



SAS Publishing



SAS/ACCESS[®] 9.1.3

Supplement for Oracle

SAS/ACCESS for Relational Databases

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2005. *SAS/ACCESS® 9.1.3 Supplement for Oracle (SAS/ACCESS for Relational Databases)*. Cary, NC: SAS Institute Inc.

SAS/ACCESS® 9.1.3 Supplement for Oracle (SAS/ACCESS for Relational Databases)

Copyright © 2005, SAS Institute Inc., Cary, NC, USA

ISBN 1-59047-803-7

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, June 2005


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

Chapter 1	△ SAS/ACCESS for Oracle	1
Introduction to the SAS/ACCESS Interface to Oracle		2
LIBNAME Statement Specifics for Oracle		2
Data Set Options for Oracle		4
Pass-Through Facility Specifics for Oracle		6
Autopartitioning Scheme for Oracle		8
Temporary Table Support for Oracle		10
ACCESS Procedure Specifics for Oracle		11
DBLOAD Procedure Specifics for Oracle		13
Maximizing Oracle Performance		15
Passing SAS Functions to Oracle		15
Passing Joins to Oracle		16
Oracle Bulk Loading		17
Locking in the Oracle Interface		19
Naming Conventions for Oracle		20
Data Types for Oracle Servers		20
Appendix 1	△ Recommended Reading	25
Recommended Reading		25
Glossary		27
Index		33



CHAPTER

1

SAS/ACCESS for Oracle

<i>Introduction to the SAS/ACCESS Interface to Oracle</i>	2
<i>LIBNAME Statement Specifics for Oracle</i>	2
Arguments	2
Oracle LIBNAME Statement Examples	4
<i>Data Set Options for Oracle</i>	4
<i>Pass-Through Facility Specifics for Oracle</i>	6
Examples	7
<i>Autopartitioning Scheme for Oracle</i>	8
Overview	8
Partitioned Oracle Tables	8
Non-partitioned Oracle Tables	9
Performance Summary	10
<i>Temporary Table Support for Oracle</i>	10
Establishing a Temporary Table	10
Syntax	11
Terminating a Temporary Table	11
Examples	11
<i>ACCESS Procedure Specifics for Oracle</i>	11
Examples	12
<i>DBLOAD Procedure Specifics for Oracle</i>	13
Examples	14
<i>Maximizing Oracle Performance</i>	15
<i>Passing SAS Functions to Oracle</i>	15
<i>Passing Joins to Oracle</i>	16
<i>Oracle Bulk Loading</i>	17
Oracle Bulk Loading: Interactions with Other Options	17
Oracle Bulk Loading: z/OS Specifics	18
Oracle Bulk Loading: Example	18
<i>Locking in the Oracle Interface</i>	19
<i>Naming Conventions for Oracle</i>	20
<i>Data Types for Oracle Servers</i>	20
Character Data	21
Numeric Data	21
Other Data Types	21
Oracle Null and Default Values	22
LIBNAME Statement Data Conversions	22
ACCESS Procedure Data Conversions	23
DBLOAD Procedure Data Conversions	24

Introduction to the SAS/ACCESS Interface to Oracle

This document includes details *only* about the SAS/ACCESS Interface to Oracle. It should be used as a supplement to the main SAS/ACCESS documentation, *SAS/ACCESS for Relational Databases: Reference*.

LIBNAME Statement Specifics for Oracle

This section describes the LIBNAME statement as supported in the SAS/ACCESS interface to Oracle. For a complete description of this feature, see the LIBNAME statement section in *SAS/ACCESS for Relational Databases: Reference*. The Oracle specific syntax for the LIBNAME statement is as follows:

```
LIBNAME libref oracle <connection-options> <LIBNAME-options>;
```

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.

oracle

is the SAS/ACCESS engine name for the interface to Oracle.

connection-options

provide connection information and control how SAS manages the timing and concurrence of the connection to the DBMS. The connection options for the interface to Oracle are as follows:

```
USER=<'>Oracle-user-name<'>
```

specifies an optional Oracle user name. If the user name contains blanks or national characters, enclose it in quotation marks. If you omit an Oracle user name and password, the default Oracle user ID OPS\$sysid is used, if it is enabled. **USER=** must be used with **PASSWORD=**.

```
PASSWORD=<'>Oracle-password<'>
```

specifies an optional Oracle password that is associated with the Oracle user name. If you omit **PASSWORD=**, the password for the default Oracle user ID OPS\$sysid is used, if it is enabled. **PASSWORD=** must be used with **USER=**.

```
PATH=<'>Oracle-database-specification<'>
```

specifies the Oracle driver, node, and database. Aliases are required if you are using SQL*Net Version 2.0 or later. In some operating environments, you can enter the information that is required by the **PATH=** statement before invoking SAS.

SAS/ACCESS uses the same Oracle path designation that you use to connect to Oracle directly. See your database administrator to determine the databases that have been set up in your operating environment, and to determine the default values if you do not specify a database. On UNIX systems, the TWO_TASK environment variable is used, if set. If neither the **PATH=** nor the TWO_TASK values have been set, the default value is the local driver.

If you specify the appropriate system options or environment variables for Oracle, you can often omit the connection options from your LIBNAME statements. See your Oracle documentation for details.

LIBNAME-options

define how DBMS objects are processed by SAS. Some LIBNAME options can enhance performance; others determine locking or naming behavior. The following table describes LIBNAME options that are supported for Oracle, and presents default values where applicable. See the section about the SAS/ACCESS LIBNAME statement in *SAS/ACCESS for Relational Databases: Reference* for detailed information about these options.

Table 1.1 SAS/ACCESS LIBNAME Options for Oracle

Option	Default Value
ACCESS=	none
CONNECTION=	SHAREDREAD
CONNECTION_GROUP=	none
DBCOMMIT=	1000 when inserting rows; 0 when updating rows, deleting rows, or appending rows to an existing table
DBCONINIT=	none
DBCONTERM=	none
DBCREATE_TABLE_OPTS=	none
DBGEN_NAME=	DBMS
DBINDEX=	NO Use this option only when the object is a TABLE, not a VIEW. Use DBKEY when you do not know whether the object is a TABLE.
DBLIBINIT=	none
DBLIBTERM=	none
DBLINK=	the local database
DBMAX_TEXT=	1024
DBNULLKEYS=	YES
DBPROMPT=	NO
DBSLICEPARM=	THREADED_APPS,2
DEFER=	NO
DIRECT_EXE=	none
DIRECT_SQL=	YES
INSERTBUFF=	10
LOCKWAIT=	YES
MULTI_DATASRC_OPT=	NONE
ORACLE_UPD_NOWHERE=	YES
PRESERVE_COL_NAMES=	NO
PRESERVE_TAB_NAMES=	NO
READBUFF=	250
READ_ISOLATION_LEVEL=	see “Locking in the Oracle Interface” on page 19

Option	Default Value
READ_LOCK_TYPE=	NOLOCK
REREAD_EXPOSURE=	NO
SCHEMA=	SAS accesses objects in the default and public schemas
SHOW_SYNONYMS=	NO
SPOOL=	YES
UPDATE_ISOLATION_LEVEL=	see “Locking in the Oracle Interface” on page 19
UPDATE_LOCK_TYPE=	NOLOCK
UPDATEBUFF=	1
UTILCONN_TRANSIENT=	NO

Oracle LIBNAME Statement Examples

In the following example, the connection is made using default settings for the connection options. If you specify the appropriate system options or environment variables for Oracle, you can often omit the connection options from your LIBNAME statements. See your Oracle documentation for details.

```
libname myoralib oracle;
```

In the following example, the libref MYDBLIB uses the SAS/ACCESS interface to Oracle to connect to an Oracle database. The SAS/ACCESS connection options are USER=, PASSWORD=, and PATH=. PATH= specifies an alias for the database specification (as required by SQL*Net).

```
libname mydblib oracle user=testuser password=testpass path=hrdept_002;
```

```
proc print data=mydblib.employees;
  where dept='CSR010';
run;
```

Data Set Options for Oracle

The following table describes all of the data set options that are supported for the Oracle interface. Default values are provided where applicable. See the section about data set options in *SAS/ACCESS for Relational Databases: Reference* for general information about these options.

Table 1.2 Data Set Options for Oracle

Option	Default Value
BL_BADFILE=	creates a file in the current directory or with the default file specifications
BL_CONTROL=	creates a file in the current directory or with the default file specifications

Option	Default Value
BL_DATAFILE=	creates a file in the current directory or with the default file specifications
BL_DELETE_DATAFILE=	YES
BL_DIRECT_PATH=	YES
BL_DISCARDFILE=	creates a file in the current directory or with the default file specifications
BL_INDEX_OPTIONS=	the current SQL*Loader Index options with bulk-loading
BL_LOAD_METHOD=	When loading an empty table, the default value is INSERT; when loading a table that contains data, the default value is APPEND.
BL_LOG=	If there is no pre-existing log file, the default action is to create a log file in the current directory or the default file specifications. If there is already a log file, the Oracle bulk loader reuses the file, replacing the contents with information from the new load.
BL_OPTIONS=	ERRORS=1000000
BL_PRESERVE_BLANKS=	NO
BL_RECOVERABLE=	YES
BL_SQLLDR_PATH=	sqlldr
BL_SUPPRESS_NULLIF=	NO
BULKLOAD=	NO
DBCOMMIT=	the current LIBNAME option setting
DBCONDITION=	none
DBCREATE_TABLE_OPTS=	the current LIBNAME option setting
DBFORCE=	NO
DBGEN_NAME=	DBMS
DBINDEX=	the current LIBNAME option setting
DBKEY=	none
DBLABEL=	NO
DBLINK=	the current LIBNAME option setting
DBMASTER=	none
DBMAX_TEXT=	1024
DBNULL=	YES
DBNULLKEYS=	the current LIBNAME option setting
DBPROMPT=	the current LIBNAME option setting
DBSASTYPE=	see “Data Types for Oracle Servers” on page 20
DBSLICE=	none
DBSLICEPARM=	THREADED_APPS,2

Option	Default Value
DBTYPE=	see “LIBNAME Statement Data Conversions” on page 22
ERRLIMIT=	1
INSERTBUFF=	the current LIBNAME option setting
NULLCHAR=	SAS
NULLCHARVAL=	a blank character
OR_PARTITION=	an Oracle table partition name
OR_UPD_NOWHERE=	the current LIBNAME option setting
ORHINTS=	no hints
PRESERVE_COL_NAMES=	current LIBNAME option setting
READ_ISOLATION_LEVEL=	the current LIBNAME option setting
READ_LOCK_TYPE=	the current LIBNAME option setting
READBUFF=	the current LIBNAME option setting
SASDATEFORMAT=	DATE/TIME20.0
SCHEMA=	the current LIBNAME option setting
UPDATE_ISOLATION_LEVEL=	the current LIBNAME option setting
UPDATE_LOCK_TYPE=	the current LIBNAME option setting
UPDATEBUFF=	the current LIBNAME option setting

Pass-Through Facility Specifics for Oracle

See the section about the Pass-Through Facility in *SAS/ACCESS for Relational Databases: Reference* for general information about this feature.

The Pass-Through Facility specifics for Oracle are as follows:

- The *dbms-name* is **oracle**.
- The CONNECT statement is optional. If you omit the CONNECT statement, an implicit connection is made with your OPS\$*sysid*, if it is enabled. When you omit a CONNECT statement, an implicit connection is performed when the first EXECUTE statement or CONNECTION TO component is passed to Oracle. In this case you must use the default DBMS name **oracle**.
- The interface to Oracle can connect to multiple databases (both local and remote) and to multiple user IDs. If you use multiple simultaneous connections, you must use an *alias* argument to identify each connection. If you do not specify an alias, the default alias, **oracle**, is used.
- The *database-connection-arguments* for the CONNECT statement are as follows:

USER=<'>Oracle-user-name<'>

specifies an optional Oracle user name. If you specify USER=, you must also specify PASSWORD=.

PASSWORD= <'>Oracle-password<'>

specifies an optional Oracle password that is associated with the Oracle user name. If you omit an Oracle password, the default Oracle user ID OPS\$*sysid* is used, if it is enabled. If you specify PASSWORD=, you must also specify USER=.

ORAPW= is an alias for this option.

Note: If you do not wish to enter your Oracle password in uncoded text, see PROC PWENCODE for a method to encode it. △

BUFSIZE=*number-of-rows*

specifies the number of rows to retrieve from an Oracle table or view with each fetch. Using this argument can improve the performance of any query to Oracle.

By setting the value of the BUFSIZE= argument in your SAS programs, you can find the optimal number of rows for a given query on a given table. The default buffer size is 250 rows per fetch. The limit is 32,767 rows per fetch, although a practical limit for most applications is less, depending on the available memory.

PRESERVE_COMMENTS

enables you to pass additional information (called *hints*) to Oracle for processing. These hints might direct the Oracle query optimizer to choose the best processing method based on your hint.

You specify PRESERVE_COMMENTS as an argument in the CONNECT statement. Then you specify the hints in the CONNECTION TO component's Oracle SQL query. The hints are entered as comments in the SQL query and are passed to and processed by Oracle.

PATH=<'>Oracle-database-specification<'>

specifies the Oracle driver, node, and database. Aliases are required if you are using SQL*Net Version 2.0 or later. In some operating environments, you can enter the information that is required by the PATH= statement before invoking SAS.

SAS/ACCESS uses the same Oracle path designation that you use to connect to Oracle directly. See your database administrator to determine the path designations that have been set up in your operating environment, and to determine the default value if you do not specify a path designation. On UNIX systems, the TWO_TASK environment variable is used, if set. If neither PATH= nor TWO_TASK have been set, the default value is the local driver.

Examples

The following example uses the alias DBCON for the DBMS connection (the connection alias is optional):

```
proc sql;
  connect to oracle as dbcon
    (user=testuser password=testpass bufsize=100
     path='myorapath');
quit;
```

The following example connects to Oracle and sends it two EXECUTE statements to process.

```
proc sql;
  connect to oracle (user=testuser password=testpass);
  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;

```

The following example performs a query, shown in highlighted text, on the Oracle table CUSTOMERS:

```

proc sql;
connect to oracle (user=testuser password=testpass);
select *
    from connection to oracle
        (select * from customers
         where customer like '1%');
    disconnect from oracle;
quit;

```

In this example, the PRESERVE_COMMENTS argument is specified after the USER= and PASSWORD= arguments. The Oracle SQL query is enclosed in the required parentheses. The SQL INDX command identifies the index for the Oracle query optimizer to use in processing the query. Note that multiple hints are separated with blanks.

```

proc sql;
connect to oracle as mycon(user=testuser
    password=testpass preserve_comments);
select *
    from connection to mycon
        (select /* +indx(empid) all_rows */
         count(*) from employees);
quit;

```

Autopartitioning Scheme for Oracle

See the section about threaded reads *SAS/ACCESS for Relational Databases: Reference* for general information about this feature.

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

Overview

In the absence of user specified partitioning from the DBSLICE= option, the SAS/ACCESS interface to Oracle attempts to use its own partitioning techniques. The partitioning technique it chooses depends on whether the table is physically partitioned on the Oracle server.

Partitioned Oracle Tables

If you are dealing with a partitioned Oracle table, it is recommended that you allow the Oracle engine to partition the table for you. The Oracle engine will gather all of the partition information needed to do a threaded read on the table.

A partitioned Oracle table is a good candidate for a threaded read, because each of the partitions in the table can be read in parallel without much contention for disk

resources. If the Oracle engine determines that the table is partitioned, it makes the same number of connections to the server as there are partitions. Each connection retrieves rows from a single partition.

For example, assume a SALES table was created in Oracle as follows:

```
CREATE TABLE SALES (acct_no NUMBER(5),
  acct_name CHAR(30), amount_of_sale NUMBER(6), qtr_no INTEGER)
PARTITION BY RANGE (qtr_no)
(PARTITION sales1 VALUES LESS THAN (2) TABLESPACE ts0,
PARTITION sales2 VALUES LESS THAN (2) TABLESPACE ts1,
PARTITION sales3 VALUES LESS THAN (2) TABLESPACE ts2,
PARTITION sales4 VALUES LESS THAN (2) TABLESPACE ts3)
```

Performing a threaded read on this table causes SAS to make four separate connections to the Oracle server. Each connection reads from each of the partitions. Turning SASTRACE on shows you the SQL that is generated for each connection:

```
libname x oracle user=testuser path=oraserver;
data new;
set x.SALES (DBSLICEPARAM=ALL);
run;

ORACLE: SELECT "ACCT_NO","ACCT_NAME", "AMOUNT_OF_SALE", "QTR_NO" FROM SALES
partition (SALES2)
ORACLE: SELECT "ACCT_NO","ACCT_NAME", "AMOUNT_OF_SALE", "QTR_NO" FROM SALES
partition (SALES3)
ORACLE: SELECT "ACCT_NO","ACCT_NAME", "AMOUNT_OF_SALE", "QTR_NO" FROM SALES
partition (SALES1)
ORACLE: SELECT "ACCT_NO","ACCT_NAME", "AMOUNT_OF_SALE", "QTR_NO" FROM SALES
partition (SALES4)
```

The number of threads used to read the table in parallel is determined by the second parameter of the DBSLICEPARAM= option. The number of connections made to the Oracle server for retrieving rows from the table is determined by the number of partitions on the table.

Non-partitioned Oracle Tables

If the table is not partitioned, and the DBSLICE= option is not specified, Oracle resorts to the MOD function discussed in the section on threaded reads in *SAS/ACCESS for Relational Databases: Reference*. With this technique, the engine makes N connections, and each connection retrieves rows based on a WHERE clause as follows:

```
WHERE ABS(MOD(ModColumn,N))=R
```

- ModColumn is a column in the table which is of type integer and is not used in any user specified WHERE clauses. (The engine selects this column. If you do not think this is the ideal partitioning column, you can use DBSLICE= to override this default behavior.)
- R varies from 0 to $(N-1)$ for each of the N WHERE clauses.
- N defaults to 2, and N can be overridden with the second parameter in DBSLICEPARAM=.

The Oracle engine selects the ModColumn to use in this technique. Any numeric column with zero scale value can qualify as the ModColumn. However, if a primary key

column is present, it is preferred over all others. Generally, values in the primary key column are in a serial order and will yield an equal number of rows for each connection.

An example illustrates this point:

```
create table employee (empno number(10) primary key,
  empname varchar2(20), hiredate date,
  salary number(8,2), gender char(1));
```

Performing a threaded read on this table causes Oracle to make two separate connections to the Oracle server. SAS tracing shows the SQL generated for each connection:

```
data new;
set x.EMPLOYEE(DBSLICPARM=ALL);
run;
ORACLE: SELECT "EMPNO", "EMPNAME", "HIREDATE", "SALARY", "GENDER"
FROM EMPLOYEE WHERE ABS(MOD("EMPNO",2))=0
ORACLE: SELECT "EMPNO", "EMPNAME", "HIREDATE", "SALARY", "GENDER"
FROM EMPLOYEE WHERE ABS(MOD("EMPNO",2))=1
```

EMPNO, the primary key, is selected as the MOD column.

The success of MOD depends on the distribution of the values within the selected ModColumn and the value of N . Ideally, the rows will be distributed evenly among the threads.

Performance Summary

There are times you might not see an improvement in performance with the MOD technique. It is possible that the engine might not be able to find a column that qualifies as a good MOD column. In these situations, you can explicitly specify DBSLICE= to force a threaded read and improve performance.

It is a good policy to let the engine autopartition and intervene with DBSLICE= only when necessary.

Temporary Table Support for Oracle

See the section on the temporary table support in *SAS/ACCESS for Relational Databases: Reference* for general information about this feature.

Establishing a Temporary Table

A temporary table in Oracle persists just like a regular table, but contains either session specific or transaction specific data. Whether the data is session or transaction specific is determined by what is specified with the ON COMMIT keyword when you create the temporary table.

In the SAS context, you must use the LIBNAME option CONNECTION=SHARED in order for the data in a temporary table to persist over procedure and DATA step boundaries. Without this option, the temporary table will persist but the data within it will not.

If you have a SAS data set and you want to join it with an Oracle table to generate a report, the join normally occurs in SAS. However, using a temporary table you can also have the join occur on the Oracle server.

Syntax

The syntax to create a temporary table whose data is transaction specific (default) is as follows:

```
CREATE GLOBAL TEMPORARY TABLE table name ON COMMIT DELETE
ROWS
```

The syntax to create a temporary table whose data is session specific is as follows:

```
CREATE GLOBAL TEMPORARY TABLE table name ON COMMIT PRESERVE
ROWS
```

Terminating a Temporary Table

You can drop a temporary table at any time, or allow it to be implicitly dropped when the connection is terminated. Temporary tables do not persist beyond the scope of a single connection.

Examples

In the following example, a temporary table, TEMPTRANS, is created in Oracle to match the TRANS SAS data set (using the Pass-Through Facility):

```
proc sql;
  connect to oracle (user=scott pw=tiger path=oracle9);
  execute (create global temporary table TEMPTRANS
          (empid number, salary number)) by oracle;
quit;

libname ora oracle user=scott pw=tiger path=oracle9 connection=shared;

/* load the data from the TRANS table into the Oracle temporary table */
proc append base=ora.TEMPTRANS set TRANS;
run;

proc sql;
/* do the join on the DBMS server */
  select lastname, firstname, salary from ora.EMPLOYEES T1, ora.TEMPTRANS T2
         where T1.empno=T2.empno;
quit;
```

ACCESS Procedure Specifics for Oracle

See the section about the ACCESS procedure in *SAS/ACCESS for Relational Databases: Reference* for general information about this feature.

The Oracle interface supports all of the ACCESS procedure statements. The ACCESS procedure specifics for Oracle are as follows:

- The PROC ACCESS step DBMS= value is **Oracle**.
- The *database-description-statements* used by PROC ACCESS are as follows:

USER=<'>*Oracle-user-name*<'>

specifies an optional Oracle user name. If you omit an Oracle password and user name, the default Oracle user ID OPS\$sysid is used if it is enabled. If you specify USER=, you must also specify ORAPW=.

ORAPW= <'>*Oracle-password*<'>

specifies an optional Oracle password that is associated with the Oracle user name. If you omit ORAPW=, the password for the default Oracle user ID OPS\$sysid is used, if it is enabled. If you specify ORAPW=, you must also specify USER=.

PATH=<'>*Oracle-database-specification*<'>

specifies the Oracle driver, node, and database. Aliases are required if you are using SQL*Net Version 2.0 or later. In some operating environments, you can enter the information that is required by the PATH= statement before invoking SAS.

SAS/ACCESS uses the same Oracle path designation that you use to connect to Oracle directly. See your database administrator to determine the path designations that have been set up in your operating environment, and to determine the default value if you do not specify a path designation. On UNIX systems, the TWO_TASK environment variable is used, if set. If neither PATH= nor TWO_TASK have been set, the default value is the local driver.

- The PROC ACCESS step TABLE= statement is as follows:

TABLE= <'><*Oracle-table-name*><'>;

specifies the name of the Oracle table or Oracle view on which the access descriptor is based. This statement is required. The *Oracle-table-name* argument can be up to 30 characters long and must be a valid Oracle table name. If the table name contains blanks or national characters, enclose it in quotation marks.

Examples

The following example creates an access descriptor and a view descriptor based on Oracle data.

```
options linesize=80;

libname adlib 'SAS-data-library';
libname vlib 'SAS-data-library';

proc access dbms=oracle;

/* create access descriptor */

create adlib.customer.access;
user=testuser;
orapw=testpass;
table=customers;
path='myorapath';
```



```

assign=yes;
rename customer=custnum;
format firstorder date9.;
list all;

/* create view descriptor */

create vlib.usacust.view;
select customer state zipcode name
       firstorder;
subset where customer like '1%';
run;

```

The following example creates another view descriptor that is based on the ADLIB.CUSTOMER access descriptor. The view is then printed.

```

/* create socust view */

proc access dbms=oracle accdesc=adlib.customer;
create vlib.socust.view;
select customer state name contact;
subset where state in ('NC', 'VA', 'TX');
run;

/* print socust view */

proc print data=vlib.socust;
title 'Customers in Southern States';
run;

```

DBLOAD Procedure Specifics for Oracle

See the section about the DBLOAD procedure in *SAS/ACCESS for Relational Databases: Reference* for general information about this feature.

The Oracle interface supports all of the DBLOAD procedure statements. The DBLOAD procedure specifics for Oracle are as follows:

- The PROC DBLOAD step DBMS= value is **Oracle**.
- The *database-description-statements* used by PROC DBLOAD are as follows:

USER=<'>Oracle-user-name<'>

specifies an optional Oracle user name. If you omit an Oracle password and user name, the default Oracle user ID OPS\$sysid is used if it is enabled. If you specify USER=, you must also specify ORAPW=.

ORAPW= <'>Oracle-password<'>

specifies an optional Oracle password that is associated with the Oracle user name. If you omit ORAPW=, the password for the default Oracle user ID OPS\$sysid is used, if it is enabled. If you specify ORAPW=, you must also specify USER=.

PATH=<'>Oracle-database-specification<'>

specifies the Oracle driver, node, and database. Aliases are required if you are using SQL*Net Version 2.0 or later. In some operating environments, you can enter the information that is required by the PATH= statement before invoking SAS.

SAS/ACCESS uses the same Oracle path designation that you use to connect to Oracle directly. See your database administrator to determine the path designations that have been set up in your operating environment, and to determine the default value if you do not specify a path designation. On UNIX systems, the TWO_TASK environment variable is used, if set. If neither PATH= nor TWO_TASK have been set, the default value is the local driver.

TABLESPACE= <'>Oracle-tablespace-name<';

specifies the name of the Oracle tablespace where you want to store the new table. The *Oracle-tablespace-name* argument can be up to 18 characters long and must be a valid Oracle tablespace name. If the name contains blanks or national characters, enclose the entire name in quotation marks.

If TABLESPACE= is omitted, the table is created in your default tablespace that is defined by the Oracle database administrator at your site.

- The PROC DBLOAD step TABLE= statement is as follows:

TABLE= <'><Oracle-table-name><';

specifies the name of the Oracle table or Oracle view on which the access descriptor is based. This statement is required. The *Oracle-table-name* argument can be up to 30 characters long and must be a valid Oracle table name. If the table name contains blanks or national characters, enclose the name in quotation marks.

Examples

The following example creates a new Oracle table, EXCHANGE, from the DLIB.RATEOFEX data file. An access descriptor, ADLIB.EXCHANGE, based on the new table, is also created. The PATH= statement uses an alias to connect to a remote Oracle7 Server database.

The SQL statement in the second DBLOAD procedure sends an SQL GRANT statement to Oracle. You must be granted Oracle privileges to create new Oracle tables or to grant privileges to other users. The SQL statement is in a separate procedure because you cannot create a DBMS table and reference it within the same DBLOAD step. The new table is not created until the RUN statement is processed at the end of the first DBLOAD step.

Note: The DLIB.RATEOFEX data set is included in the sample data shipped with your software. Δ

```
libname adlib 'SAS-data-library';
libname dlib 'SAS-data-library';

proc dbload dbms=oracle data=dlib.rateofex;
  user=testuser;
  orapw=testpass;
  path='myorapath';
  table=exchange;
  accdesc=adlib.exchange;
  rename fgnindol=fgnindolar 4=dolrsinfgn;
  nulls updated=n fgnindol=n 4=n country=n;
  load;
run;

proc dbload dbms=oracle;
```

```

user=testuser;
orapw=testpass;
path='myorapath';
sql grant select on testuser.exchange to pham;
run;

```

The next example uses the APPEND option to append rows from the INVDATA data set to an existing Oracle table named INVOICE.

```

proc dbload dbms=oracle data=invdata append;
  user=testuser;
  orapw=testpass;
  path='myorapath';
  table=invoice;
  load;
run;

```

Note: This example uses a previously created data set, INVDATA. Δ

Maximizing Oracle Performance

There are several measures you can take to optimize performance when using the SAS/ACCESS interface to Oracle. See the section about performance considerations in *SAS/ACCESS for Relational Databases: Reference* for general information about improving performance when using SAS/ACCESS engines.

The SAS/ACCESS interface to Oracle has several options that you can use to further improve performance. See the INSERTBUFF=, UPDATEBUFF= and READBUFF= LIBNAME options for tips on multi-row processing. See “Oracle Bulk Loading” on page 17 for instructions on using Oracle’s SQL*Loader to increase performance when loading rows of data into Oracle tables.

Note: If you choose the transactional inserting of rows (specify BULKLOAD=NO), you can improve performance by inserting multiple rows at a time. This performance enhancement is comparable to using the Oracle SQL*Loader Conventional Path Load. For more information about inserting multiple rows, see the INSERTBUFF= option. Δ

Passing SAS Functions to Oracle

The interface to Oracle passes the following SAS functions to Oracle for processing. Where the Oracle function name is different than the SAS function name, the Oracle name appears in parentheses. See the section about optimizing SQL usage in *SAS/ACCESS for Relational Databases: Reference* for information.

```

ABS
ARCOS (ACOS)
ARSIN (ASIN)
ATAN
CEIL
COS

```

COSH
DATETIME
(SYSDATE)
EXP
FLOOR
LOG
LOG10 (LOG)
LOG2 (LOG)
LOWCASE
(LOWER)
SIGN
SIN
SINH
SOUNDEX
SQRT
STRIP (TRIM)
TAN
TANH
TRIMN (RTRIM)
TRANSLATE
UPCASE (UPPER)
SUM
COUNT
AVE
MIN
MAX

Passing Joins to Oracle

In order for a join to pass to Oracle, each of the following components of the LIBNAME statements must match exactly:

- user ID
- password
- path

See the section about performance considerations in *SAS/ACCESS for Relational Databases: Reference* for more information about when and how SAS/ACCESS passes joins to the DBMS.

Oracle Bulk Loading

The SAS/ACCESS interface to Oracle can call the Oracle SQL*Loader (SQLLDR) when you set the data set option BULKLOAD=YES. The support for Oracle's bulk loader provides superior load performance, enabling you to rapidly move data from a SAS file into an Oracle table. In future releases, SAS/ACCESS software will continue to make use of powerful Oracle tools to improve its loading performance.

The bulk loading data set options for Oracle are listed below. See the section about data set options in *SAS/ACCESS for Relational Databases: Reference* for additional information about these options. Note that they all begin with BL_ (for BULKLOAD):

BL_BADFILE=
 BL_CONTROL=
 BL_DATAFILE=
 BL_DELETE_DATAFILE=
 BL_DIRECT_PATH=
 BL_DISCARDFILE=
 BL_INDEX_OPTIONS=
 BL_LOAD_METHOD=
 BL_LOG=
 BL_OPTIONS=
 BL_PARFILE=
 BL_PRESERVE_BLANKS=
 BL_RECOVERABLE=
 BL_SQLLDR_PATH=
 BL_SUPPRESS_NULLIF=
 BULKLOAD=

BULKLOAD= invokes Oracle's bulk loader, enabling the Oracle engine to move data from a SAS file into an Oracle table using SQL*Loader (SQLLDR).

Note: SQL*Loader Direct path load has a number of limitations. Please refer to the Oracle utilities manual for details.

When using bulk load, you should consult the SQL*Loader log file (rather than the SAS log) for information about the load. Δ

Oracle Bulk Loading: Interactions with Other Options

When BULKLOAD=YES, the following is true:

- The DBCOMMIT=, DBFORCE=, ERRLIMIT=, and INSERTBUFF= options are ignored.
- If NULLCHAR=SAS, and the NULLCHARVAL value is blank, then the SQL*Loader attempts to insert a NULL instead of a NULLCHARVAL value.
- If NULLCHAR=NO, and the NULLCHARVAL value is blank, then the SQL*Loader attempts to insert a NULL even if the DBMS does not allow NULL.

To avoid this result, set BL_PRESERVE_BLANKS=YES or set NULLCHARVAL to a non-blank value (and then replace the non-blank value with blanks after processing, if necessary).

Oracle Bulk Loading: z/OS Specifics

When you are using bulk loading in the z/OS operating environment, the files used by the SQL*Loader must conform to z/OS data set standards. The data sets can be either sequential data sets or partitioned data sets. Each of the file names supplied to the SQL*Loader are subject to extension and FNA processing.

If you do not specify file names using data set options, then default names in the form of *userid.SAS.data-set-extension* apply. The *userid* is the TSO prefix when running under TSO, and it is the PROFILE PREFIX in batch. The *data-set-extensions* are:

- BAD for the bad file
- CTL for the control file
- DAT for the data file
- DSC for the discard file
- LOG for the log file

If you want to specify file names using data set options, then you must use one of the following forms:

- /DD/ddname*
- /DD/ddname(membername)*
- Name*

For detailed information about these forms, refer to the SQL*Loader chapter in the Oracle user's guide for z/OS.

The SQL*Loader is executed by the Oracle engine by issuing a host-system command from within your SAS session. The data set where the SQLLDR executable resides must be available to your TSO session or allocated to your batch job. Check with your system administrator if you do not know the name or availability of the data set that contains the SQLLDR executable.

On z/OS, the bad file and the discard file are, by default, not created in the same format as the data file. This makes it difficult to load the contents of these files after making corrections. Refer to the section on SQL*Loader file attributes in the SQL*Loader section in the Oracle user's guide for z/OS for information about overcoming this limitation.

Oracle Bulk Loading: Example

The following example shows how to create a SAS data set and use it to create and load to a large Oracle table, FLIGHTS98. This load uses the SQL*Loader direct path method because you specified BULKLOAD=YES. BL_OPTIONS= passes the specified SQL*Loader options to SQL*Loader when it is invoked. In this example, the ERRORS= option enables you to have 899 errors in the load before the load terminates, and the LOAD= option loads the first 5,000 rows of the input data set, SASFLT.FLT98.

```
options yearcutoff=1925;    /* included for Year 2000 compliance */

libname sasflt 'SAS-Data-Library';
libname ora_air oracle user=testuser password=testpass
    path='ora8_flt' schema=statsdiv;

data sasflt.flt98;
    input flight $3. +5 dates date7. +3 depart time5. +2 orig $3.
```

```

+3 dest $3. +7 miles +6 boarded +6 capacity;
format dates date9. depart time5.;
informat dates date7. depart time5.;
datalines;
114    01JAN98    7:10  LGA  LAX      2475      172      210
202    01JAN98   10:43 LGA  ORD       740      151      210
219    01JAN98    9:31  LGA  LON     3442      198      250

<...10,000 more observations...>

proc sql;
create table ora_air.flights98
(BULKLOAD=YES BL_OPTIONS='ERRORS=899,LOAD=5000') as
  select * from sasflt.flt98;
quit;

```

During a load, certain SQL*Loader files are created, such as the data, log, and control files. Unless otherwise specified, they are given a default name and written to the current directory. For this example, the default names would be **bl_flights98.dat**, **bl_flights98.log**, and **bl_flights98ctl**.

Locking in the Oracle Interface

The SAS/ACCESS interface to Oracle supports the following locking options as both LIBNAME and data set options. See the section about the LIBNAME statement in *SAS/ACCESS for Relational Databases: Reference* for additional information about these options.

READ_LOCK_TYPE= NOLOCK | ROW | TABLE

The default value is NOLOCK. The valid values for this option are as follows:

- NOLOCK — table locking is not used during the reading of tables and views.
- ROW — the Oracle ROW SHARE table lock is used during the reading of tables and views.
- TABLE — the Oracle SHARE table lock is used during the reading of tables and views.

Note: If you set READ_LOCK_TYPE= to either TABLE or ROW, you must also set the CONNECTION= option to UNIQUE. If not, an error occurs. Δ

UPDATE_LOCK_TYPE= NOLOCK | ROW | TABLE

The default value is NOLOCK. The valid values for this option are as follows:

- ROW — the Oracle ROW SHARE table lock is used during the reading of tables and views for update.
- TABLE — the Oracle EXCLUSIVE table lock is used during the reading of tables and views for update.
- NOLOCK — table locking is not used during the reading of tables and views for update.
 - If OR_UPD_NOWHERE=YES, updates are performed using serializable transactions.
 - If OR_UPD_NOWHERE=NO, updates are performed using an extra WHERE clause to ensure that the row has not been updated since it was first read. Updates might fail under these conditions, because other users might modify a row after the row was read for update.

READ_ISOLATION_LEVEL= READCOMMITTED | SERIALIZABLE

Oracle supports the READCOMMITTED and SERIALIZABLE read isolation levels, as defined in the following table. The SPOOL= option overrides the READ_ISOLATION_LEVEL= option. The READ_ISOLATION_LEVEL= option should be rarely needed because the SAS/ACCESS engine chooses the appropriate isolation level based on other locking options.

Table 1.3 Isolation Levels for Oracle

Isolation Level	Definition
SERIALIZABLE	Does not allow dirty reads, non-repeatable reads, or phantom reads.
READCOMMITTED	Does not allow dirty reads; does allow non-repeatable reads and phantom reads

UPDATE_ISOLATION_LEVEL= READCOMMITTED | SERIALIZABLE

Oracle supports the READCOMMITTED and SERIALIZABLE isolation levels, as defined in the preceding table, for updates.

This option should be rarely needed because the SAS/ACCESS engine chooses the appropriate isolation level based on other locking options.

Naming Conventions for Oracle

The PRESERVE_COL_NAMES= and PRESERVE_TAB_NAMES= options determine how the interface to Oracle handles case-sensitivity, spaces, and special characters. See the section about the LIBNAME statement in *SAS/ACCESS for Relational Databases: Reference* for information about these options.

Oracle objects that can be named include tables, views, columns, and indexes. For the Oracle7 Server, objects also include database triggers, procedures, and stored functions. Use the following Oracle naming conventions:

- A name must start with a letter. However, if the name appears within double quotation marks, it may start with any character.
- A name must be from 1 to 30 characters long, except for database names, which are limited to 8 characters, and link names, which are limited to 128 characters.
- A name may contain the letters A through Z, the digits 0 through 9, the underscore (_), \$, and #. If the name appears within double quotation marks, it may contain any characters, except double quotation marks.
- A name is not case-sensitive. For example, CUSTOMER is the same as customer. However, if the name of the object appears within double quotation marks when it is used, then it is case-sensitive.
- A name cannot be an Oracle reserved word.
- A name cannot be the same name as another Oracle object in the same schema.

Data Types for Oracle Servers

Every column in a table has a name and a data type. The data type tells Oracle how much physical storage to set aside for the column and the form in which the data is stored.

Note: The SAS/ACCESS interface to Oracle does not support the following Oracle data types: MLSLABEL and ROWID. Δ

Character Data

CHAR (*n*)

contains fixed-length character string data with a length of *n*, where *n* must be at least 1 and cannot exceed 255 characters. (The limit is 2,000 characters with an Oracle8 Server.) Note that the Oracle7 Server CHAR data type is not equivalent to the Oracle Version 6 CHAR data type. The Oracle7 Server CHAR data type is new with the Oracle7 Server and uses blank-padded comparison semantics.

LONG

contains varying-length character string data that is similar to type VARCHAR2. Type LONG is character data of variable length with a maximum length of 2 gigabytes. You can define only one LONG column per table. Available memory considerations might also limit the size of a LONG data type.

VARCHAR2(*n*)

contains character string data with a length of *n*, where *n* must be at least 1 and cannot exceed 2000 characters. (The limit is 4,000 characters with an Oracle8 Server.) The VARCHAR2 data type is equivalent to the Oracle Version 6 CHAR data type except for the difference in maximum lengths. The VARCHAR2 data type uses nonpadded comparison semantics.

Numeric Data

NUMBER(*p,s*)

specifies a fixed-point number with an implicit decimal point, where *p* is the total number of digits (precision) and can range from 1 to 38, and *s* is the number of digits to the right of the decimal point (scale) and can range from -84 to 127.

NUMBER(*p*)

specifies an integer of precision *p* that can range from 1 to 38 and a scale of 0.

NUMBER

specifies a floating-point number with a precision of 38. A floating-point value can either specify a decimal point anywhere from the first to the last digit or omit the decimal point. A scale value does not apply to floating-point numbers since there is no restriction on the number of digits that can appear after the decimal point.

Other Data Types

DATE

contains date values. Valid dates are from January 1, 4712 BC to December 31, 4712 AD. The default format is DD-MON-YY, for example '05-OCT-98'.

LONG RAW

contains raw binary data of variable length up to 2 gigabytes. Values entered into columns of this type must be inserted as character strings in hexadecimal notation.

RAW(*n*)

contains raw binary data where *n* must be at least 1 and cannot exceed 255 bytes. (In Oracle Version 8, the limit is 2,000 bytes.) Values entered into columns of this type must be inserted as character strings in hexadecimal notation. You must specify *n* for this data type.

Note: For compatibility with other DBMSs, Oracle supports the syntax for a wide variety of numeric data types, including DECIMAL, INTEGER, REAL, DOUBLE-PRECISION, and SMALLINT. All forms of numeric data types are actually stored in the same internal Oracle NUMBER format. The additional numeric data types are variations of precision and scale. A null scale implies a floating-point number, and a non-null scale implies a fixed-point number. Δ

Oracle Null and Default Values

Oracle has a special value called NULL. An Oracle NULL value means an absence of information and is analogous to a SAS missing value. When SAS/ACCESS reads an Oracle NULL value, it interprets it as a SAS missing value.

By default, Oracle columns accept NULL values. However, you can define columns so that they cannot contain NULL data. NOT NULL tells Oracle not to add a row to the table unless the row has a value for that column. When creating an Oracle table with SAS/ACCESS, you can use the DBNULL= data set option to indicate whether NULL is a valid value for specified columns.

For more information about how SAS handles NULL values, see in *SAS/ACCESS for Relational Databases: Reference*.

Note: To control how SAS missing character values are handled by Oracle, use the NULLCHAR= and NULLCHARVAL= data set options. Δ

LIBNAME Statement Data Conversions

The following table shows the default SAS variable formats that SAS/ACCESS assigns to Oracle data types during input operations when you use the LIBNAME statement.

Table 1.4 LIBNAME Statement: Default SAS Formats for Oracle Data Types

Oracle Data Type	Default SAS Format
CHAR(<i>n</i>)	$\$n.$
VARCHAR2(<i>n</i>)	$\$n.$
NUMBER	none (BEST. on OS/390)
NUMBER(<i>p</i>)	$w.$ (BEST. on OS/390)
NUMBER(<i>p</i> , <i>s</i>)	$w.d$
DATE	DATE TIME20.
LONG	$\$1024.$
RAW(<i>n</i>)	$\$HEXw.$
LONG RAW	$\$HEX2048.$

Oracle data types that are omitted from this table are not supported by SAS/ACCESS.

If Oracle data falls outside valid SAS data ranges, the values are usually counted as missing.

Note: SAS automatically converts Oracle NUMBER types to SAS number formats by using an algorithm that determines the correct scale and precision. When the scale and precision cannot be determined, SAS/ACCESS allows the procedure or application to determine the format.

You can also convert numeric data to character data by using the Pass-Through facility with the Oracle TO_CHAR function. See your Oracle documentation for more details. △

The following table shows the default Oracle data types that SAS/ACCESS assigns to SAS variable formats during output operations when you use the LIBNAME statement.

Table 1.5 LIBNAME Statement: Default Oracle Data Types for SAS Formats

SAS Variable Format	Oracle Data Type
$\$w.$	VARCHAR2(n)***
$w.$ with SAS format name of NULL	NUMBER(p)
$w.d$ with SAS format name of NULL	NUMBER(p,s)
all other numerics *	NUMBER (NUMBER(38,10) on OS/390)
datetime $w.d$	DATE
date $w.$	DATE
time. **	DATE

* Includes all SAS numeric formats, such as BINARY8 and E10.0.

** Includes all SAS time formats, such as TOD w,d and HHMM w,d .

***If the SAS char variable format is greater than \$4000, the default Oracle data type is LONG. If you use Oracle7 and the CHAR variable format is between \$2000. and \$4000. then use the DBTYPE= option to change the default VARCHAR2 type to LONG.

To override these data types, use the DBTYPE= data set option during output processing.

ACCESS Procedure Data Conversions

The following table shows the default SAS variable formats that SAS/ACCESS assigns to Oracle data types when you use the ACCESS procedure.

Table 1.6 PROC ACCESS: Default SAS Formats for Oracle Data Types

Oracle Data Type	Default SAS Format
CHAR(n)	$\$n.$ ($n \leq 200$) \$200. ($n > 200$)
VARCHAR2(n)	$\$n.$ ($n \leq 200$) \$200. ($n > 200$)
FLOAT	BEST22.
NUMBER	BEST22.
NUMBER(p)	$w.$
NUMBER(p, s)	$w.d$
DATE	DATETIME16.
LONG	\$200.

Oracle Data Type	Default SAS Format
RAW(<i>n</i>)	$\$n$. ($n < 200$) $\$200$. ($n > 200$)
LONG RAW	$\$200$.

Oracle data types that are omitted from this table are not supported by SAS/ACCESS. If Oracle data falls outside valid SAS data ranges, the values are usually counted as missing.

The following table shows the correlation between the Oracle NUMBER data types and the default SAS formats that are created from that data type.

Table 1.7 Default SAS Formats for Oracle NUMBER Data Types

Oracle NUMBER Data Type	Rules	Default SAS Format
NUMBER(<i>p</i>)	$0 < p \leq 32$	$(p + 1).0$
NUMBER(<i>p,s</i>)	$p > 0, s < 0, s < p$	$(p + s + 1).0$
NUMBER(<i>p,s</i>)	$p > 0, s < 0, s \geq p$	$(p + s + 1).0$
NUMBER(<i>p,s</i>)	$p > 0, s > 0, s < p$	$(p + 2).s$
NUMBER(<i>p,s</i>)	$p > 0, s > 0, s \geq p$	$(s + 3).s$
NUMBER(<i>p</i>)	$p > 32$	BEST22. SAS selects format
NUMBER	p, s unspecified	BEST22. SAS selects format

Note: The general form of an Oracle number is NUMBER(*p,s*) where *p* is the precision and *s* is the scale of the number. Oracle defines precision as the total number of digits, with a valid range of -84 to 127. However, a negative scale means that the number is rounded to the specified number of places to the left of the decimal. For example, if the number 1,234.56 is specified as data type NUMBER(8,-2), it is rounded to the nearest hundred and stored as 1,200. Δ

DBLOAD Procedure Data Conversions


The following table shows the default Oracle data types that SAS/ACCESS assigns to SAS variable formats when you use the DBLOAD procedure.

Table 1.8 PROC DBLOAD: Default Oracle Data Types for SAS Formats

SAS Variable Format	Oracle Data Type
$\$w$.	CHAR(<i>n</i>)
<i>w</i> .	NUMBER(<i>p</i>)
<i>w.d</i>	NUMBER(<i>p,s</i>)
all other numerics *	NUMBER
datetime <i>w.d</i>	DATE
date <i>w</i> .	DATE
time. **	NUMBER

* Includes all SAS numeric formats, such as BINARY8 and E10.0.

** Includes all SAS time formats, such as TOD*w,d* and HHMM*w,d*.



APPENDIX

1

Recommended Reading

Recommended Reading 25

Recommended Reading

Here is the recommended reading list for this title:

- SAS/ACCESS for Relational Databases: Reference*
- SAS Language Reference: Concepts*
- SAS Language Reference: Dictionary*
- Base SAS Procedures Guide*
- SAS Companion that is specific to your operating environment

For a complete list of SAS publications, see the current *SAS Publishing Catalog*. To order the most current publications or to receive a free copy of the catalog, contact a SAS representative at

SAS Publishing Sales
SAS Campus Drive
Cary, NC 27513
Telephone: (800) 727-3228*
Fax: (919) 677-8166
E-mail: sasbook@sas.com

Web address: support.sas.com/pubs

* For other SAS Institute business, call (919) 677-8000.

Customers outside the United States should contact their local SAS office.

Glossary

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**
- abstract data
 - Oracle data types 21
 - access descriptors
 - Oracle data, based on (example) 19
 - ACCESS procedure
 - Oracle specifics 11, 23
 - autopartitioning
 - Oracle specifics 8
- B**
- BL_ bulk loading options
 - Oracle interface 17
 - BUFSIZE= option, SQL CONNECT statement 7
 - bulk loading
 - Oracle 17, 18
 - BULKLOAD= option, LIBNAME statement
 - Oracle interface 17, 18
- C**
- CHAR data type
 - Oracle 21
 - character data
 - Oracle data types 21
 - CONNECT statement, SQL procedure
 - options, Oracle 6
- D**
- data set options
 - Oracle interface 4
 - data types
 - Oracle servers 20
 - date and time data
 - Oracle data types 21
 - DATE data type
 - Oracle 21
 - DBLOAD procedure
 - Oracle specifics 13, 24
- J**
- joins
 - passing to Oracle 16
- L**
- LIBNAME statement
 - Oracle specifics 2, 22
 - locking data, handling
 - Oracle interface 19
 - LONG data type 21
 - LONG RAW data type 21
- M**
- MOD function
 - nonpartitioned Oracle tables 9
- N**
- naming conventions
 - Oracle interface 20
 - NULL values
 - Oracle 22
 - NUMBER data type 21
 - numeric data
 - Oracle data types 21
- O**
- Oracle, interface to 2
 - ACCESS procedure 11, 23
 - autopartitioning scheme 8
 - bulk loading 17, 18
 - data set options 4
 - data types 20
 - DBLOAD procedure 13, 24
 - LIBNAME statement 2, 22
 - locking in 19
 - naming conventions 20
 - Pass-Through Facility 6
 - passing joins to 16
 - passing SAS functions to 15
 - Oracle SQL*Loader 17
 - OS/390 operating environment 18
 - ORAPW= option
 - PROC ACCESS statement 12
 - PROC DBLOAD statement 13
 - OS/390, interface to
 - Oracle bulk loading 18
- P**
- partitioning
 - Oracle tables 8
 - Pass-Through Facility
 - Oracle interface 6
 - PASSWORD= option, LIBNAME statement
 - Oracle interface 2
 - PASSWORD= option, SQL CONNECT statement
 - Oracle interface 6
 - PATH= option
 - LIBNAME statement 2
 - PROC ACCESS statement 12
 - PROC DBLOAD statement 13
 - SQL CONNECT statement 7
 - performance
 - Oracle, SAS/ACCESS for 20
 - Oracle partitioned tables 10
 - PRESERVE_COL_NAMES= option
 - Oracle interface 20
 - PRESERVE_COMMENTS option, SQL CONNECT statement 7
 - PRESERVE_TAB_NAMES= option, LIBNAME statement
 - Oracle interface 20
- R**
- RAW data type 21
 - READ_ISOLATION_LEVEL= option
 - Oracle interface 20
 - READ_LOCK_TYPE= option
 - Oracle interface 19
- S**
- SAS/ACCESS data set options
 - Oracle interface 4

SAS SQL functions
 passing to Oracle 15

T

TABLE= option, DBLOAD procedure
 Oracle 14
TABLE= statement, ACCESS procedure
 Oracle interface 19
TABLESPACE= option, PROC DBLOAD state-
 ment 14

threaded reads
 Oracle interface 8

U

UPDATE_ISOLATION_LEVEL= option
 Oracle interface 20
UPDATE_LOCK_TYPE= option
 Oracle interface 19

USER= option
 PROC ACCESS statement 12
 PROC DBLOAD statement 13
 SQL CONNECT statement 6
USER= option, LIBNAME statement
 Oracle interface 2

V

VARCHAR2 data type 21
view descriptors
 Oracle data, based on (example) 19

Your Turn

If you have comments or suggestions about *SAS/ACCESS 9.1.3 Supplement for Oracle (SAS/ACCESS for Relational Databases)*, please send them to us on a photocopy of this page, or send us electronic mail.

For comments about this book, please return the photocopy to

SAS Publishing
SAS Campus Drive
Cary, NC 27513
E-mail: yourturn@sas.com

For suggestions about the software, please return the photocopy to

SAS Institute Inc.
Technical Support Division
SAS Campus Drive
Cary, NC 27513
E-mail: suggest@sas.com

