



SAS Publishing



# **SAS/ACCESS<sup>®</sup> 9.1**

# **Supplement for ODBC**

SAS/ACCESS for Relational Databases

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2004. *SAS/ACCESS® 9.1 Supplement for ODBC (SAS/ACCESS for Relational Databases)*. Cary, NC: SAS Institute Inc.

**SAS/ACCESS® 9.1 Supplement for ODBC (SAS/ACCESS for Relational Databases)**

Copyright © 2004, SAS Institute Inc., Cary, NC, USA

ISBN 1-59047-249-7

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

**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, January 2004

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](http://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

<b>Chapter 1</b>	<b>△ SAS/ACCESS for ODBC</b>	<b>1</b>
Introduction to the SAS/ACCESS Interface to ODBC		1
LIBNAME Statement Specifics for ODBC		3
Data Set Options for ODBC		7
Pass-Through Facility Specifics for ODBC		8
Autopartitioning Scheme for ODBC		15
DBLOAD Procedure Specifics for ODBC		19
Passing SAS Functions to ODBC		21
Passing Joins to ODBC		22
Temporary Table Support for ODBC		22
ODBC Bulk Loading		24
Locking in the ODBC Interface		24
Naming Conventions for ODBC		25
Data Types for ODBC		26
<b>Appendix 1</b>	<b>△ Recommended Reading</b>	<b>29</b>
Recommended Reading		29
<b>Glossary</b>		<b>31</b>
<b>Index</b>		<b>37</b>





## CHAPTER

## 1

# SAS/ACCESS for ODBC

---

<i>Introduction to the SAS/ACCESS Interface to ODBC</i>	1
<i>Overview of ODBC</i>	2
<i>LIBNAME Statement Specifics for ODBC</i>	3
<i>Arguments</i>	3
<i>ODBC LIBNAME Statement Examples</i>	7
<i>Data Set Options for ODBC</i>	7
<i>Pass-Through Facility Specifics for ODBC</i>	8
<i>CONNECT Statement Examples</i>	9
<i>Pass-Through Views</i>	10
<i>IBM AS/400 Specifics</i>	10
<i>Microsoft SQL Server Specifics</i>	12
<i>Connection To Component Examples</i>	12
<i>Special ODBC Queries</i>	13
<i>Autopartitioning Scheme for ODBC</i>	15
<i>Overview</i>	15
<i>Autopartitioning Restrictions</i>	15
<i>Nullable Columns</i>	15
<i>Using WHERE Clauses</i>	16
<i>Using DBSLICEPARM=</i>	16
<i>Using DBSLICE=</i>	16
<i>Configuring SQL Server Partitioned Views for Use with DBSLICE=</i>	17
<i>DBLOAD Procedure Specifics for ODBC</i>	19
<i>Examples</i>	20
<i>Passing SAS Functions to ODBC</i>	21
<i>Passing Joins to ODBC</i>	22
<i>Temporary Table Support for ODBC</i>	22
<i>Establishing a Temporary Table</i>	22
<i>Terminating a Temporary Table</i>	22
<i>Examples</i>	22
<i>ODBC Bulk Loading</i>	24
<i>Locking in the ODBC Interface</i>	24
<i>Naming Conventions for ODBC</i>	25
<i>Data Types for ODBC</i>	26
<i>ODBC Null Values</i>	27

---

## Introduction to the SAS/ACCESS Interface to ODBC

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

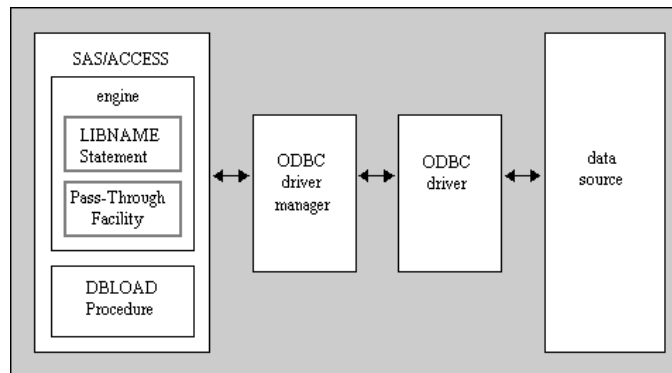
## Overview of ODBC

Open database connectivity (ODBC) standards provide a common interface to a variety of databases, including AS/400, dBASE, Microsoft Access, Oracle, Paradox, and Microsoft SQL Server databases. Specifically, ODBC standards define application programming interfaces (APIs) that enable an application to access a database if both the application and the database adhere to the specification. ODBC also provides a mechanism to enable dynamic selection of a database that an application is accessing, so end users have the flexibility of selecting databases other than those that are specified by the application developer.

The basic components and features of ODBC include the following:

- ODBC functionality is provided by three components: the client interface, the ODBC driver manager, and the ODBC driver. SAS provides the SAS/ACCESS interface to ODBC, which is the client interface. For PC platforms, Microsoft developed the ODBC Administrator, which is used from the Windows Control Panel to perform software administration and maintenance activities. The ODBC driver manager also manages the interaction between the client interface and the ODBC driver. Other software vendors provide the ODBC manager with their ODBC drivers, which process requests for external data. These drivers also either directly manipulate and retrieve the data or they pass the request to a native library for the specific DBMS. The ODBC interface to SAS is illustrated in the figure below.

**Figure 1.1** The ODBC Interface to SAS



- The ODBC administrator defines a *data source* as the data that is used in an application and the operating system and network that are used to access the data. You create a data source by using the ODBC administrator in the Windows Control Panel, selecting an ODBC driver, and providing the information (for example, data source name, user ID, password, description, server name) required by the driver to make a connection to the desired data. The driver displays dialog boxes in which you enter this information. During operation, a client application usually requests a connection to a named data source, not just to a specific ODBC driver. In a UNIX environment such as HP-UX, AIX, or Solaris, no ODBC Administrator exists. During an install, the driver creates a generic .odbc.ini file that can be edited to create your own data source names.

For more information about customizing your SAS application, refer to your vendor-specific documentation.

- ODBC uses SQL syntax for queries and statement execution (or for statements that are executed as commands). However, all databases that support ODBC are

not necessarily SQL databases. For example, many databases do not have system tables, and the term *table* may be used to describe a variety of items, including a file, a part of a file, a group of files, a typical SQL table, generated data, or any potential source of data. This distinction is important because although all ODBC data sources respond to a base set of SQL statements such as SELECT, INSERT, UPDATE, DELETE, CREATE, and DROP in their simplest forms, some databases do not support other statements and more complex forms of the SQL statements.

- The ODBC standard allows for various levels of conformance, generally categorized as low, medium, and high. As mentioned previously, the level of SQL syntax that is supported varies. There are also many programming interfaces that might not be supported by some drivers. The SAS/ACCESS interface to ODBC is designed to work with API calls that conform to the lowest level of ODBC compliance, Level 1. However, the interface does use some Level 2 API calls if they are available.

It is the responsibility of the SAS programmer or end user to ensure that the SQL syntax that is used is supported by the particular driver that is being used. If the ODBC driver supports a higher level of API conformance, some of the advanced features are made available through the PROC SQL CONNECT statement and special queries supported by the SAS/ACCESS interface to ODBC. For more information, see “Special ODBC Queries” on page 13.

- The ODBC manager and drivers return standard operation states and custom text for any warnings or errors. The state variables and their associated text are available through the SAS macro variables SYSDBRC and SYSDBMSG.
- There are three types of data source names that can be specified. A user DSN is specific to an individual user and is available only to the user who creates it. A system DSN can be used by anyone who has permission to access the data source. A file DSN can be shared among users even though it is created locally. Since it is file based, it contains all the information that is required to connect to a data source.
- In addition to the information provided in this documentation, you need to refer to the documentation provided with your ODBC driver. Most ODBC drivers supply a help file that you can access online. In the Windows Control Panel, double click the ODBC icon to start the ODBC Administrator application. Within the ODBC Data Source Administrator, double-click the data source name from the User DSN, System DSN, or File DSN tabbed dialog box. This brings up the ODBC driver setup dialog box for your specific ODBC driver. Clicking the Help button provides the information that you need to configure the ODBC data source for your driver.

---

## LIBNAME Statement Specifics for ODBC

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

```
LIBNAME libref odbc <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.

odbc

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

*connection-options*

provide connection information and control how SAS manages the timing and concurrence of the connection to the DBMS. There are multiple ways that you can connect to ODBC when you use the LIBNAME statement. The methods are mutually exclusive, so you must use only one of the following for each connection:

- specify USER=, PASSWORD=, and DATASRC=
- specify COMPLETE=
- specify NOPROMPT=
- specify PROMPT=
- specify REQUIRED=.

These connection options are defined as follows:

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

enables you to connect to an ODBC database, such as Microsoft SQL Server or AS/400, with a user ID that is different from the default ID.

USER= is optional. UID= is an alias for this option.

PASSWORD=<'>*password*<'>

specifies the ODBC password that is associated with your user ID.

PASSWORD= is optional. PWD is an alias for this option.

*Note:* If you do not wish to enter your ODBC password in clear text on this statement, see PROC PWENCODE for a method to encode it.  $\triangle$

DATASRC=<'>*ODBC-data-source*<'>

specifies the ODBC data source to which you want to connect. For PC platforms, data sources must be configured by using the ODBC icon in the Windows Control Panel. For UNIX platforms, data sources must be configured by modifying the .odbc.ini file.

DSN= is an alias for this option that indicates that the connection is attempted using the ODBC SQLConnect API, which requires a data source name. Optionally, a user ID and password can be used in conjunction with DSN=.

If you want to use an ODBC file DSN, then instead of supplying DATASRC=<'>*ODBC-data-source*<'>, use the PROMPT= or NOPROMPT= option followed by "*filedsn=(name-of-your-file-dsn)*";. For example:

```
libname mydblib odbc noprompt="filedsn=d:\share\msafiledsn.dsn";
```

COMPLETE=<'>*ODBC-connection-options*<'>

specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connection string is returned in the SYSDBMSG macro variable.

If you do not specify enough correct connection options, you are prompted with a dialog box that displays the values from the COMPLETE= connection string. You can edit any field before you connect to the data source.

This option is not supported on UNIX platforms. See your ODBC driver documentation for more details.

NOPROMPT=<'>*ODBC-connection-options*<'>

specifies connection options for your data source or database. Separate multiple options with a semicolon.

If you do not specify enough correct connection options, an error is returned. No dialog box is displayed to help you complete the connection string.



PROMPT=<'>*ODBC-connection-information*<'>

specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connection string is returned in the SYSDBMSG macro variable.

PROMPT= does not immediately attempt to connect to the DBMS. Instead, it displays a dialog box that contains the values that you entered in the PROMPT= connection string. You can edit values or enter additional values in any field before you connect to the data source.

This option is not supported on UNIX platforms.

REQUIRED=<'>*ODBC-connection-options*<'>

specifies connection options for your data source or database. Separate multiple options with a semicolon. When a successful connection is made, the complete connection string is returned in the SYSDBMSG macro variable.

If you do not specify enough correct connection options, a dialog box prompts you for the connection options. REQUIRED= allows you to modify only required fields in the dialog box.

This option is not supported on UNIX platforms.

*Note:* See your ODBC driver documentation for a list of the ODBC connection options that your ODBC driver supports. △

The following ODBC connection options are not supported on UNIX:

REQUIRED=

BULKCOPY=

PROMPT=

COMPLETE=

#### *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 the LIBNAME options that are supported for ODBC, 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 ODBC

Option	Default Value
ACCESS=	none
AUTOCOMMIT=	data source specific
BL_LOG=	none
BL_OPTIONS=	none
BULKLOAD=	NO
CONNECTION=	data source specific
CONNECTION_GROUP=	none
CURSOR_TYPE=	DYNAMIC
DBCOMMIT=	1000 (inserting) or 0 (updating)
DBCONINIT=	none
DBCONTERM=	none

Option	Default Value
DB_CREATE_TABLE_OPTS=	none
DBGEN_NAME=	DBMS
DBINDEX=	YES
DBLIBINIT=	none
DBLIBTERM=	none
DBMAX_TEXT=	1024
DBNULLKEYS=	YES
DBPROMPT=	NO
DBSLICEPARM=	THREADED_APPS,2 or 3
DEFER=	NO
DELETE_MULT_ROWS=	NO
DIRECT_EXE=	none
DIRECT_SQL=	YES
IGNORE_READ_ONLY_COLUMNS=	NO
INSERT_SQL=	data source specific
INSERTBUFF=	1
KEYSET_SIZE=	0
MULTI_DATASRC_OPT=	NONE
PRESERVE_COL_NAMES=	see “Naming Conventions for ODBC” on page 25
PRESERVE_TAB_NAMES =	see “Naming Conventions for ODBC” on page 25
QUALIFIER=	none
QUERY_TIMEOUT=	0
QUOTE_CHAR=	none
READBUFF=	0
READ_ISOLATION_LEVEL=	RC (see “Locking in the ODBC Interface” on page 24)
READ_LOCK_TYPE=	ROW
REREAD_EXPOSURE=	NO
SCHEMA=	none
SPOOL=	YES
SQL_FUNCTIONS=	NONE
STRINGDATES=	NO
TRACE=	NO
TRACEFILE=	none
UPDATE_ISOLATION_LEVEL=	RC (see “Locking in the ODBC Interface” on page 24)
UPDATE_LOCK_TYPE=	ROW
UPDATE_MULT_ROWS=	NO

Option	Default Value
UPDATE_SQL=	driver specific
USE_ODBC_CL =	NO
UTILCONN_TRANSIENT=	YES

## ODBC LIBNAME Statement Examples

In the following example, USER=, PASSWORD=, and DATASRC= are connection options.

```
libname mydblib odbc user=testuser password=testpass datasrc=mydatasource;
```

In the following example, the libref MYLIB uses the ODBC engine to connect to an AS/400 database. The connection options are USER=, PASSWORD=, and DATASRC=.

```
libname mydblib odbc datasrc=as400 user=testuser
password=testpass;
```

```
proc print data=mydblib.customers;
  where state='CA';
run;
```

In the following example, the libref MYDBLIB uses the ODBC engine to connect to a Microsoft SQL Server database. The connection option is NOPROMPT=.

```
libname mydblib odbc
  noprompt="uid=testuser;pwd=testpass;dsn=sqlservr;"
  stringdates=yes;
```

```
proc print data=mydblib.customers;
  where state='CA';
run;
```

## Data Set Options for ODBC

The following table describes the data set options that are supported for ODBC, and provides default values where applicable. See the section about data set options in *SAS/ACCESS for Relational Databases: Reference* for detailed information about these options.

**Table 1.2** SAS/ACCESS Data Set Options

Option	Default Value
CURSOR_TYPE=	LIBNAME option setting
DBCOMMIT=	LIBNAME option setting
DBCONDITION=	none
DBCREATE_TABLE_OPTS=	LIBNAME option setting
DBFORCE=	NO
DBGEN_NAME=	DBMS

Option	Default Value
DBINDEX=	LIBNAME option setting
DBKEY=	none
DBLABEL=	NO
DBMASTER=	none
DBMAX_TEXT=	1024
DBNULL=	YES
DBNULLKEYS=	LIBNAME option setting
DBPROMPT=	LIBNAME option setting
DBSASTYPE=	see “Data Types for ODBC” on page 26
DBSLICE=	none
DBSLICEPARM=	THREADED_APPS,2 or 3
DBTYPE=	see “Data Types for ODBC” on page 26
ERRLIMIT=	1
IGNORE_READ_ONLY_COLUMNS=	NO
INSERT_SQL=	LIBNAME option setting
INSERTBUFF=	LIBNAME option setting
KEYSET_SIZE=	LIBNAME option setting
NULLCHAR=	SAS
NULLCHARVAL=	a blank character
PRESERVE_COL_NAMES=	LIBNAME option setting
QUALIFIER=	LIBNAME option setting
QUERY_TIMEOUT=	LIBNAME option setting
READBUFF=	LIBNAME option setting
READ_ISOLATION_LEVEL=	LIBNAME option setting
READ_LOCK_TYPE=	LIBNAME option setting
SASDATEFMT=	none
SCHEMA=	LIBNAME option setting
UPDATE_ISOLATION_LEVEL=	LIBNAME option setting
UPDATE_LOCK_TYPE=	LIBNAME option setting
UPDATE_SQL=	LIBNAME option setting

---

## Pass-Through Facility Specifics for ODBC

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 the ODBC interface are as follows:

- The *dbms-name* is **ODBC**.
- The CONNECT statement is required.

- PROC SQL supports multiple connections to ODBC. If you use multiple simultaneous connections, you must use the *alias* argument to identify the different connections. If you do not specify an alias, the default alias, **odbc**, is used. The functionality of multiple connections to the same ODBC data source might be limited by the particular data source's driver.
- The CONNECT statement *database-connection-arguments* are identical to its LIBNAME connection-options. Not all of these arguments are supported by all ODBC drivers. Refer to your driver documentation for more information.
- On some DBMSs, the *DBMS-SQL-query* argument can be a DBMS-specific SQL EXECUTE statement that executes a DBMS stored procedure. However, if the stored procedure contains more than one query, only the first query is processed.
- The following LIBNAME options are available with the CONNECT statement:
  - AUTOCOMMIT=
  - CURSOR\_TYPE=
  - KEYSET\_SIZE=
  - QUERY\_TIMEOUT=
  - READBUFF=
  - READ\_ISOLATION\_LEVEL=
  - TRACE=
  - TRACEFILE=
  - USE\_ODBC\_CL=

See the section about the LIBNAME statement in *SAS/ACCESS for Relational Databases: Reference* for information about these options.

---

## CONNECT Statement Examples

The following examples use ODBC to connect to a data source that is configured under the data source name **User's Data** using the alias USER1. The first example uses the connection method that is guaranteed to be present at the lowest level of ODBC conformance. Note that DATASRC= names can contain quotation marks and spaces.

```
proc sql;
  connect to ODBC as user1
  (datasrc="User's Data" user=testuser password=testpass);
```

The following example uses the connection method that represents a more advanced level of ODBC conformance. It uses the input dialog box that is provided by the driver. The DATASRC= and USER= arguments are within the connection string and, therefore, are not parsed by the Pass-Through Facility but instead are passed to the ODBC manager.

```
proc sql;
  connect to odbc as user1
  (required = "dsn=User's Data;uid=testuser");
```

The following ODBC example enables you to select any data source that is configured on your machine. The example uses the connection method that represents a more advanced level of ODBC conformance, Level 1. When a successful connection is made, the connection string is returned in the SQLXMSG and SYSDBMSG macro variables and can be stored if this method is used to configure a connection for later use.

```
proc sql;
  connect to odbc (required);
```

The following ODBC example prompts you to specify the information that is required to make a connection to the DBMS. You are prompted to supply the data source name, user ID, and password in the dialog boxes that are displayed.

```
proc sql;
  connect to odbc (prompt);
```

---

## Pass-Through Views

Version 6 SQL views do not need to be updated to be used in Version 7, Version 8, or SAS 9. The ODBC interface and DBMS client must be available and ready to connect. In order for any truncated variable names to be correctly interpreted by the ODBC driver, you must specify the VALIDVARNAME=V6.

The following example, must use the SAS option VALIDVARNAME=V6 in order to successfully process a Version 6 SQL view. See the section about the LIBNAME statement in *SAS/ACCESS for Relational Databases: Reference* for more information about this option.

```
options validvarname=v6;
proc sql;
  describe view as4sql.invoice4;
run;

/* NOTE: SQL view AS4SQL.INVOICE4 is defined as: */
select
  INVOICEN as INVOICE,
  AMTBILLE as AMOUNT format=DOLLAR20.2,
  BILLEDON
from connection to AS400
/* dbms=AS400, connect options=() */
(select invoicenum, amtbilled, billedon
  from sasdemo/invoice
  where paidon = '18OCT1998');
```

*Note:* If a view cannot be processed, or if you want to see what a view is, use the DESCRIBE VIEW statement to see what the existing view is. Then you can use the PROC SQL statements to create a new view for the ODBC connection.  $\Delta$

In Version 6, the AS/400 column name INVOICENUM is mapped to the SAS variable INVOICEN, and AMTBILLED is mapped to AMTBILLE. If you do not specify option VALIDVARNAME=V6, you get the following error because the ODBC driver attempts to find the truncated column names in the DBMS table:

```
ERROR: The following columns were not found in the
       contributing tables: AMTBILLE, INVOICEN.
```

---

## IBM AS/400 Specifics

To run your SQL views for IBM AS/400, you must do the following:

- create a data source name first by using the ODBC administrator. Refer to the installation instructions for the SAS/ACCESS interface to ODBC for more information.
- set the environment variable AS400DSN (located in your SASV9.cfg) to the data source name that you assigned. Quotation marks are required if the name includes blanks or special characters.

In this example,

```
CONNECT TO AS400 AS market;
```

is converted to

```
CONNECT TO ODBC AS market
  (NOPROMPT="DATASRC=IBM AS/400 Database;
    USER=TESTUSER; PASSWORD=TESTPASS;
    NAM=1"
  )
;
```

This example demonstrates a problem in which AS/400 short alias names cannot be returned by the AS/400 ODBC driver. This problem causes you to get an error, for example, if you have specified the short alias names in your selection list before the CONNECTION TO component, but have not specified the short alias names in the selection list that defines the view. If you encounter this problem with your Version 6 SQL views, you need to re-create the views.

This example creates an AS/400 table named TEST5 with the columns CUSTOMER\_FIRST\_NAME and CUSTOMER\_LAST\_NAME. The short name alias for CUSTOMER\_FIRST\_NAME is FNAME and the short name alias for CUSTOMER\_LAST\_NAME is LNAME.

```
options validvarname=v6;
%let name=test5;
proc sql;
  describe view as4sql.&name;
  /* NOTE: SQL view AS4SQL.TEST5 is defined as: */
  select FNAME, LNAME from connection to AS400
  /* dbms=AS400, connect options=() */
  ( select * from sasdemo/test
    where lname = 'Ju' );
quit;
proc print data=as4sql.&name;
run;
```

This example generates the following errors:

```
ERROR: The following columns were not found in
       the contributing tables: FNAME, LNAME.
ERROR: SQL View AS4SQL.TEST5 could not be processed.
```

The following two examples work successfully because the short alias names are specified in the SELECT statement that defines the view.

```
create view as4sql.&name as
  select FNAME, LNAME from connection to AS400
  /* dbms=AS400, connect options=() */
  (select FNAME, LNAME from sasdemo/test
    where lname = 'Ju' );

create view as4sql.&name as
  select * from connection to AS400
  /* dbms=AS400, connect options=() */
  (select fname, lname from sasdemo/test
    where lname = 'Ju' );
```

---

## Microsoft SQL Server Specifics

To run your SQL views for Microsoft SQL Server, you are encouraged, but not required, to create a data source name. You can use the ODBC administrator to create it. Refer to the installation instructions for this interface for more information. If you do create a data source name, you must set the environment variable MSSQLDSN to be the *'data-source-name'*. Quotation marks are required if the name includes blanks or special characters.

In this example,

```
CONNECT TO SQLSERVER AS finance
  (user=testuser password=testpass
   server='dbipcl.pc.sas.com'
   database='sample'
  )
;
```

is converted to

```
CONNECT TO ODBC AS finance
  (NOPROMPT="DATASRC=Microsoft SQL Server Database;
   SERVER=dbipcl.pc.sas.com;
   USER=testuser; PASSWORD=testpass;
   DATABASE=sample"
  )
;
```

---

## Connection To Component Examples

The following example sends an Oracle SQL query (presented in highlighted text) to the Oracle database for processing. The results from the query serve as a virtual table for the PROC SQL FROM clause. In this example, MYCON is a connection alias.

```
proc sql;
connect to odbc as mycon
  (datasrc=ora7 user=testuser password=testpass);

select *
  from connection to mycon
    (select empid, lastname, firstname,
      hiredate, salary
     from sasdemo.employees
     where hiredate>='31.12.1988' ) ;

disconnect from mycon;
quit;
```

The following example gives the previous query a name and stores it as the SQL view Samples.Hires88. The CREATE VIEW statement appears highlighted.

```
libname samples 'SAS-data-library';

proc sql;
connect to odbc as mycon
  (datasrc=ora7 user=testuser password=testpass);
```



```

create view samples.hires88 as
  select *
    from connection to mycon
      (select empid, lastname, firstname,
         hiredate, salary from sasdemo.employees
         where hiredate>='31.12.1988');

disconnect from mycon;
quit;

```

The following example connects to Microsoft Access 7 and creates a view NEWORDERS from all the columns in the ORDERS table.

```

proc sql;
  connect to odbc as mydb
    (datasrc=access7);
  create view neworders as
    select * from connection to mydb
      (select * from orders);
disconnect from mydb;
quit;

```

The following example sends an SQL query to Microsoft SQL Server 6.5, configured under the data source name **SQL Server**, for processing. The results from the query serve as a virtual table for the PROC SQL FROM clause.

```

proc sql;
  connect to odbc as mydb
    (datasrc="SQL Server" user=testuser password=testpass);
  select * from connection to mydb
    (select CUSTOMER, NAME, COUNTRY
     from CUSTOMERS
     where COUNTRY <> 'USA');
quit;

```

The following example returns a list of the columns in the CUSTOMERS table.

```

proc sql;
  connect to odbc as mydb
    (datasrc="SQL Server" user=testuser password=testpass);
  select * from connection to mydb
    (ODBC::SQLColumns ( , , "CUSTOMERS"));
quit;

```

---

## Special ODBC Queries

The following special queries are supported by the SAS/ACCESS interface to ODBC. Many databases provide or use system tables that allow queries to return the list of available tables, columns, procedures, and other useful information. In ODBC, much of this functionality is provided through special APIs (application programming interfaces) in order to accommodate databases that do not follow the SQL table structure. You can use these special queries on non-SQL and SQL databases. The general format of the special queries is as follows:

```
ODBC::SQLAPI "parameter 1", "parameter n"
```

where

ODBC::

is required to distinguish special queries from regular queries.

SQLAPI

is the specific API that is being called. Both ODBC:: and SQLAPI are case sensitive.

"parameter n"

is a quoted string that is delimited by commas.

Within the quoted string, two characters are universally recognized: the percent sign (%) and the underscore (\_). The percent sign matches any sequence of zero or more characters; the underscore represents any single character. Each driver also has an escape character that can be used to place characters within the string. Consult the driver's documentation to determine the valid escape character.

The values for the special query arguments are DBMS specific. For example, you supply the fully qualified table name for a "Catalog" argument. In dBase, the value of "Catalog" might be `c:\dbase\tst.dbf` and in SQL Server, the value might be `test.customer`. In addition, depending on the DBMS that you are using, valid values for a "Schema" argument might be a user ID, a database name, or a library. All arguments are optional. If you specify some but not all the arguments within a parameter, use a comma to indicate the omitted arguments. If you do not specify any parameters, commas are not necessary.

*Note:* These special queries might not be available for all ODBC drivers.  $\Delta$

The following special queries are supported:

ODBC::SQLTables <"Catalog", "Schema", "Table-name", "Type">

returns a list of all the tables that match the specified arguments. If no arguments are specified, all accessible table names and information are returned.

ODBC::SQLColumns <"Catalog", "Schema", "Table-name", "Column-name">

returns a list of all the columns that match the specified arguments. If no arguments are specified, all accessible column names and information are returned.

ODBC::SQLColumnPrivileges <"Catalog", "Schema", "Table-name", "Column-name">

returns a list of all the column privileges that match the specified arguments. If no arguments are specified, all accessible column names and privilege information are returned.

ODBC::SQLForeignKeys <"PK-catalog", "PK-schema", "PK-table-name", "FK-catalog", "FK-schema", "FK-table-name">

returns a list of all the columns that comprise foreign keys that match the specified arguments. If no arguments are specified, all accessible foreign key columns and information are returned.

ODBC::SQLPrimaryKeys <"Catalog", "Schema", "Table-name">

returns a list of all the columns that compose the primary key that matches the specified table. A primary key can be composed of one or more columns. If no table name is specified, this special query fails.

ODBC::SQLProcedureColumns <"Catalog", "Schema", "Procedure-name", "Column-name">

returns a list of all the procedure columns that match the specified arguments. If no arguments are specified, all accessible procedure columns are returned.

ODBC::SQLProcedures <"Catalog", "Schema", "Procedure-name">

returns a list of all the procedures that match the specified arguments. If no arguments are specified, all accessible procedures are returned.

ODBC::SQLSpecialColumns <"Identifier-type", "Catalog-name", "Schema-name", "Table-name", "Scope", "Nullable">  
 returns a list of the optimal set of columns that uniquely identify a row in the specified table.

ODBC::SQLStatistics <"Catalog", "Schema", "Table-name">  
 returns a list of the statistics for the specified table name, with options of SQL\_INDEX\_ALL and SQL\_ENSURE set in the SQLStatistics API call. If the table name argument is not specified, this special query fails.

ODBC::SQLTablePrivileges <"Catalog", "Schema", "Table-name">  
 returns a list of all the tables and associated privileges that match the specified arguments. If no arguments are specified, all accessible table names and associated privileges are returned.

ODBC::SQLGetTypeInfo  
 returns information about the data types that are supported in the data source.

---

## Autopartitioning Scheme for ODBC

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

---

### Overview

The autopartitioning method available for SAS/ACCESS to ODBC is modeled after the MOD function method as described in the section about autopartitioning techniques in *SAS/ACCESS for Relational Databases: Reference*.

---

### Autopartitioning Restrictions

SAS/ACCESS to ODBC places additional restrictions on which columns can be used for the partitioning column during the autopartitioning phase. Columns are partitioned as follows:

- SQL\_INTEGER, SQL\_BIT, SQL\_SMALLINT, and SQL\_TINYINT columns are given preference.
- SQL\_DECIMAL, SQL\_DOUBLE, SQL\_FLOAT, SQL\_NUMERIC, and SQL\_REAL columns might be used for partitioning, provided the following conditions are met:
  - The ODBC driver supports converting these types to SQL\_INTEGER via the INTEGER cast function.
  - The precision minus the scale of the column is greater than 0 but less than 10, that is,  $0 < (\text{precision} - \text{scale}) < 10$ .

The exception to the above rule is for Oracle SQL\_DECIMAL columns. As long as the scale of the SQL\_DECIMAL column is 0, the column can be used as the partitioning column.

---

### Nullable Columns

If a nullable column is selected for autopartitioning, then the SQL statement "OR<column-name>IS NULL" will be appended to the end of the SQL code that is

generated for the threaded read to ensure that any possible NULL values are returned in the result set. In addition, if the column to be used for the partitioning is SQL\_BIT, then the number of threads will automatically be changed to two, regardless of the setting of the DBSLICEPARM= option.

---

## Using WHERE Clauses

Autopartitioning does not select a column to be the partitioning column if it appears in the WHERE clause. For instance, the following data step would not be able to use a threaded read to retrieve the data since all of the numeric columns in the table are in the WHERE clause:

```
data work.locemp;
set trlib.MYEMPS;
where EMPNUM<=30 and ISTENURE=0 and
    SALARY<=35000 and NUMCLASS>2;
run;
```

---

## Using DBSLICEPARM=

When using autopartitioning, and DBSLICEPARM= does not specify a maximum number of threads to use for the threaded read, SAS/ACCESS to ODBC defaults to three threads.

---

## Using DBSLICE=

You might achieve the best possible performance when using threaded reads by specifying an ODBC-specific DBSLICE= option in your SAS operation. This is especially true if your DBMS supports multiple database partitions and provides a mechanism to allow connections to individual partitions. If your DBMS supports this concept, you can configure an ODBC datasource for each partition and use the DBSLICE= clause to specify both the datasource and the WHERE clause for each partition, as shown in the following example:

```
proc print data=trilib.MYEMPS(DBSLICE=(DSN1="EMPNUM BETWEEN 1 AND 33"
DSN2="EMPNUM BETWEEN 34 AND 66"
DSN3="EMPNUM BETWEEN 67 AND 100"));
run;
```

Consult your DBMS or ODBC driver documentation for more information about configuring for multiple partition access. You can also refer to “Configuring SQL Server Partitioned Views for Use with DBSLICE=” on page 17 for an example of configuring multiple partition access to a table.

Using the DATASOURCE= syntax is not required in order to use DBSLICE= with threaded reads for SAS/ACCESS to ODBC. The methods and examples described in DBSLICE= work well in instances where the table you want to read is not stored in multiple partitions in your DBMS. These methods also give you flexibility in column selection. For example, if you know that the STATE column in your employee table only contains a few distinct values, you can tailor your DBSLICE= clause accordingly:

```
datawork.locemp;
set trlib2.MYEMP(DBSLICE=("STATE='FL'" "STATE='GA'"
"STATE='SC'" "STATE='VA'" "STATE='NC'"));
```

```
where EMPNUM<=30 and ISTEMURE=0 and SALARY<=35000 and NUMCLASS>2;
run;
```

---

## Configuring SQL Server Partitioned Views for Use with DBSLICE=

Microsoft SQL Server implements multiple partitioning by creating a global view across multiple instances of a Microsoft SQL Server database. For this example, assume that Microsoft SQL Server has been installed on three separate machines (SERVER1, SERVER2, SERVER3), and three ODBC datasources (SSPART1, SSPART2, SSPART3) have been configured against these servers. Also, a linked server definition for each of these servers has been defined. This example uses SAS to create the tables and associated views, but this can be accomplished outside of the SAS environment.

First create a local SAS table to build the Microsoft SQL Server tables:

```
data work.MYEMPS;
format HIREDATE mmddyy 0. SALARY 9.2
      NUMCLASS 6. GENDER $1. STATE $2. EMPNUM 10.;
do EMPNUM=1 to 100;
  morf=mod(EMPNUM,2)+1;
  if(morf eq 1) then
    GENDER='F';
  else
    GENDER='M';
  SALARY=(ranuni(0)*5000);
  HIREDATE=int(ranuni(13131)*3650);
  whatstate=int(EMPNUM/5);
  if(whatstate eq 1) then
    STATE='FL';
  if(whatstate eq 2) then
    STATE='GA';
  if(whatstate eq 3) then
    STATE='SC';
  if(whatstate eq 4) then
    STATE='VA';
  else
    state='NC';
  ISTEMURE=mod(EMPNUM,2);
  NUMCLASS=int(EMPNUM/5)+2;
  output;
end;
run;
```

Next, create a table on each of the SQL server databases with the same table structure, and insert 1/3 of the overall data into each table. These table definitions also use CHECK constraints to enforce the distribution of the data on each of the subtables of the target view.

```
libname trlib odbc user=ssuser pw=sspwd dsn=sspart1;
proc delete data=trlib.MYEMPS1;
run;
data trlib.MYEMPS1(drop=morf whatstate
  DBTYPE=(HIREDATE="datetime" SALARY="numeric(8,2)"
  NUMCLASS="smallint" GENDER="char(1)" ISTEMURE="bit" STATE="char(2)"
  EMPNUM="int NOT NULL Primary Key CHECK (EMPNUM BETWEEN 0 AND 33)));
set work.MYEMPS;
```

```

where (EMPNUM BETWEEN 0 AND 33);
run;

libname trlib odbc user=ssuer pw=sspwd dsn=sspart2;
proc delete data=trlib.MYEMPS2;
run;
data trlib.MYEMPS2(drop=morf whatstate
  DBTYPE=(HIREDATE="datetime" SALARY="numeric(8,2)"
  NUMCLASS="smallint" GENDER="char(1)" ISTENURE="bit" STATE="char(2)"
  EMPNUM="int NOT NULL Primary Key CHECK (EMPNUM BETWEEN 34 AND 66)));
set work.MYEMPS;
where (EMPNUM BETWEEN 34 AND 66);
run;

libname trlib odbc user=ssuer pw=sspwd dsn=sspart3;
proc delete data=trlib.MYEMPS3;
run;
data trlib.MYEMPS3(drop=morf whatstate
  DBTYPE=(HIREDATE="datetime" SALARY="numeric(8,2)"
  NUMCLASS="smallint" GENDER="char(1)" ISTENURE="bit" STATE="char(2)"
  EMPNUM="int NOT NULL Primary Key CHECK (EMPNUM BETWEEN 67 AND 100)));
set work.MYEMPS;
where (EMPNUM BETWEEN 67 AND 100);
run;

```

Next, create a view using the UNION ALL construct on each Microsoft SQL Server instance which references the other two tables. This creates a global view that references the entire data set.

```

/*SERVER1,SSPART1*/
proc sql noerrorstop;
connect to odbc (UID=ssuser PWD=sspwd DSN=SSPART1);
execute (drop view MYEMPS) by odbc;
execute (create view MYEMPS AS
  SELECT * FROM users.ssuser.MYEMPS1
  UNION ALL
  SELECT * FROM SERVER2.users.ssuser.MYEMPS2
  UNION ALL
  SELECT * FROM SERVER3.users.ssuser.MYEMPS3) by odbc;
quit;

/*SERVER2,SSPART2*/
proc sql noerrorstop;
connect to odbc (UID=ssuser PWD=sspwd DSN=SSPART2);
execute (drop view MYEMPS) by odbc;
execute (create view MYEMPS AS
  SELECT * FROM users.ssuser.MYEMPS2
  UNION ALL
  SELECT * FROM SERVER1.users.ssuser.MYEMPS1
  UNION ALL
  SELECT * FROM SERVER3.users.ssuser.MYEMPS3) by odbc;
quit;

/*SERVER3,SSPART3*/
proc sql noerrorstop;

```

```

connect to odbc (UID=ssuser PWD=sspwd DSN=SSPART3);
execute (drop view MYEMPS) by odbc;
execute (create view MYEMPS AS
        SELECT * FROM users.ssuser.MYEMPS3
        UNION ALL
        SELECT * FROM SERVER2.users.ssuser.MYEMPS2
        UNION ALL
        SELECT * FROM SERVER1.users.ssuser.MYEMPS1) by odbc;
quit;

```

Finally, set up your SAS operation to perform the threaded read. The DBSLICE option contains the Microsoft SQL Server partitioning information.

```

proc print data=trlib.MYEMPS(DBLICE=(sspart1="EMPNUM BETWEEN 1 AND 33"
sspart2="EMPNUM BETWEEN 34 AND 66"
sspart3="EMPNUM BETWEEN 67 AND 100"));
run;

```

This configuration enables SAS/ACCESS to ODBC to access the data for the MYEMPS view directly from each subtable on the corresponding Microsoft SQL Server instance. The data is inserted directly into each subtable, but this process can also be accomplished by using the global view to divide up the data. For instance, you can create empty tables and then create the view as seen in the example with the UNION ALL construct. You can then insert the data into the view MYEMPS. The CHECK constraints will allow the Microsoft SQL Server query processor to determine which subtables should receive the data.

There are other tuning options available when configuring Microsoft SQL Server to use partitioned data. For more information, see the "Creating a Partitioned View" and "Using Views with Partitioned Data" sections in the *SQL Server On-line Guide*.

---

## DBLOAD Procedure Specifics for ODBC

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

The ODBC interface supports all of the DBLOAD procedure statements (except ACCDESC=) in batch mode. The DBLOAD procedure specifics for ODBC are as follows:

- The DBLOAD step DBMS= value is **ODBC**.
- PROC DBLOAD uses the following database description statements:

**DSN=** <'>*ODBC-data-source*<'>;

specifies the name of the data source in which you want to store the new ODBC table. The *data-source* is limited to eight characters.

The data source that you specify must already exist. If the data source name contains the following special characters (.,\$,@,#), you must enclose it in quotation marks. However, the ODBC standard recommends against using special characters in data source names.

**USER=** <'>*username*<'>;

enables you to connect to an ODBC database, such as Microsoft SQL Server or AS/400, with a user ID that is different from the default ID.

USER= is optional in ODBC. If you specify USER=, you must also specify PASSWORD=. If USER= is omitted, your default user ID is used.

**PASSWORD=**<'>*password*<'>;

specifies the ODBC password that is associated with your user ID.

PASSWORD= is optional in ODBC because users have default user IDs. If you specify USER=, you must specify PASSWORD=.

*Note:* If you do not wish to enter your ODBC password in uncoded text on this statement, see PROC PWENCODE for a method to encode it.  $\Delta$

BULKCOPY= YES|NO;

determines whether SAS uses the Microsoft Bulk Copy facility to insert data into a DBMS table (Microsoft SQL Server only). The default value is NO.

BCP is Microsoft's Bulk Copy facility, and it enables you to efficiently insert rows of data into a DBMS table as a unit. As SAS/ACCESS sends each row of data to BCP, the data is written to an input buffer. When you have inserted all the rows, or the buffer reaches a certain size (as determined by the DBCOMMIT= data set option), all of the rows are inserted as a unit into the table, and the data is committed to the table.

Alternatively, you can set the DBCOMMIT=*n* option to commit rows after every *n* insertions.

If an error occurs, a message is written to the SAS log, and any rows that have been inserted in the table before the error are rolled back.

*Note:* To use BULKCOPY=, your installation of Microsoft SQL Server must include the ODBCBCP.DLL, which is currently only supported by Microsoft SQL Server 7.0. BULKCOPY= is not supported on UNIX.  $\Delta$

- The TABLE= statement is as follows:

TABLE= <authorization-id.>table-name;

identifies the table or view that you want to use to create an access descriptor. The TABLE= statement is required.

The *authorization-id* is a user ID or group ID that is associated with the table.

- The NULLS statement is as follows:

NULLS *variable-identifier-1* =Y|N|D < . . . *variable-identifier-n* =Y|N|D >;

enables you to specify whether the columns that are associated with the listed SAS variables allow NULL values. By default, all columns accept NULL values.

The NULLS statement accepts any one of these three values:

Y – specifies that the column accepts NULL values. This is the default.

N – specifies that the column does not accept NULL values.

D – specifies that the column is defined as NOT NULL WITH DEFAULT

## Examples

The following example creates a new ODBC table, TESTUSER.EXCHANGE, from the DLIB.RATEOFEX data file. You must be granted the appropriate privileges in order to create new ODBC tables or views.

```
proc dbload dbms=odbc data=dlib.rateofex;
  dsn=sample;
  user='testuser';
  password='testpass';
  table=exchange;
  rename fgnindol=fgnindollars
         4=dollarsinfgn;
  nulls updated=n fgnindollars=n
        dollarsinfgn=n country=n;
```



```
load;  
run;
```

The following example only sends an ODBC SQL GRANT statement to the SAMPLE database and does not create a new table. Therefore, the TABLE= and LOAD statements are omitted.

```
proc dbload dbms=odbc;  
  user='testuser';  
  password='testpass';  
  dsn=sample;  
  sql grant select on testuser.exchange  
    to dbitest;  
run;
```

---

## Passing SAS Functions to ODBC

The interface to ODBC passes the following SAS functions to the data source for processing (if the DBMS server supports the function). See the section about optimizing SQL usage in *SAS/ACCESS for Relational Databases: Reference* for information.

ABS  
ARCOS  
ARSIN  
ATAN  
AVG  
CEIL  
COS  
EXP  
FLOOR  
LOG  
LOG10  
LOWCASE  
MAX  
MIN  
SIGN  
SIN  
SQRT  
TAN  
UPCASE  
SUM  
COUNT

---

## Passing Joins to ODBC

In order for a multiple libref join to pass to ODBC, all of the following components of the LIBNAME statements must match exactly:

- user ID
- password
- datasource
- catalog
- UPDATE\_ISOLATION\_LEVEL=  
(if specified)
- READ\_ISOLATION\_LEVEL=  
(if specified)
- PROMPT=  
must *not* be  
specified

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.

---

## Temporary Table Support for ODBC

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

When you want to use temporary tables that persist across SAS procedures and DATA steps with ODBC, you must use the CONNECTION=SHARED LIBNAME option. In doing so, the temporary table is available for processing until the libref is closed.

---

### 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

Using the Internat sample table, the following example creates a temporary table, #LONDON, with Microsoft SQL Server that contains information about flights that flew to London. This table is then joined with a larger SQL Server table that lists all the flights, March, but matched only on flights that flew to London.

```
libname samples odbc dsn=lupinss uid=dbitest pwd=dbigrpl connection=shared;
```

```

data samples. '#LONDON' n;
  set work.internat;
  where dest='LON';
run;

proc sql;
  select b.flight, b.dates, b.depart, b.orig
         from samples. '#LONDON' n a, samples.march b
         where a.dest=b.dest;
quit;

```

In the following example a temporary table called New is created with Microsoft SQL Server. The data from this table is then appended to an existing SQL Server table named Inventory.

```

libname samples odbc dsn=lupinss uid=dbitest pwd=dbigrpl connection=shared;

data samples.inventory(DBTYPE=(itemnum='char(5)' item='varchar(30)'
                              quantity='numeric'));
  itemnum='12001';
  item='screwdriver';
  quantity=15;
  output;
  itemnum='12002';
  item='hammer';
  quantity=25;
  output;
  itemnum='12003';
  item='sledge hammer';
  quantity=10;
  output;
  itemnum='12004';
  item='saw';
  quantity=50;
  output;
  itemnum='12005';
  item='shovel';
  quantity=120;
  output;
run;

data samples. '#new' n(DBTYPE=(itemnum='char(5)' item='varchar(30)'
                              quantity='numeric'));
  itemnum='12006';
  item='snow shovel';
  quantity=5;
  output;
  itemnum='12007';
  item='nails';
  quantity=500;
  output;
run;

proc append base=samples.inventory data=samples. '#new' n;
run;

```

```
proc print data=samples.inventory;
run;
```

The following example demonstrates the use of a temporary table using the Pass-Through Facility.

```
proc sql;
  connect to odbc as test (dsn=lupinss uid=dbitest
                        pwd=dbigrpl connection=shared);
  execute (create table #FRANCE (flight char(3), dates datetime,
                               dest char(3))) by test;

  execute (insert #FRANCE select flight, dates, dest from internat
          where dest like '%FRA%') by test;
  select * from connection to test (select * from #FRANCE);
quit;
```

---

## ODBC Bulk Loading

The LIBNAME option BULKLOAD= calls the Bulk Copy facility (BCP), which enables you to efficiently insert rows of data into a DBMS table as a unit. BCP= is an alias for this option.

*Note:* The Bulk Copy facility is available only when you are accessing Microsoft SQL Server data on Windows platforms. To use this facility, your installation of Microsoft SQL Server must include the ODBCBCP.DLL file. BULKCOPY= is not available on UNIX.  $\Delta$

As SAS/ACCESS sends rows of data to the Bulk Copy facility, the data is written to an input buffer. When you have sent all of the rows or when the buffer reaches a certain size (as determined by the DBCOMMIT= option), all of the rows are inserted as a unit into the table and the data is committed to the table. You can set the DBCOMMIT= option to commit rows after a specified number of rows are inserted.

If an error occurs, a message is written to the SAS log, and any rows that were inserted before the error are rolled back.

---

## Locking in the ODBC Interface

The following LIBNAME and data set options enable you to control how the interface to ODBC handles locking. See the section about the LIBNAME statement in *SAS/ACCESS for Relational Databases: Reference* for additional information about these options.

```
READ_LOCK_TYPE= ROW | TABLE | NOLOCK
```

```
UPDATE_LOCK_TYPE= ROW | TABLE | NOLOCK
```

```
READ_ISOLATION_LEVEL= S | RR | RC | RU | V
```

The ODBC driver manager supports the S, RR, RC, RU, and V isolation levels defined in the following table.

**Table 1.3** Isolation Levels for ODBC

Isolation Level	Definition
S (serializable)	Does not allow dirty reads, nonrepeatable reads, or phantom reads.
RR (repeatable read)	Does not allow dirty reads or nonrepeatable reads; does allow phantom reads.
RC (read committed)	Does not allow dirty reads or nonrepeatable reads; does allow phantom reads.
RU (read uncommitted)	Allows dirty reads, nonrepeatable reads and phantom reads.
V (versioning)	Does not allow dirty reads, nonrepeatable reads, or phantom reads. These transactions are serializable but higher concurrency is possible than with the serializable isolation level. Typically, a nonlocking protocol is used.

The terms in the table are defined as follows:

- *Dirty read* — A transaction that exhibits this phenomenon has very minimal isolation from concurrent transactions. In fact, it can see changes that are made by those concurrent transactions even before they commit.

For example, suppose that transaction T1 performs an update on a row, transaction T2 then retrieves that row, and transaction T1 then terminates with rollback. Transaction T2 has then seen a row that no longer exists.

- *Nonrepeatable read* — If a transaction exhibits this phenomenon, it is possible that it might read a row once and if it attempts to read that row again later in the course of the same transaction, the row might have been changed or even deleted by another concurrent transaction. Therefore, the read is not (necessarily) repeatable.

For example, suppose that transaction T1 retrieves a row, transaction T2 then updates that row, and transaction T1 then retrieves the same row again. Transaction T1 has now retrieved the same row twice but has seen two different values for it.

- *Phantom reads* — When a transaction exhibits this phenomenon, a set of rows that it reads once might be a different set of rows if the transaction attempts to read them again.

For example, suppose that transaction T1 retrieves the set of all rows that satisfy some condition. Suppose that transaction T2 then inserts a new row that satisfies that same condition. If transaction T1 now repeats its retrieval request, it sees a row that did not previously exist, a phantom.

UPDATE\_ISOLATION\_LEVEL= S | RR | RC | V

The ODBC driver manager supports the S, RR, RC, and V isolation levels defined in the preceding table.

---

## Naming Conventions for ODBC

Since ODBC is not a database but rather is an application programming interface, or API, table names and column names are determined at run time. Beginning in Version 7 of SAS, table or column names can be up to 32 characters long. The ODBC engine supports table and column names up to 32 characters long. If the DBMS column names

are longer than 32 characters, they are truncated to 32 characters. If truncating a column's name results in identical names, then SAS generates unique names by replacing the last character with a number. DBMS table names must be 32 characters or less, since SAS will *not* truncate a longer name. If you already have a table name greater than 32 characters, it is recommended that you create a table view.

The PRESERVE\_COL\_NAMES= and PRESERVE\_TAB\_NAMES= options determine how the interface to ODBC handles case sensitivity, spaces, and special characters. The default value for both options is YES for Microsoft Access, Microsoft Excel, and Microsoft SQL Server; and NO for all others.

The following example specifies SYBASE as the DBMS.

```
libname mydblib odbc user=TESTUSER password=testpass
      database=sybase;

data mydblib.a;
  x=1;
  y=2;
run;
```

SYBASE is generally case sensitive. Therefore, this example would produce a SYBASE table named **a** and columns named **x** and **y**.

If the DBMS being accessed was Oracle, which is not case sensitive, the example would produce an Oracle table named **A** and columns named **x** and **y**. The object names would be normalized to uppercase.

---

## Data Types for ODBC

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

The following table shows all of the data types and default SAS formats that are supported by the SAS/ACCESS interface to ODBC. This table does not explicitly define the data types as they exist for each DBMS. It lists the SQL types that each DBMS data type would map to. For example, a CHAR data type under DB2 would map to an ODBC data type of SQL\_CHAR. There are no unsupported data types.

**Table 1.4** ODBC Data Types and Default SAS Formats

ODBC Data Type	Default SAS Format
SQL_CHAR	$\$n$
SQL_VARCHAR	$\$n$
SQL_LONGVARCHAR	$\$n$
SQL_BINARY	$\$n.*$
SQL_VARBINARY	$\$n.*$
SQL_LONGVARBINARY	$\$n.*$
SQL_DECIMAL	$m$ or $m.n$ or none if $m$ and $n$ are not specified
SQL_NUMERIC	$m$ or $m.n$ or none if $m$ and $n$ are not specified
SQL_INTEGER	11.
SQL_SMALLINT	6.

ODBC Data Type	Default SAS Format
SQL_TINYINT	4.
SQL_BIT	1.
SQL_REAL	none
SQL_FLOAT	none
SQL_DOUBLE	none
SQL_BIGINT	20.
SQL_INTERVAL	\$n
SQL_GUID	\$n
SQL_TYPE_DATE	DATE9.
SQL_TYPE_TIME	TIME8. ODBC cannot support fractions of seconds for time values
SQL_TYPE_TIMESTAMP	DATETIME $m.n$ where $m$ and $n$ depend on precision

\* Because the ODBC driver does the conversion, this field is displayed as though the \$HEX $n$ . format were applied.

The following table shows the default data types that the SAS/ACCESS interface to ODBC uses when creating tables.

**Table 1.5** Default ODBC Output Data Types

SAS Variable Format	Default ODBC Data Type
$m.n$	SQL_DOUBLE or SQL_NUMERIC using $m.n$ if the DBMS allows it
\$ $n$ .	SQL_VARCHAR using $n$
datetime formats	SQL_TIMESTAMP
date formats	SQL_DATE
time formats	SQL_TIME

The interface to ODBC allows nondefault data types to be specified with the DBTYPE= data set option.

## ODBC Null Values

Many relational database management systems have a special value called NULL. A DBMS NULL value means an absence of information and is analogous to a SAS missing value. When SAS/ACCESS reads a DBMS NULL value, it interprets it as a SAS missing value.


In most relational databases, columns can be defined as NOT NULL so that they require data (they cannot contain NULL values). When a column is defined as NOT NULL, the DBMS will not add a row to the table unless the row has a value for that column. When creating a DBMS table with SAS/ACCESS, you can use the DBNULL= data set option to indicate whether NULL is a valid value for specified columns.

ODBC mirrors the behavior of the underlying DBMS with regard to NULL values. Refer to the documentation for your DBMS for information about how it handles NULL values.

For more information about how SAS handles NULL values, see “Potential Result Set Differences When Processing Null Data” in *SAS/ACCESS for Relational Databases: Reference*.

*Note:* To control how SAS missing character values are handled by the DBMS, use the NULLCHAR= and NULLCHARVAL= data set options.  $\triangle$





## APPENDIX

## 1

## Recommended Reading

---

*Recommended Reading* 29

---

### 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](mailto:sasbook@sas.com)

Web address: [support.sas.com/pubs](http://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

autopartitioning  
ODBC specifics 15

## B

BCP (Bulk Copy facility) 20, 24  
Bulk Copy facility (BCP) 20, 24  
BULKCOPY= option, PROC DBLOAD statement 20  
BULKLOAD= option, LIBNAME statement  
ODBC interface 24

## C

COMPLETE= option, LIBNAME statement  
ODBC interface 4  
conformance, ODBC 3  
CONNECT statement, SQL procedure  
options, ODBC 9  
CONNECTION TO component, SQL SELECT  
statement  
ODBC interface 12

## D

data set options  
ODBC interface 7  
data source, ODBC 2  
data types  
ODBC interface 26  
DATASRC= option, LIBNAME statement  
ODBC interface 4  
DBLOAD procedure  
ODBC specifics 19  
DBSLICE= option  
autopartitioning, ODBC 16  
DBSLICEPARAM= option, LIBNAME statement  
autopartitioning, ODBC 16  
dirty reads 25  
DSN= option  
LIBNAME statement 4  
PROC DBLOAD statement 19

## I

IBM AS/400, SQL views for 10

## J

joins  
passing to ODBC 22

## L

LIBNAME statement  
ODBC specifics 3, 7  
locking data, handling  
ODBC interface 24

## M

Microsoft Bulk Copy facility 20, 24  
Microsoft SQL Server, interface to  
multiple partitioning, configuring 17  
SQL views 12

## N

naming conventions  
ODBC interface 25  
nonrepeatable reads 25  
NOPROMPT= option, LIBNAME statement  
ODBC interface 4  
NULL values  
ODBC 27  
nullable columns  
autopartitioning, ODBC 15  
NULLS option, PROC DBLOAD statement  
ODBC interface 20

## O

ODBC, interface to 1  
autopartitioning scheme 15  
data set options 7  
data types 26  
DBLOAD procedure 19

LIBNAME statement 3, 7  
locking in 24  
naming conventions 25  
ODBC standards basics 2  
Pass-Through Facility 8  
passing joins to 22  
passing SAS functions to 21  
ODBC:: special queries 13

## P

partitioning  
partitioned views, Microsoft SQL Server 17  
Pass-Through Facility  
ODBC specifics 8  
PASSWORD= option, LIBNAME statement  
ODBC interface 4  
PASSWORD= option, PROC DBLOAD statement  
ODBC 19  
phantom reads 25  
PROMPT= option, LIBNAME statement  
ODBC interface 5

## Q

queries, SQL  
ODBC special queries 13

## R

READ\_ISOLATION\_LEVEL= option  
ODBC interface 24  
READ\_LOCK\_TYPE= option  
ODBC interface 24  
REQUIRED= option, LIBNAME statement  
ODBC interface 5

## S

SAS/ACCESS data set options  
ODBC interface 7  
SAS SQL functions  
passing to ODBC 21  
SELECT statement (SQL)  
CONNECTION TO component, ODBC 12

SQL\_ data types  
    ODBC interface 26  
SQL SELECT statement  
    CONNECTION TO component, ODBC 12

## **T**

TABLE= option, DBLOAD procedure  
    ODBC interface 20

threaded reads  
    ODBC interface 15

## **U**

UPDATE\_ISOLATION\_LEVEL= option 25  
UPDATE\_LOCK\_TYPE= option  
    ODBC interface 24

USER= option  
    PROC DBLOAD statement 19  
USER= option, LIBNAME statement  
    ODBC interface 4

## **V**

VALIDVARNAME= option, SQL CONNECT  
    statement 10  
views, SQL  
    Pass-Through Facility, ODBC interface 10

# Your Turn

---

If you have comments or suggestions about *SAS/ACCESS® 9.1 Supplement for ODBC*, 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](mailto: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](mailto:suggest@sas.com)

