



SAS Publishing



SAS/ACCESS[®] 9.1

Supplement for SYBASE

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 SYBASE*. Cary, NC: SAS Institute Inc.

SAS/ACCESS® 9.1 Supplement for SYBASE

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

ISBN 1-59047-253-5

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

2nd printing, October 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 Sybase	1
Introduction to the SAS/ACCESS Interface to Sybase		2
LIBNAME Statement Specifics for Sybase		2
Data Set Options for Sybase		4
Pass-Through Facility Specifics for Sybase		5
Autopartitioning Scheme for Sybase		7
Temporary Table Support for Sybase		9
ACCESS Procedure Specifics for Sybase		10
DBLOAD Procedure Specifics for Sybase		11
Passing SAS Functions to Sybase		13
Passing Joins to Sybase		14
Reading Multiple Sybase Tables		14
Locking in the Sybase Interface		15
Naming Conventions for Sybase		16
Case-Sensitivity in Sybase		17
Data Types for Sybase		17
Inserting TEXT into Sybase from SAS		22
National Language Support for Sybase		23
Appendix 1	△ Recommended Reading	25
Recommended Reading		25
Glossary		27
Index		33



CHAPTER

1

SAS/ACCESS for Sybase

<i>Introduction to the SAS/ACCESS Interface to Sybase</i>	2
<i>LIBNAME Statement Specifics for Sybase</i>	2
Arguments	2
Sybase LIBNAME Statement Example	4
<i>Data Set Options for Sybase</i>	4
<i>Pass-Through Facility Specifics for Sybase</i>	5
Example	7
<i>Autopartitioning Scheme for Sybase</i>	7
Overview	7
Indexes	8
Partitioning Criteria	8
Data types	8
Examples	8
<i>Temporary Table Support for Sybase</i>	9
Establishing a Temporary Table	9
Terminating a Temporary Table	9
Example	9
<i>ACCESS Procedure Specifics for Sybase</i>	10
Example	10
<i>DBLOAD Procedure Specifics for Sybase</i>	11
Example	13
<i>Passing SAS Functions to Sybase</i>	13
<i>Passing Joins to Sybase</i>	14
<i>Reading Multiple Sybase Tables</i>	14
<i>Locking in the Sybase Interface</i>	15
Understanding Sybase Update Rules	16
<i>Naming Conventions for Sybase</i>	16
<i>Case-Sensitivity in Sybase</i>	17
<i>Data Types for Sybase</i>	17
Character Data	17
Numeric Data	18
Abstract Data	18
User-Defined Data Types	19
Sybase Null Values	19
LIBNAME Statement Data Conversions	20
ACCESS Procedure Data Conversions	21
DBLOAD Procedure Data Conversions	21
Data Returned as SAS Binary Data with Default Format \$HEX	22
Data Returned as SAS Character Data	22
<i>Inserting TEXT into Sybase from SAS</i>	22
<i>National Language Support for Sybase</i>	23

Introduction to the SAS/ACCESS Interface to Sybase

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

LIBNAME Statement Specifics for Sybase

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

```
LIBNAME libref sybase <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.

sybase

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

connection-options

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

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

specifies the Sybase user name (also called the login name) that you use to connect to your database. If the user name contains spaces or non-alphanumeric characters, you must enclose it in quotation marks.

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

specifies the password that is associated with your Sybase user name.

If you omit the password, a default password of NULL is used. If the password contains spaces or non-alphanumeric characters, you must enclose it in quotation marks.

PASSWORD= can also be specified with the SYBPW=, PASS=, and PW= aliases.

```
DATABASE=<'>database-name<'>
```

specifies the name of the Sybase database that contains the tables and views that you want to access.

If the database name contains spaces or non-alphanumeric characters, you must enclose it in quotation marks. If you omit DATABASE=, the default database for your Sybase user name is used.

DATABASE= can also be specified with the DB= alias.

```
SERVER=<'>server-name<'>
```

specifies the server that you want to connect to. This server accesses the database that contains the tables and views that you want to access.

If the server name contains lowercase, spaces, or non-alphanumeric characters, you must enclose it in quotation marks.

If you omit `SERVER=`, the default action for your operating system occurs. On UNIX systems, the value of the environment variable `DSQUERY` is used if it has been set.

`IP_CURSOR= YES | NO`

specifies whether Implicit Proc SQL pass-through processes multiple result sets simultaneously.

`IP_CURSOR` is set to `NO` by default. Setting it to `YES` allows this type of extended processing, but will decrease performance, since cursors are being used, and not result sets. Do not set to `YES` unless needed.

If you specify the appropriate system options or environment variables for your database, you can often omit the connection options. See your Sybase 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 which `LIBNAME` options are supported for Sybase, 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 Sybase

Option	Default Value
<code>ACCESS=</code>	none
<code>AUTOCOMMIT=</code>	YES
<code>CONNECTION=</code>	SHAREDREAD
<code>CONNECTION_GROUP=</code>	none
<code>DBCOMMIT=</code>	1000 (inserting) or 0 (updating)
<code>DBCONINIT=</code>	none
<code>DBCONTERM=</code>	none
<code>DBCREATE_TABLE_OPTS=</code>	none
<code>DBGEN_NAME=</code>	DBMS
<code>DBINDEX=</code>	NO
<code>DBLIBINIT=</code>	none
<code>DBLIBTERM=</code>	none
<code>DBLINK=</code>	the local database
<code>DBMAX_TEXT=</code>	1024
<code>DBPROMPT=</code>	NO
<code>DBSASLABEL=</code>	COMPAT
<code>DBSLICEPARM=</code>	THREADED_APPS,2 or 3
<code>DEFER=</code>	NO
<code>DIRECT_EXE=</code>	none
<code>DIRECT_SQL=</code>	YES
<code>ENABLE_BULK=</code>	YES

Option	Default Value
INTERFACE=	none
MAX_CONNECTS=	25
MULTI_DATASRC_OPT=	NONE
PACKETSIZE=	server setting
QUOTED_IDENTIFIER=	NO
READBUFF=	100
READ_ISOLATION_LEVEL=	1
READ_LOCK_TYPE=	NOLOCK
REREAD_EXPOSURE=	NO
SCHEMA=	none
SPOOL=	YES
UPDATE_ISOLATION_LEVEL=	1
UPDATE_LOCK_TYPE=	PAGE
UTILCONN_TRANSIENT=	NO

Sybase LIBNAME Statement Example

In the following example, the libref MYDBLIB uses the Sybase engine to connect to a Sybase database. USER= and PASSWORD= are connection options.

```
libname mydblib sybase user=testuser password=testpass;
```

If you specify the appropriate system options or environment variables for your database, you can often omit the connection options. See your Sybase documentation for details.

Data Set Options for Sybase

The following table describes the data set options that are supported for Sybase, 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 for Sybase

Option	Default Value
AUTOCOMMIT=	LIBNAME option setting
BULK_BUFFER=	100
BULKLOAD=	NO
DBCOMMIT=	LIBNAME setting

Option	Default Value
DBCONDITION=	none
DBCREATE_TABLE_OPTS=	LIBNAME setting
DBFORCE=	NO
DBGEN_NAME=	LIBNAME option setting
DBINDEX=	LIBNAME option setting
DBKEY=	none
DBLABEL=	NO
DBLINK=	LIBNAME option setting
DBMASTER=	none
DBMAX_TEXT=	LIBNAME option setting
DBNULL=	_ALL_YES
DBPROMPT=	LIBNAME option setting
DBSASLABEL=	COMPAT
DBSLICE=	none
DBSLICEPARM=	THREADED_APPS,2 or 3
DBTYPE=	see “Data Types for Sybase” on page 17
ERRLIMIT=	1
NULLCHAR=	SAS
NULLCHARVAL=	a blank character
READBUFF=	LIBNAME option setting
READ_ISOLATION_LEVEL=	LIBNAME option setting
READ_LOCK_TYPE=	LIBNAME option setting
SASDATEFMT=	DATETIME22.3
SCHEMA=	LIBNAME option setting
SEGMENT_NAME=	none
UPDATE_ISOLATION_LEVEL=	LIBNAME option setting
UPDATE_LOCK_TYPE=	LIBNAME option setting

Pass-Through Facility Specifics for Sybase

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 Sybase are as follows:

- The *dbms-name* is **SYBASE**.
- The CONNECT statement is optional. If you omit the CONNECT statement, an implicit connection is made using the default values for all of the connection options.
- The interface can connect multiple times to one or more servers.

- The *database-connection-arguments* for the CONNECT statement are as follows:

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

specifies the Sybase user name (also called the login name) that you use to connect to your database. If the user name contains spaces or non-alphanumeric characters, you must enclose it in quotation marks.

PASSWORD=<'>SYBASE-password<'>

specifies the password that is associated with the Sybase user name.

If you omit the password, a default password of NULL is used. If the password contains spaces or non-alphanumeric characters, you must enclose it in quotation marks.

PASSWORD= can also be specified with the SYBPW=, PASS=, and PW= aliases.

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

DATABASE=<'>database-name<'>

specifies the name of the Sybase database that contains the tables and views that you want to access.

If the database name contains spaces or non-alphanumeric characters, you must enclose it in quotation marks. If you omit DATABASE=, the default database for your Sybase user name is used.

DATABASE= can also be specified with the DB= alias.

SERVER=<'>server-name<'>

specifies the server you want to connect to. This server accesses the database that contains the tables and views that you want to access.

If the server name contains lowercase, spaces, or non-alphanumeric characters, you must enclose it in quotation marks.

If you omit SERVER=, the default action for your operating system occurs. On UNIX systems, the value of the environment variable DSQUERY is used if it has been set.

INTERFACE=*filename*

specifies the name and location of the Sybase interfaces file. The interfaces file contains the names and network addresses of all of the available servers on the network.

If you omit this statement, the default action for your operating system occurs. INTERFACE= is not used in some operating environments. Contact your database administrator to determine whether it applies to your operating environment.

SYBBUFSZ=*number-of-rows*

specifies the number of rows of DBMS data to write to the buffer. If this statement is used, the SAS/ACCESS interface view engine creates a buffer that is large enough to hold the specified number of rows. This buffer is created when the associated database table is read. The interface view engine uses SYBBUFSZ= to improve performance.

If you omit this statement, no data is written to the buffer.

Note: Connection options for Sybase are all case-sensitive. They are passed to Sybase exactly as you type them. △

- The following LIBNAME options are available with the CONNECT statement:

```
DBMAX_TEXT=
MAX_CONNECTS=
READBUFF=
PACKETSIZE=
```

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

Example

The following example retrieves a subset of rows from the Sybase INVOICE table. Because the WHERE clause is specified in the DBMS query (the inner SELECT statement), the DBMS processes the WHERE expression and returns a subset of rows to SAS.

```
proc sql;
connect to sybase(server=SERVER1
                  database=INVENTORY
                  user=testuser password=testpass);
%put &sqlxmsg;

select * from connection to sybase
      (select * from INVOICE where BILLEDY=457232);
%put &sqlxmsg;
```

Note: The SELECT statement that is enclosed in parentheses is sent directly to the database and therefore must be specified using valid database variable names and syntax. Δ

Autopartitioning Scheme for Sybase

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

Overview

Sybase autopartitioning utilizes the Sybase MOD function (%) to create multiple SELECT statements with WHERE clauses, which, in the optimum scenario, divide the result set into equal chunks; one chunk per thread. For example, assume that your original SQL statement was **SELECT * FROM DBTAB**, and assume that DBTAB has a primary key column PKCOL of type integer and that you want it partitioned into three threads. The autopartitioning scheme would break up the table into three SQL statements as follows:

```
select * from DBTAB where (abs(PKCOL))%3=0
select * from DBTAB where (abs(PKCOL))%3=1
select * from DBTAB where (abs(PKCOL))%3=2
```

Since PKCOL is a primary key column, you should get a fairly even distribution among the three partitions, which is the primary goal.

Indexes

An index on a SAS partitioning column increases performance of the threaded read. If a primary key is not defined for the table, an index should be placed on the partitioning column in order to attain similar benefits. Understanding and following *Sybase ASE Performance and Tuning Guide* documentation recommendations with respect to index creation and usage is essential in order to achieve optimum database performance. The order of column selection for the partitioning column is as follows:

- 1 Identity column
- 2 Primary key column (integer or numeric)
- 3 integer, numeric, or bit; not nullable
- 4 integer, numeric, or bit; nullable

Note: If the column selected is a bit type, there will be only two partitions created since the only values are 0 and 1. Δ

Partitioning Criteria

The most efficient partitioning column is an Identity column, which is usually identified as a primary key column. Identity columns usually lead to evenly partitioned result sets because of the sequential values they store.

The least efficient partitioning column is a numeric, decimal, or float column that is NULLABLE, and does not have an index defined.

Given equivalent selection criteria, columns defined at the beginning of the table definition that meet the selection criteria will take precedence over columns defined toward the end of the table definition.

Data types

The following data types are supported in partitioning column selection:

integer
tinyint
smallint
numeric
decimal
float
bit.

Examples

The following are examples of generated SELECT statements involving various column data types:

COL1 is numeric, decimal, or float. This example uses three threads (the default) and COL1 is NOT NULL.

```
select * from DBTAB where (abs(convert(INTEGER, COL1)))%3=0
select * from DBTAB where (abs(convert(INTEGER, COL1)))%3=1
select * from DBTAB where (abs(convert(INTEGER, COL1)))%3=2
```

COL1 is bit, integer, smallint, or tinyint. This example uses two threads (the default) and COL1 is NOT NULL.

```
select * from DBTAB where (abs(COL1))%3=0
select * from DBTAB where (abs(COL1))%3=1
```

COL1 is and integer and is nullable.

```
select * from DBTAB where (abs(COL1))%3=0 OR COL1 IS NULL
select * from DBTAB where (abs(COL1))%3=1
```

Temporary Table Support for Sybase

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 specify CONNECTION=GLOBAL, you can reference a temporary table throughout a SAS session, in both DATA steps and procedures. The name of the table MUST start with the character '#'. To reference it, use the SAS convention of an *n* literal, as in mylib.'#foo'n.

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.

Example

The following example demonstrates how to use temporary tables:

```
/* clear any connection */
libname x clear;

libname x sybase user=test pass=test connection=global;

/* create the temp table. You can even use bulk copy */
/* Notice how the name is specified: '#mytemp'n */
data x.'#mytemp'n (bulk=yes);
  x=55;
  output;
  x=44;
  output;
run;

/* print it */
proc print data=x.'#mytemp'n;
run ;

/* The same temp table persists in PROC SQL, */
```

```

/* with the global connection specified */
proc sql;
    connect to sybase (user=austin pass=austin connection=global);
    select * from connection to sybase (select * from #mytemp);
quit;

/* use the temp table again in a procedure */
proc means data=x.'#mytemp'n;
run;

/* drop the connection, the temp table is automatically dropped */
libname x clear;

/* to convince yourself it's gone, try to access it */
libname x sybase user=austin password=austin connection=global;

/* it's not there */
proc print data=x.'#mytemp'n;
run;

```

ACCESS Procedure Specifics for Sybase

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

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

- The DBMS= value for PROC ACCESS is **SYBASE**.
- The *database-description-statements* used by PROC ACCESS are identical to the Pass-Through Facility's CONNECT statement database-connection-arguments on page 6.
- The TABLE= statement for PROC ACCESS is:

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

specifies the name of the Sybase table or Sybase view on which the access descriptor is based.

Example

The following example creates access descriptors and view descriptors for the EMPLOYEES and INVOICE tables. These tables have different owners and are stored in PERSONNEL and INVENTORY databases that reside on different machines. The USER= and PASSWORD= statements identify the owners of the Sybase tables and their passwords.

```

libname vlib 'sas-data-library';

proc access dbms=sybase;
    create work.employee.access;
        server='server1';
        database='personnel';
        user='testuser1';

```

```

        password='testpass1';
        table=EMPLOYEES;
    create vlib.emp_acc.view;
        select all;
        format empid 6.;
        subset where DEPT like 'ACC%';
run;

proc access dbms=sybase;
    create work.invoice.access;
        server='server2';
        database='inventory';
        user='testuser2';
        password='testpass2';
        table=INVOICE;
        rename invoicenum=invnum;
        format invoicenum 6. billedon date9.
            paidon date9.;
    create vlib.sainv.view;
        select all;
        subset where COUNTRY in ('Argentina','Brazil');
run;

options linesize=120;
title 'South American Invoices and
      Who Submitted Them';

proc sql;
    select invnum, country, billedon, paidon,
           billedby, lastname, firstnam
    from vlib.emp_acc, vlib.sainv
    where emp_acc.empid=sainv.billedby;

```

Sybase is a case-sensitive database. The PROC ACCESS database identification statements and the Sybase column names in all of the statements except SUBSET are converted to uppercase unless the names are enclosed in quotation marks. The SUBSET statements are passed to Sybase exactly as you type them, so you must use the correct case for the Sybase column names.

DBLOAD Procedure Specifics for Sybase

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

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

- The DBMS= value for PROC DBLOAD is **SYBASE**.
- The TABLE= statement for PROC DBLOAD is:

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

- PROC DBLOAD uses the following *database-description-statements*:

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

specifies the Sybase user name (also called the login name) that you use to connect to your database. If the user name contains spaces or non-alphanumeric characters, you must enclose it in quotation marks.

PASSWORD=<'>SYBASE-password<'>

specifies the password that is associated with the Sybase user name.

If you omit the password, a default password of NULL is used. If the password contains spaces or non-alphanumeric characters, you must enclose it in quotation marks.

PASSWORD= can also be specified with the SYBPW=, PASS=, and PW= aliases.

DATABASE=<'>database-name<'>

specifies the name of the Sybase database that contains the tables and views that you want to access.

If the database name contains spaces or non-alphanumeric characters, you must enclose it in quotation marks. If you omit DATABASE=, the default database for your Sybase user name is used.

DATABASE= can also be specified with the DB= alias.

SERVER=<'>server-name<'>

specifies the server that you want to connect to. This server accesses the database that contains the tables and views that you want to access.

If the server name contains lowercase, spaces, or non-alphanumeric characters, you must enclose it in quotation marks.

If you omit SERVER=, the default action for your operating system occurs. On UNIX systems, the value of the environment variable DSQUERY is used if it has been set.

INTERFACE=filename

specifies the name and location of the Sybase interfaces file. The interfaces file contains the names and network addresses of all of the available servers on the network.

If you omit this statement, the default action for your operating system occurs. INTERFACE= is not used in some operating environments. Contact your database administrator to determine whether it applies to your operating environment.

BULKCOPY= Y|N;

uses the Sybase bulk copy utility to insert rows into a Sybase table. The default value is N.

If you specify BULKCOPY=Y, BULKCOPY= calls the Sybase bulk copy utility in order to load data into a Sybase table. This utility groups rows so that they are inserted as a unit into the new table. Using the bulk copy utility can improve performance.

You use the COMMIT= statement to specify the number of rows in each group (this argument must be a positive integer). After each group of rows is inserted, the rows are permanently saved in the table. While each group is being inserted, if one row in the group is rejected, then all of the rows in that group are rejected.

If you specify BULKCOPY=N, rows are inserted into the new table using Transact-SQL INSERT statements. Refer to your Sybase documentation for more information about the bulk copy utility.

Example

The following example creates a new Sybase table, EXCHANGE, from the DLIB.RATEOFEX data file. An access descriptor ADLIB.EXCHANGE is also created, and it is based on the new table. The DBLOAD procedure sends a Transact-SQL GRANT statement to Sybase. You must be granted Sybase privileges to create new Sybase tables or to grant privileges to other users.

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

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

proc dbload dbms=sybase data=dlib.rateofex;
  server='server1';
  database='testdb';
  user='testuser';
  password='testpass';
  table=EXCHANGE;
  accdesc=adlib.exchange;
  rename fgnindol=fgnindolar 4=dolrsinfgn;
  nulls updated=n fgnindol=n 4=n country=n;
  load;
run;
```

Passing SAS Functions to Sybase

The interface to Sybase passes the following SAS functions to Sybase for processing. See the section about optimizing SQL usage in *SAS/ACCESS for Relational Databases: Reference* for information.

ABS
ARCOS
ARSIN
ATAN
AVG
CEIL
COS
DATETIME
EXP
FLOOR
LOG
MAX
MIN
SIGN

SIN
 SQRT
 TAN
 SUM
 COUNT

Passing Joins to Sybase

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

user ID
 password
 database
 server.

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.

Reading Multiple Sybase Tables

SAS opens multiple Sybase tables for simultaneous reading in the following situations:

- When you are using PROC COMPARE. For example:

```
proc compare base=syb.data1 compare=syb.data2;
```

- When you are running an SCL program that reads from more than one Sybase table simultaneously.
- When you are joining Sybase tables in SAS (when implicit pass through is not used (DIRECT_SQL=NO)). For example:

```
proc sql ;
  select * from syb.table1, syb.table2 where table1.x=table2.x;
```

or

```
proc sql;
  select * from syb.table1 where table1.x = (select x from syb.table2
  where y = 33);
```

or

```
proc sql;
  select empname from syb.employee where empyears > all (select empyears
  from syb.employee where emptitle = 'salesrep');
```

or

```
proc sql ;
  create view myview as
```

```

select * from employee where empyears > all (select empyears from
syb.employee where emptitle = 'salesrep');
proc print data=myview ;

```

In order to read two or more Sybase tables simultaneously, you must specify either the LIBNAME option CONNECTION=UNIQUE or the LIBNAME option READLOCK_TYPE=PAGE. Because READLOCK_TYPE=PAGE can degrade performance, it is generally recommended that you use CONNECTION=UNIQUE (unless there is a concern about the number of connections that are opened on the database).

Locking in the Sybase Interface

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

READ_LOCK_TYPE= PAGE | NOLOCK

The default value for Sybase is NOLOCK.

UPDATE_LOCK_TYPE= PAGE | NOLOCK

PAGE

SAS/ACCESS uses a cursor that can be updated. When you use this setting, it is recommended that the table have a defined primary key. PAGE is the default value for Sybase.

NOLOCK

SAS/ACCESS uses Sybase browse mode updating, in which the table that is being updated must have a primary key and timestamp.

READ_ISOLATION_LEVEL= 1 | 2 | 3

For reads, Sybase supports isolation levels 1, 2, and 3, as defined in the following table. Refer to your Sybase documentation for more information.

Table 1.3 Isolation Levels for Sybase

Isolation Level	Definition
1	Prevents dirty reads. This is the default transaction isolation level.
2	Uses serialized reads.
3	Also uses serialized reads.

UPDATE_ISOLATION_LEVEL = 1 | 3

Sybase uses a shared or update lock on base table pages that contain rows representing a current cursor position. This option applies to updates only when UPDATE_LOCK_TYPE=PAGE because cursor updating is in effect. It does not apply when UPDATE_LOCK_TYPE=NOLOCK.

For updates, Sybase supports isolation levels 1 and 3, as defined in the preceding table. Refer to your Sybase documentation for more information.

Understanding Sybase Update Rules

To avoid data integrity problems when updating and deleting data in Sybase tables, take the following precautionary measures:

- Always define a primary key.
- If the updates are not taking place via cursor processing, define a timestamp column as well.

It is not always obvious whether or not updates are utilizing cursor processing. Cursor processing is *never* used for LIBNAME statement updates if UPDATE_LOCK_TYPE=NOLOCK. Cursor processing is *always* used in the following situations:

- Updates using the LIBNAME statement with UPDATE_LOCK_TYPE=PAGE. *Note that this is the default setting for this option.*
- Updates using PROC SQL views.
- Updates using PROC ACCESS view descriptors.

Naming Conventions for Sybase

Sybase database objects that can be named include tables, views, columns, indexes, and database procedures. Use the following Sybase naming conventions:

- Database names must be unique. For each owner within a database, names of database objects must be unique. Column names and index names must be unique within a table.
- A name must be from 1 to 30 characters long (or 28 characters if quoted).
- A name must start with an alphabetic character or an underscore (_), unless the name is enclosed in quotation marks.
- After the first character, a name may contain the letters A through Z (in uppercase or lowercase), the digits 0 through 9, the underscore (_), the dollar sign (\$), the pound sign (#), the at sign (@), the yen sign (¥), and the monetary pound sign (£).
- Embedded spaces are not permitted unless the name is enclosed in quotation marks.
- A name cannot be a Sybase reserved word unless the name is enclosed in quotation marks. See your Sybase documentation for more information about reserved words.
- Embedded quotation marks are not permitted.
- Case-sensitivity is set when a server is installed. By default, the names of database objects are case-sensitive. On a case-sensitive server, the names **CUSTOMER** and **customer** are different.

Note: By default, column and table names are not quoted in the SAS/ACCESS interface to Sybase. To quote the table and column names, you must use the LIBNAME statement QUOTED_IDENTIFIER= option when you assign a libref. \triangle

When you use the DATASETS procedure to list your Sybase tables, the table names appear exactly as they exist in the Sybase data dictionary. If you specified the LIBNAME option SCHEMA=, SAS/ACCESS lists the tables for the specified schema user name.

To reference a table or other named object that you own, or for the specified schema, refer to the table name (for example, CUSTOMERS). If you use the LIBNAME statement DBLINK= option, all references to the libref refer to the specified database.

Case-Sensitivity in Sybase

SAS names can be entered in either uppercase or lowercase. When you reference Sybase objects through the SAS/ACCESS interface, objects are case-sensitive and require no quotation marks.

However, Sybase is generally set for case-sensitivity, and special consideration should be given to the names of objects (such as tables and columns) when they are to be used in SAS by the ACCESS or DBLOAD procedures. The ACCESS procedure converts Sybase object names to uppercase unless they are enclosed in quotation marks. Any Sybase objects that were given lowercase names, or whose names contain national or special characters, must be enclosed in quotation marks. The only exceptions are the SUBSET statement in the ACCESS procedure and the SQL statement in the DBLOAD procedure. Arguments or values from these statements are passed to Sybase exactly as you type them, with the case preserved.

In the Pass-Through Facility, all Sybase object names are case-sensitive. The names are passed to Sybase exactly as they are typed.

For more information about case-sensitivity and Sybase names, see “Naming Conventions for Sybase” on page 16.

Data Types for Sybase

Every column in a table has a name and a data type. The data type indicates to the DBMS how much physical storage to reserve for the column and the format in which the data is stored.

Note: SAS/ACCESS does not support the following Sybase data types: BINARY, VARBINARY, IMAGE, NCHAR(n), and NVARCHAR(n). SAS/ACCESS provides an error message when it attempts to read a table that has at least one column that uses an unsupported data type. Δ

Character Data

You must enclose all character data in single or double quotation marks.

CHAR(n)

CHAR(n) is a character string that can have 1 to 255 letters, symbols, and numbers. You specify the maximum length of the string with n . Storage size is also n , regardless of the actual entry length.

VARCHAR(n)

VARCHAR(n) is a varying-length character string that can have 1 to 255 letters, symbols, and numbers. You specify the maximum length of the string with n . Storage size is the actual entry length.

TEXT

TEXT stores character data of variable length up to two gigabytes. SAS supports the TEXT data type provided in Sybase. However, SAS only allows a maximum of 32,767 bytes of character data.

Numeric Data

NUMERIC(*p,s*), DECIMAL(*p,s*)

Exact numeric values have specified degrees of precision (*p*) and scale (*s*).

NUMERIC data can have a precision of 1 to 38 and scale of 0 to 38, where the value of *s* must be less or equal to than the value of *p*. The DECIMAL data type is identical to the NUMERIC data type. The default precision and scale are (18,0) for the DECIMAL data type.

REAL, FLOAT

Floating-point values consist of an integer part, a decimal point, and a fraction part, or scientific notation. The exact format for REAL and FLOAT data depends on the number of significant digits and the precision that your machine supports. You can use all arithmetic operations and aggregate functions with REAL and FLOAT except modulus. The REAL (4 byte) range is approximately 3.4E-38 to 3.4E+38, with 7-digit precision. The FLOAT (8 byte) range is approximately 1.7E-308 to 1.7E+308, with 15-digit precision.

TINYINT, SMALLINT, INT

Integers contain no fractional part. The three integer data types are TINYINT (1 byte), which has a range of 0 to 255; SMALLINT (2 bytes), which has a range of -32,768 to +32,767; and INT (4 bytes), which has a range of -2,147,483,648 to +2,147,483,647.

BIT

BIT data has a storage size of one bit and holds either a 0 or a 1. Other integer values are accepted but are interpreted as 1. BIT data cannot be NULL and cannot have indexes defined on it.

Abstract Data

Sybase date and money data types are abstract data types and are described in this section. Refer to your documentation on Transact-SQL for more information about abstract data types.

SMALLDATETIME

SMALLDATETIME data is 4 bytes long and consists of one small integer that represents the number of days after January 1, 1900, and one small integer that represents the number of minutes past midnight. The date range is from January 1, 1900, to December 31, 2079.

DATETIME

DATETIME data has two 4-byte integers. The first integer represents the number of days after January 1, 1900, and the second integer represents the number of milliseconds past midnight. Values can range from January 1, 1753 to December 31, 9999.

DATETIME values are input as quoted character strings in various alphabetic or numeric formats. Time data must be entered in the prescribed order (hours; minutes; seconds; milliseconds; AM, am, PM, pm) and must include either a colon or an AM/PM designator. Case is ignored, and spaces can be inserted anywhere within the value.

When you input DATETIME values, the national language setting determines how the date values are interpreted. You can change the default date order with the SET DATEFORMAT statement. See your Transact-SQL documentation for more information.

You can use Sybase built-in date functions to perform some arithmetic calculations on DATETIME values.

TIMESTAMP

TIMESTAMP data is used by SAS in UPDATE mode. If you select a column that contains TIMESTAMP data for input into SAS, the values are displayed in hex format.

SMALLMONEY

SMALLMONEY data is 4 bytes long and can range from -214,748.3648 to 214,748.3647. When displayed, it is rounded up to two places.

MONEY

MONEY data is 8 bytes long and can range from -922,337,203,685,477.5808 to 922,337,203,685,477.5807. When input, a dollar sign (\$) must appear before the MONEY value. For negative values, the minus sign must follow the dollar sign. Commas are not allowed.

MONEY values are accurate to a ten-thousandth of a monetary unit. However, when they are displayed, the dollar sign is omitted and MONEY values are rounded up to two places. A comma is inserted after every three digits.

You can store values for currencies other than U.S. dollars, but no form of conversion is provided.

User-Defined Data Types

You can supplement the Sybase system data types by defining your own data types with the Sybase system procedure `sp_addtype`. When you define your own data type for a column, you can specify a default value (other than NULL) for the column and define a range of allowable values for the column.

Sybase Null Values

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

By default, Sybase columns are defined as NOT NULL. NOT NULL tells Sybase not to add a row to the table unless the row has a value for the specified column.

If you want a column to accept NULL values, you must explicitly define it as NULL. Here is an example of a CREATE TABLE statement that defines all of the columns for a table to be NULL except for CUSTOMER. In this case, Sybase only accepts a row if it contains a value for CUSTOMER.

```
create table CUSTOMERS
  (CUSTOMER      char(8)      not null,
   STATE        char(2)      null,
   ZIPCODE      char(5)      null,
   COUNTRY      char(20)     null,
   TELEPHONE    char(12)     null,
   NAME         char(60)     null,
   CONTACT      char(30)     null,
   STREETADDRESS char(40)    null,
   CITY         char(25)     null,
   FIRSTORDERDATE datetime   null);
```

When creating a Sybase 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 Sybase, 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 Sybase data types during input operations when you use the LIBNAME statement.

Table 1.4 LIBNAME Statement: Default SAS Formats for Sybase Server Data Types

Sybase Column Type	SAS Data Type	Default SAS Format
CHAR(<i>n</i>)	character	\$ <i>n</i> . (<i>n</i> <= 255) \$255. (<i>n</i> > 255)
VARCHAR(<i>n</i>)	character	\$ <i>n</i> . (<i>n</i> <= 255) \$255. (<i>n</i> > 255)
TEXT	character	\$ <i>n</i> . (where <i>n</i> is the value of the DBMAX_TEXT= option)
BIT	numeric	1.0
TINYINT	numeric	4.0
SMALLINT	numeric	6.0
INT	numeric	11.0
NUMERIC	numeric	<i>w</i> , <i>w.d</i> (if possible)
DECIMAL	numeric	<i>w</i> , <i>w.d</i> (if possible)
FLOAT	numeric	
REAL	numeric	
SMALLMONEY	numeric	DOLLAR12.2
MONEY	numeric	DOLLAR24.2
SMALLDATETIME	numeric	DATETIME22.3
DATETIME	numeric	DATETIME22.3
TIMESTAMP	hex	\$HEX <i>w</i>

The following table shows the default Sybase 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 Sybase Data Types for SAS Variable Formats

SAS Variable Format	Sybase Data Type
\$ <i>w</i> ., \$CHAR <i>w</i> ., \$VARYING <i>w</i> ., \$HEX <i>w</i> .	VARCHAR(<i>w</i>)
DOLLAR <i>w.d</i>	SMALLMONEY (where <i>w</i> < 6) MONEY (where <i>w</i> >= 6)

SAS Variable Format	Sybase Data Type
any datetime, date, or time format	DATETIME
any numeric with a SAS format name of $w.d$ (where $d > 0$ and $w > 10$) or w .	NUMERIC(p,s)
any numeric with a SAS format name of $w.d$ (where $d = 0$ and $w < 10$)	TINYINT (where $w < 3$) SMALLINT (where $w < 5$) INT (where $w < 10$)
any other numeric	FLOAT

You can override these default data types by using the DBTYPE= data set option.

ACCESS Procedure Data Conversions

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

Table 1.6 PROC ACCESS: Default SAS Formats for Sybase Server Data Types

Sybase Column Type	SAS Data Type	Default SAS Format
CHAR(n)	character	$\$n.$ ($n \leq 200$) $\$200.$ ($n > 200$)
VARCHAR(n)	character	$\$n.$ ($n \leq 200$) $\$200.$ ($n > 200$)
BIT	numeric	1.0
TINYINT	numeric	4.0
SMALLINT	numeric	6.0
INT	numeric	11.0
FLOAT	numeric	BEST22.
REAL	numeric	BEST11.
SMALLMONEY	numeric	DOLLAR12.2
MONEY	numeric	DOLLAR24.2
SMALLDATETIME	numeric	DATETIME21.2
DATETIME	numeric	DATETIME21.2

The ACCESS procedure also supports Sybase user-defined data types. The ACCESS procedure uses the Sybase data type on which a user-defined data type is based in order to assign a default SAS format for columns.

Note: The DECIMAL, NUMERIC, and TEXT data types are not supported in PROC ACCESS. The TIMESTAMP data type is not displayed in PROC ACCESS. Δ

DBLOAD Procedure Data Conversions

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

Table 1.7 PROC DBLOAD: Default Sybase Data Types for SAS Variable Formats

SAS Variable Format	Sybase Data Type
$\$w.$, $\$CHARw.$, $\$VARYINGw.$, $\$HEXw.$	CHAR(w)
$w.$	TINYINT
$w.$	SMALLINT
$w.$	INT
$w.$	FLOAT
$w.d$	FLOAT
IB $w.d$, PIB $w.d$	INT
FRACT, E format, and other numeric formats	FLOAT
DOLLAR $w.d$, $w \leq 12$	SMALLMONEY
DOLLAR $w.d$, $w > 12$	MONEY
any datetime, date, or time format	DATETIME

The DBLOAD procedure also supports Sybase user-defined data types. Use the TYPE= statement to specify a user-defined data type.

Data Returned as SAS Binary Data with Default Format \$HEX

BINARY
VARBINARY
IMAGE

Data Returned as SAS Character Data

NCHAR
NVARCHAR

Inserting TEXT into Sybase from SAS

TEXT data can only be inserted into a Sybase table by using the BULKLOAD= data set option, as in the following example:


```
data yourlib.newtable(bulkload=yes);
  set work.sasbigtext;
run;
```

If the BULKLOAD= option is not used, you receive the following error message:

```
ERROR: Object not found in database. Error Code: -2782
An untyped variable in the PREPARE statement 'S401bcf78'
is being resolved to a TEXT or IMAGE type.
This is illegal in a dynamic PREPARE statement.
```

National Language Support for Sybase

To support output and update processing from SAS into Sybase in languages other than English, special setup steps are required so that date, time, and datetime values can be processed correctly. In SAS, you must ensure that the DFLANG= system option is set to the correct language. This can be globally set by the system administrator or set by a user within a single SAS session. In Sybase, the default client language, set in the *locales.dat* file, must match the language that is used in SAS.



APPENDIX

1

Recommended Reading

Recommended Reading 25

Recommended Reading

Here is the recommended reading list for this title:

- *SAS/ACCESS Supplement for DB2 under z/OS (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for DB2 under UNIX and PC Hosts (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for Informix (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for Microsoft SQL Server (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for ODBC (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for OLE DB (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for Oracle (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for SYBASE (SAS/ACCESS for Relational Databases)*
- *SAS/ACCESS Supplement for Teradata*
- *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
 SYBASE data types 18
 ACCESS procedure
 SYBASE specifics 10, 21
 autopartitioning
 SYBASE specifics 7

B

BIT data type 18
 BULKCOPY= option, PROC DBLOAD statement 12
 BULKLOAD= data set option
 inserting TEXT into SYBASE from SAS 22

C

capitalization and case sensitivity
 SYBASE objects 17
 CHAR data type
 SYBASE 17
 character data
 SYBASE data types 17
 CONNECT statement, SQL procedure
 options, SYBASE 5

D

data set options
 SYBASE interface 4
 data types
 SYBASE interface 17
 DATABASE= option
 LIBNAME statement 2
 PROC DBLOAD statement 12
 SQL CONNECT statement 6
 date and time data
 SYBASE data types 18
 DATETIME data type 18
 DBLOAD procedure
 SYBASE specifics 11, 21
 DBMS= option
 PROC ACCESS statement 10
 PROC DBLOAD statement 11

F

FLOAT data type
 SYBASE 18

I

indexes
 SYBASE interface 8
 INT data type 18
 INTERFACE= option
 PROC DBLOAD statement 12
 SQL CONNECT statement 6
 internationalization 23

J

joins
 passing to SYBASE 14

L

language support, SYBASE 23
 LIBNAME statement
 SYBASE specifics 2, 20
 locking data, handling
 SYBASE interface 15

M

MONEY data type
 SYBASE 19

N

naming conventions
 SYBASE interface 16
 NULL values
 SYBASE 19
 numeric data
 SYBASE data types 18
 NUMERIC data type
 SYBASE 18

P

Pass-Through Facility
 SYBASE interface 5
 PASSWORD= option, LIBNAME statement
 SYBASE interface 2
 PASSWORD= option, PROC DBLOAD statement
 SYBASE 12
 PASSWORD= option, SQL CONNECT statement
 SYBASE interface 6

R

READ_ISOLATION_LEVEL= option
 SYBASE interface 15
 READ_LOCK_TYPE= option
 SYBASE interface 15
 REAL data type
 SYBASE 18

S

SAS/ACCESS data set options
 SYBASE interface 4
 SAS SQL functions
 passing to SYBASE 13
 SERVER= option
 LIBNAME statement 2
 PROC DBLOAD statement 12
 SQL CONNECT statement 6
 SMALLDATETIME data type 18
 SMALLINT data type
 SYBASE 18
 SMALLMONEY data type 19
 SYBASE, interface to 2
 ACCESS procedure 10, 21
 autopartitioning scheme 7
 data set options 4
 data types 17
 DBLOAD procedure 11, 21
 inserting TEXT data from SAS 22
 LIBNAME statement 2, 20
 locking in 15
 naming conventions 16
 national language support 23
 Pass-Through Facility 5
 passing joins to 14
 passing SAS functions to 13

reading multiple tables 14
SYBBUFSZ= option, SQL CONNECT statement 6

T

TABLE= statement, ACCESS procedure
SYBASE 10
TEXT data type
inserting into SYBASE from SAS 22
SYBASE 17

threaded reads
SYBASE interface 7
TIMESTAMP data type
SYBASE 19
TINYINT data type 18

U

UPDATE_ISOLATION_LEVEL= option
SYBASE interface 15

UPDATE_LOCK_TYPE= option
SYBASE interface 15
USER= option
PROC DBLOAD statement 12
SQL CONNECT statement 6
USER= option, LIBNAME statement
SYBASE interface 2

V

VARCHAR data type
SYBASE 17

Your Turn

If you have comments or suggestions about *SAS/ACCESS 9.1 Supplement for SYBASE*, 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

SAS Publishing gives you the tools to flourish in any environment with SAS®!

Whether you are new to the workforce or an experienced professional, you need a way to distinguish yourself in this rapidly changing and competitive job market. SAS Publishing provides you with a wide range of resources, from software to online training to publications to set yourself apart.

Build Your SAS Skills with SAS Learning Edition

SAS Learning Edition is your personal learning version of the world's leading business intelligence and analytic software. It provides a unique opportunity to gain hands-on experience and learn how SAS gives you the power to perform.

support.sas.com/LE

Personalize Your Training with SAS Self-Paced e-Learning

You are in complete control of your learning environment with SAS Self-Paced e-Learning! Gain immediate 24/7 access to SAS training directly from your desktop, using only a standard Web browser. If you do not have SAS installed, you can use SAS Learning Edition for all Base SAS e-learning.

support.sas.com/selfpaced

Expand Your Knowledge with Books from SAS Publishing

SAS Press offers user-friendly books for all skill levels, covering such topics as univariate and multivariate statistics, linear models, mixed models, fixed effects regression and more. View our complete catalog and get free access to the latest reference documentation by visiting us online.

support.sas.com/pubs



SAS Publishing

