is 6 minutes. The value specified for SLEEP must be greater than 0.

SLEEP= is used in conjunction with TENACITY=, which specifies the time in hours that MultiLoad must continue to retry the logon operation. SLEEP= and TENACITY= function very much like the SLEEP and TENACITY run-time options of the native Teradata MultiLoad utility.

See Also

See TENACITY= for information about specifying how long to continue to retry a logon operation.

TENACITY= Data Set Option

Specifies how many hours MultiLoad continues to retry logging on to Teradata if the maximum number of Teradata utilities are already running.

Valid in: DATA and PROC steps (when creating and appending to DBMS tables using SAS/ACCESS software)

DBMS support: Teradata

Default value: 4

Syntax

TENACITY=*number-of-hours*

Syntax Description

number-of-hours

the number of hours to continue to retry logging on to Teradata.

Details

Use the TENACITY= data set option to indicate to MultiLoad how long to continue retrying a logon operation when the maximum number of utilities are already running. (The maximum number of Teradata utilities that can run concurrently varies from 5 to 15, depending upon the database server setting.) The default value for TENACITY= is four hours. The value specified for TENACITY= must be greater than zero.

Use TENACITY= with SLEEP=, which specifies the number of minutes that MultiLoad waits before it retries logging on to Teradata. SLEEP= and TENACITY= function very much like the SLEEP and TENACITY run-time options of the native Teradata MultiLoad utility.

The following message is written to the SAS log if the time period that TENACITY= specifies is exceeded.

```
ERROR: MultiLoad failed unexpectedly with returncode 12
```

Note: Check the MultiLoad log for more information about the cause of the MultiLoad failure. SAS does not receive any informational messages from Teradata in either of these situations:

- when the currently run MultiLoad process waits because the maximum number of utilities are already running
- if MultiLoad is terminated because the time limit that TENACITY= specifies has been exceeded

The native Teradata MultiLoad utility sends messages associated with SLEEP and TENACITY only to the MultiLoad log. So nothing is written to the SAS log. Δ

See Also

See SLEEP= for information about specifying how long to wait before retrying a logon operation.

TRAP151= Data Set Option

Enables columns that cannot be updated to be removed from a FOR UPDATE OF clause so updating of columns can proceed as normal

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 z/OS

Default value: NO

Syntax

TRAP151=YES | NO

Syntax Description

YES

removes the non-updatable column that is designated in the error-151 and reprepares the statement for processing. This process is repeated until all columns that cannot be updated are removed, and all remaining columns can be updated.

NO

disables TRAP151=. TRAP151= is disabled by default. It is not necessary to specify NO.

Examples

In the following example, DB2DEBUG is turned on so that you can see what occurs when TRAP151=YES:

Output 10.1 SAS Log for TRAP151=YES

```
proc fsedit data=x.v4(trap151=yes);
run;
SELECT * FROM V4 FOR FETCH ONLY
SELECT * FROM V4 FOR FETCH ONLY
SELECT "A","X","Y","B","Z","C" FROM V4 FOR UPDATE OF "A","X","Y","B","Z","C"
DB2 SQL Error, sqlca->sqlcode=-151
WARNING: SQLCODE -151: reparing SELECT as:
  SELECT "A","X","Y","B","Z","C" FROM V4 FOR UPDATE OF "A","Y","B","Z","C"
DB2 SQL Error, sqlca->sqlcode=-151
WARNING: SQLCODE -151: reparing SELECT as:
  SELECT "A","X","Y","B","Z","C" FROM V4 FOR UPDATE OF "A","B","Z","C"
DB2 SQL Error, sqlca->sqlcode=-151
WARNING: SQLCODE -151: reparing SELECT as:
  SELECT "A","X","Y","B","Z","C" FROM V4 FOR UPDATE OF "A","B","C"
COMMIT WORK
NOTE: The PROCEDURE FSEDIT used 0.13 CPU seconds and 14367K.
```

The following example features the same code with TRAP151 turned off:

Output 10.2 SAS Log for TRAP151=NO

```
proc fsedit data=x.v4(trap151=no);
run;
SELECT * FROM V4 FOR FETCH ONLY
SELECT * FROM V4 FOR FETCH ONLY
SELECT "A","X","Y","B","Z","C" FROM V4 FOR UPDATE OF "A","X","Y","B","Z","C"
DB2 SQL Error, sqlca->sqlcode=-151
ERROR: DB2 prepare error; DSNT4081 SQLCODE= ---151, ERROR;
      THE UPDATE STATEMENT IS INVALID BECAUSE THE CATALOG DESCRIPTION OF COLUMN C
      INDICATES THAT IT CANNOT BE UPDATED.
COMMIT WORK
NOTE: The SAS System stopped processing this step because of errors.
NOTE: The PROCEDURE FSEDIT used 0.08 CPU seconds and 14367K.
```

UPDATE_ISOLATION_LEVEL= Data Set Option

Defines the degree of isolation of the current application process from other concurrently running application processes

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, Microsoft SQL Server, MySQL, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: LIBNAME setting

Syntax

UPDATE_ISOLATION_LEVEL=*DBMS-specific-value*

Syntax Description

dbms-specific-value

See the documentation for your SAS/ACCESS interface for the values for your DBMS.

Details

The degree of isolation identifies the following:

- the degree to which rows that are read and updated by the current application are available to other concurrently executing applications
- the degree to which update activity of other concurrently executing application processes can affect the current application.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “UPDATE_ISOLATION_LEVEL= LIBNAME Option” on page 150.

UPDATE_LOCK_TYPE= Data Set Option

Specifies how data in a DBMS table is locked during an update transaction

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: DB2 UNIX/PC, DB2 z/OS, Microsoft SQL Server, ODBC, OLE DB, Oracle, Sybase, Teradata

Default value: LIBNAME setting

Syntax

UPDATE_LOCK_TYPE=ROW | PAGE | TABLE | NOLOCK | VIEW

Syntax Description

Not all values are valid for every interface. See details below.

ROW

locks a row if any of its columns are going to be updated. (This is valid in the DB2 under UNIX and PC hosts, Microsoft SQL Server, ODBC, OLE DB, and Oracle interfaces.)

PAGE

locks a page of data, which is a DBMS-specific number of bytes. (This is valid in the Sybase interface.)

TABLE

locks the entire DBMS table. (This is valid in the DB2 under z/OS, DB2 under UNIX and PC hosts, ODBC, Oracle, Microsoft SQL Server, and Teradata interfaces.)

NOLOCK

does not lock the DBMS table, page, or any rows when reading them for update. (This is valid in the Microsoft SQL Server, ODBC, Oracle, and Sybase interfaces.)

VIEW

locks the entire DBMS view. (This is valid in the Teradata interface.)

Details

If you omit UPDATE_LOCK_TYPE=, you get either the default action for the DBMS that you are using, or a lock for the DBMS that was set with the LIBNAME statement. You can set a lock for one DBMS table by using the data set option or for a group of DBMS tables by using the LIBNAME option.

For more information, see the documentation for your SAS/ACCESS interface.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “UPDATE_LOCK_TYPE= LIBNAME Option” on page 151.

UPDATE_MODE_WAIT= Data Set Option

Specifies during SAS/ACCESS update operations whether the DBMS waits to acquire a lock or fails your request when the DBMS resource is locked by a different user

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Teradata

Default value: LIBNAME setting

Syntax

UPDATE_MODE_WAIT= YES|NO

Syntax Description

YES

specifies that Teradata waits to acquire the lock, so SAS/ACCESS waits indefinitely until it can acquire the lock.

NO

specifies that Teradata fails the lock request if the specified DBMS resource is locked.

Details

If you specify UPDATE_MODE_WAIT=NO and if a different user holds a *restrictive* lock, then your SAS step fails and SAS/ACCESS continues the job by processing the next step. If you specify UPDATE_MODE_WAIT=YES, SAS/ACCESS waits indefinitely until it can acquire the lock.

A *restrictive* lock means that a different user is holding a lock that prevents you from obtaining your desired lock. Until the other user releases the restrictive lock, you cannot obtain your lock. For example, another user's table-level WRITE lock prevents you from obtaining a READ lock on the table.

Use SAS/ACCESS locking options only when Teradata's standard locking is undesirable.

For more information, see the documentation for the interface to Teradata.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option "UPDATE_MODE_WAIT= LIBNAME Option" on page 152.

UPDATE_SQL= Data Set Option

Determines the method that is used to update and delete rows in a data source

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Microsoft SQL Server, ODBC

Default value: LIBNAME setting

Syntax

UPDATE_SQL=YES | NO

Syntax Description

YES

specifies that SAS/ACCESS uses Current-of-Cursor SQL to update or delete rows in a table.

NO

specifies that SAS/ACCESS uses the SQLSetPos() API to update or delete rows in a table.

Details

This is the update and delete equivalent of the INSERT_SQL= option.

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “UPDATE_SQL= LIBNAME Option” on page 154.

UPDATEBUFF= Data Set Option

Specifies the number of rows that are processed in a single DBMS update or delete operation

Valid in: DATA and PROC steps (when accessing DBMS data using SAS/ACCESS software)

DBMS support: Oracle

Default value: LIBNAME setting

Syntax

UPDATEBUFF=*positive-integer*

Syntax Description

positive-integer


is the maximum value that is allowed by the DBMS.

Details

When updating with the VIEWTABLE window or PROC FSVIEW, use UPDATEBUFF=1 to prevent the DBMS interface from trying to update multiple rows. By default, these features update only one observation at a time (since by default they use record-level locking, they lock only the observation that is currently being edited).

See Also

To assign this option to a *group* of relational DBMS tables or views, see the LIBNAME option “UPDATEBUFF= LIBNAME Option” on page 154.



CHAPTER

11

Macro Variables and System Options for Relational Databases

Introduction to Macro Variables and System Options for Relational Databases 269

Macro Variables for Relational Databases 269

System Options for Relational Databases 271

Introduction to Macro Variables and System Options for Relational Databases

This section describes the system options and macro variables that are available to use with SAS/ACCESS software. It describes only those components of the macro facility that depend on the SAS/ACCESS engines. Most features of the SAS macro facility are portable. For more information, refer to the *SAS Macro Language: Reference* and the SAS Help for the macro facility.

Macro Variables for Relational Databases

SYSDBMSG, SYSDBRC, SQLXMSG, and SQLXRC are automatic SAS macro variables. The SAS/ACCESS engine and your DBMS determine their values. Initially, SYSDBMSG and SQLXMSG are blank, and SYSDBRC and SQLXRC are set to 0.

SAS/ACCESS generates several return codes and error messages while it processes your programs. This information is available to you through these SAS macro variables:

SYSDBMSG

contains DBMS-specific error messages that are generated when you use SAS/ACCESS software to access your DBMS data.

SYSDBRC

contains DBMS-specific error codes that are generated when you use SAS/ACCESS software to access your DBMS data. Error codes that are returned are text, not numbers.

You can use these variables anywhere while you are accessing DBMS data. Only one set of macro variables is provided. It is possible that, if tables from two different DBMSs are accessed, it might not be clear from which DBMS the error message originated. To address this problem, the name of the DBMS is inserted at the beginning of the SYSDBMSG macro variable message or value. The contents of the SYSDBMSG and SYSDBRC macro variables can be printed in the SAS log by using the %PUT macro. They are reset after each SAS/ACCESS LIBNAME statement, DATA step, or procedure is executed. In the following statement, %SUPERQ masks special characters

such as `&`, `%`, and any unbalanced parentheses or quotation marks that might exist in the text stored in the `SYSDBMSG` macro.

```
%put %superq(SYSDBMSG)
```

These special characters can cause unpredictable results if you use this statement:

```
%put &SYSDBMSG
```

It is more advantageous to use `%SUPERQ`.

If you try to connect to Oracle and use the incorrect password, you receive the messages shown in this output.

Output 11.1 SAS Log for an Oracle Error

```
2? libname mydblib oracle user=pierre pass=paris path="orav7";
ERROR: Oracle error trying to establish connection. Oracle error is
      ORA-01017: invalid username/password; logon denied
ERROR: Error in the LIBNAME or FILENAME statement.
3? %put %superq(sysdbmsg);

Oracle: ORA-01017: invalid username/passsword; logon denied
4? %put &sysdbrc;

-1017
5?
```

You can also use `SYMGET` to retrieve error messages:

```
msg=symget("SYSDBMSG");
```

For example:

```
data_null_;
msg=symget("SYSDBMSG");
put msg;
run;
```

The Pass-Through Facility generates return codes and error messages that are available to you through these SAS macro variables:

SQLXMSG

contains DBMS-specific error messages.

SQLXRC

contains DBMS-specific error codes.

You can use `SQLXMSG` and `SQLXRC` only through explicit pass-through with the Pass-Through Facility. See Return Codes“Return Codes” on page 286.

You can print the contents of `SQLXMSG` and `SQLXRC` in the SAS log by using the `%PUT` macro. `SQLXMSG` is reset to a blank string, and `SQLXRC` is reset to 0 when any Pass-Through Facility statement is executed.

System Options for Relational Databases

DBSRVTP= and DBSLICEPARM= are SAS system options for databases. SASTRACE=, SASTRACELOC=, and VALIDVARNAME= are SAS system options that have applications specific to SAS/ACCESS.

Note: The SAS system option REPLACE= is not supported by the SAS/ACCESS interfaces. Δ

DBSRVTP= System Option

Specifies whether SAS/ACCESS engines put a hold (or block) on the originating client while making performance-critical calls to the database. This option applies when SAS is invoked as a server responding to multiple clients

Valid in: SAS invocation

Default value: NONE

Syntax

DBSRVTP= 'ALL' | 'NONE' | '(engine-name(s))

Syntax Description

ALL

indicates that SAS does not use any blocking operations for all underlying SAS/ACCESS products that support this option.

NONE

indicates that SAS uses standard blocking operations for all SAS/ACCESS products.

engine-name(s)

indicates that SAS does not use any blocking operations for the specified SAS/ACCESS interface(s). You can specify one or more of the following engine names. If you specify more than one name, separate them with blank spaces and enclose the list in parentheses.

db2 (only supported under UNIX and PC hosts)

informix

odbc (indicates that SAS uses nonblocking operations for the SAS/ACCESS interfaces to ODBC and Microsoft SQL Server).

oledb

Oracle

sybase

teradata (not supported on OS/390)

Details

This option can be used to help throughput of the SAS server by supporting multiple simultaneous execution streams, if the server uses certain SAS/ACCESS interfaces.

Better throughput is accomplished when the underlying SAS/ACCESS engine does not hold or block the originating client, such that any one client using a SAS/ACCESS product does not keep the SAS server from responding to other clients' requests. SAS/SHARE software and SAS Integration Technologies are two ways of invoking SAS as a server.

This option is a system invocation option, which means the value is set when SAS is invoked. Because the DBSRVTP= option uses multiple native threads, enabling this option uses the underlying DBMS's threading support. Some databases handle threading better than others, so you might want to invoke DBSRVTP= for some DBMSs and not others. Refer to your documentation for your DBMS for more information.

The option accepts a string where values are the engine name of a SAS/ACCESS product, ALL, or NONE. If multiple values are specified, enclose the values in quotation marks and parentheses, and separate the values with a space.

This option is applicable on all Windows platforms, AIX, SLX, and MVS (Oracle only). On some of these hosts, SAS can be invoked with the -SETJMP system option. Setting -SETJMP disables the DBSRVTP= option.

Examples

The following examples invoke SAS from the UNIX command line:

```
sas -dbsrvtp all
```

```
sas -dbsrvtp '(oracle db2)'
```

```
sas -dbsrvtp teradata
```

```
sas -dbsrvtp '(SYBASE informix odbc oledb)'
```

```
sas -dbsrvtp none
```

SASTRACE= System Option

Generates trace information from a DBMS engine

Valid in: OPTIONS statement, configuration file, SAS invocation

Default value: NONE

Syntax

SASTRACE= ',,d' | ',,d,' | 'd,' | ',,,s' | ',,,sa' | ',,t,'

Syntax Description

',,d'

specifies that all SQL statements sent to the DBMS are sent to the log. These statements include the following:

SELECT
 CREATE
 DROP
 INSERT
 UPDATE
 DELETE
 SYSTEM CATALOG
 COMMIT
 ROLLBACK

For those engines that do not generate SQL statements, the API calls, including all parameters, are sent to the log.

'd'

specifies that all routine calls are sent to the log. When this option is selected, all function enters and exits, as well as pertinent parameters and return codes, are traced. The information, however, varies from engine to engine.

This option is most useful if you are having a problem and need to send a SAS log to technical support for troubleshooting.

'd'

specifies that all DBMS calls, such as API and Client calls, connection information, column bindings, column error information, and row processing are sent to the log. However, this information varies from engine to engine.

This option is most useful if you are having a problem and need to send a SAS log to technical support for troubleshooting.

',,s'

specifies that a summary of timing information for calls made to the DBMS is sent to the log.

',,sa'

specifies that timing information for each call made to the DBMS, along with a summary, is sent to the log.

',,t'

specifies that all threading information is sent to the log. This information includes:

- the number of threads spawned
- the number of observations each thread contains
- the exit code of the thread, should it fail.

SAS/ACCESS Specific Details

The SASTRACE= options have behavior that is specific to SAS/ACCESS software. SASTRACE= is a very powerful tool to use when you want to see the commands that are sent to your DBMS by the SAS/ACCESS engine. SASTRACE= output is DBMS-specific. However, most SAS/ACCESS engines show statements like SELECT or COMMIT as the DBMS processes them for the SAS application. The following details can help you manage SASTRACE= output in your DBMS:

- When using SASTRACE= on PC platforms, you must also specify SASTRACELOC=.
- In order to turn SAS tracing off, you can specify the following option:

```
options sastrace=off;
```

- Log output is much easier to read if you specify NOSTSUFFIX.

Note: NOSTSUFFIX is not supported on MVS. Δ

The following code is entered without specifying the option, and the resulting log is longer and harder to decipher.

```
options sastrace=',,,d' sastraceloc=saslog;
proc print data=mydblib.snow_birthdays;
run;
```

The resulting log is as follows:

```
0 1349792597 sastb_next 2930 PRINT
ORACLE_5: Prepared: 1 1349792597 sastb_next 2930 PRINT
SELECT * FROM scott.SNOW_BIRTHDAYS 2 1349792597 sastb_next 2930 PRINT
3 1349792597 sastb_next 2930 PRINT
16 proce print data=mydblib.snow_birthdays; run;

4 1349792597 sastb_next 2930 PRINT
ORACLE_6: Executed: 5 1349792597 sastb_next 2930 PRINT
Prepared statement ORACLE_5 6 1349792597 sastb_next 2930 PRINT
7 1349792597 sastb_next 2930 PRINT
```

However, by using NOSTSUFFIX, the log file becomes much easier to read:

```
options sastrace=',,,d' sastraceloc=saslog nostsuffix;
proc print data=mydblib.snow_birthdays;
run;
```

The resulting log is as follows:

```
ORACLE_1: Prepared:
SELECT * FROM scott.SNOW_BIRTHDAYS

12 proc print data=mydblib.snow_birthdays; run;

ORACLE_2: Executed:
Prepared statement ORACLE_1
```

Examples

The examples in this section are based on the following data set, and are shown using NOSTSUFFIX and SASTRACELOC=SASLOG.

```
data work.winter_birthdays;
  input empid birthdat date9. lastname $18.;
  format birthdat date9.;
datalines;
678999 28DEC1966 PAVEO          JULIANA          3451
456788 12JAN1977 SHIPTON       TIFFANY          3468
890123 20FEB1973 THORSTAD       EDVARD           3329
;
run;
```

Example 1: SQL Trace ',,,d'

```
options sastrace=',,,d' sastraceloc=saslog nostsuffix;
libname mydblib oracle user=scott password=tiger schema=bday_data;

data mydblib.snow_birthdays;
  set work.winter_birthdays;
run;

libname mydblib clear;
```

The output is written to the SAS log, as specified by the SASTRACELOC=SASLOG option.

Output 11.2 SAS Log Output from the SASTRACE= ',,,d' System Option

```
30 data work.winter_birthdays;
31   input empid birthdat date9. lastname $18.;
32   format birthdat date9.;
33 datalines;

NOTE: The data set WORK.WINTER_BIRTHDAYS has 3 observations and 3 variables.
NOTE: DATA statement used (Total process time):
      real time           0.03 seconds
      cpu time            0.04 seconds
37 ;
38 run;
39 options sastrace=',,,d' sastraceloc=saslog nostsuffix;

40 libname mydblib oracle user=scott password=XXXXX schema=bday_data;
NOTE: Libref MYDBLIB was successfully assigned as follows:
      Engine:           ORACLE
      Physical Name:

41 proc delete data=mydblib.snow_birthdays; run;

ORACLE_1: Prepared:
SELECT * FROM SNOW_BIRTHDAYS
```

```

ORACLE_2: Executed:
DROP TABLE SNOW_BIRTHDAYS

NOTE: Deleting MYDBLIB.SNOW_BIRTHDAYS (memtype=DATA).
NOTE: PROCEDURE DELETE used (Total process time):
      real time          0.26 seconds
      cpu time           0.12 seconds

42  data mydblib.snow_birthdays;
43      set work.winter_birthdays;
44  run;

ORACLE_3: Prepared:
SELECT * FROM SNOW_BIRTHDAYS

NOTE: SAS variable labels, formats, and lengths are not written to DBMS tables.

ORACLE_4: Executed:
CREATE TABLE SNOW_BIRTHDAYS(empid NUMBER ,birthdat DATE,lastname VARCHAR2 (18))

ORACLE_5: Prepared:
INSERT INTO SNOW_BIRTHDAYS (empid,birthdat,lastname) VALUES
(:empid,TO_DATE(:birthdat,'DDMONYYYY','NLS_DATE_LANGUAGE=American'),:lastname)

NOTE: There were 3 observations read from the data set WORK.WINTER_BIRTHDAYS.

ORACLE_6: Executed:
Prepared statement ORACLE_5

ORACLE: *--*--*--*--* COMMIT *--*--*--*--*
NOTE: The data set MYDBLIB.SNOW_BIRTHDAYS has 3 observations and 3 variables.
ORACLE: *--*--*--*--* COMMIT *--*--*--*--*
NOTE: DATA statement used (Total process time):
      real time          0.47 seconds
      cpu time           0.13 seconds

ORACLE_7: Prepared:
SELECT * FROM SNOW_BIRTHDAYS

45  proc print data=mydblib.snow_birthdays; run;

ORACLE_8: Executed:
Prepared statement ORACLE_7

NOTE: There were 3 observations read from the data set MYDBLIB.SNOW_BIRTHDAYS.
NOTE: PROCEDURE PRINT used (Total process time):
      real time          0.04 seconds
      cpu time           0.04 seconds

46
47  libname mydblib clear;
NOTE: Libref MYDBLIB has been deassigned.

```

Example 2: Log Trace ',,d'

```

options sastrace=',,d,' sastraceloc=saslog nostsuffix;
libname mydblib oracle user=scott password=tiger schema=bday_data;

data mydblib.snow_birthdays;
    set work.winter_birthdays;
run;

libname mydblib clear;

```

The output is written to the SAS log, as specified by the SASTRACELOC=SASLOG option.

Output 11.3 SAS Log Output from the SASTRACE= ',,d,' System Option

```

84  options sastrace=',,d,' sastraceloc=saslog nostsuffix;

ACCESS ENGINE: Entering DBICON
ACCESS ENGINE: Number of connections is 1
ORACLE: orcon()
ACCESS ENGINE: Successful physical conn id 1
ACCESS ENGINE: Exiting DBICON, Physical Connect id = 1, with rc=0X00000000
85  libname mydblib oracle user=dbitest password=XXXXX schema=bday_data;
ACCESS ENGINE: CONNECTION= SHAREDREAD
NOTE: Libref MYDBLIB was successfully assigned as follows:
      Engine:          ORACLE
      Physical Name:  lupin

86  data mydblib.snow_birthdays;
87      set work.winter_birthdays;
88  run;

ACCESS ENGINE: Entering yoeopen
ACCESS ENGINE: Entering dbiopen
ORACLE: oropen()
ACCESS ENGINE: Successful dbiopen, open id 0, connect id 1
ACCESS ENGINE: Exit dbiopen with rc=0X00000000
ORACLE: orqall()
ORACLE: orprep()
ACCESS ENGINE: Entering dbiclose
ORACLE: orclose()
ACCESS ENGINE: DBICLOSE open_id 0, connect_id 1
ACCESS ENGINE: Exiting dbiclos with rc=0X00000000
ACCESS ENGINE: Access Mode is XO_OUTPUT
ACCESS ENGINE: Access Mode is XO_SEQ
ACCESS ENGINE: Shr flag is XHSHRMEM
ACCESS ENGINE: Entering DBICON
ACCESS ENGINE: CONNECTION= SHAREDREAD
ACCESS ENGINE: Number of connections is 2
ORACLE: orcon()
ACCESS ENGINE: Successful physical conn id 2
ACCESS ENGINE: Exiting DBICON, Physical Connect id = 2, with rc=0X00000000
ACCESS ENGINE: Entering dbiopen
ORACLE: oropen()
ACCESS ENGINE: Successful dbiopen, open id 0, connect id 2
ACCESS ENGINE: Exit dbiopen with rc=0X00000000
ACCESS ENGINE: Exit yoeopen with SUCCESS.
ACCESS ENGINE: Begin yoeinfo
ACCESS ENGINE: Exit yoeinfo with SUCCESS.
ORACLE: orovar()
NOTE: SAS variable labels, formats, and lengths are not written to DBMS tables.
ORACLE: oroload()
ACCESS ENGINE: Entering dbrload with SQL Statement set to
      CREATE TABLE SNOW_BIRTHDAYS(empid NUMBER ,birthdat DATE,lastname VARCHAR2 (18))

ORACLE: orexec()
ORACLE: orexec() END
ORACLE: orins()
ORACLE: orubuf()
ORACLE: orubuf()
ORACLE: SAS date : 28DEC1966
ORACLE: orins()
ORACLE: SAS date : 12JAN1977
ORACLE: orins()
ORACLE: SAS date : 20FEB1973
NOTE: There were 3 observations read from the data set WORK.WINTER_BIRTHDAYS.
ORACLE: orforc()
ORACLE: orflush()
NOTE: The data set MYDBLIB.SNOW_BIRTHDAYS has 3 observations and 3 variables.
ACCESS ENGINE: Enter yoeclos
ACCESS ENGINE: Entering dbiclose
ORACLE: orclose()
ORACLE: orforc()
ORACLE: orflush()

```

```

ACCESS ENGINE: DBICLOSE open_id 0, connect_id 2
ACCESS ENGINE: Exiting dbiclos with rc=0X00000000
ACCESS ENGINE: Entering DBIDCON
ORACLE: ordcon
ACCESS ENGINE: Physical disconnect on id = 2
ACCESS ENGINE: Exiting DBIDCON with rc=0X00000000, rc2=0X00000000
ACCESS ENGINE: Exit yocelos with rc=0x00000000
NOTE: DATA statement used (Total process time):
      real time          0.21 seconds
      cpu time           0.06 seconds

ACCESS ENGINE: Entering DBIDCON
ORACLE: ordcon
ACCESS ENGINE: Physical disconnect on id = 1
ACCESS ENGINE: Exiting DBIDCON with rc=0X00000000, rc2=0X00000000
89  libname mydblib clear;
NOTE: Libref MYDBLIB has been deassigned.

```

Example 3: DBMS Trace 'd,'

```

options sastrace='d,' sastraceloc=saslog nostsuffix;
libname mydblib oracle user=scott password=tiger schema=bday_data;

data mydblib.snow_birthdays;
    set work.winter_birthdays;
run;

libname mydblib clear;

```

The output is written to the SAS log, as specified by the SASTRACELOC=SASLOG option.

Output 11.4 SAS Log Output from the SASTRACE= 'd,' System Option

```

ORACLE: PHYSICAL connect successful.
ORACLE: USER=scott
ORACLE: PATH=lupin
ORACLE: SCHEMA=bday_data
110 libname mydblib oracle user=dbitest password=XXXXX path=lupin schema=bday_data;

NOTE: Libref MYDBLIB was successfully assigned as follows:
      Engine:          ORACLE
      Physical Name:  lupin

111 data mydblib.snow_birthdays;
112   set work.winter_birthdays;
113 run;

ORACLE: PHYSICAL connect successful.
ORACLE: USER=scott
ORACLE: PATH=lupin
ORACLE: SCHEMA=bday_data
NOTE: SAS variable labels, formats, and lengths are not written to DBMS tables.
ORACLE: INSERTBUFF option value set to 10.
NOTE: There were 3 observations read from the data set WORK.WINTER_BIRTHDAYS.
ORACLE: Rows processed: 3
ORACLE: Rows failed   : 0
NOTE: The data set MYDBLIB.SNOW_BIRTHDAYS has 3 observations and 3 variables.
ORACLE: Successfully disconnected.
ORACLE: USER=scott
ORACLE: PATH=lupin
NOTE: DATA statement used (Total process time):
      real time          0.21 seconds
      cpu time           0.04 seconds

```

```
ORACLE: Successfully disconnected.
ORACLE: USER=scott
ORACLE: PATH=lupin
114 libname mydblib clear;
NOTE: Libref MYDBLIB has been deassigned.
```

Example 4: Time Trace ',,,s'

```
options sastrace=',,,s' sastraceloc=saslog nostsuffix;
libname mydblib oracle user=scott password=tiger schema=bday_data;

data mydblib.snow_birthdays;
    set work.winter_birthdays;
run;

libname mydblib clear;
```

The output is written to the SAS log, as specified by the SASTRACELOC=SASLOG option.

Output 11.5 SAS Log Output from the SASTRACE= ',,,s' System Option

```
118 options sastrace=',,,s' sastraceloc=saslog nostsuffix;

119 libname mydblib oracle user=dbitest password=XXXXX schema=bday_data;

NOTE: Libref MYDBLIB was successfully assigned as follows:
      Engine:          ORACLE
      Physical Name:  lupin

120 data mydblib.snow_birthdays;
121     set work.winter_birthdays;
122 run;

NOTE: SAS variable labels, formats, and lengths are not written to DBMS tables.
NOTE: There were 3 observations read from the data set WORK.WINTER_BIRTHDAYS.
NOTE: The data set MYDBLIB.SNOW_BIRTHDAYS has 3 observations and 3 variables.

Summary Statistics for ORACLE are:
Total SQL execution seconds were:          0.127079
Total SQL prepare seconds were:           0.004404
Total SQL row insert seconds were:        0.004735
Total seconds used by the ORACLE ACCESS engine were 0.141860

NOTE: DATA statement used (Total process time):
      real time          0.21 seconds
      cpu time           0.04 seconds

123 libname mydblib clear;
NOTE: Libref MYDBLIB has been deassigned.
```

Example 5: Time All Trace ',,sa'

```

options sastrace=',,sa' sastraceloc=saslog nostsuffix;
libname mydblib oracle user=scott password=tiger schema=bday_data;

data mydblib.snow_birthdays;
    set work.winter_birthdays;
run;

libname mydblib clear;

```

The output is written to the SAS log, as specified by the SASTRACELOC=SASLOG option.

Output 11.6 SAS Log Output from the SASTRACE= ',,sa' System Option

```

146 options sastrace=',,sa' sastraceloc=saslog nostsuffix;
147
148 libname mydblib oracle user=dbitest password=XXXXX path=lupin schema=dbitest insertbuff=1;

NOTE: Libref MYDBLIB was successfully assigned as follows:
      Engine:          ORACLE
      Physical Name:   lupin

149 data mydblib.snow_birthdays;
150     set work.winter_birthdays;
151 run;

NOTE: SAS variable labels, formats, and lengths are not written to DBMS tables.
ORACLE: The insert time in seconds is    0.004120
ORACLE: The insert time in seconds is    0.001056
ORACLE: The insert time in seconds is    0.000988
NOTE: There were 3 observations read from the data set WORK.WINTER_BIRTHDAYS.
NOTE: The data set MYDBLIB.SNOW_BIRTHDAYS has 3 observations and 3 variables.

Summary Statistics for ORACLE are:
Total SQL execution seconds were:          0.130448
Total SQL prepare seconds were:           0.004525
Total SQL row insert seconds were:         0.006158
Total seconds used by the ORACLE ACCESS engine were    0.147355

NOTE: DATA statement used (Total process time):
      real time          0.20 seconds
      cpu time           0.00 seconds

152
153 libname mydblib clear;
NOTE: Libref MYDBLIB has been deassigned.

```

Example 6: Threaded Trace ',,t'

```

options sastrace=',,t,' sastraceloc=saslog nostsuffix;
libname mydblib oracle user=scott password=tiger schema=bday_data;

data mydblib.snow_birthdays(DBTYPE=(empid'number(10')));
    set work.winter_birthdays;
run;

proc print data=mydblib.snow_birthdays(dbsliceparm=(all,3));
run;

```

The output is written to the SAS log, as specified by the SASTRACELOC=SASLOG option.

Output 11.7 SAS Log Output from the SASTRACE= ',,t,' System Option

```

165 options sastrace=',,t,' sastraceloc=saslog nostsuffix;
166 data mydblib.snow_birthdays(DBTYPE=(empid='number(10)'));
167     set work.winter_birthdays;
168 run;

NOTE: SAS variable labels, formats, and lengths are not written to DBMS tables.
NOTE: There were 3 observations read from the data set WORK.WINTER_BIRTHDAYS.
NOTE: The data set MYDBLIB.SNOW_BIRTHDAYS has 3 observations and 3 variables.
NOTE: DATA statement used (Total process time):
      real time          0.21 seconds
      cpu time           0.06 seconds

169 proc print data=mydblib.snow_birthdays(dbsliceparm=(all,3));
170 run;

ORACLE: DBSLICEPARAM option set and 3 threads were requested
ORACLE: No application input on number of threads.
ORACLE: Thread 1 contains 1 obs.
ORACLE: Thread 2 contains 0 obs.
ORACLE: Thread 3 contains 2 obs.
ORACLE: Threaded read enabled. Number of threads created: 3
NOTE: There were 3 observations read from the data set MYDBLIB.SNOW_BIRTHDAYS.
NOTE: PROCEDURE PRINT used (Total process time):
      real time          1.12 seconds
      cpu time           0.17 seconds

```

For more information about tracing threaded reads, refer to “Generating Trace Information for Threaded Reads” on page 50.

Note: The SASTRACE= options can also be used in conjunction with each other. For example, SASTRACE=’,,d,d’. Δ

SASTRACELOC= System Option

Prints SASTRACE information to a specified location

Valid in: OPTIONS statement, configuration file, SAS invocation

Default value: stdout

Syntax

SASTRACELOC=stdout | SASLOG | FILE 'path-and-filename'

Details

SASTRACELOC= enables you to specify where to put the trace messages that are generated by SASTRACE=. By default, the output goes to the default output location for your operating environment. You can send the output to a SAS log by specifying SASTRACELOC=SASLOG.

Note: This option and its values might differ for each host. Δ

Example

The following example on a PC platform writes the trace information to the TRACE.LOG file in the work directory on the C drive.

```
options sastrace=',,,d' sastraceloc=file 'c:\work\trace.log';
```

VALIDVARNAME= System Option

Controls the type of SAS variable names that can be used or created during a SAS session

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window

Default value: V7

Syntax

VALIDVARNAME= V7 | V6 | UPCASE | ANY

SAS/ACCESS Specific Details

VALIDVARNAME= enables you to control which rules apply for SAS variable names. For more information about the VALIDVARNAME= system option, see the *SAS Language Reference: Dictionary*. The settings are as follows:

VALIDVARNAME=V7

indicates that a DBMS column name is changed to a valid SAS name by using the following rules:

- Up to 32 mixed-case alphanumeric characters are allowed.
- Names must begin with an alphabetic character or an underscore.
- Invalid characters are changed to underscores.
- Any column name that is not unique when it is normalized is made unique by appending a counter (0,1,2,...) to the name.

This is the default value for SAS Version 7 and later.

VALIDVARNAME=V6

indicates that only those variable names that are considered valid in Version 6 are considered valid SAS variable names. When V6 is specified in Pass-Through Facility code, the DBMS engine truncates column names to eight characters as it did in Version 6. If required, numbers are appended to the end of the truncated name to make it unique.

VALIDVARNAME=UPCASE

indicates that a DBMS column name is changed to a valid SAS name as described in VALIDVARNAME=V7 except that variable names are in uppercase.

VALIDVARNAME=ANY

allows any characters in DBMS column names to appear as valid characters in SAS variable names. Symbols, such as the equal sign (=) and the asterisk (*), must be contained in a *'variable-name'*n construct. You must use ANY whenever you want to read DBMS column names that do not follow the SAS naming conventions.

Example

The following example shows how the Pass-Through Facility works with VALIDVARNAME=V6.


```
options validvarname=v6;
proc sql;
  connect to oracle (user=testuser pass=testpass);
  create view myview as
    select amount_b, amount_s
      from connection to oracle
        (select "Amount Budgeted$", "Amount Spent$"
          from mytable);
quit;

proc contents data=myview;
run;
```

The output from this example would show that "Amount Budgeted\$" becomes AMOUNT_B and "Amount Spent\$" becomes AMOUNT_S.

See Also

“SAS Names and Support for DBMS Names” in *SAS/ACCESS for Relational Databases: Reference*



CHAPTER

12

The Pass-Through Facility for Relational Databases

<i>Overview of the SQL Procedure's Interactions with SAS/ACCESS</i>	285
<i>Overview of the Pass-Through Facility</i>	285
<i>Syntax for the Pass-Through Facility for Relational Databases</i>	286
<i>Return Codes</i>	286

Overview of the SQL Procedure's Interactions with SAS/ACCESS

The SQL procedure implements structured query language (SQL) for SAS software. See the *Base SAS Procedures Guide* for information about PROC SQL. SAS/ACCESS software for relational databases interacts with PROC SQL in the following ways:

- You can assign a libref to a DBMS using the SAS/ACCESS LIBNAME statement and then reference the new libref in a PROC SQL statement to query, update, or delete DBMS data. See the Chapter 9, “The LIBNAME Statement for Relational Databases,” on page 73 for more information about this feature.
- You can embed LIBNAME information in a PROC SQL view and then automatically connect to the DBMS every time the PROC SQL view is processed. See “SQL Views with Embedded LIBNAME Statements” on page 76 for more information.
- You can send DBMS-specific SQL statements directly to a DBMS using an extension to PROC SQL called the Pass-Through Facility. See “Syntax for the Pass-Through Facility for Relational Databases” on page 286 for the syntax for this feature.

Overview of the Pass-Through Facility

The Pass-Through Facility uses SAS/ACCESS to connect to a DBMS and to send statements directly to the DBMS for execution. This facility is an alternative to the SAS/ACCESS LIBNAME statement. It enables you to use the SQL syntax of your DBMS, and it supports any non-ANSI standard SQL that is supported by your DBMS.

Note: Not all SAS/ACCESS interfaces support this feature. See Chapter 8, “SAS/ACCESS Features by Host,” on page 65 to determine whether this feature is available in your environment. △

The Pass-Through Facility enables you to complete the following tasks:

- establish and terminate connections with a DBMS using the facilities CONNECT and DISCONNECT statements
- send dynamic, non-query, DBMS-specific SQL statements to a DBMS using the facility's EXECUTE statement

- retrieve data directly from a DBMS using the facilities CONNECTION TO component in the FROM clause of a PROC SQL SELECT statement.

You can use Pass-Through Facility statements in a PROC SQL query or you can store them in an SQL view. When you create an SQL view, any arguments that you specify in the CONNECT statement are stored with the view. Therefore, when the view is used in a SAS program, SAS can establish the appropriate connection to the DBMS.

Syntax for the Pass-Through Facility for Relational Databases

This section presents the syntax for the Pass-Through Facility statements and the CONNECTION TO component. For DBMS-specific details, see the documentation for your SAS/ACCESS interface.

```
PROC SQL <option(s)>;
CONNECT TO dbms-name <AS alias> <(<database-connection-arguments>
    <connect-statement-arguments> )>;
DISCONNECT FROM dbms-name | alias;
EXECUTE (dbms-specific-SQL-statement) BY dbms-name | alias;
SELECT column-list FROM CONNECTION TO dbms-name | alias (dbms-query)
```

Return Codes

As you use the PROC SQL statements that are available in the Pass-Through Facility, any error return codes and error messages are written to the SAS log. These codes and messages are available to you through the following two SAS macro variables:

```
SQLXRC
    contains the DBMS return code that identifies the DBMS error.
SQLXMSG
    contains descriptive information about the DBMS error that is generated by the
    DBMS.
```

The contents of the SQLXRC and SQLXMSG macro variables are printed in the SAS log using the %PUT macro. They are reset after each Pass-Through Facility statement has been executed.

See “Macro Variables for Relational Databases” on page 269 for more information about these return codes.

CONNECT Statement

Establishes a connection with the DBMS

Valid in: PROC SQL steps (when accessing DBMS data using SAS/ACCESS software)

Syntax

```
CONNECT TO dbms-name <AS alias> <(<database-connection-arguments>
    <connect-statement-arguments> )>;
```

The CONNECT statement establishes a connection with the DBMS. You establish a connection to send DBMS-specific SQL statements to the DBMS or to retrieve DBMS data. The connection remains in effect until you issue a DISCONNECT statement or terminate the SQL procedure.

To connect to a DBMS using the Pass-Through Facility, complete the following steps:

- 1 Initiate a PROC SQL step.
- 2 Use the Pass-Through Facility's CONNECT statement, identify the DBMS (such as Oracle or DB2), and (optionally) assign an alias.
- 3 Specify any attributes for the connection (such as multiple, shared, unique).
- 4 Specify any arguments that are needed to connect to the database.

The CONNECT statement is optional for some DBMSs. However, if it is not specified, the default values for all of the database connection arguments are used.

Any return code or message that is generated by the DBMS is available in the macro variables SQLXRC and SQLXMSG after the statement executes. See "Macro Variables for Relational Databases" on page 269 for more information about these macro variables.

Arguments

You use the following arguments with the CONNECT statement:

dbms-name

identifies the database management system to which you want to connect. You must specify the DBMS name for your SAS/ACCESS interface. You can also specify an optional alias.

alias

specifies for the connection an optional alias that has 1 to 32 characters. If you specify an alias, the keyword AS must appear before the alias. If an alias is not specified, the DBMS name is used as the name of the Pass-Through connection.

database-connection-arguments

specifies the DBMS-specific arguments that are needed by PROC SQL to connect to the DBMS. These arguments are optional for most databases, but if they are included, they must be enclosed in parentheses. See the documentation for your SAS/ACCESS interface for information about these arguments.

connect-statement-arguments

specifies arguments that indicate whether you can make multiple connections, shared or unique connections, and so on, to the database. These arguments enable the Pass-Through Facility to use some of the LIBNAME statement's connection management features. These arguments are optional, but if they are included, they must be enclosed in parentheses.

CONNECTION= SHARED | GLOBAL

indicates whether multiple CONNECT statements for a DBMS can use the same connection.

The CONNECTION= option enables you to control the number of connections, and therefore transactions, that your SAS/ACCESS engine executes and supports for each Pass-Through CONNECT statement.

When CONNECTION=GLOBAL, multiple CONNECT statements that use identical values for CONNECTION=, CONNECTION_GROUP=, DBCONINIT=, DBCONTERM=, and any database connection arguments can share the same connection to the DBMS.

When CONNECTION=SHARED, the CONNECT statement makes one connection to the DBMS. Only Pass-Through statements that use this alias share the connection. SHARED is the default value for CONNECTION=.

In the following example, the two CONNECT statements share the same connection to the DBMS because CONNECTION=GLOBAL. Only the first CONNECT statement actually makes the connection to the DBMS, while the last DISCONNECT statement is the only statement that disconnects from the DBMS.

```
proc sql;

/*...SQL Pass-Through statements referring to mydbone...*/

connect to oracle as mydbone
  (user=testuser pw=testpass
   path='myorapath'
   connection=global);

/*...SQL Pass-Through statements referring to mydbtwo...*/

connect to oracle as mydbtwo
  (user=testuser pw=testpass
   path='myorapath'
   connection=global);

disconnect from mydbone;
disconnect from mydbtwo;
quit;
```

CONNECTION_GROUP=*connection-group-name*

specifies a connection that can be shared among several CONNECT statements in the Pass-Through Facility.

Default value: none

By specifying the name of a connection group, you can share one DBMS connection among several CONNECT statements. The connection to the DBMS can be shared only if each CONNECT statement specifies the same CONNECTION_GROUP= value and specifies identical DBMS connection arguments.

When CONNECTION_GROUP= is specified, it implies that the value of the CONNECTION= option is GLOBAL.

DBCONINIT=*<'>DBMS-user-command<'>*

specifies a user-defined initialization command to be executed immediately after the connection to the DBMS.

You can specify any DBMS command that can be passed by the SAS/ACCESS engine to the DBMS and that does not return a result set or output parameters. The command executes immediately after the DBMS connection is established successfully. If the command fails, a disconnect occurs, and the CONNECT statement fails. You must specify the command as a single, quoted string, unless it is an environment variable.

DBCONTERM=*'DBMS-user-command'*

specifies a user-defined termination command to be executed before the disconnect from the DBMS that occurs with the DISCONNECT statement.

Default value: none

The termination command that you select can be a script, stored procedure, or any DBMS SQL language statement to provide additional control over the

interaction between the SAS/ACCESS engine and the DBMS. You can specify any valid DBMS command that can be passed by the SAS/ACCESS engine to the DBMS and that does not return a result set or output parameters. The command executes immediately before SAS terminates each connection to the DBMS. If the command fails, SAS provides a warning message but the disconnect still occurs. You must specify the command as a quoted string.

DBGEN_NAME= DBMS | SAS

specifies whether to automatically rename DBMS columns containing characters that SAS does not allow, such as \$, to valid SAS variable names. See “DBGEN_NAME= LIBNAME Option” on page 100 for further information.

DBMAX_TEXT=*integer*

determines the length of any very long DBMS character data type that is read into SAS or written from SAS when using a SAS/ACCESS engine. This option applies to reading, appending, and updating rows in an existing table. It does not apply when you are creating a table.

Examples of a long DBMS data type are the SYBASE TEXT data type or the Oracle LONG RAW data type.

DBPROMPT=YES | NO

specifies whether SAS displays a window that prompts the user to enter DBMS connection information before connecting to the DBMS.

Default value: NO

If you specify DBPROMPT=YES, SAS displays a window that interactively prompts you for the DBMS connection arguments when the CONNECT statement is executed. Therefore, it is not necessary to provide connection arguments with the CONNECT statement. If you do specify connection arguments with the CONNECT statement and you specify DBPROMPT=YES, the connection argument values are displayed in the window. These values can be overridden interactively.

If you specify DBPROMPT=NO, SAS does not display the prompting window.

The DBPROMPT= option interacts with the DEFER= option to determine when the prompt window appears. If DEFER=NO, the DBPROMPT window opens when the CONNECT statement is executed. If DEFER=YES, the DBPROMPT window opens the first time a Pass-Through statement is executed. The DEFER= option normally defaults to NO, but defaults to YES if DBPROMPT=YES. You can override this default by explicitly setting DEFER=NO.

DEFER=NO | YES

determines when the connection to the DBMS occurs.

Default value: NO

If DEFER=YES, the connection to the DBMS occurs when the first Pass-Through statement is executed. If DEFER=NO, the connection to the DBMS occurs when the CONNECT statement occurs.

VALIDVARNAME=V6

indicates that only those variable names considered valid SAS variable names in Version 6 of SAS software are considered valid. Specify this connection argument if you want Pass-Through to operate in Version 6 compatibility mode.

By default DBMS column names are changed to valid SAS names by using the following rules:

- Up to 32 mixed-case alphanumeric characters are allowed.

- Names must begin with an alphabetic character or an underscore.
- Characters that are not permitted are changed to underscores.
- Any column name that is not unique when it is normalized is made unique by appending a counter (0,1,2,...) to the name.

When VALIDVARNAME=V6 is specified, the SAS/ACCESS engine for the DBMS truncates column names to eight characters, as it did in Version 6. If required, numbers are appended to the ends of the truncated names to make them unique. Setting this option overrides the value of the SAS system option VALIDVARNAME= during (and only during) the Pass-Through connection.

The following example shows how the Pass-Through Facility uses VALIDVARNAME=V6 as a connection argument. Using this option causes the output to show the DBMS column "Amount Budgeted\$" as AMOUNT_B and "Amount Spent\$" as AMOUNT_S.

```
proc sql;
connect to oracle (user=gloria password=teacher
                  validvarname=v6)
create view budget2000 as
  select amount_b, amount_s
  from connection to oracle
      (select "Amount Budgeted$", "Amount Spent$"
       from annual_budget);
quit;
proc contents data=budget2000;
run;
```

For this example, if you *omit* VALIDVARNAME=V6 as a connection argument, you must add it in an OPTIONS= statement in order for PROC CONTENTS to work:

```
options validvarname=v6;
proc contents data=budget2000;
run;
```

Thus, using it as a connection argument saves you coding later.

Note: In addition to the arguments listed here, several other LIBNAME options are available for use with the CONNECT statement. See the section about the Pass-Through Facility in the documentation for your SAS/ACCESS interface to determine which LIBNAME options are available in the Pass-Through Facility for your DBMS. When used with the Pass-Through Facility CONNECT statement, these options have the same effect as they do in a LIBNAME statement. Δ

CONNECT Statement Example

The following example connects to a Sybase server and assigns the alias SYBCON1 to it. Sybase is a case-sensitive database. Therefore, the database objects are in uppercase, as they were created.

```
proc sql;
connect to sybase as sybcon1
  (server=SERVER1 database=PERSONNEL
   user=testuser password=testpass
   connection=global);
%put &sqlxmsg &sqlxrc;
```

Note: You might be able to omit the CONNECT statement and implicitly connect to a database using default settings. See the documentation for your SAS/ACCESS interface for more information. △

DISCONNECT Statement

Terminates the connection to the DBMS

Valid in: PROC SQL steps (when accessing DBMS data using SAS/ACCESS software)

Syntax

DISCONNECT FROM *dbms-name* | *alias*

The DISCONNECT statement ends the connection with the DBMS. If the DISCONNECT statement is omitted, an implicit DISCONNECT is performed when PROC SQL terminates. The SQL procedure continues to execute until you submit a QUIT statement, another SAS procedure, or a DATA step.

Any return code or message that is generated by the DBMS is available in the macro variables SQLXRC and SQLXMSG after the statement executes. See “Macro Variables for Relational Databases” on page 269 for more information about these macro variables.

Arguments

Use one of the following arguments with the DISCONNECT statement:

dbms-name

specifies the database management system from which you want to disconnect. You must specify the DBMS name for your SAS/ACCESS interface, or use an alias in the DISCONNECT statement.

Note: If you used the CONNECT statement to connect to the DBMS, the DISCONNECT statement’s DBMS name or alias must match the name or alias specified in the CONNECT statement. △

alias

specifies an alias that was defined in the CONNECT statement.

Example

To exit the Pass-Through Facility, use the facilities DISCONNECT statement and then QUIT the PROC SQL statement. The following example disconnects the user from a DB2 database with the alias DBCON1 and terminates the SQL procedure:

```
proc sql;
connect to db2 as dbcon1 (ssid=db2a);
...more SAS statements...
disconnect from dbcon1;
quit;
```

EXECUTE Statement

Sends DBMS-specific, non-query SQL statements to the DBMS

Valid in: PROC SQL steps (when accessing DBMS data using SAS/ACCESS software)

Syntax

EXECUTE (*dbms-specific-sql-statement*) BY *dbms-name* | *alias*;

The EXECUTE statement sends dynamic non-query, DBMS-specific SQL statements to the DBMS and processes those statements.

In some SAS/ACCESS interfaces, you can issue an EXECUTE statement directly without first explicitly connecting to a DBMS (see “CONNECT Statement” on page 286). If you omit the CONNECT statement, an implicit connection is performed (by using default values for all database connection arguments) when the first EXECUTE statement is passed to the DBMS. See the documentation for your SAS/ACCESS interface for details.

The EXECUTE statement cannot be stored as part of an Pass-Through Facility query in a PROC SQL view.

Arguments

(dbms-specific-sql-statement)

a dynamic non-query, DBMS-specific SQL statement. This argument is required and must be enclosed in parentheses. However, the SQL statement cannot contain a semicolon because a semicolon represents the end of a statement in SAS. The SQL statement might be case-sensitive, depending on your DBMS, and it is passed to the DBMS exactly as you type it.

On some DBMSs, this argument can be a DBMS stored procedure. However, stored procedures with output parameters are not supported in the Pass-Through Facility. Furthermore, if the stored procedure contains more than one query, only the first query is processed.

Any return code or message that is generated by the DBMS is available in the macro variables SQLXRC and SQLXMSG after the statement executes. See “Macro Variables for Relational Databases” on page 269 for more information about these macro variables.

dbms-name

identifies the database management system to which you direct the DBMS-specific SQL statement. The keyword BY must appear before the *dbms-name* argument. You must specify either the DBMS name for your SAS/ACCESS interface or an alias.

alias

specifies an alias that was defined in the CONNECT statement. (You cannot use an alias if the CONNECT statement was omitted.)

Useful Statements to Include in EXECUTE Statements

You can pass the following statements to the DBMS by using the Pass-Through Facility’s EXECUTE statement.

CREATE

creates a DBMS table, view, index, or other DBMS object, depending on how the statement is specified.

DELETE

deletes rows from a DBMS table.

DROP

deletes a DBMS table, view, or other DBMS object, depending on how the statement is specified.

GRANT

gives users the authority to access or modify objects such as tables or views.

INSERT

adds rows to a DBMS table.

REVOKE

revokes the access or modification privileges that were given to users by the GRANT statement.

UPDATE

modifies the data in one column of a row in a DBMS table.

For more information and restrictions on these and other SQL statements, see the SQL documentation for your DBMS.

CONNECTION TO Component

Retrieves and uses DBMS data in a PROC SQL query or view

Valid in: PROC SQL step SELECT statements (when accessing DBMS data using SAS/ACCESS software)

Syntax

CONNECTION TO *dbms-name* | *alias* (*dbms-query*)

The CONNECTION TO component specifies the DBMS connection that you want to use or that you want to create (if you have omitted the CONNECT statement). CONNECTION TO then enables you to retrieve DBMS data directly through a PROC SQL query.

You use the CONNECTION TO component in the FROM clause of a PROC SQL SELECT statement:

```
PROC SQL;
  SELECT column-list
  FROM CONNECTION TO dbms-name (dbms-query)
  other optional PROC SQL clauses
QUIT;
```

CONNECTION TO can be used in any FROM clause, including those in nested queries (that is, subqueries).

You can store a Pass-Through Facility query in an SQL view and then use that view in SAS programs. When you create an SQL view, any options that you specify in the corresponding CONNECT statement are stored too. Thus, when the SQL view is used in a SAS program, SAS can establish the appropriate connection to the DBMS.

On many relational databases, you can issue a CONNECTION TO component in a PROC SQL SELECT statement directly without first connecting to a DBMS (see “CONNECT Statement” on page 286). If you omit the CONNECT statement, an implicit connection is performed when the first PROC SQL SELECT statement that contains a CONNECTION TO component is passed to the DBMS. Default values are used for all DBMS connection arguments. See the documentation for your SAS/ACCESS interface for details.

Because relational databases and SAS have different naming conventions, some DBMS column names might be changed when you retrieve DBMS data through the CONNECTION TO component. See Chapter 2, “SAS Names and Support for DBMS Names,” on page 7 for more information.

Arguments

dbms-name

identifies the database management system to which you direct the DBMS-specific SQL statement. See the documentation for your SAS/ACCESS interface for the name for your DBMS.

alias

specifies an alias, if one was defined in the CONNECT statement.

(dbms-query)

specifies the query that you are sending to the DBMS. The query can use any DBMS-specific SQL statement or syntax that is valid for the DBMS. However, the query cannot contain a semicolon because a semicolon represents the end of a statement in SAS.

You must specify a query argument in the CONNECTION TO component, and the query must be enclosed in parentheses. The query is passed to the DBMS exactly as you type it. Therefore, if your DBMS is case-sensitive, you must use the correct case for DBMS object names.

On some DBMSs, the *dbms-query* argument can be a DBMS stored procedure. However, stored procedures with output parameters are not supported in the Pass-Through Facility. Furthermore, if the stored procedure contains more than one query, only the first query is processed.

Example

After you connect (explicitly using the CONNECT statement or implicitly using default settings) to a DBMS, you can send a DBMS-specific SQL query to the DBMS using the facilities CONNECTION TO component. You issue a SELECT statement to indicate which columns you want to retrieve, identify your DBMS (such as Oracle or DB2), and issue your query using the SQL syntax of your DBMS.

The following example sends an Oracle SQL query, shown by highlighting, to the Oracle database for processing. The results from the Oracle SQL query serve as a virtual table for the PROC SQL FROM clause. In this example, MYCON is a connection alias.

```
proc sql;
  connect to oracle as mycon (user=testuser
```

```

password=testpass path='myorapath');
%put &sqlxmsg;

select *
  from connection to mycon
    (select empid, lastname, firstname,
        hiredate, salary
     from employees where
        hiredate>='31-DEC-88');
%put &sqlxmsg;

disconnect from mycon;
quit;

```

The SAS %PUT macro displays the &SQLXMSG macro variable for error codes and information from the DBMS. See “Macro Variables for Relational Databases” on page 269 for more information.

The following example gives the query a name and stores it as the SQL view samples.HIRES88:

```

libname samples 'SAS-data-library';


proc sql;
connect to oracle as mycon (user=testuser
    password=testpass path='myorapath');
%put &sqlxmsg;

create view samples.hires88 as
  select *
    from connection to mycon
      (select empid, lastname, firstname,
          hiredate, salary
       from employees where
          hiredate>='31-DEC-88');
%put &sqlxmsg;

disconnect from mycon;


quit;

```

Converting SAS/ACCESS Descriptors to PROC SQL Views

Chapter 13 **The CV2VIEW Procedure** 299



CHAPTER

13

The CV2VIEW Procedure

<i>Overview of the CV2VIEW Procedure</i>	299
<i>Procedure Syntax</i>	300
<i>PROC CV2VIEW Statement</i>	300
<i>FROM_VIEW= Statement</i>	300
<i>FROM_LIBREF= Statement</i>	301
<i>REPLACE= Statement</i>	301
<i>SAVEAS= Statement</i>	302
<i>SUBMIT Statement</i>	302
<i>TO_VIEW= Statement</i>	303
<i>TO_LIBREF= Statement</i>	303
<i>TYPE= Statement</i>	304
<i>CV2VIEW Procedure Examples</i>	304
<i>Example 1: Converting an Individual View Descriptor</i>	304
<i>Example 2: Converting a Library of View Descriptors for a Single DBMS</i>	306
<i>Example 3: Converting a Library of View Descriptors for All Supported DBMSs</i>	307

Overview of the CV2VIEW Procedure

The CV2VIEW procedure converts SAS/ACCESS view descriptors and access descriptors into SQL views. We recommend that you consider converting your descriptors for the following reasons:

- Descriptors are no longer the recommended method for accessing relational database data. By converting to SQL views, you can use the LIBNAME statement, which is the preferred method. The LIBNAME statement provides greater control over DBMS operations such as locking, spooling, and data type conversions. The LIBNAME statement can also handle long field names, whereas descriptors cannot.
- SQL views are platform-independent. SAS/ACCESS descriptors are not.

The CV2VIEW procedure in SAS 9.1 can convert both of these descriptors:

- 64-bit SAS/ACCESS view descriptors that were created in either 64-bit SAS Version 8 or 64-bit SAS 9.1
- 32-bit SAS/ACCESS access descriptors that were created in 32-bit SAS Version 6 and Version 8

If the descriptor that you want to convert is READ, WRITE, or ALTER protected, then those values are applied to the output SQL view. For security reasons, these values do not appear if you save the generated SQL to a file. The PASSWORD portion of the LIBNAME statement is also not visible. This prevents the generated SQL statements from being manually submitted without modification.

Procedure Syntax

The syntax for the CV2VIEW procedure is as follows:

```
PROC CV2VIEW DBMS= dbms-name | ALL;
FROM_VIEW= libref.input-descriptor;
FROM_LIBREF= input-library;
TO_VIEW= libref.output-view;
TO_LIBREF= output-library;
TYPE= SQL | VIEW | ACCESS;
SAVEAS= external-filename;
SUBMIT;
REPLACE ALL | VIEW | FILE;
```

PROC CV2VIEW Statement

```
PROC CV2VIEW DBMS= dbms-name | ALL;
```

Arguments

dbms-name

specifies the name of a supported database from which you want to obtain descriptors. Valid values for *dbms-name* are **DB2**, **Oracle**, and **SYBASE**.

ALL

specifies that you want the descriptors from all supported databases.

FROM_VIEW= Statement

Specifies the name of the view descriptor or access descriptor that you want to convert

Requirement: You must specify either the FROM_VIEW= statement or the FROM_LIBREF= statement.

Requirement: The FROM_VIEW= and TO_VIEW= statements are always used together.

Restriction: If you specify DBMS=ALL, then you cannot use the FROM_VIEW= statement.

```
FROM_VIEW=libref.input-descriptor;
```


Arguments

libref

specifies the libref that contains the view descriptor or access descriptor that you want to convert.

input-descriptor

specifies the view descriptor or access descriptor that you want to convert.

FROM_LIBREF= Statement

Specifies the library that contains the view descriptors or access descriptors that you want to convert

Requirement: You must specify either the FROM_VIEW= statement or the FROM_LIBREF= statement.

Requirement: The FROM_LIBREF= and TO_LIBREF= statements are always used together.

```
FROM_LIBREF= input-library;
```

Argument

input-library

specifies a previously assigned library that contains the view descriptors or access descriptors that you want to convert. All descriptors that are in the specified library and that access data in the specified DBMS are converted into SQL views. If you specify DBMS=ALL, then all descriptors that are in the specified library and that access any supported DBMS are converted.

REPLACE= Statement

Specifies whether existing views and files are replaced

```
REPLACE= ALL | FILE | VIEW ;
```

Arguments

ALL

replaces the TO_VIEW= file if it already exists and replaces the SAVEAS= file if it already exists.

FILE

replaces the SAVEAS= file if it already exists. If the file already exists, and if REPLACE=FILE or REPLACE=ALL are not specified, the generated PROC SQL code is appended to the file.

VIEW

replaces the TO_VIEW= file if it already exists.

SAVEAS= Statement

Saves the generated PROC SQL statements to a file

Interaction: If you specify the SAVEAS= statement, the generated SQL is not automatically submitted, so you must use the SUBMIT statement.

SAVEAS=*external-filename*;

Argument

external-filename

lets you save the PROC SQL statements that are generated by PROC CV2VIEW to an external file. You can modify this file and submit it on another platform.

Details

PROC CV2VIEW inserts comments in the generated SQL to replace any statements that contain passwords. For example, if a view descriptor is READ, WRITE, or ALTER protected, then the output view has the same level of security. However, the file that contains the SQL statements does not show the password values. This is also the case for the password in the LIBNAME statement.

SUBMIT Statement

Causes PROC CV2VIEW to submit the generated PROC SQL statements when you specify the SAVEAS= statement

Tip: If you do not use the SAVEAS= statement, PROC CV2VIEW automatically submits the generated SQL, so you do not need to specify the SUBMIT statement.

SUBMIT;

TO_VIEW= Statement

Specifies the name of the new (converted) SQL view

Requirement: You must specify either the TO_VIEW= statement or the TO_LIBREF= statement.

Requirement: The FROM_VIEW= and TO_VIEW= statements are always used together.

Restriction: If you specify DBMS=ALL, then you cannot use the TO_VIEW= statement.

Interaction: Use the REPLACE= statement to control whether the output file is overwritten or appended if it already exists.

```
TO_VIEW=libref.output-view;
```

Arguments

libref

specifies the libref where you want to store the new SQL view.

output-view

specifies the name for the new SQL view that you want to create.

TO_LIBREF= Statement

Specifies the library that contains the new (converted) SQL views

Requirement: You must specify either the TO_VIEW= statement or the TO_LIBREF= statement.

Requirement: The FROM_LIBREF= and TO_LIBREF= statements are always used together.

Interaction: Use the REPLACE= statement if a file with the name of one of your output views already exists. If a file with the name of one of your output views already exists and you do not specify the REPLACE statement, PROC CV2VIEW does not convert that view.

```
TO_LIBREF= output-library;
```

Argument

output-library

specifies the name of a previously assigned library where you want to store the new SQL views.

Details

The names of the input view descriptors or access descriptors are used as the output view names. In order to individually name your output views, use the FROM_VIEW= statement and the TO_VIEW= statement.

TYPE= Statement

Specifies what type of conversion should occur

TYPE= SQL | VIEW | ACCESS;

Arguments

SQL

specifies that PROC CV2VIEW converts descriptors to SQL views. This is the default behavior.

VIEW

specifies that PROC CV2VIEW converts descriptors to native view descriptor format. This is most useful in the 32-bit to 64-bit case. This does not convert view descriptors across different operating systems.

ACCESS

specifies that PROC CV2VIEW converts access descriptors to native access descriptor format. This is most useful in the 32-bit to 64-bit case. This does not convert access descriptors across different operating systems.

Details

When TYPE=VIEW or TYPE=ACCESS, SAVEAS=, SUBMIT, and REPLACE= or REPLACE_FILE= are not valid options.

CV2VIEW Procedure Examples

Example 1: Converting an Individual View Descriptor

In this example, PROC CV2VIEW converts the MYVIEW view descriptor to the SQL view NEWVIEW. The MYVIEW view descriptor is ALTER, READ, and WRITE protected. The PROC SQL statements that are generated by PROC CV2VIEW are submitted and saved to an external file named SQL.SAS.

```
libname input '/username/descriptors/';
libname output '/username/sqlviews/';

proc cv2view dbms=oracle;
from_view = input.myview (alter=apwd);
to_view = output.newview;
saveas = '/username/vsql/sql.sas';
submit;
replace file;
run;
```

PROC CV2VIEW generates the following PROC SQL statements.

```

/* SOURCE DESCRIPTOR: MYVIEW */
PROC SQL DQUOTE=ANSI;
  CREATE VIEW OUTPUT.NEWVIEW
  (
/* READ= */
/* WRITE= */
/* ALTER= */
  LABEL=EMPLINFO
  )
  AS SELECT
    "EMPLOYEE " AS EMPLOYEE INFORMAT= 5.0 FORMAT= 5.0
      LABEL= 'EMPLOYEE ' ,
    "LASTNAME " AS LASTNAME INFORMAT= $10. FORMAT= $10.
      LABEL= 'LASTNAME ' ,
    "SEX " AS SEX INFORMAT= $6. FORMAT= $6.
      LABEL= 'SEX ' ,
    "STATUS " AS STATUS INFORMAT= $9. FORMAT= $9.
      LABEL= 'STATUS ' ,
    "DEPARTMENT" AS DEPARTME INFORMAT= 7.0 FORMAT= 7.0
      LABEL= 'DEPARTMENT' ,
    "CITYSTATE " AS CITYSTAT INFORMAT= $15. FORMAT= $15.
      LABEL= 'CITYSTATE '
  FROM _CVLIB_."EMPLINFO"
  USING LIBNAME _CVLIB_
  Oracle
/* PW= */
  USER=ordevxx PATH=OracleV8 PRESERVE_TAB_NAMES=YES;
  QUIT;

```

The REPLACE FILE statement causes an existing file named SQL.SAS to be overwritten. Without this statement the text would be appended to SQL.SAS if the user has the appropriate privileges.

The LABEL value of **EMPLINFO** is the name of the underlying database table that is referenced by the view descriptor.

If the underlying DBMS is Oracle or DB2, the CV2VIEW procedure adds the PRESERVE_TAB_NAMES= option to the embedded LIBNAME statement. You can then use CV2VIEW to access those tables with mixed-case or embedded-blank table names.

Note: This SQL syntax fails if you try to submit it because the PW field of the LIBNAME statement is replaced with a comment in order to protect the password. The ALTER, READ, and WRITE protection is commented out for the same reason. You can add the passwords to the code and then submit the SQL to re-create the view. Δ

Example 2: Converting a Library of View Descriptors for a Single DBMS

In this example, PROC CV2VIEW converts all of the Oracle view descriptors in the input library into SQL views. If an error occurs during the conversion of a view descriptor, the procedure moves to the next view. The PROC SQL statements that are generated by PROC CV2VIEW are both submitted and saved to an external file named SQL.SAS.

```
libname input '/username/descriptors/';
libname output '/username/sqlviews/';
proc cv2view dbms=oracle;
from_libref = input;
to_libref = output;
saveas = '/username/vsql/manyview.sas';
submit;
run;
```

PROC CV2VIEW generates the following PROC SQL statements for one of the views.

```
/* SOURCE DESCRIPTOR: PPCV2R */
PROC SQL DQUOTE=ANSI;
CREATE VIEW OUTPUT.PPCV2R
(
LABEL=EMPLOYEES
)
AS SELECT
"EMPID" " AS EMPID INFORMAT= BEST22. FORMAT= BEST22.
LABEL= 'EMPID' ,
"HIREDATE " AS HIREDATE INFORMAT= DATETIME16. FORMAT= DATETIME16.
LABEL= 'HIREDATE' ,
"JOBCODE " AS JOBCODE INFORMAT= BEST22. FORMAT= BEST22.
LABEL= 'JOBCODE' ,
"SEX " AS SEX INFORMAT= $1. FORMAT= $1.
LABEL= 'SEX'
FROM _CVLIB_."EMPLOYEES" (
SASDATEFMT = ( "HIREDATE"= DATETIME16. ) )
USING LIBNAME _CVLIB_
Oracle
/* PW= */
USER=ordevxx PATH=OracleV8 PRESERVE_TAB_NAMES=YES;
QUIT;
```

The SAVEAS= statement causes all of the generated SQL for all of the Oracle view descriptors to be stored in the MANYVIEW.SAS file.


If the underlying DBMS is Oracle or DB2, the CV2VIEW procedure adds the PRESERVE_TAB_NAMES= option to the embedded LIBNAME statement. You can then use CV2VIEW to access those tables with mixed-case or embedded-blank table names.

Example 3: Converting a Library of View Descriptors for All Supported DBMSs

In this example, PROC CV2VIEW converts all of the view descriptors that are in the input library and that access data in any supported DBMS. If an error occurs during the conversion of a view descriptor, then the procedure moves to the next view. The PROC SQL statements that are generated by PROC CV2VIEW are automatically submitted but are not saved to an external file (because the SAVEAS= statement is not used).


```
libname input '/username/descriptors/';  
libname output '/username/sqlviews/';
```

```
proc cv2view dbms=all;  
from_libref = input;  
to_libref = output;  
run;
```

Appendixes

<i>Appendix 1</i>	The ACCESS Procedure for Relational Databases	<i>311</i>
<i>Appendix 2</i>	The DBLOAD Procedure for Relational Databases	<i>329</i>
<i>Appendix 3</i>	Recommended Reading	<i>343</i>



APPENDIX

1

The ACCESS Procedure for Relational Databases

<i>Overview of the ACCESS Procedure for Relational Databases</i>	311
<i>Accessing DBMS Data</i>	311
<i>About ACCESS Procedure Statements</i>	312
<i>Procedure Syntax</i>	313
<i>PROC ACCESS Statement</i>	314
<i>Database Connection Statements</i>	314
<i>ASSIGN Statement</i>	315
<i>CREATE Statement</i>	315
<i>DROP Statement</i>	317
<i>FORMAT Statement</i>	317
<i>LIST Statement</i>	318
<i>QUIT Statement</i>	319
<i>RENAME Statement</i>	319
<i>RESET Statement</i>	320
<i>SELECT Statement</i>	321
<i>SUBSET Statement</i>	322
<i>TABLE= Statement</i>	323
<i>UNIQUE Statement</i>	323
<i>UPDATE Statement</i>	324
<i>Using Descriptors with the ACCESS Procedure</i>	325
<i>What Are Descriptors?</i>	325
<i>Access Descriptors</i>	325
<i>View Descriptors</i>	326
<i>Accessing Data Sets and Descriptors</i>	326
<i>Examples of Using the ACCESS Procedure</i>	327
<i>Example 1: Updating an Access Descriptor</i>	327
<i>Example 2: Creating a View Descriptor</i>	328

Overview of the ACCESS Procedure for Relational Databases

Accessing DBMS Data

The ACCESS procedure is still supported for the database systems and environments on which it was available in Version 6. However, it is no longer the recommended method for accessing relational DBMS data. It is recommended that you access your DBMS data more directly, using the LIBNAME statement or Pass-Through Facility .

Note: Not all SAS/ACCESS interfaces support this feature. See Chapter 8, “SAS/ACCESS Features by Host,” on page 65 to determine whether this feature is available in your environment. \triangle

This section provides general reference information for the ACCESS procedure; see the SAS/ACCESS documentation for your DBMS for DBMS-specific details.

The ACCESS procedure, along with the DBLOAD procedure and an interface view engine, creates an interface between SAS and data in other vendors’ databases. You can use the ACCESS procedure to create and update descriptors.

About ACCESS Procedure Statements

The ACCESS procedure has several types of statements:

- Database connection statements* are used to connect to your DBMS. For details, see the SAS/ACCESS documentation for your DBMS.
- Creating and updating statements* are CREATE and UPDATE.
- Table and editing statements* include ASSIGN, DROP, FORMAT, LIST, QUIT, RENAME, RESET, SELECT, SUBSET, TABLE, and UNIQUE.

The following table summarizes the PROC ACCESS options and statements that are required to accomplish common tasks.

Table A1.1 Statement Sequence for Accomplishing Tasks with the ACCESS Procedure

To do this...	Use these statements and options
Create an access descriptor	<pre>PROC ACCESS <i>statement-options</i>; CREATE <i>libref.member-name.ACCESS</i>; <i>database-connection-statements</i>; <i>editing-statements</i>; RUN;</pre>
Create an access descriptor and a view descriptor	<pre>PROC ACCESS <i>statement-options</i>; CREATE <i>libref.member-name.ACCESS</i>; <i>database-connection-statements</i>; <i>editing-statements</i>; CREATE <i>libref.member-name.VIEW</i>; SELECT <i>column-list</i>; <i>editing-statements</i>; RUN;</pre>
Create a view descriptor from an existing access descriptor	<pre>PROC ACCESS <i>statement-options</i>, including <i>ACCDESC=libref.access-descriptor</i>; CREATE <i>libref.member-name.VIEW</i>; SELECT <i>column-list</i>; <i>editing-statements</i>; RUN;</pre>
Update an access descriptor	<pre>PROC ACCESS <i>statement-options</i>; UPDATE <i>libref.member-name.ACCESS</i>; <i>database-connection-statements</i>; <i>editing-statements</i>; RUN;</pre>

To do this...	Use these statements and options
Update an access descriptor and a view descriptor	<pre> PROC ACCESS <i>statement-options</i>; UPDATE <i>libref.member-name.ACCESS</i>; <i>database-connection-statements</i>; <i>editing-statements</i>; UPDATE <i>libref.member-name.VIEW</i>; <i>editing-statements</i>; RUN; </pre>
Update an access descriptor and create a view descriptor	<pre> PROC ACCESS <i>statement-options</i>; UPDATE <i>libref.member-name.ACCESS</i>; <i>database-connection-statements</i>; <i>editing-statements</i>; CREATE <i>libref.member-name.VIEW</i>; SELECT <i>column-list</i>; <i>editing-statements</i>; RUN; </pre>
Update a view descriptor from an existing access descriptor	<pre> PROC ACCESS <i>statement-options</i>, including <i>ACCDESC=libref.access-descriptor</i>; UPDATE <i>libref.member-name.VIEW</i>; <i>editing-statements</i>; RUN; </pre>
Create a SAS data set from a view descriptor	<pre> PROC ACCESS <i>statement-options</i>, including <i>DBMS=dbms-name</i>; <i>VIEWDESC=libref.member</i>; <i>OUT=libref.member</i>; RUN; </pre>

Procedure Syntax

The general syntax for the ACCESS procedure is presented here. See the SAS/ACCESS documentation for your DBMS for DBMS-specific details.

```

PROC ACCESS<options>;
           database-connection-statements;
CREATE libref.member-name.ACCESS | VIEW <password-option>;
UPDATE libref.member-name.ACCESS | VIEW <password-option>;
TABLE= <'>table-name<'>;
ASSIGN <=>YES | NO | Y | N;
DROP <'>column-identifier-1<'> <...<'>column-identifier-n<'>>;
FORMAT <'>column-identifier-1<'> <=> SAS-format-name-1
           <...<'>column-identifier-n<'> <=> SAS-format-name-n>;
LIST <ALL | VIEW | <'>column-identifier<'>>;
QUIT;

```

```

RENAME <'>column-identifier-1<'> <=> SAS-variable-name-1
          <...>'>column-identifier-n<'> <=> SAS-variable-name-n>;
RESET ALL | <'>column-identifier-1<'> <...>'>column-identifier-n<'>>;
SELECT ALL | <'>column-identifier-1<'> <...>'>column-identifier-n<'>>;
SUBSET selection-criteria;
UNIQUE <=> YES | NO | Y | N;

RUN;

```

PROC ACCESS Statement

```
PROC ACCESS <options>;
```

Options

ACCDESC=libref.access-descriptor

specifies an access descriptor. ACCDESC= is used with the DBMS= option to create or update a view descriptor that is based on the specified access descriptor. You can use a SAS data set option on the ACCDESC= option to specify any passwords that have been assigned to the access descriptor.

Note: The ODBC interface does not support this option. Δ

DBMS=database-management-system

specifies which database management system you want to use. This DBMS-specific option is required. See the SAS/ACCESS documentation for your DBMS.

OUT=libref.member-name

specifies the SAS data file to which DBMS data is output.

VIEWDESC=libref.view-descriptor

specifies a view descriptor through which you extract the DBMS data.

Database Connection Statements

Provide DBMS-specific connection information

```
database-connection-statements;
```

Database connection statements are used to connect to your DBMS. For the statements to use with your DBMS, see the documentation for your SAS/ACCESS interface.

ASSIGN Statement

Indicates whether SAS variable names and formats are generated

Applies to: access descriptor

Interacts with: FORMAT, RENAME, RESET, UNIQUE

Default: NO

ASSIGN <=>YES | NO | Y | N;

YES

generates unique SAS variable names from the first eight characters of the DBMS column names. If you specify **YES**, you cannot specify the RENAME, FORMAT, RESET, or UNIQUE statements when you create view descriptors that are based on the access descriptor.

NO

lets you modify SAS variable names and formats when you create an access descriptor and when you create view descriptors that are based on this access descriptor.

Details

The ASSIGN statement indicates how SAS variable names and formats are assigned:

- SAS automatically generates SAS variable names.
- You can change SAS variable names and formats in the view descriptors that are created from the access descriptor.

Each time the SAS/ACCESS interface encounters a CREATE statement to create an access descriptor, the ASSIGN statement is reset to the default **NO** value.

When you create an access descriptor, use the RENAME statement to change SAS variable names and the FORMAT statement to change SAS formats.

When you specify **YES**, SAS generates names according to these rules:

- You can change the SAS variable names only in the access descriptor.
- SAS variable names that are saved in an access descriptor are *always* used when view descriptors are created from the access descriptor. You cannot change them in the view descriptors.
- The ACCESS procedure allows names only up to eight characters.

CREATE Statement

Creates a SAS/ACCESS descriptor file

Applies to: access descriptor or view descriptor

CREATE *libref.member-name*.ACCESS | VIEW <password-option>;

libref.member-name

identifies the libref of the SAS library where you want to store the descriptor and identifies the descriptor name.

ACCESS

specifies an access descriptor.

VIEW

specifies a view descriptor.

password-option

specifies a password.

Details

The CREATE statement is required. It names the access descriptor or view descriptor that you are creating. Use a three-level name:

- The first level identifies the libref of the SAS library where you want to store the descriptor.
- The second level is the descriptor name,
- The third level specifies the type of SAS file (specify **ACCESS** for an access descriptor or **VIEW** for a view descriptor).

See Statement Sequence for Accomplishing Tasks with the ACCESS ProcedureTable A1.1 on page 312 for the appropriate sequence of statements for creating access and view descriptors.

Example

The following example creates an access descriptor AdLib.Employ on the Oracle table Employees, and a view descriptor Vlib.Emp1204 based on AdLib.Employ, in the same PROC ACCESS step.

```
proc access dbms=oracle;

    /* create access descriptor */

    create adlib.employ.access;
    database='qa:[dubois]textile';
    table=employees;
    assign=no;
    list all;

    /* create view descriptor */

    create vlib.empl204.view;
    select empid lastname hiredate salary dept
    gender birthdate;
    format empid 6.
           salary dollar12.2
           jobcode 5.
           hiredate datetime9.
           birthdate datetime9.;
    subset where jobcode=1204;
run;
```

DROP Statement

Drops a column so that it cannot be selected in a view descriptor

Applies to: access and view descriptors

Interacts with: RESET, SELECT

DROP <'>*column-identifier-1*<'> <...<'>*column-identifier-n*<'>>;

column-identifier

specifies the column name or the positional equivalent from the LIST statement, which is the number that represents the column's place in the access descriptor. For example, to drop the third and fifth columns, submit the following statement:

```
drop 3 5;
```

Details

The DROP statement drops one or more specified columns from a descriptor. You can drop a column when creating or updating an access descriptor; you can also drop a column when updating a view descriptor. If you drop a column when creating an access descriptor, you cannot select that column when creating a view descriptor that is based on the access descriptor. The underlying DBMS table is unaffected by this statement.

To display a column that was previously dropped, specify that column name in the RESET statement. However, doing so also resets all the column's attributes (such as SAS variable name, format, and so on) to their default values.

FORMAT Statement

Changes a SAS format for a DBMS column

Applies to: access descriptor or view descriptor

Interacts with: ASSIGN, DROP, RESET

FORMAT <'>*column-identifier-1*<'> <=>*SAS-format-name-1*
<...<'>*column-identifier-n*<'> <=> *SAS-format-name-n*>;

column-identifier

specifies the column name or the positional equivalent from the LIST statement, which is the number that represents the column's place in the access descriptor. If the column name contains lowercase characters, special characters, or national characters, enclose the name in quotation marks.

SAS-format-name

specifies the SAS format to be used.

Details

The FORMAT statement changes SAS variable formats from their default formats. The default SAS variable format is based on the data type of the DBMS column. See the SAS/ACCESS documentation for your DBMS for information about default formats that SAS assigns to your DBMS data types.

You can use the FORMAT statement with a view descriptor only if the ASSIGN statement that was used when creating the access descriptor was specified with the **NO** value. When you use the FORMAT statement with access descriptors, the FORMAT statement also reselects columns that were previously dropped with the DROP statement.

For example, to associate the DATE9. format with the BIRTHDATE column and with the second column in the access descriptor, submit the following statement:

```
format 2=date9. birthdate=date9.;
```

The equal sign (=) is optional. For example, you can use the FORMAT statement to specify new SAS variable formats for four DBMS table columns:

```
format productid 4.
      weight      e16.9
      fibersize   e20.13
      width       e16.9;
```

LIST Statement

Lists columns in the descriptor and gives information about them

Applies to: access descriptor or view descriptor

Default: ALL

LIST <ALL | VIEW | <'>column-identifier<'>>;

ALL

lists all the DBMS columns in the table, the positional equivalents, the SAS variable names, and the SAS variable formats that are available for a descriptor.

VIEW

lists all the DBMS columns that are selected for a view descriptor, their positional equivalents, their SAS names and formats, and any subsetting clauses.

column-identifier

lists information about a specified DBMS column, including its name, positional equivalent, SAS variable name and format, and whether it has been selected. If the column name contains lowercase characters, special characters, or national characters, enclose the name in quotation marks.

The *column-identifier* argument can be either the column name or the positional equivalent, which is the number that represents the column's place in the descriptor. For example, to list information about the fifth column in the descriptor, submit the following statement:

```
list 5;
```

Details

The LIST statement lists columns in the descriptor, along with information about the columns. The LIST statement can be used only when *creating* an access descriptor or a view descriptor. The LIST information is written to your SAS log.

Note: To review the contents of an existing view descriptor, use the CONTENTS procedure. △

When you use LIST for an access descriptor, ***NON-DISPLAY*** appears next to the column description for any column that has been dropped; ***UNSUPPORTED*** appears next to any column whose data type is not supported by your DBMS interface view engine. When you use LIST for a view descriptor, ***SELECTED*** appears next to the column description for columns that you have selected for the view.

Specify LIST last in your PROC ACCESS code in order to see the entire descriptor. If you create or update multiple descriptors, specify LIST before each CREATE or UPDATE statement to list information about all descriptors that you are creating or updating.

QUIT Statement

Terminates the procedure

Applies to: access descriptor or view descriptor

QUIT;

Details

The QUIT statement terminates the ACCESS procedure without any further descriptor creation. Changes made since the last CREATE, UPDATE, or RUN statement are not saved; changes are saved only when a new CREATE, UPDATE, or RUN statement is submitted.

RENAME Statement

Modifies the SAS variable name

Applies to: access descriptor or view descriptor

Interacts with: ASSIGN, RESET

RENAME <'>column-identifier-1<'> <=> SAS-variable-name-1
<...<'>column-identifier-n<'> <=> SAS-variable-name-n>;

column-identifier

specifies the DBMS column name or the positional equivalent from the LIST statement, which is the number that represents the column's place in the descriptor. If the column name contains lowercase characters, special characters, or national characters, enclose the name in quotation marks. The equal sign (=) is optional.

SAS-variable-name

specifies a SAS variable name.

Details

The RENAME statement sets or modifies the SAS variable name that is associated with a DBMS column.

Two factors affect the use of the RENAME statement: whether you specify the ASSIGN statement when you are creating an access descriptor, and the kind of descriptor you are creating.

- If you omit the ASSIGN statement or specify it with a **NO** value, the renamed SAS variable names that you specify in the access descriptor are retained when an ACCESS procedure executes. For example, if you rename the CUSTOMER column to CUSTNUM when you create an access descriptor, the column is still named CUSTNUM when you select it in a view descriptor unless you specify another RESET statement or RENAME statement.

When you create a view descriptor that is based on this access descriptor, you can specify the RESET statement or another RENAME statement to rename the variable. However, the new name applies only in that view. When you create other view descriptors, the SAS variable names are derived from the access descriptor.

- If you specify the **YES** value in the ASSIGN statement, you can use the RENAME statement to change SAS variable names only while creating an access descriptor. Remember that SAS variable names and formats that are saved in an access descriptor are always used when creating view descriptors that are based on the access descriptor.

For example, to rename the SAS variable names that are associated with the seventh column and the nine-character FIRSTNAME column in a descriptor, submit the following statement:

```
rename
7 birthdy 'firstname'=fname;
```

Note: When you are creating a view descriptor, the RENAME statement automatically selects the renamed column for the view. That is, if you rename the SAS variable associated with a DBMS column, you do not have to issue a SELECT statement for that column. \triangle

RESET Statement

Resets DBMS columns to their default settings

Applies to: access descriptor or view descriptor

Interacts with: ASSIGN, DROP, FORMAT, RENAME, SELECT

RESET ALL | <'>column-identifier-1<'> <...<'>column-identifier-n<'>>;

ALL

resets all columns in an access descriptor to their default names and formats and reselects any dropped columns. **ALL** deselects all columns in a view descriptor so that no columns are selected for the view.

column-identifier

can be either the DBMS column name or the positional equivalent from the **LIST** statement, which is the number that represents the column's place in the access descriptor. If the column name contains lowercase characters, special characters, or national characters, enclose the name in quotation marks. For example, to reset the SAS variable name and format associated with the third column, submit the following statement:

```
reset
3;
```

For access descriptors, the specified column is reset to its default name and format settings. For view descriptors, the specified column is no longer selected for the view.

Details

The **RESET** statement resets column attributes to their default values. This statement has different effects on access and view descriptors.

For *access descriptors*, the **RESET** statement resets the specified column names to the default names that are generated by the **ACCESS** procedure. The **RESET** statement also changes the current SAS variable format to the default SAS format. Any previously-dropped columns that are specified in the **RESET** statement become available.

Note: When creating an access descriptor, if you omit the **ASSIGN** statement or set it to **NO**, the default SAS variable names are blanks. If you set **ASSIGN=YES**, default names are the first eight characters of each DBMS column name. △

For *view descriptors*, the **RESET** statement clears (deselects) any columns that were included in the **SELECT** statement. When you create a view descriptor that is based on an access descriptor that is created without an **ASSIGN** statement or with **ASSIGN=NO**, resetting and then reselecting (within the same procedure execution) a SAS variable changes the SAS variable names and formats to their default values. When you create a view descriptor that is based on an access descriptor created with **ASSIGN=YES**, the **RESET** statement does not have this effect.

SELECT Statement

Selects DBMS columns for the view descriptor

Applies to: view descriptor

Interacts with: **RESET**

SELECT ALL | <'>column-identifier-1<'> <...<'>column-identifier-n <'>>;

ALL

includes in the view descriptor all the columns that were defined in the access descriptor and that were not dropped.

column-identifier

can be either the DBMS column name or the positional equivalent from the LIST statement. The positional equivalent is the number that represents the where the column is located in the access descriptor on which the view is based. For example, to select the first three columns, submit the following statement:

```
select 1 2 3;
```

If the column name contains lowercase characters, special characters, or national characters, enclose the name in quotation marks.

Details

The SELECT statement is required. The SELECT statement specifies which DBMS columns in an access descriptor to include in a view descriptor.

SELECT statements are cumulative within a view creation. That is, if you submit the following two SELECT statements, columns 1, 5, and 6 are selected:

```
select 1;
select 5 6;
```

To clear your current selections when creating a view descriptor, use the **RESET ALL** statement.

SUBSET Statement

Adds or modifies selection criteria for a view descriptor

Applies to: view descriptor

SUBSET *selection-criteria*;

selection-criteria

one or more DBMS-specific SQL expressions that are accepted by your DBMS, such as WHERE, ORDER BY, HAVING, and GROUP BY. Use DBMS column names, not SAS variable names, in your selection criteria.

Details

You can use the SUBSET statement to specify selection criteria when you create a view descriptor. This statement is optional; if you omit it, the view retrieves all the data (that is, all the rows) in the DBMS table.

For example, for a view descriptor that retrieves rows from a DBMS table, you could submit the following SUBSET statement:

```
subset where firstorder is not null;
```

If you have multiple selection criteria, enter them all in one SUBSET statement, as in the following example:

```
subset where firstorder is not null
       and country = 'USA'
       order by country;
```

Unlike other ACCESS procedure statements, the SUBSET statement is case-sensitive. The SQL statement is sent to the DBMS exactly as you type it. Therefore, you must use the correct case for any DBMS object names. See the SAS/ACCESS documentation for your DBMS for details.

SAS does not check the SUBSET statement for errors. The statement is verified only when the view descriptor is used in a SAS program.

If you specify more than one SUBSET statement per view descriptor, the last SUBSET overwrites the earlier SUBSETs. To delete the selection criteria, submit a SUBSET statement without any arguments.

TABLE= Statement

Identifies the DBMS table on which the access descriptor is based

Applies to: access descriptor

TABLE= <'>*table-name*<'>;

table-name

a valid DBMS table name. If it contains lowercase characters, special characters, or national characters, you must enclose it in quotation marks. See the SAS/ACCESS documentation for your DBMS for details on the TABLE= statement.

Details

This statement is required with the CREATE statement and optional with the UPDATE statement.

UNIQUE Statement

Generates SAS variable names based on DBMS column names

Applies to: view descriptor

Interacts with: ASSIGN

UNIQUE <=> YES | NO | Y | N;

YES

causes the SAS/ACCESS interface to append numbers to any duplicate SAS variable names, thus making each variable name unique.

NO

causes the SAS/ACCESS interface to continue to allow duplicate SAS variable names to exist. You must resolve these duplicate names before saving (and thereby creating) the view descriptor.

Details

The UNIQUE statement specifies whether the SAS/ACCESS interface should generate unique SAS variable names for DBMS columns for which SAS variable names have not been entered.

The UNIQUE statement is affected by whether you specified the ASSIGN statement when you created the access descriptor on which the view is based:

- If you specified the ASSIGN=YES statement, you cannot specify UNIQUE when creating a view descriptor. **YES** causes SAS to generate unique names, so UNIQUE is not necessary.
- If you omitted the ASSIGN statement or specified ASSIGN=NO, you must resolve any duplicate SAS variable names in the view descriptor. You can use UNIQUE to generate unique names automatically, or you can use the RENAME statement to resolve duplicate names yourself. See “RENAME Statement” on page 319 for information.

If duplicate SAS variable names exist in the access descriptor on which you are creating a view descriptor, you can specify UNIQUE to resolve the duplication.

Note: It is recommended that you use the UNIQUE statement and specify UNIQUE=YES. If you omit the UNIQUE statement or specify UNIQUE=NO and SAS encounters duplicate SAS variable names in a view descriptor, your job fails. Δ

The equal sign (=) is optional in the UNIQUE statement.

UPDATE Statement

Updates a SAS/ACCESS descriptor file

Applies to: access descriptor or view descriptor

UPDATE *libref.member-name*.ACCESS | VIEW <password-option>;

libref.member-name

identifies the libref of the SAS library where you want to store the descriptor and identifies the descriptor name.

ACCESS

specifies an access descriptor.

VIEW

specifies a view descriptor.

password-option

specifies a password.

Details

The UPDATE statement identifies an existing access descriptor or view descriptor that you want to update. UPDATE is normally used to update database connection information, such as user IDs and passwords. If your descriptor requires many changes, it might be easier to use the CREATE statement to overwrite the old descriptor with a new one.

Note: Altering a DBMS table might invalidate descriptor files that are based on the DBMS table, or it might cause these files to be out of date. If you re-create a table, add a new column to a table, or delete an existing column from a table, use the UPDATE statement to modify your descriptors so that they use the new information. △

Rules that apply to the CREATE statement also apply to the UPDATE statement. For example, the SUBSET statement is valid only for updating view descriptors.

Note: The following statements are not supported when using the UPDATE statement: ASSIGN, RESET, SELECT, and UNIQUE. △

See Table A1.1 on page 312 for the appropriate sequence of statements for updating descriptors.

Using Descriptors with the ACCESS Procedure

What Are Descriptors?

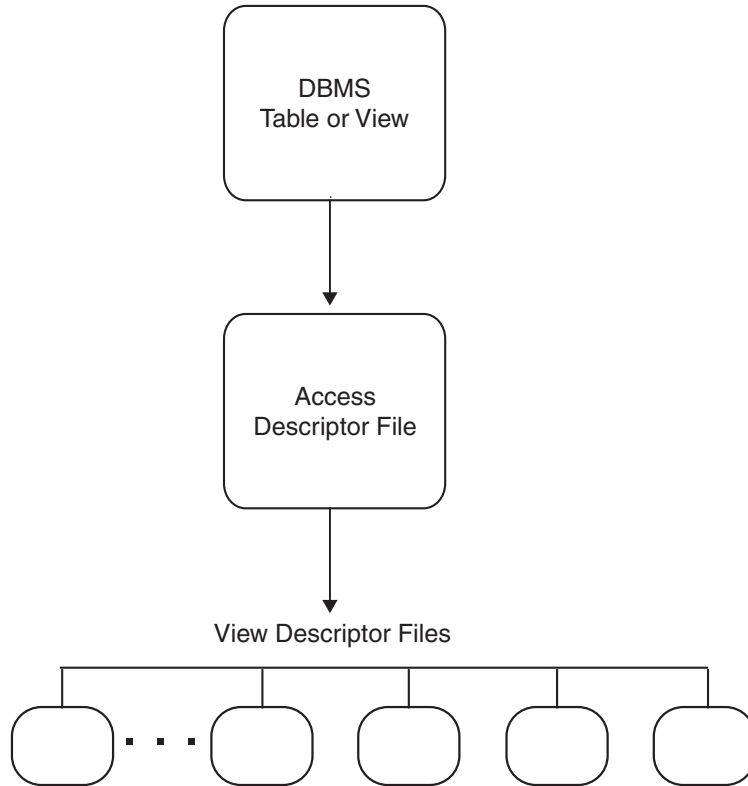
Descriptors work with the ACCESS procedure by providing information about DBMS objects to SAS, enabling you to access and update DBMS data from within a SAS session or program.

There are two types of descriptors, *access descriptors* and *view descriptors*. Access descriptors provide SAS with information about the structure and attributes of a DBMS table or view. An access descriptor, in turn, is used to create one or more view descriptors, or SAS data views, of the DBMS data.

Access Descriptors

Typically, each DBMS table or view has a single access descriptor that provides connection information, data type information, and names for databases, tables, and columns.

You use an access descriptor to create one or more view descriptors. When creating a view descriptor, you select the columns and specify criteria for the rows you want to retrieve. The figure below illustrates the descriptor creation process. Note that an access descriptor, which contains the metadata of the DBMS table, must be created before view descriptors can be created.

Figure 14.1 Creating an Access Descriptor and View Descriptors for a DBMS Table

View Descriptors

You use a view descriptor in a SAS program much as you would any SAS data set. For example, you can specify a view descriptor in the `DATA=` statement of a SAS procedure or in the `SET` statement of a `DATA` step.

You can also use a view descriptor to copy DBMS data into a SAS data file, which is called *extracting* the data. When you need to use DBMS data in several procedures or `DATA` steps, you might use fewer resources by extracting the data into a SAS data file instead of repeatedly accessing the data directly.

Note: The SAS/ACCESS interface view engine usually tries to pass `WHERE` conditions to the DBMS for processing. In most cases it is more efficient for a DBMS to process `WHERE` conditions than for SAS to do the processing. \triangle

Accessing Data Sets and Descriptors

SAS let you control access to SAS data sets and access descriptors by associating one or more SAS passwords with them. When you create an access descriptor, the connection information that you provide is stored in the access descriptor and in any view descriptors based on that access descriptor. The password is stored in an encrypted form. When these descriptors are accessed, the connection information that was stored is also used to access the DBMS table or view. To ensure data security, you might want to change the protection on the descriptors to prevent others from seeing the connection information stored in the descriptors.

When you create or update view descriptors, you can use a SAS data set option after the `ACCDESC=` option to specify the access descriptor password, if one exists. In this

case, you are *not* assigning a password to the view descriptor that is being created or updated. Instead, using the password grants you permission to use the access descriptor to create or update the view descriptor. Here is an example:

```
proc access dbms=sybase accdesc=adlib.customer
            (alter=rouge);
  create vlib.customer.view;
  select all;
run;
```

By specifying the ALTER level of password, you can read the AdLib.Customer access descriptor and create the Vlib.Customer view descriptor.

Examples of Using the ACCESS Procedure

Example 1: Updating an Access Descriptor

The following example updates an access descriptor AdLib.Employ on the Oracle table Employees. The original access descriptor includes all of the columns in the table. The updated access descriptor omits the Salary and BirthDate columns.

```
proc access dbms=oracle ad=adlaib.employ;

  /* update access descriptor */

  update adlib.employ.access;
  drop salary birthdate;
  list all;
run;
```

You can use the LIST statement to write all variables to the SAS log so that you can see the complete access descriptor before you update it.

Example 2: Creating a View Descriptor


The following example re-creates a view descriptor, VLIB.EMP1204, which is based on an access descriptor, ADLIB.EMPLOY, which was previously updated.

```
proc access dbms=oracle;

    /* re-create view descriptor */

    create vlib.emp1204.view;
    select empid hiredate dept jobcode gender
           lastname firstname middlename phone;
    format empid 6.
           jobcode 5.
           hiredate datetime9.;
    subset where jobcode=1204;
run;
```

Because **SELECT** and **RESET** are not supported when **UPDATE** is used, the view descriptor Vlib.Emp1204 must be re-created to omit the Salary and BirthDate columns.



APPENDIX

2

The DBLOAD Procedure for Relational Databases

<i>Overview of the DBLOAD Procedure for Relational Databases</i>	329
<i>Sending Data from SAS to a DBMS</i>	329
<i>Properties of the DBLOAD Procedure</i>	330
<i>About DBLOAD Procedure Statements</i>	330
<i>Procedure Syntax</i>	331
<i>PROC DBLOAD Statement</i>	332
<i>Database Connection Statements</i>	333
<i>ACCDDESC= Statement</i>	333
<i>COMMIT= Statement</i>	334
<i>DELETE Statement</i>	334
<i>ERRLIMIT= Statement</i>	334
<i>LABEL Statement</i>	335
<i>LIMIT= Statement</i>	335
<i>LIST Statement</i>	336
<i>LOAD Statement</i>	336
<i>NULLS Statement</i>	337
<i>QUIT Statement</i>	338
<i>RENAME Statement</i>	338
<i>RESET Statement</i>	339
<i>SQL Statement</i>	340
<i>TABLE= Statement</i>	340
<i>TYPE Statement</i>	341
<i>WHERE Statement</i>	341
<i>Example of Using the DBLOAD Procedure</i>	342
<i>Example 1: Appending a Data Set to a DBMS Table</i>	342

Overview of the DBLOAD Procedure for Relational Databases

Sending Data from SAS to a DBMS

The DBLOAD procedure is still supported for the database systems and environments on which it was available in Version 6. However, it is no longer the recommended method for sending data from SAS to a DBMS. It is recommended that you access your DBMS data more directly, using the LIBNAME statement or the Pass-Through Facility.

Note: Not all SAS/ACCESS interfaces support this feature. See Chapter 8, “SAS/ACCESS Features by Host,” on page 65 to determine whether this feature is available in your environment. \triangle

Properties of the DBLOAD Procedure

This section provides general reference information for the DBLOAD procedure; see the SAS/ACCESS documentation for your DBMS for DBMS-specific details.

The DBLOAD procedure, along with the ACCESS procedure and an interface view engine, creates an interface between SAS and data in other vendors' databases.

The DBLOAD procedure enables you to create and load a DBMS table, append rows to an existing table, and submit non-query DBMS-specific SQL statements to the DBMS for processing. The procedure constructs DBMS-specific SQL statements to create and load, or append, to a DBMS table by using one of the following:

- a SAS data file
- an SQL view or DATA step view
- a view descriptor that was created with the SAS/ACCESS interface to your DBMS or with another SAS/ACCESS interface product
- another DBMS table referenced by a SAS libref that was created with the SAS/ACCESS LIBNAME statement.

The DBLOAD procedure associates each SAS variable with a DBMS column and assigns a default name and data type to each column. It also specifies whether each column accepts NULL values. You can use the default information or change it as necessary. When you are finished customizing the columns, the procedure creates the DBMS table and loads or appends the input data.

About DBLOAD Procedure Statements

There are several types of DBLOAD statements:

- Database connection statements* are used to connect to your DBMS. See the SAS/ACCESS documentation for your DBMS for details.
- Creating and loading statements* are LOAD and RUN.
- Table and editing statements* are used to specify how a table is populated.

The following table summarizes the PROC DBLOAD options and statements required to accomplish common tasks.

Table A2.1 Statement Sequence for Accomplishing Common Tasks with the DBLOAD Procedure

To do this...	Use these options and statements
Create and load a DBMS table	PROC DBLOAD <i>statement-options;</i> <i>database-connection-options;</i> TABLE= <'>table-name<'>; LOAD; RUN;
Submit a dynamic, non-query DBMS-SQL statement to DBMS (without creating a table)	PROC DBLOAD <i>statement-options;</i> <i>database-connection-options;</i> SQL <i>DBMS-specific-SQL-statements;</i> RUN;

Note: LOAD must appear before RUN to create and load a table or append data to a table. △

Procedure Syntax

The general syntax for the DBLOAD procedure is presented here; see the SAS/ACCESS documentation for your DBMS for DBMS-specific details.

```

PROC DBLOAD <options>;
database connection statements;
TABLE= <'>table-name<'>;
ACCDESC= <libref.>access-descriptor;
COMMIT= commit-frequency;
DELETE variable-identifier-1
    <...variable-identifier-n>;
ERRLIMIT= error-limit;
LABEL;
LIMIT= load-limit;
LIST <ALL | COLUMN | variable-identifier>;
NULLS variable-identifier-1 = Y | N
    <...variable-identifier-n = Y | N>;
QUIT;
RENAME variable-identifier-1 = <'>column-name-1<'>
    <...variable-identifier-n = <'>column-name-n<'>>;
RESET ALL | variable-identifier-1 <...variable-identifier-n>;
SQL DBMS-specific-SQL-statement;
TYPE variable-identifier-1 = 'column-type-1' <...variable-identifier-n = 'column-type-n'>;

```

```

WHERE SAS-where-expression;
LOAD;

RUN;

```

PROC DBLOAD Statement

```
PROC DBLOAD <options>;
```

Options

DBMS=*database-management-system*

specifies which database management system you want to access. This DBMS-specific option is required. See the SAS/ACCESS documentation for your DBMS.

DATA=<*libref*>*SAS-data-set*

specifies the input data set. You can retrieve input data from a SAS data file, an SQL view, a DATA step view, a SAS/ACCESS view descriptor, or another DBMS table to which a SAS/ACCESS libref points. If the SAS data set is permanent, you must use its two-level name, *libref.SAS-data-set*. If you omit the DATA= option, the default is the last SAS data set that was created.

APPEND

appends data to an existing DBMS table that you identify by using the TABLE= statement. When you specify APPEND, the input data specified with the DATA= option is inserted into the existing DBMS table. Your input data can be in the form of a SAS data set, SQL view, or SAS/ACCESS view (view descriptor).

CAUTION:

When you use APPEND, you must ensure that your input data corresponds exactly to the columns in the DBMS table. If your input data does not include values for all columns in the DBMS table, you might corrupt your DBMS table by inserting data into the wrong columns. Use the COMMIT, ERRLIMIT, and LIMIT statements to help safeguard against data corruption. Use the DELETE and RENAME statements to drop and rename SAS input variables that do not have corresponding DBMS columns. \triangle

All PROC DBLOAD statements and options can be used with APPEND, except for the NULLS and TYPE statements, which have no effect when used with APPEND. The LOAD statement is required.

The following example appends new employee data from the SAS data set NEWEMP to the DBMS table EMPLOYEES. The COMMIT statement causes a DBMS commit to be issued after every 100 rows are inserted. The ERRLIMIT statement causes processing to stop after five errors occur.

```
proc dbload dbms=oracle data=newemp append;
  user=testuser;
  password=testpass;
  path='myorapath';
  table=employees;
  commit=100;
  errlimit=5;
  load;
run;
```

Note: By omitting the APPEND option from the DBLOAD statement, you can use the PROC DBLOAD SQL statements to create a DBMS table and append to it in the same PROC DBLOAD step. △

Database Connection Statements

Provide DBMS connection information

database-connection-statements

These statements are used to connect to your DBMS and vary depending on which SAS/ACCESS interface you are using. See the documentation for your SAS/ACCESS interface for details. Examples include USER=, PASSWORD=, and DATABASE=.

ACCDESC= Statement

Creates an access descriptor based on the new DBMS table

ACCDESC=<libref.>access-descriptor;

Details

The ACCDESC= statement creates an access descriptor based on the DBMS table that you are creating and loading. If you specify ACCDESC=, the access descriptor is automatically created after the new table is created and loaded. You must specify an access descriptor if it does not already exist.

COMMIT= Statement

Issues a commit or saves rows after a specified number of inserts

Default: 1000

COMMIT=*commit-frequency*;

Details

The COMMIT= statement issues a commit (that is, generates a DBMS-specific SQL COMMIT statement) after the specified number of rows has been inserted.

Using this statement might improve performance by releasing DBMS resources each time the specified number of rows has been inserted.

If you omit the COMMIT= statement, a commit is issued (or a group of rows is saved) after each 1,000 rows are inserted and after the last row is inserted.

The *commit-frequency* argument must be a non-negative integer.

DELETE Statement

Does not load specified variables into the new table

DELETE *variable-identifier-1* <...*variable-identifier-n*>;

Details

The DELETE statement drops the specified SAS variables before the DBMS table is created. The *variable-identifier* argument can be either the SAS variable name or the positional equivalent from the LIST statement. The positional equivalent is the number that represents the variable's place in the data set. For example, if you want to drop the third variable, submit the following statement:

```
delete 3;
```

When you drop a variable, the positional equivalents of the variables do not change. For example, if you drop the second variable, the third variable is still referenced by the number 3, not 2. If you drop more than one variable, separate the identifiers with spaces, not commas.

ERRLIMIT= Statement

Stops the loading of data after a specified number of errors

Default: 100; see the SAS/ACCESS documentation for your DBMS for possible exceptions

ERRLIMIT=*error-limit*;

Details

The **ERRLIMIT**= statement stops the loading of data after the specified number of DBMS SQL errors has occurred. Errors include observations that fail to be inserted and commits that fail to execute. The **ERRLIMIT**= statement defaults to 10 when used with **APPEND**.

The *error-limit* argument must be a non-negative integer. To allow an unlimited number of DBMS SQL errors to occur, specify **ERRLIMIT**=0. If the SQL **CREATE TABLE** statement that is generated by the procedure fails, the procedure terminates.

LABEL Statement

Causes DBMS column names to default to SAS labels

Interacts with: RESET

Default: DBMS column names default to SAS variable names

LABEL;

Details

The **LABEL** statement causes the DBMS column names to default to the SAS variable labels when the new table is created. If a SAS variable has no label, the variable name is used. If the label is too long to be a valid DBMS column name, the label is truncated.

You must use the **RESET** statement *after* the **LABEL** statement for the **LABEL** statement to take effect.

LIMIT= Statement

Limits the number of observations that are loaded

Default: 5000

LIMIT=*load-limit*;

Details

The **LIMIT**= statement places a limit on the number of observations that can be loaded into the new DBMS table. The *load-limit* argument must be a non-negative integer. To load all the observations from your input data set, specify **LIMIT**=0.

LIST Statement

Lists information about the variables to be loaded

Default: ALL

LIST <ALL | FIELD | *variable-identifier*>;

Details

The LIST statement lists information about some or all of the SAS variables to be loaded into the new DBMS table. By default, the list is sent to the SAS log.

The LIST statement can take the following arguments:

ALL

lists information about all variables in the input SAS data set, despite whether those variables are selected for the load.

FIELD

lists information about only the input SAS variables that are selected for the load.

variable-identifier

lists information about only the specified variable. The *variable-identifier* argument can be either the SAS variable name or the positional equivalent. The positional equivalent is the number that represents the variable's position in the data set. For example, if you want to list information for the column associated with the third SAS variable, submit the following statement:

```
list 3;
```

You can specify LIST as many times as you want while creating a DBMS table; specify LIST before the LOAD statement to see the entire table.

LOAD Statement

Creates and loads the new DBMS table

Valid: in the DBLOAD procedure (required statement for loading or appending data)

LOAD;

Details

The LOAD statement informs the DBLOAD procedure to execute the action that you request, including loading or appending data. This statement is required to create and load a new DBMS table or to append data to an existing table.

When you create and load a DBMS table, you must place statements or groups of statements in a certain order after the PROC DBLOAD statement and its options, as listed in Table A2.1 on page 331 .

Example

The following example creates the SummerTemps table in Oracle based on the DLib.TempEmps data file.

```
proc dbload dbms=oracle data=dlib.tempemps;
  user=testuser; password=testpass;
  path='testpath';
  table=summertemps;
  rename firstnam=firstname
         middlena=middlename;
  type hiredate 'date'
         empid 'number(6,0)'
         familyid 'number(6,0)';
  nulls l=n;
  list;
  load;
run;
```

NULLS Statement

Specifies whether DBMS columns accept NULL values

Default: Y

NULLS *variable-identifier-1* = Y | N <...*variable-identifier-n* = Y | N>;

Details

Note: Some DBMSs have three valid values for this statement, Y, N, and D. See the SAS/ACCESS documentation for your DBMS for further details. △

The NULLS statement specifies whether the DBMS columns that are associated with the listed input SAS variables allow NULL values. Specify Y to accept NULL values. Specify N to reject NULL values and to require data in that column.

If you specify N for a numeric column, no observations that contain missing values in the corresponding SAS variable are loaded into the table. A message is written to the SAS log, and the current error count increases by one for each observation that is not loaded. See “ERRLIMIT= Statement” on page 334 for more information.

If a character column contains blanks (the SAS missing value) and you have specified N for the DBMS column, then blanks are inserted. If you specify Y, NULL values are inserted.

The *variable-identifier* argument can be either the SAS variable name or the positional equivalent from the LIST statement. The positional equivalent is the number that represents the variable’s place in the data set. For example, if you want the

column that is associated with the third SAS variable to accept NULL values, submit the following statement:

```
nulls 3=y;
```

If you omit the NULLS statement, the DBMS default action occurs. You can list as many variables as you want in one NULLS statement. If you have previously defined a column as NULLS=N, you can use the NULLS statement to redefine it to accept NULL values.

QUIT Statement

Terminates the procedure

Valid: in the DBLOAD procedure (control statement)

```
QUIT;
```

Details

The QUIT statement terminates the DBLOAD procedure without further processing.

RENAME Statement

Renames DBMS columns

Interacts with: DELETE, LABEL, RESET

```
RENAME variable-identifier-1 = <'>column-name-1<'> <...variable-identifier-n =  
<'>column-name-n<'>>;
```

Details

The RENAME statement changes the names of the DBMS columns that are associated with the listed SAS variables. If you omit the RENAME statement, all the DBMS column names default to the corresponding SAS variable names (unless the LABEL statement is specified).

The *variable-identifier* argument can be either the SAS variable name or the positional equivalent from the LIST statement. The positional equivalent is the number that represents the variable's place in the data set. For example, if you want to rename the column associated with the third SAS variable, submit the following statement:

```
rename 3=employeename;
```

Note: The *column-name* argument must be a valid DBMS column name. If the column name includes lowercase characters, special characters, or national characters, you must enclose the column name in single or double quotation marks. If no quotation marks are used, the DBMS column name is created in uppercase. To preserve case, use the following syntax: `rename 3="employee"` △

The RENAME statement enables you to include variables that you have previously deleted. For example, suppose you submit the following statements:

```
delete 3;
rename 3=empname;
```

The DELETE statement drops the third variable. The RENAME statement includes the third variable and assigns the name EMPNAME and the default column type to it.

You can list as many variables as you want in one RENAME statement. The RENAME statement overrides the LABEL statement for columns that are renamed. COLUMN is an alias for the RENAME statement.

RESET Statement

Resets column names and data types to their default values

Interacts with: DELETE, LABEL, RENAME, TYPE

RESET ALL | *variable-identifier-1* <...*variable-identifier-n*>;

Details

The RESET statement resets columns that are associated with the listed SAS variables to default values for the DBMS column name, column data type, and ability to accept NULL values. If you specify ALL, all columns are reset to their default values, and any dropped columns are restored with their default values. Here are the default values:

- column name
 - defaults to the SAS variable name, or to the SAS variable label (if you have used the LABEL statement).
- column type
 - is generated from the SAS variable format.
- nulls
 - uses the DBMS default value.

The *variable-identifier* argument can be either the SAS variable name or the positional equivalent from the LIST statement. The positional equivalent is the number that represents the variable's place in the data set. For example, if you want to reset the column associated with the third SAS variable, submit the following statement:

```
reset 3;
```

You must use the RESET statement *after* the LABEL statement for the LABEL statement to take effect.

SQL Statement

Submits a DBMS-specific SQL statement to the DBMS

SQL *DBMS-specific-SQL-statement*;

Details

The SQL statement submits a dynamic, non-query, DBMS-specific SQL statement to the DBMS. You can use the DBLOAD statement to submit these DBMS-specific SQL statements despite whether you create and load a DBMS table.

You must enter the keyword SQL before each DBMS-specific SQL statement that you submit. The *SQL-statement* argument can be any valid dynamic DBMS-specific SQL statement except the SELECT statement. However, you can enter a SELECT statement as a substatement within another statement, such as in a CREATE VIEW statement. You must use DBMS-specific SQL object names and syntax in the DBLOAD SQL statement.

You cannot create a DBMS table and reference it in your DBMS-specific SQL statements within the same PROC DBLOAD step. The new table is not created until the RUN statement is processed.

To submit dynamic, non-query DBMS-specific SQL statements to the DBMS *without* creating a DBMS table, you use the DBMS= option, any database connection statements, and the SQL statement.

Example

The following PROC DBLOAD example grants UPDATE privileges to user MARURI on the DB2 SasDemo.Orders table.

```
proc dbload dbms=db2;
  in sample;
  sql grant update on sasdemo.orders to maruri;
run;
```

TABLE= Statement

Names the DBMS table to be created and loaded

TABLE= <'>*DBMS-specific-syntax*<'>;

Details

When you create and load or append to a DBMS table, the TABLE= statement is required. It must follow other database connection statements such as DATABASE= or USER=. The TABLE= statement specifies the name of the DBMS table to be created and loaded into a DBMS database. The table name must be a valid table name for the

DBMS. (See the SAS/ACCESS documentation for your DBMS for the syntax.) If your table name contains lowercase characters, special characters, or national characters, it must be enclosed in quotation marks.

In addition, you must specify a table name that does not already exist. If a table by that name exists, an error message is written to the SAS log, and the table specified in this statement is not loaded.

When you are submitting dynamic DBMS-specific SQL statements to the DBMS without creating and loading a table, do not use this statement.

TYPE Statement

Changes default DBMS data types in the new table

TYPE *variable-identifier-1* = '*column-type-1*' <...*variable-identifier-n* = '*column-type-n*'>;

Details

The TYPE statement changes the default DBMS column data types that are associated with the corresponding SAS variables.

The *variable-identifier* argument can be either the SAS variable name or the positional equivalent from the LIST statement. The positional equivalent is the number that represents the variable's place in the data set. For example, if you want to change the data type of the DBMS column associated with the third SAS variable, submit the following statement:

```
type 3='char(17)';
```

The argument *column-type* must be a valid data type for the DBMS and must be enclosed in quotation marks.

If you omit the TYPE statement, the column data types are generated with default DBMS data types that are based on the SAS variable formats. You can change as many data types as you want in one TYPE statement. See the documentation for your SAS/ACCESS interface for a complete list of the default conversion data types for the DBLOAD procedure.

WHERE Statement

Loads a subset of data into the new table

WHERE *SAS-where-expression*;

Details

The WHERE statement causes a subset of observations to be loaded into the new DBMS table. The *SAS-where-expression* must be a valid SAS WHERE statement that

uses SAS variable names (not DBMS column names) as defined in the input data set. The following example loads only the observations in which the SAS variable COUNTRY has the value **BRAZIL**:

```
where country='Brazil';
```

For more information about the syntax of the SAS WHERE statement, see *SAS Language Reference: Dictionary*.


Example of Using the DBLOAD Procedure

Example 1: Appending a Data Set to a DBMS Table

The following example appends new employee data from the SAS data set NewEmp to the DBMS table Employees. The COMMIT statement causes a DBMS commit to be issued after every 100 rows are inserted. The ERRLIMIT statement causes processing to stop after 10 errors occur.

```
proc dbload dbms=oracle data=newemp append;
  user=testuser;
  password=testpass;
  path='myorapath';
  table=employees;
  commit=100;
  errlimit=10;
  load;
run;
```

Note: By omitting the APPEND option from the DBLOAD statement, you can use the PROC DBLOAD SQL statements to create a DBMS table and append to it in the same PROC DBLOAD step. \triangle



APPENDIX

3

Recommended Reading

Recommended Reading 343

Recommended Reading

Here is the recommended reading list for this title:

- SAS/ACCESS Supplement for DB2 under UNIX and PC Hosts*
- SAS/ACCESS Supplement for DB2 under z/OS*
- SAS/ACCESS Supplement for HP Neoview (SAS/ACCESS for Relational Databases)*
- SAS/ACCESS Supplement for Informix*
- SAS/ACCESS Supplement for Microsoft SQL Server*
- SAS/ACCESS Supplement for MySQL*
- SAS/ACCESS Supplement for Netezza (SAS/ACCESS for Relational Databases)*
- SAS/ACCESS Supplement for ODBC*
- SAS/ACCESS Supplement for OLE DB*
- SAS/ACCESS Supplement for Oracle*
- SAS/ACCESS Supplement for Sybase*
- SAS/ACCESS Supplement for Teradata*
- SAS/ACCESS Interface to PC Files: Reference*
- Base SAS Procedures Guide*
- SAS Language Reference: Concepts*
- SAS Language Reference: Dictionary*
- SAS Companion that is specific to your operating environment

For a complete list of SAS publications, go to support.sas.com/bookstore. If you have questions about which titles you need, please contact a SAS Publishing Sales Representative at:

SAS Publishing Sales
 SAS Campus Drive
 Cary, NC 27513
 Telephone: 1-800-727-3228
 Fax: 1-919-531-9439
 E-mail: sasbook@sas.com
 Web address: support.sas.com/bookstore

Customers outside the United States and Canada, please contact your local SAS office for assistance.

Glossary

This glossary defines SAS software terms that are used in this document as well as terms that relate specifically to SAS/ACCESS software.

access descriptor

a SAS/ACCESS file that describes data that is managed by a data management system. After creating an access descriptor, you can use it as the basis for creating one or more view descriptors. See also view and view descriptor.

browsing data

the process of viewing the contents of a file. Depending on how the file is accessed, you can view SAS data either one observation (row) at a time or as a group in a tabular format. You cannot update data that you are browsing.

bulk load

to load large amounts of data into a database object, using methods that are specific to a particular DBMS. Bulk loading enables you to rapidly and efficiently add multiple rows of data to a table as a single unit.

client

(1) a computer or application that requests services, data, or other resources from a server. (2) in the X Window System, an application program that interacts with the X server and can perform tasks such as terminal emulation or window management. For example, SAS is a client because it requests windows to be created, results to be displayed, and so on.

column

in relational databases, a vertical component of a table. Each column has a unique name, contains data of a specific type, and has certain attributes. A column is analogous to a variable in SAS terminology.

column function

an operation that is performed for each value in the column that is named as an argument of the function. For example, AVG(SALARY) is a column function.

commit

the process that ends a transaction and makes permanent any changes to the database that the user made during the transaction. When the commit process occurs, locks on the database are released so that other applications can access the changed data. The SQL COMMIT statement initiates the commit process.

DATA step view

a type of SAS data set that consists of a stored DATA step program. Like other SAS data views, a DATA step view contains a definition of data that is stored elsewhere; the view does not contain the physical data. The view's input data can come from one or more sources, including external files and other SAS data sets. Because a DATA step view only reads (opens for input) other files, you cannot update the view's underlying data.

data type

a unit of character or numeric information in a SAS data set. A data value represents one variable in an observation.

data value

in SAS, a unit of character or numeric information in a SAS data set. A data value represents one variable in an observation.

database

an organized collection of related data. A database usually contains named files, named objects, or other named entities such as tables, views, and indexes

database management system (DBMS)

an organized collection of related data. A database usually contains named files, named objects, or other named entities such as tables, views, and indexes

editing data

the process of viewing the contents of a file with the intent and the ability to change those contents. Depending on how the file is accessed, you can view the data either one observation at a time or in a tabular format.

engine

a component of SAS software that reads from or writes to a file. Each engine enables SAS to access files that are in a particular format. There are several types of engines.

file

a collection of related records that are treated as a unit. SAS files are processed and controlled by SAS and are stored in SAS data libraries.

format

a collection of related records that are treated as a unit. SAS files are processed and controlled by SAS and are stored in SAS data libraries. In SAS/ACCESS software, the default formats vary according to the interface product.

index

(1) in SAS software, a component of a SAS data set that enables SAS to access observations in the SAS data set quickly and efficiently. The purpose of SAS indexes is to optimize WHERE-clause processing and to facilitate BY-group processing. (2) in other software vendors' databases, a named object that directs the DBMS to the storage location of a particular data value for a particular column. Some DBMSs have additional specifications. These indexes are also used to optimize the processing of WHERE clauses and joins. Depending on the SAS interface to a database product and how selection criteria are specified, SAS may or may not be able to use the indexes of the DBMS to speed data retrieval.

Depending on how selection criteria are specified, SAS might use DBMS indices to speed data retrieval.

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.

interface view engine

a SAS engine that is used by SAS/ACCESS software to retrieve data from files that have been formatted by another vendor's software. Each SAS/ACCESS interface has its own interface view engine, which reads the interface product data and returns the data in a form that SAS can understand (that is, in a SAS data set). SAS automatically uses an interface view engine; the engine name is stored in SAS/ACCESS descriptor files so that you do not need to specify the engine name in a LIBNAME statement.

libref

a name that is temporarily associated with a SAS data library. The complete name of a SAS file consists of two words, separated by a period. The libref, which is the first word, indicates the library. The second word is the name of the specific SAS file. For example, in VLIB.NEWBDAY, the libref VLIB tells SAS which library contains the file NEWBDAY. You assign a libref with a LIBNAME statement or with an operating system command.

member

a SAS file in a SAS data library.

member name

a name that is given to a SAS file in a SAS data library.

member type

a SAS name that identifies the type of information that is stored in a SAS file. Member types include ACCESS, DATA, CATALOG, PROGRAM, and VIEW.

missing value

in SAS, a term that describes the contents of a variable that contains no data for a particular row or observation. By default, SAS prints or displays a missing numeric value as a single period, and it prints or displays a missing character value as a blank space.

observation

a row in a SAS data set. All of the data values in an observation are associated with a single entity such as a customer or a state. Each observation contains one data value for each variable. In a database product table, an observation is analogous to a row. Unlike rows in a database product table or file, observations in a SAS data file have an inherent order.

Pass-Through Facility

a group of SQL procedure statements that send and receive data directly between a relational database management system and SAS. The Pass-Through Facility includes the CONNECT, DISCONNECT, and EXECUTE statements, and the CONNECTION TO component. SAS/ACCESS software is required in order to use the Pass-Through Facility.

PROC SQL view

a SAS data set (of type VIEW) that is created by the SQL procedure. A PROC SQL view contains no data. Instead, it stores information that enables it to read data values from other files, which can include SAS data files, SAS/ACCESS views, DATA step views, or other PROC SQL views. A PROC SQL view's output can be either a subset or a superset of one or more files.

query

a set of instructions that requests particular information from one or more data sources.

referential integrity

a set of rules that a DBMS uses to ensure that whenever a data value in one table is changed, the appropriate change is also made to any related values in other tables or in the same table. Referential integrity is also used to ensure that related data is not deleted or changed accidentally.

relational database management system

a database management system that organizes and accesses data according to relationships between data items. Oracle and DB2 are examples of relational database management systems.

rollback

in most databases, the process that restores the database to its state when changes were last committed, voiding any recent changes. The SQL ROLLBACK statement initiates the rollback processes. See also commit.

row

in relational database management systems, the horizontal component of a table. A row is analogous to a SAS observation.

SAS data file

a type of SAS data set that contains data values as well as descriptor information that is associated with the data. The descriptor information includes information such as the data types and lengths of the variables, as well as the name of the engine that was used to create the data. A PROC SQL table is a SAS data file. SAS data files are of member type DATA.

SAS data library

a collection of one or more SAS files that are recognized by SAS and that are referenced and stored as a unit. Each file is a member of the library.

SAS data set

a file whose contents are in one of the native SAS file formats. There are two types of SAS data sets: SAS data files and SAS data views. SAS data files contain data values in addition to descriptor information that is associated with the data. SAS data views contain only the descriptor information plus other information that is required for retrieving data values from other SAS data sets or from files whose contents are in other software vendors' file formats.

SAS data view

a file whose contents are in one of the native SAS file formats. There are two types of SAS data sets: SAS data files and SAS data views. SAS data files contain data values in addition to descriptor information that is associated with the data. SAS data views contain only the descriptor information plus other information that is required for retrieving data values from other SAS data sets or from files whose contents are in other software vendors' file formats.

SAS/ACCESS views

See view descriptor and SAS data view.

server

in a network, a computer that is reserved for servicing other computers in the network. Servers can provide several different types of services, such as file services and communication services. Servers can also enable users to access shared resources such as disks, data, and modems.

Structured Query Language (SQL)

the standardized, high-level query language that is used in relational database management systems to create and manipulate database management system objects. SAS implements SQL through the SQL procedure.

table

a two-dimensional representation of data, in which the data values are arranged in rows and columns.

trigger

a type of user-defined stored procedure that is executed whenever a user issues a data-modification command such as INSERT, DELETE, or UPDATE for a specified table or column. Triggers can be used to implement referential integrity or to maintain business constraints.

variable

a column in a SAS data set. A variable is a set of data values that describe a given characteristic across all observations.

view

a definition of a virtual data set. The definition is named and stored for later use. A view contains no data; it merely describes or defines data that is stored elsewhere. SAS data views can be created by the ACCESS and SQL procedures.

view descriptor

a file created by SAS/ACCESS software that defines part or all of the database management system (DBMS) data or PC file data that is described by an access descriptor. The access descriptor describes the data in a single DBMS table, DBMS view, or PC file.

wildcard

a file created by SAS/ACCESS software that defines part or all of the database management system (DBMS) data or PC file data that is described by an access descriptor. The access descriptor describes the data in a single DBMS table, DBMS view, or PC file.

Index

A

ACCDESC= option
 PROC ACCESS statement 314, 326
 ACCDESC= statement
 DBLOAD procedure 333
 access descriptors
 ACCESS procedure with 325
 converting into SQL views 299, 300, 304
 creating 316, 333
 data set and descriptor access 326
 identifying DBMS table for 323
 listing columns in, with information 318
 resetting columns to default settings 321
 updating 324, 327
 ACCESS= LIBNAME option 79
 access methods
 relational DBMS data 3
 selecting 4
 ACCESS procedure, relational databases 313
 accessing DBMS data 311
 descriptors with 325
 examples 327
 names and 9
 overview 60, 311
 syntax 313
 accessing DBMS data
 ACCESS procedure 311
 methods for 3
 repeated accessing 33
 acquisition error tables 232, 233
 aggregate functions
 passing to DBMS 38
 AIX (RS/6000) 65
 Informix 67
 Microsoft SQL Server 67
 MySQL 68
 ODBC 69
 Sybase 71
 Teradata 72
 ALL option
 LIST statement 318
 PROC CV2VIEW statement 300
 RESET statement 321
 SELECT statement 322
 APPEND option
 PROC DBLOAD statement 332
 ASSIGN statement
 ACCESS procedure 315
 AUTHID= data set option 158

AUTHID= LIBNAME option 80
 authorization ID 158
 autocommit capability 159
 AUTOCOMMIT= data set option 159
 AUTOCOMMIT= LIBNAME option 80
 automatic COMMIT 200
 autopartitioning 50
 DBSLICE= option 54

B

BL_BADFILE= data set option 159
 BL_CODEPAGE= data set option 160
 BL_CONTROL= data set option 161
 BL_COPY_LOCATION= data set option 162
 BL_DATAFILE= data set option 162
 BL_DB2CURSOR= data set option 163
 BL_DB2DEVT_PERM= data set option 164
 BL_DB2DEVT_TEMP= data set option 164
 BL_DB2DISC= data set option 165
 BL_DB2ERR= data set option 165
 BL_DB2IN= data set option 166
 BL_DB2LDCT1= data set option 166
 BL_DB2LDCT2= data set option 167
 BL_DB2LDTEXT= data set option 167
 BL_DB2MAP= data set option 168
 BL_DB2PRINT= data set option 169
 BL_DB2PRNLOG= data set option 169
 BL_DB2REC= data set option 170
 BL_DB2RECS= data set option 170
 BL_DB2RSTRT= data set option 171
 BL_DB2SPC_PERM= data set option 171
 BL_DB2SPC_TEMP= data set option 172
 BL_DB2TBLXST= data set option 172
 BL_DB2UTID= data set option 173
 BL_DELETE_DATAFILE= data set option 173
 BL_DIRECT_PATH= data set option 175
 BL_DISCARDFILE= data set option 176
 BL_INDEXING_MODE= data set option 178
 BL_INDEX_OPTIONS= data set option 177
 BL_KEEPIENTITY= data set option 179
 BL_KEEPIENTITY= LIBNAME option 81
 BL_KEEPNULLS= data set option 180
 BL_KEEPNULLS= LIBNAME option 82
 BL_LOAD_METHOD= data set option 181
 BL_LOAD_REPLACE= data set option 181
 BL_LOG= data set option 182
 BL_LOG= LIBNAME option 83
 BL_METHOD= data set option 183

BL_OPTIONS= data set option 183
 BL_OPTIONS= LIBNAME option 83
 BL_PARFILE= data set option 185
 BL_PRESERVE_BLANKS= data set option 186
 BL_RECOVERABLE= data set option 187
 BL_REMOTE_FILE= data set option 188
 BL_SERVER_DATAFILE= data set option 189
 BL_SQLLDR_PATH= data set option 189
 BL_SUPPRESS_NULLIF= data set option 190
 BL_WARNING_COUNT= data set option 192
 buffering bulk rows 193
 buffers
 reading DBMS data 140
 reading rows of DBMS data 256
 BUFFERS= data set option 192
 bulk loading 183
 appending vs. replacing rows 181
 codepage for converting character data 160
 data file as seen by DB2 server instance 189
 data file for 162
 DB2 method 183
 DB2 SELECT statement 163
 file containing SQLLDR control statements 161
 filtered out records 176
 generic device type for permanent data sets 164
 identity column 179
 loading rows of data as one unit 194
 log file for 182
 MultiLoad 238
 NULL values in Microsoft SQL Server columns 180
 Oracle method 181
 rejected records 159
 saving copy of loaded data 162
 SQL*Loader Index options 177
 unit address for permanent data sets 164
 warning count 192
 bulk rows
 buffering for output 193
 BULK_BUFFER= data set option 193
 BULKLOAD= data set option 194
 BULKLOAD= LIBNAME option 84
 BYTEINT data type 85

C

case sensitivity 8
 DBMS column names 252
 CAST= data set option 196
 CAST= LIBNAME option 85
 CAST_OVERHEAD_MAXPERCENT= data set option 197
 CAST_OVERHEAD_MAXPERCENT= LIBNAME option 86
 CELLPROP= LIBNAME option 87
 character data
 codepage for converting during bulk load 160
 length of 211
 checkpoints
 interval between 231
 restart table 236
 codepage 160
 column labels
 returned by engine 215
 column names
 preserving 14

columns
 limiting retrieval 31
 NULL as valid value 212
 renaming because of disallowed characters 204
 command timeout 198
 COMMAND_TIMEOUT= data set option 198
 COMMAND_TIMEOUT= LIBNAME option 88
 COMMIT, automatic 200
 COMMIT= statement
 DBLOAD procedure 334
 Compaq Tru64
 Informix 67
 ODBC 69
 Oracle 70
 Sybase 71
 connect exits
 customizing 27
 CONNECT statement
 SQL procedure 287
 connection information
 protecting 24
 CONNECTION= LIBNAME option 89
 CONNECTION TO component 293
 CONNECTION_GROUP= LIBNAME option 93
 CREATE statement
 ACCESS procedure 316
 SQL procedure 293
 CREATE TABLE statement
 SQL procedure 202
 CREATE TABLE statement (SQL) 62
 currency control 26
 cursor type 198
 CURSOR_TYPE= data set option 198
 CURSOR_TYPE= LIBNAME option 95
 CV2VIEW procedure 300
 examples 304
 overview 299
 syntax 300
 cylinders
 LOAD utility 171, 172

D

data buffers
 MultiLoad 238
 transferring data to Teradata 192, 231
 data conversions
 overhead limit 197
 Teradata DBMS server vs. SAS 196
 DATA= option
 PROC DBLOAD statement 332
 data security 21
 privileges 21
 protecting connection information 24
 SAS security 22
 triggers 22
 data set options 157
 data set tables
 updating 153
 data sets
 appending to DBMS table 342
 controlling access to 326
 creating DBMS tables 20
 extracting DBMS data to 24
 data sources
 schemas 259

- DATA step
 - DBKEY= option with 45
- DATA step views 6
- data types
 - DBMS columns 339, 341
 - overriding SAS defaults 216
 - specifying 221
- database administrator (DBA) 21
- database links 209
- database tables
 - large 55
- DATASETS procedure
 - passwords for access descriptors 23
 - reading Oracle table names 133
- DATE data type
 - casting 85
- DB2
 - appending vs. replacing rows during bulk loading 181
 - bulk loading data file as seen by server instance 189
 - saving copy of loaded data 162
- DB2 LOAD
 - base filename and location of temporary files 188
- DB2 load utility
 - index maintenance 178
- DB2 LOAD utility
 - execution mode 167
 - unique identifier for a given run 173
- DB2 SELECT statement 163
- DB2 server data file 189
- DB2 under UNIX and PC hosts
 - features 65
- DB2 under z/OS
 - features 66
- DBA (database administrator) 21
- DBCOMMIT= data set option 200
- DBCOMMIT= LIBNAME option 96
- DBCONDITION= data set option 201
- DBCONINIT= LIBNAME option 97
- DBCONTERM= LIBNAME option 99
- DBCREATE_TABLE_OPTS= data set option 202
- DBCREATE_TABLE_OPTS= LIBNAME option 100
- DBCREATE_TABLE_OPTS= option 202
- DBFORCE= data set option 203
- DBGEN_NAME= data set option 204
- DBGEN_NAME= LIBNAME option 100
- DBINDEX= data set option 45, 205
 - passing joins to DBMS 43
 - replacing missing values 245
- DBINDEX= LIBNAME option 101
- DBKEY= data set option 43, 207
 - replacing missing values 245
- DBLABEL= data set option 208
- DBLIBINIT= LIBNAME option 102
- DBLIBTERM= LIBNAME option 103
- DBLINK= data set option 209
- DBLINK= LIBNAME option 104
- DBLOAD procedure, relational databases 331
 - example 342
 - how it works 62
 - names and 10
 - overview 329
 - sending data from SAS to DBMS 329
 - syntax 331
- DBMASTER= data set option 210
- DBMAX_TEXT= data set option 211
- DBMAX_TEXT= LIBNAME option 105
- DBMS
 - assigning libref to remote DBMS 78
 - passing DISTINCT processing to 42
 - passing functions to 38
 - passing functions to, with WHERE clauses 43
 - passing joins to 38, 40
 - passing UNION processing to 42
 - submitting SQL statements to 340
- DBMS, SQL statements with
 - connecting to DBMS 287
 - disconnecting from 291
 - sending non-query SQL statements 292
 - specifying DBMS connection 293
- DBMS autocommit capability 159
- DBMS bulk-load facility
 - deleting the data file 173
- DBMS columns
 - basing variable names on column names 324
 - changing column formats 317
 - names defaulting to labels 335
 - naming during output 208
 - null values accepted in 337
 - preserving names 251
 - renaming 11, 338, 339
 - resetting to default settings 321
 - selecting 322
- DBMS connections
 - controlling 25
- DBMS data
 - access methods 3
 - accessing/extracting 311
 - extracting to data set 24
 - pushing updates 36
 - renaming 10
 - repeatedly accessing 33
 - SAS views of 6
 - sorting 33
 - subsetting and ordering 201
- DBMS data types
 - changing default 341
 - resetting to default 339
- DBMS engine
 - codepage for converting character data 160
- DBMS objects
 - naming 12
- DBMS= option
 - PROC ACCESS statement 314
 - PROC DBLOAD statement 332
- DBMS tables
 - access descriptors based on 333
 - appending SAS data sets to 342
 - committing or saving after inserts 334
 - creating and loading 336, 340
 - creating with data sets 20
 - creating with DBMS data 19
 - dropping variables before creating 334
 - limiting observations loaded to 335
 - loading data subsets into 341
 - naming 340
 - preserving column names 251
 - preserving names 15
 - renaming 10
 - verifying indexes 205
- DBMS variables
 - renaming 11
- DBNULL= data set option 212

DBNULLKEYS= data set option 213
 DBNULLKEYS= LIBNAME option 106
 DBPROMPT= data set option 214
 DBPROMPT= LIBNAME option 106
 DBSASLABEL= data set option 215
 DBSASLABEL= LIBNAME option 108
 DBSASTYPE= data set option 216
 DBSLICE= data set option 49, 217
 DBSLICEPARM= data set option 49, 54, 219
 DBSLICEPARM= LIBNAME option 109
 DBSRVTP= system option 271, 272
 DBTYPE= data set option 221
 DECIMAL data type
 casting 85
 DEFER= LIBNAME option 111
 DEGREE= LIBNAME option 112
 DELETE statement
 DBLOAD procedure 334
 passing to empty a table 40
 SQL procedure 61, 293
 DELETE_MULT_ROWS= LIBNAME option 112
 descriptor files
 ACCESS procedure with 325
 converting into SQL views 300, 304
 creating descriptor files 316
 defined 325
 listing columns in, with information 318
 resetting columns to default settings 321
 updating 324
 DIRECT option
 SQL*Loader 175
 DIRECT_EXE= LIBNAME option 113
 DIRECT_SQL= LIBNAME option 114
 disconnect exits
 customizing 27
 DISCONNECT statement
 SQL procedure 291
 DISTINCT operator
 pass processing to DBMS 42
 DQUOTE= option
 SQL procedure 11, 17
 DROP= data set option
 limiting retrieval 32
 DROP statement
 ACCESS procedure 317
 SQL procedure 293

E

ENABLE_BULK= LIBNAME option 116
 ERRLIMIT= data set option 223
 ERRLIMIT= LIBNAME option 116
 ERRLIMIT= statement
 DBLOAD procedure 335
 error limits with rollbacks 223
 error tracking
 acquisition error tables 232, 233
 EXECUTE statement
 SQL procedure 60, 292
 extracting DBMS data 24
 ACCESS procedure 311

F

FORMAT statement
 ACCESS procedure 317

FROM_LIBREF= statement
 CV2VIEW procedure 301
 FROM_VIEW= statement
 CV2VIEW procedure 301
 functions
 LIBNAME statement and 74
 passing to DBMS with WHERE clauses 43

G

GRANT statement
 SQL procedure 293
 group ID 158

H

heterogeneous joins
 pushing to DBMS 36
 HP-UX
 DB2 under UNIX and PC hosts 65
 Informix 67
 Microsoft SQL Server 67
 MySQL 68
 ODBC 69
 Oracle 70
 Sybase 71
 Teradata 72
 HP-UX for Itanium Processor Family Architecture
 Informix 67
 ODBC 69
 Sybase 71
 Teradata 72

I

identity column 179
 IGNORE_READ_ONLY_COLUMNS= data set option 224
 IGNORE_READ_ONLY_COLUMNS= LIBNAME option 117
 IN= data set option 226
 IN= LIBNAME option 119
 indexes 44
 checking with the DBKEY= option 45
 maintenance, DB2 load utility 178
 Informix
 features 67
 input processing
 overriding default SAS data types 216
 insert processing
 forcing truncation of data 203
 INSERT statement
 SQL procedure 61, 62, 293
 INSERTBUFF= data set option 227
 INSERTBUFF= LIBNAME option 120
 inserting data
 appending data sets to DBMS tables 342
 limiting observations loaded 335
 loading data subsets into DBMS tables 341
 saving DBMS table after inserts 334
 INSERT_SQL= data set option 226
 INSERT_SQL= LIBNAME option 119
 installing SAS/ACCESS 57
 INTEGER data type
 casting 85
 INTERFACE= LIBNAME option 121
 isolation levels 264

J

- joins
 - determining larger table 210
 - passing to DBMS 38, 40, 43
 - processing 59
 - pushing heterogeneous joins 36

K

- KEEP= data set option
 - limiting retrieval 32
- key column for DBMS retrieval 207
- KEYSET_SIZE= data set option 228
- KEYSET_SIZE= LIBNAME option 122

L

- LABEL statement
 - DBLOAD procedure 335
- labels
 - DBMS column names defaulting to 335
- LIBNAME statement, relational databases 58, 73
 - accessing data from DBMS objects 58
 - arguments 75
 - assigning librefs 73, 77
 - assigning librefs interactively 74
 - data from a DBMS 76
 - disassociating librefs 76
 - functions and 74
 - options 78
 - prompting window 78
 - sorting data 73
 - SQL views embedded with 76
 - syntax 75
 - writing library attributes to log 76
- libraries
 - disassociating librefs 76
 - writing attributes to log 76
- librefs
 - assigning 77
 - assigning interactively 74
 - assigning to remote DBMS 78
 - disassociating 76
- LIMIT= statement
 - DBLOAD procedure 335
- links
 - database links 209
- Linux for Intel Architecture
 - DB2 under UNIX and PC hosts 65
 - MySQL 68
 - ODBC 69
 - Oracle 70
 - Sybase 71
 - Teradata 72
- Linux for Itanium-based Systems
 - DB2 under UNIX and PC hosts 65
 - Informix 67
 - Microsoft SQL Server 67
 - MySQL 68
 - ODBC 69
 - Oracle 70
 - Sybase 71
- LIST statement
 - ACCESS procedure 318
 - DBLOAD procedure 336

- LOAD process
 - recoverability of 187
 - LOAD statement
 - DBLOAD procedure 336
 - LOAD utility
 - restarts 171
 - running against existing tables 172
 - SYSDISC data set name 165
 - SYSERR data set name for 165
 - SYSIN data set name 166
 - SYSMAP data set name 168
 - SYSPRINT data set name 169
 - SYSREC data set name 170
 - temporary data sets 164
 - LOCATION= data set option 229
 - LOCATION= LIBNAME option 123
 - locking 26, 256, 266
 - DBMS resources 152
 - during read isolation 254
 - during read transactions 138, 255
 - during update transaction 265
 - shared locks 230
 - LOCKTABLE= data set option 230
 - LOCKTABLE= LIBNAME option 123
 - LOCKTIME= LIBNAME option 124
 - LOCKWAIT= LIBNAME option 124
 - log
 - SQL statements 273
 - writing library attributes to 76
 - log files
 - for bulk loading 182
 - LOGDB= LIBNAME option 125
-
- M**
 - macro variables 269
 - MAX_CONNECTS= LIBNAME option 126
 - MBUFFSIZE= data set option 231
 - Microsoft SQL Server
 - features 67
 - NULL values during bulk loading 180
 - missing values 27
 - replacing character values 245
 - ML_CHECKPOINT= data set option 231
 - ML_ERROR1= data set option 232
 - ML_ERROR2= data set option 233
 - ML_LOG= data set option 235
 - ML_RESTART= data set option 236
 - ML_WORK= data set option 237
 - MOD function 53
 - MULTI_DATASCR_OPT= option
 - LIBNAME statement 43
 - MULTI_DATASRC_OPT= LIBNAME option 126
 - MultiLoad
 - acquisition error tables 232, 233
 - bulk loading 238
 - data buffers 238
 - enabling/disabling 238
 - examples 243
 - prefix for temporary table names 235
 - restart table 236
 - restarting 239
 - retries for logging in to Teradata 261, 262
 - storing intermediate data 237
 - temporary tables 238
 - work table 237

MULTILOAD= data set option 238
 MySQL
 features 68

N

name literals 9, 18
 names
 ACCESS procedure and 9
 case sensitivity 8
 creating DBMS objects 12
 DBLOAD procedure and 10
 DBMS columns 338, 339
 DBMS tables 340
 length of 8
 modification 9
 options affecting 11
 replacing unsupported characters 13
 retrieving DBMS data 12
 SAS 8
 SAS/ACCESS 7
 truncation 9
 NULL
 as valid value when tables are created 212
 null values 27
 accepted in DBMS columns 337
 NULL values
 in Microsoft SQL Server columns 180
 NULLCHAR= data set option 45, 244
 NULLCHARVAL= data set option 45, 245
 NULLIF clause
 suppressing 190
 NULLS statement
 DBLOAD procedure 337

O

observations 335
 ODBC
 features 69
 OLE DB
 features 70
 OpenVMS Alpha 70
 optimizing SQL usage 37
 Oracle
 bulk loading method 181
 database links 209
 features 70
 hints 250
 ordering
 DBMS data 201
 ORHINTS= data set option 250
 OR_PARTITION= data set option 246
 OR_UPD_NOWHERE= data set option 249
 OR_UPD_NOWHERE= LIBNAME option 128
 OUT= option
 PROC ACCESS statement 314

P

PACKETSIZE= LIBNAME option 129
 partitioning
 queries for threaded reads 217
 Pass-Through Facility 285
 advantages of 5
 how it works 59

 macro variables and 270
 relational databases 286
 syntax 286
 passwords
 assigning 22
 data set and descriptor access 326
 performance
 limiting retrieval 31
 processing queries, joins, and data functions 59
 reducing table read time 47
 repeatedly accessing data 33
 SAS data files 43
 SAS server throughput 31
 sorting DBMS data 33
 SQL usage 37
 temporary table support 34
 threaded reads 50
 PREFETCH= LIBNAME option 129
 PRESERVE_COL_NAMES= data set option 251
 PRESERVE_COL_NAMES= LIBNAME option 130
 PRESERVE_COL_NAMES= option
 SAS/ACCESS LIBNAME statement 11
 PRESERVE_TAB_NAMES= LIBNAME option 132
 PRESERVE_TAB_NAMES= option
 SAS/ACCESS LIBNAME statement 11
 privileges
 DBMS security 21
 prompting window 78
 prompts
 DBMS connections and 214
 pushing heterogeneous joins 36
 pushing updates 36

Q

QUALIFIER= data set option 253
 QUALIFIER= LIBNAME option 134
 QUALIFY_ROWS= LIBNAME option 134
 queries
 partitioning for threaded reads 217
 QUERY_TIMEOUT= data set option 253
 QUERY_TIMEOUT= LIBNAME option 135
 QUIT statement
 ACCESS procedure 319
 DBLOAD procedure 338
 QUOTE_CHAR= LIBNAME option 136
 QUOTED_IDENTIFIER= LIBNAME option 137

R

READBUFF= data set option 256
 READBUFF= LIBNAME option 140
 READ_ISOLATION_LEVEL= data set option 254
 READ_ISOLATION_LEVEL= LIBNAME option 137
 READ_LOCK_TYPE= data set option 255
 READ_LOCK_TYPE= LIBNAME option 138
 READ_MODE_WAIT= data set option 256
 READ_MODE_WAIT= LIBNAME option 139
 relational databases
 data set options for 157
 Pass-Through Facility 286
 remote DBMS
 assigning libref to 78
 REMOTE_DBTYPE= LIBNAME option 141
 RENAME statement
 ACCESS procedure 320

- DBLOAD procedure 338
 - renaming
 - columns, because of disallowed characters 204
 - DBMS columns 338
 - DBMS data 10
 - SAS variables 320
 - REPLACE= data set option 157
 - REPLACE= statement
 - CV2VIEW procedure 301
 - REREAD_EXPOSURE= LIBNAME option 142
 - RESET statement
 - ACCESS procedure 321
 - DBLOAD procedure 339
 - restart table 236
 - retrieving DBMS data
 - ACCESS procedure 311
 - return codes
 - PROC SQL statement 286
 - REVOKE statement
 - SQL procedure 293
 - rollbacks
 - error limits and 223
 - row warnings 192
 - rows
 - number in single DBMS insert 227
 - number in single insert operation 120
 - reading into buffers 256
- S**
- SAS/ACCESS
 - features 5
 - interactions with SQL procedure 285
 - task table 5
 - SAS/ACCESS engine
 - buffering bulk rows for output 193
 - SAS/ACCESS LIBNAME statement
 - advantages of 4
 - SAS/ACCESS views 6
 - SAS data views 6
 - creating 6
 - SAS server
 - increasing throughput 31
 - SASDATEFMT= data set option 257
 - SASTRACE= system option 273
 - SASTRACELOC= system option 49, 53, 281
 - SAVEAS= statement
 - CV2VIEW procedure 302
 - SCHEMA= data set option 259
 - SCHEMA= LIBNAME option 143
 - schemas
 - data security 25
 - SEGMENT_NAME= data set option 260
 - SELECT statement
 - ACCESS procedure 322
 - SQL procedure 60
 - SHOW_SYNONYMS= LIBNAME option 145
 - SLEEP= data set option 261
 - SMALLINT data type
 - casting 85
 - Solaris 2
 - DB2 under UNIX and PC hosts 65
 - Informix 67
 - Microsoft SQL Server 67
 - MySQL 68
 - ODBC 69
 - Oracle 70
 - Sybase 71
 - Teradata 72
 - sorting data 73
 - data ordering in SAS/ACCESS 54
 - DBMS data 33
 - subsetting and ordering DBMS data 201
 - spool files
 - UTILLOC= system option with 55
 - SPOOL= LIBNAME option 145
 - SQL operations
 - optimizing 37
 - SQL procedure
 - creating DBMS tables 19
 - interactions with SAS/ACCESS 285
 - passing functions to DBMS 38
 - passing joins to DBMS 38
 - SQL statement
 - DBLOAD procedure 340
 - SQL views 6
 - converting descriptors to 299, 300, 304
 - embedded LIBNAME statements in 76
 - SQL_FUNCTIONS= LIBNAME option 146
 - SQLLDR control statements
 - file containing 161
 - SQLLDR executable file
 - location specification 189
 - SQL*Loader
 - blank spaces in CHAR/VARCHAR2 columns 186
 - command line options 185
 - DIRECT option 175
 - discarded rows file 176
 - index options for bulk loading 177
 - SQLXMSG macro 270
 - SQLXRC macro 270
 - STRINGDATES= LIBNAME option 149
 - SUBMIT statement
 - CV2VIEW procedure 303
 - SUBSET statement
 - ACCESS procedure 322
 - subsetting
 - DBMS data 201
 - SUPERQ macro 270
 - Sybase
 - database links 209
 - features 71
 - SYSDBMSG macro variable 269
 - SYSDBRC macro variable 269
 - SYSDISC data set name 165
 - SYSIN data set name 166
 - SYSMAP data set name 168
 - SYSPRINT data set name 169
 - SYSPRINT output 169
 - SYSREC data set
 - number of cylinders 170
 - SYSREC data set name 170
- T**
- TABLE= statement
 - ACCESS procedure 323
 - DBLOAD procedure 340
 - tables
 - emptying with DELETE statement 40
 - temporary tables 34

temporary tables 34
 acquisition error tables 232, 233
 MultiLoad 238
 prefix for names of 235
 restart table 236
 work table 237
 TENACITY= data set option 262
 Teradata
 buffers and transferring data to 192, 231
 features 72
 locking DBMS resources 152
 MultiLoad retries for logging in to 261, 262
 threaded reads 47, 281
 controlling scope of 219
 data set options 49
 partitioning queries for 217
 scope 49
 trace information 50
 two-pass processing 54
 TO_LIBREF= statement
 CV2VIEW procedure 303
 TO_VIEW= statement
 CV2VIEW procedure 303
 TRACE= LIBNAME option 149
 TRACEFILE= LIBNAME option 150
 tracking errors
 acquisition error tables 232, 233
 transaction control 26
 TRAP151= data set option 263
 triggers 22
 truncation
 forcing during insert processing 203
 two-pass processing
 threaded reads 54
 TYPE statement
 DBLOAD procedure 341
 CV2VIEW procedure 304

U

UNION operator
 pass processing to DBMS 42
 UNIQUE statement
 ACCESS procedure 324
 UPDATE statement
 ACCESS procedure 324
 SQL procedure 61, 293
 UPDATEBUFF= data set option 155, 268
 UPDATEBUFF= LIBNAME option 155
 UPDATE_ISOLATION_LEVEL= data set option 264
 UPDATE_ISOLATION_LEVEL= LIBNAME option 151
 UPDATE_LOCK_TYPE= data set option 265
 UPDATE_LOCK_TYPE= LIBNAME option 151
 UPDATE_MODE_WAIT= data set option 266
 UPDATE_MODE_WAIT= LIBNAME option 152
 UPDATE_MULT_ROWS= LIBNAME option 153
 updates
 pushing 36
 UPDATE_SQL= data set option 267
 UPDATE_SQL= LIBNAME option 154
 updating
 access descriptors 324, 327
 method for updating rows 267
 non-updatable columns 263
 reading data 61

specifying number of rows 268
 USE_ODBC_CL= LIBNAME option 155
 user IDs 158
 UTILCONN_TRANSIENT= LIBNAME option 156
 UTILLOC= system option 55

V

VALIDVARNAME= system option 131, 282
 naming and 12
 variable labels
 as DBMS column names 208
 variable names
 as DBMS column names 208
 valid names 282
 variable names and formats 282
 basing on DBMS column names 324
 changing from default 317
 generating 315
 modifying variable names 320
 variables
 dropping before creating a table 334
 listing information about, before loading 336
 macro variables 269
 view descriptors
 ACCESS procedure with 325
 converting into SQL views 299, 300, 304
 creating 316, 328
 dropping columns to make unselectable 317
 listing columns in, with information 318
 resetting columns to default settings 321
 selecting DBMS columns 322
 selection criteria, adding or modifying 322
 updating 324
 VIEWDESC= option
 PROC ACCESS statement 314
 views
 data security 25

W

WHERE clause
 NULL values and format of 213
 WHERE clauses
 partitioning queries for threaded reads 217
 passing functions to DBMS with 43
 WHERE statement
 DBLOAD procedure 341
 Windows NT and 2000
 DB2 under UNIX and PC hosts 65
 MySQL 68
 ODBC 69
 OLE DB 70
 Oracle 70
 Sybase 71
 Teradata 72
 work table 237

Z

z/OS
 DB2 under z/OS 66
 Oracle 70
 Teradata 72

Your Turn

We want 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 gives you the tools to flourish in any environment with SAS!

Whether you are new to the workforce 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.

SAS® Press Series

Need to learn the basics? Struggling with a programming problem? You'll find the expert answers that you need in example-rich books from the SAS Press Series. Written by experienced SAS professionals from around the world, these 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: 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, and hard-copy books.

support.sas.com/publishing

SAS® Learning Edition 4.1

Get a workplace advantage, perform analytics in less time, and prepare for the SAS Base Programming exam and SAS Advanced Programming exam with SAS® Learning Edition 4.1. This inexpensive, intuitive personal learning version of SAS includes Base SAS® 9.1.3, SAS/STAT®, SAS/GRAPH®, SAS/QC®, SAS/ETS®, and SAS® Enterprise Guide® 4.1. Whether you are a professor, student, or business professional, this is a great way to learn SAS.

support.sas.com/LE



**THE
POWER
TO KNOW®**

