



SAS Publishing



SAS/ACCESS[®] 9.1

Supplement for Informix

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 Informix (SAS/ACCESS for Relational Databases)*. Cary, NC: SAS Institute Inc.

SAS/ACCESS® 9.1 Supplement for Informix (SAS/ACCESS for Relational Databases)

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

ISBN 1-59047-247-0

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 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 Informix	1
Introduction to the SAS/ACCESS Interface to Informix		1
LIBNAME Statement Specifics for Informix		2
Data Set Options for Informix		4
Pass-Through Facility Specifics for Informix		5
Autopartitioning Scheme for Informix		8
Passing SAS Functions to Informix		9
Passing Joins to Informix		10
Temporary Table Support for Informix		10
Locking in the Informix Interface		11
Naming Conventions for Informix		12
Informix Data Types		12
Overview of Informix Servers		15
Appendix 1	△ Recommended Reading	17
Recommended Reading		17
Glossary		19
Index		25



CHAPTER

1

SAS/ACCESS for Informix

<i>Introduction to the SAS/ACCESS Interface to Informix</i>	1
<i>Default Environment</i>	2
<i>LIBNAME Statement Specifics for Informix</i>	2
<i>Arguments</i>	2
<i>Informix LIBNAME Statement Example</i>	4
<i>Data Set Options for Informix</i>	4
<i>Pass-Through Facility Specifics for Informix</i>	5
<i>Stored Procedures and the Pass-Through Facility</i>	5
<i>Command Restrictions for the Pass-Through Facility</i>	6
<i>Examples</i>	6
<i>Autopartitioning Scheme for Informix</i>	8
<i>Overview</i>	8
<i>Autopartitioning Restrictions</i>	8
<i>Using WHERE Clauses</i>	8
<i>Using DBSLICEPARM=</i>	9
<i>Using DBSLICE=</i>	9
<i>Passing SAS Functions to Informix</i>	9
<i>Passing Joins to Informix</i>	10
<i>Temporary Table Support for Informix</i>	10
<i>Establishing a Temporary Table</i>	10
<i>Terminating a Temporary Table</i>	10
<i>Example</i>	11
<i>Locking in the Informix Interface</i>	11
<i>Naming Conventions for Informix</i>	12
<i>Informix Data Types</i>	12
<i>Character Data</i>	13
<i>Numeric Data</i>	13
<i>Abstract Data</i>	13
<i>Informix Null Values</i>	14
<i>LIBNAME Statement Data Conversions</i>	14
<i>Pass-Through Facility Data Conversions</i>	15
<i>Overview of Informix Servers</i>	15
<i>Informix Database Servers</i>	15
<i>Using the DBDATASRC Environment Variables</i>	16
<i>Using Fully Qualified Table Names</i>	16

Introduction to the SAS/ACCESS Interface to Informix

This document includes details *only* about the SAS/ACCESS interface to Informix. It should be used as a supplement to the generic SAS/ACCESS documentation

SAS/ACCESS for Relational Databases: Reference. See “Overview of Informix Servers” on page 15 for background information about Informix.

Default Environment

When you access Informix tables by using the SAS/ACCESS interface to Informix, the default Informix read isolation level is set for committed reads, and SAS spooling is on. Committed reads enable you to read rows unless another user or process is updating the rows. Reading in this manner does not lock the rows. SAS spooling guarantees that you get identical data each time you re-read a row because SAS buffers the rows after you read them the first time. This default environment is suitable for most users. If this default environment is unsuitable for your needs, see “Locking in the Informix Interface” on page 11.

To see the SQL statements that SAS issues to the Informix server, include the `SASTRACE=` option in your code:

```
option sastrace=',,,'d';
```

If you use quotation marks in your Informix SQL statements, your `DELIMIDENT=` environment variable should be set to `DELIMIDENT=YES`, or your statements could be rejected by Informix. Because some of the SAS options that preserve case generate SQL statements that contain quotation marks, you should set `DELIMIDENT=YES` in your environment.

LIBNAME Statement Specifics for Informix

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

```
LIBNAME libref informix <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.

`informix`

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

connection-options

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

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

specifies the Informix user name that you use to connect to the database that contains the tables and views that you want to access. If you omit the `USER=` option, your operating environment account name is used, if applicable to your operating environment.

USING=<'>*Informix-password*<'>

specifies the password that is associated with the Informix user. If you omit the password, Informix uses the password in the `/etc/passwd` file.

USING= can also be specified with the PASSWORD= and PWD= aliases.

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

specifies the ODBC data source to which you want to connect. An error occurs if the SERVER= option is not set. For UNIX platforms, the data source must be configured by modifying the `.ODBC.ini` file. See your ODBC driver manual for details.

Note: For the SAS/ACCESS 9 interface to Informix, the Informix ODBC Driver API is used to connect to Informix, and the connection options have been changed accordingly. The DATABASE= option from SAS/ACCESS Version 8 has been removed. If you need to specify a database, set it in the `.ODBC.ini` file. For SERVER= options, instead of specifying the server name, as in Version 8, you specify an ODBC data source name. Optionally, a user ID and password can be used in conjunction with SERVER=. Δ

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 Informix, and presents default values where applicable. See the section about the LIBNAME statement in *SAS/ACCESS for Relational Databases: Reference* for detailed information about these options.

Table 1.1 SAS/ACCESS LIBNAME Options for Informix

Option	Default Value
ACCESS=	none
CONNECTION=	SHAREDREAD
CONNECTION_GROUP=	none
DBCOMMIT=	1000 (insert) or 0 (update)
DBCONINIT=	none
DBCONTERM=	none
DBCREATE_TABLE_OPTS=	none
DBGEN_NAME=	DBMS
DBINDEX=	NO
DBLIBINIT=	none
DBLIBTERM=	none
DBNULLKEYS=	NO
DBPROMPT=	NO
DBSASLABEL=	COMPAT
DBSLICEPARAM=	THREADED_APPS,2 or 3
DEFER=	NO
DIRECT_EXE=	none
DIRECT_SQL=	YES

Option	Default Value
LOCKTABLE=	no locking
LOCKTIME=	none
LOCKWAIT=	not set
MULTI_DATASRC_OPT=	NONE
PRESERVE_COL_NAMES=	NO
PRESERVE_TAB_NAMES=	NO
READ_ISOLATION_LEVEL=	COMMITTED READ (see “Locking in the Informix Interface” on page 11)
REREAD_EXPOSURE=	NO
SCHEMA=	your user name
SPOOL=	YES
UTILCONN_TRANSIENT=	NO

Informix LIBNAME Statement Example

In the following example, the libref MYDBLIB uses the Informix interface to connect to an Informix database:

```
libname mydblib informix user=testuser using=testpass server=testdsn;
```

In this example, USER=, USING=, and SERVER= are connection options.

Data Set Options for Informix

The following table describes the data set options that are supported for Informix, 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
DBCOMMIT=	LIBNAME option setting
DBCONDITION=	none
DBCREATE_TABLE_OPTS=	LIBNAME option setting
DBFORCE=	NO
DBGEN_NAME=	DBMS
DBINDEX=	LIBNAME option setting
DBKEY=	none
DBLABEL=	NO
DBMASTER=	none
DBNULL=	_ALL_=YES

Option	Default Value
DBNULLKEYS=	LIBNAME option setting
DBSASLABEL=	COMPAT
DBSASTYPE=	see “Informix Data Types” on page 12
DBSLICE	none
DBSLICEPARM	THREADED_APPS,2 or 3
DBTYPE=	see “Informix Data Types” on page 12
ERRLIMIT=	1
LOCKTABLE=	LIBNAME option setting
NULLCHAR=	SAS
NULLCHARVAL=	a blank character
PRESERVE_COL_NAMES=	LIBNAME option setting
SASDATEFMT=	DATE TIME
SCHEMA=	LIBNAME option setting

Pass-Through Facility Specifics for Informix

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 Informix are as follows:

- The *dbms-name* is **informix**.
- The CONNECT statement is optional when you are connecting to an Informix database if the DBDATASRC environment variable has been set. When you omit a CONNECT statement, an implicit connection is performed when the first EXECUTE statement or CONNECTION TO component is passed to the DBMS.
- You can connect to only one Informix database at a time. However, you can specify multiple CONNECT statements if they all connect to the same Informix database. If you use multiple connections, you must use an *alias* to identify the different connections. If you omit an alias, **informix** is automatically used.
- The CONNECT statement *database-connection-arguments* are identical to its connection-options .
- If you use quotation marks in your Informix Pass-Through statements, your DELIMIDENT= environment variable must be set to DELIMIDENT=YES, or your statements are rejected by Informix.

Stored Procedures and the Pass-Through Facility

The Pass-Through Facility recognizes two types of stored procedures in Informix that perform only database functions. The methods for executing the two types of stored procedures are different.

- Procedures that return no values to the calling application:

Stored procedures that do not return values can be executed directly by using the Informix SQL EXECUTE statement. Stored procedure execution is initiated with the Informix EXECUTE PROCEDURE statement. The following example executes the stored procedure `make_table`. The stored procedure has no input parameters and returns no values.

```
execute (execute procedure make_table())
      by informix;
```

- Procedures that return values to the calling application:

Stored procedures that return values must be executed by using the PROC SQL SELECT statement with a CONNECTION TO component. The following example executes the stored procedure `read_address`, which has one parameter, `"Putnum"`.

The values that are returned by `read_address` serve as the contents of a virtual table for the PROC SQL SELECT statement.

```
select * from connection to informix
      (execute procedure read_address ("Putnum"));
```

For example, when you try to execute a stored procedure that returns values from a PROC SQL EXECUTE statement, you get the following error message:

```
execute (execute procedure read_address
      ("Putnum")) by informix;

ERROR: Informix EXECUTE Error: Procedure
      (read_address) returns too many values.
```

Command Restrictions for the Pass-Through Facility

Informix SQL contains extensions to the ANSI-89 standards. Some of these extensions, such as LOAD FROM and UNLOAD TO, are restricted from use by any applications other than the Informix DB-Access product. Specifying these extensions in the PROC SQL EXECUTE statement generates this error:

```
-201
A syntax error has occurred
```

Examples

The following example connects to Informix by using data source `testdsn`:

```
proc sql;
  connect to informix
  (user=SCOTT password=TIGER server=testdsn);
```

Note: You can use the `DBDATASRC` environment variable to specify the data source. Δ

The following example grants UPDATE and INSERT authority to user `gomez` on the Informix ORDERS table. Because the CONNECT statement is omitted, an implicit connection is made that uses a default value of `informix` as the connection alias and default values for the `DATABASE=` and `SERVER=` arguments. Informix is a case-sensitive database; therefore, the database object ORDERS is in uppercase, as it was created.

```
proc sql;
  execute (grant update, insert on ORDERS to gomez) by informix;
quit;
```

The following example connects to Informix and drops (that is, removes) the table TempData from the database. The alias Temp5 that is specified in the CONNECT statement is used in the EXECUTE statement's BY clause.

```
proc sql;
  connect to informix as temp5
  (server=testdsn);
  execute (drop table tempdata) by temp5;
  disconnect from temp5;
quit;
```

The following example sends an SQL query, shown with highlighting, to the database for processing. The results from the SQL query serve as a virtual table for the PROC SQL FROM clause. In this example, DBCON is a connection alias.

```
proc sql;
connect to informix as dbcon
  (user=testuser using=testpass
  server=testdsn);

select *
  from connection to dbcon
  (select empid, lastname, firstname,
  hiredate, salary
  from employees
  where hiredate>='31JAN88');

disconnect from dbcon;
quit;
```

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

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

proc sql;
connect to informix as mycon
  (user=testuser using=testpass
  server=testdsn);

create view samples.hires88 as
select *
  from connection to mycon
  (select empid, lastname, firstname,
  hiredate, salary from employees
  where hiredate>='31JAN88');

disconnect from mycon;
quit;
```

The following example connects to Informix and executes the stored procedure **testproc**. The **select *** clause displays the results from the stored procedure.

```
proc sql;
  connect to informix as mydb
    (server=testdsn);
  select * from connection to mydb
    (execute procedure testproc('123456'));
  disconnect from mydb;
quit;
```

Autopartitioning Scheme for Informix

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 Informix 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 Informix places additional restrictions on which columns can be used for the partitioning column during the autopartitioning phase. Columns are partitioned as follows:

- INTEGER
- SMALLINT
- BIT
- TINYINT

DECIMALS with 0 scales columns may also be used as the partitioning column. Nullable columns are the least preferable.

Using WHERE Clauses

Autopartitioning does not select a column to be the partitioning column if it appears in a WHERE clause. For instance, the following DATA step cannot to use a threaded read to retrieve the data since all of the numeric columns in the table (see the table definition as described in “Using DBSLICE=” on page 9) 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 you use autopartitioning, and DBSLICEPARM= does not specify a maximum number of threads to use for the threaded read, the SAS/ACCESS interface to Informix defaults to three threads.

The following example demonstrates the use of DBSLICEPARM=, with the maximum number of threads set to five:

```
libname x informix user=dbitest using=dbigrpl server=odbc15;
proc print data=x.dept(dbsliceparm=(ALL,5));
run;
```

Using DBSLICE=

You might achieve the best possible performance when using threaded reads by specifying an Informix specific DBSLICE= data set option in your SAS operation. The following example demonstrates the use of DBSLICE=:

```
libname x informix user=dbitest using=dbigrpl server=odbc15;
data xottest;
set x.invoice(dbslice=("amt billed<10000000" "amt billed>=10000000"));
run;
```

Passing SAS Functions to Informix

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

ABS
ARCOS
ARSIN
ATAN
ATAN2
AVG
COS
DATE
EXP
INT
LOG
LOG10
MAX
MDY
MIN
SIN

SQRT
TAN
TODAY
YEAR
MONTH
DAY
SUM
COUNT

Passing Joins to Informix

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

user ID
password
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.

Temporary Table Support for Informix

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

To establish the DBMS connection to support the creation and use of temporary tables, issue a LIBNAME statement with the connection options CONNECTION_GROUP=*connection-group* and CONNECTION=GLOBAL. This LIBNAME statement is required even if you connect to the database using the Pass-Through Facility CONNECT statement, because it establishes a connection group.

For every new PROC SQL step or LIBNAME statement, you must reissue a CONNECT statement with the CONNECTION_GROUP= option set to the same value so that the connection can be reused.

Terminating a Temporary Table

To terminate a temporary table, disassociate the libref by issuing the following statement:

```
libname libref clear;
```

Example

In the following Pass-Through example, joins are pushed to Informix:

```
libname x informix user=tester using=xxxxx server=dsn_name
          connection=global connection_group=mygroup;

proc sql;
  connect to informix (user=tester using=xxxxx server=dsn_name
                    connection=global connection_group=mygroup);
  execute (select * from t1 where (id >100)
          into scratch scrl ) by informix;
  create table count2 as select * from connection to informix
    (select count(*) as totrec from scrl);
quit;

proc print data=count2;
run;

proc sql;
  connect to informix (user=tester using=xxxxx server=dsn_name
                    connection=global connection_group=mygroup);
  execute(select t2.fname, t2.lname, scrl.dept from t2, scrl where
    (t2.id = scrl.id) into scratch t3 ) by informix;
quit;

libname x clear; /* connection closed, temp table closed */
```

Locking in the Informix Interface

In most situations, SAS spooling, which is on by default with the Informix interface, provides the data consistency you need.

The `READ_ISOLATION_LEVEL= LIBNAME` option enables you to control how the interface to Informix handles locks. This option can take the following values:

COMMITTED_READ

retrieves only committed rows. No locks are acquired, and rows can be locked exclusively for update by other users or processes. This is the default setting.

REPEATABLE_READ

gives you a shared lock on every row that is selected during the transaction. Other users or processes can also acquire a shared lock, but no other process can modify any row that is selected by your transaction. If you repeat the query during the transaction, you re-read the same information. The shared locks are released only when the transaction commits or rolls back. Another process cannot update or delete a row that is accessed by using a repeatable read.

DIRTY_READ

retrieves committed and uncommitted rows that might include phantom rows, which are rows that are created or modified by another user or process that might subsequently be rolled back. This type of read is most appropriate for tables that are not frequently updated.

CURSOR_STABILITY

gives you a shared lock on the selected row. Another user or process can acquire a shared lock on the same row, but no process can acquire an exclusive lock to modify data in the row. When you retrieve another row or close the cursor, the shared lock is released.

If you set `READ_ISOLATION_LEVEL=` to `REPEATABLE_READ` or `CURSOR_STABILITY`, it is recommended that you assign a separate libref and that you clear that libref when you have finished working with the tables. This technique minimizes the negative performance impact on other users that occurs when you lock the tables. To clear the libref, include the following code:

```
libname libref clear;
```

Note: For current Informix releases, `READ_ISOLATION_LEVEL=` is only valid when transaction logging is enabled. If transaction logging is not enabled, an error is generated when you use this option. Also, locks placed when `READ_ISOLATION_LEVEL= REPEATABLE READ` or `CURSOR_STABILITY` are *not freed* until the libref is cleared. Δ

To see the SQL locking statements that SAS issues to the Informix server, include the `SASTRACE=` option in your code:

```
option sastrace=',,,'d';
```

For more details about Informix locking, see your Informix documentation.

Naming Conventions for Informix

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

The Informix naming conventions are as follows:

- Table and column names must begin with a letter or an underscore followed by letters, numbers, or underscores. However, if the name appears within quotation marks and `PRESERVE_TAB_NAMES=YES` (when applicable), it can start with any character.
- Table and column names can contain up to 32 characters for connecting to Informix Dynamic Server 200, and 18 characters for the other servers.

Note: Informix encourages users to utilize lower case for table and column names. Several problems have been found in the Informix ODBC driver that result from using upper case or mixed case.

Currently Informix has no schedule for fixing these known problems. Δ

Informix Data Types

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

Character Data

CHAR(*n*), NCHAR(*n*)

contains character string data from 1 to 32,767 characters in length and can include tabs and spaces.

VARCHAR(*m,n*), NVARCHAR(*m,n*)

contains character string data from 1 to 255 characters in length.

TEXT

contains unlimited text data, depending on memory capacity.

BYTE

contains binary data of variable length.

Numeric Data

DECIMAL, MONEY, NUMERIC

contains numeric data with definable scale and precision. The amount of storage that is allocated depends on the size of the number.

FLOAT, DOUBLE PRECISION

contains double-precision numeric data up to 8 bytes.

INTEGER

contains an integer up to 32 bits (from -2^{31} to $2^{31}-1$).

REAL, SMALLFLOAT

contains single-precision, floating-point numbers up to 4 bytes.

SERIAL

stores sequential integers up to 32 bits.

SMALLINT

contains integers up to 2 bytes.

INT8

contains an integer up to 64 bits ($-2^{(63-1)}$ to $2^{(63-1)}$).

SERIAL8

contains sequential integers up to 64 bits.

Abstract Data

DATE

contains a calendar date in the form of a signed integer value.

DATETIME

contains a calendar date and time of day stored in 2 to 11 bytes, depending on precision.

Note: When the DATETIME column is in an uncommon format (i.e., DATETIME MINUTE TO MINUTE or DATETIME SECOND TO SECOND), the date and time values might not display correctly. Δ

INTERVAL

contains a span of time stored in 2 to 12 bytes, depending on precision.

Informix Null Values

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

If you do not indicate a default value for an Informix column, the default value is NULL. You can specify the keywords NOT NULL after the data type of the column when you create an Informix table to prevent NULL values from being stored in the column. When creating an Informix 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 *SAS/ACCESS for Relational Databases: Reference*.

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

LIBNAME Statement Data Conversions

The following table shows the default SAS variable formats that SAS/ACCESS applies to Informix data types during input operations when you use the LIBNAME statement. You can override these default data types by using the DBTYPE= data set option on a specific data set.

Table 1.3 LIBNAME Statement: Default SAS Formats for Informix Data Types

Informix Column Type	Default SAS Format
CHAR(<i>n</i>)	$\$n$
DATE	DATE9.
DATETIME**	DATETIME24.5
DECIMAL	$m+2.n$
DOUBLE PRECISION	none
FLOAT	none
INTEGER	none
INT8 [#]	none
INTERVAL	$\$n$
MONEY	none
NCHAR(<i>n</i>)	$\$n$
	NLS support required
NUMERIC	none
NVARCHAR(<i>m,n</i>)*	$\$m$
	NLS support required
REAL	none
SERIAL	none
SERIAL8 [#]	none
SMALLFLOAT	none

Informix Column Type	Default SAS Format
SMALLINT	none
TEXT*	\$n
VARCHAR(m,n)*	\$m

* Only supported by Informix-Online databases

The precision of a INT8 or SERIAL8 is 15 digit.

** If the Informix field qualifier specifies either HOUR, MINUTE, SECOND, or FRACTION as the largest unit, the value is converted to a SAS TIME value. All others, such as YEAR, MONTH, or DAY, are converted to a SAS DATETIME value.

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

Table 1.4 LIBNAME Statement: Default Informix Data Types for SAS Variable Formats

SAS Variable Format	Informix Data Type
\$w.	CHAR(w).
w. with SAS format name of NULL	DOUBLE
w.d with SAS format name of NULL	DOUBLE
all other numerics	DOUBLE
datetimew.d	DATETIME YEAR TO FRACTION(5)
datew.	DATE
time.	DATETIME HOUR TO SECOND

Pass-Through Facility Data Conversions

The Pass-Through Facility uses the same default conversion formats as the LIBNAME statement. See “LIBNAME Statement Data Conversions” on page 14 for the conversion tables.

Overview of Informix Servers

Informix Database Servers

There are two types of Informix database servers, the Informix-Online and Informix-SE servers. *Informix-Online database servers* can support many users and provide tools that ensure high availability, high reliability, and that support critical applications. *Informix-SE database servers* are designed to manage relatively small databases that are used privately by individuals or shared among a small number of users.

Using the DBDATASRC Environment Variables

The Pass-Through Facility supports the environment variable DBDATASRC, which is an extension to the Informix environment variable. If you set DBDATASRC, you can omit the CONNECT statement. The value of DBDATASRC is used instead of the SERVER= argument in the CONNECT statement. The syntax for setting DBDATASRC is like the syntax of the SERVER= argument:

Bourne shell:

```
DBDATABASE='testdsn' export DBDATASRC
```

C shell:

```
setenv DBDATASRC testdsn
```

If you set DBDATASRC, you can issue a PROC SQL SELECT or EXECUTE statement without first connecting to Informix with the CONNECT statement.

If you omit the CONNECT statement, an implicit connection is performed when the SELECT or EXECUTE statement is passed to Informix.

If you create an SQL view without an explicit CONNECT statement, the view can dynamically connect to different databases, depending on the value of the DBDATASRC environment variable.

Using Fully Qualified Table Names

Informix supports a connection to only one database. If you have data that spans multiple databases, you must use fully qualified table names to work within the Informix single-connection constraints.

In the following example, the tables Tab1 and Tab2 reside in different databases, MyDB1 and MyDB2, respectively.


```
proc sql;
  connect to informix
  (database=corpdb server=online);

  create view tab1v as
  select * from connection
  to informix
  (select * from mydb1.tab1);

  create view tab2v as
  select * from connection
  to informix
  (select * from mydb2.tab2);
quit;

data getboth;
  merge tab1v tab2v;
  by common;
run;
```

Because the tables reside in separate databases, you cannot connect to each database with a PROC SQL CONNECT statement and then retrieve the data in a single step. Using the fully qualified table name (that is, *database.table*) enables you to use any Informix database in the CONNECT statement and access Informix tables in the *same* or *different* databases in a single SAS procedure or DATA step.



Recommended Reading

Recommended Reading 17

Recommended Reading

Here is the recommended reading list for this title:

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

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

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

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

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

Glossary

This glossary defines SAS software terms that are used in this document as well as terms that relate specifically to SAS/ACCESS software.

access descriptor

a SAS/ACCESS file that describes data that is managed by a data management system. After creating an access descriptor, you can use it as the basis for creating one or more view descriptors. See also view and view descriptor.

browsing data

the process of viewing the contents of a file. Depending on how the file is accessed, you can view SAS data either one observation (row) at a time or as a group in a tabular format. You cannot update data that you are browsing.

bulk load

to load large amounts of data into a database object, using methods that are specific to a particular DBMS. Bulk loading enables you to rapidly and efficiently add multiple rows of data to a table as a single unit.

client

(1) a computer or application that requests services, data, or other resources from a server. (2) in the X Window System, an application program that interacts with the X server and can perform tasks such as terminal emulation or window management. For example, SAS is a client because it requests windows to be created, results to be displayed, and so on.

column

in relational databases, a vertical component of a table. Each column has a unique name, contains data of a specific type, and has certain attributes. A column is analogous to a variable in SAS terminology.

column function

an operation that is performed for each value in the column that is named as an argument of the function. For example, AVG(SALARY) is a column function.

commit

the process that ends a transaction and makes permanent any changes to the database that the user made during the transaction. When the commit process occurs, locks on the database are released so that other applications can access the changed data. The SQL COMMIT statement initiates the commit process.

DATA step view

a type of SAS data set that consists of a stored DATA step program. Like other SAS data views, a DATA step view contains a definition of data that is stored elsewhere; the view does not contain the physical data. The view's input data can come from one or more sources, including external files and other SAS data sets. Because a DATA step view only reads (opens for input) other files, you cannot update the view's underlying data.

data type

a unit of character or numeric information in a SAS data set. A data value represents one variable in an observation.

data value

in SAS, a unit of character or numeric information in a SAS data set. A data value represents one variable in an observation.

database

an organized collection of related data. A database usually contains named files, named objects, or other named entities such as tables, views, and indexes

database management system (DBMS)

an organized collection of related data. A database usually contains named files, named objects, or other named entities such as tables, views, and indexes

editing data

the process of viewing the contents of a file with the intent and the ability to change those contents. Depending on how the file is accessed, you can view the data either one observation at a time or in a tabular format.

engine

a component of SAS software that reads from or writes to a file. Each engine enables SAS to access files that are in a particular format. There are several types of engines.

file

a collection of related records that are treated as a unit. SAS files are processed and controlled by SAS and are stored in SAS data libraries.

format

a collection of related records that are treated as a unit. SAS files are processed and controlled by SAS and are stored in SAS data libraries. In SAS/ACCESS software, the default formats vary according to the interface product.

index

(1) in SAS software, a component of a SAS data set that enables SAS to access observations in the SAS data set quickly and efficiently. The purpose of SAS indexes is to optimize WHERE-clause processing and to facilitate BY-group processing. (2) in other software vendors' databases, a named object that directs the DBMS to the storage location of a particular data value for a particular column. Some DBMSs have additional specifications. These indexes are also used to optimize the processing of WHERE clauses and joins. Depending on the SAS interface to a database product and how selection criteria are specified, SAS may or may not be able to use the indexes of the DBMS to speed data retrieval.

Depending on how selection criteria are specified, SAS might use DBMS indices to speed data retrieval.

informat

a pattern or set of instructions that SAS uses to determine how data values in an input file should be interpreted. SAS provides a set of standard informats and also enables you to define your own informats.

interface view engine

a SAS engine that is used by SAS/ACCESS software to retrieve data from files that have been formatted by another vendor's software. Each SAS/ACCESS interface has its own interface view engine, which reads the interface product data and returns the data in a form that SAS can understand (that is, in a SAS data set). SAS automatically uses an interface view engine; the engine name is stored in SAS/ACCESS descriptor files so that you do not need to specify the engine name in a LIBNAME statement.

libref

a name that is temporarily associated with a SAS data library. The complete name of a SAS file consists of two words, separated by a period. The libref, which is the first word, indicates the library. The second word is the name of the specific SAS file. For example, in VLIB.NEWBDAY, the libref VLIB tells SAS which library contains the file NEWBDAY. You assign a libref with a LIBNAME statement or with an operating system command.

member

a SAS file in a SAS data library.

member name

a name that is given to a SAS file in a SAS data library.

member type

a SAS name that identifies the type of information that is stored in a SAS file. Member types include ACCESS, DATA, CATALOG, PROGRAM, and VIEW.

missing value

in SAS, a term that describes the contents of a variable that contains no data for a particular row or observation. By default, SAS prints or displays a missing numeric value as a single period, and it prints or displays a missing character value as a blank space.

observation

a row in a SAS data set. All of the data values in an observation are associated with a single entity such as a customer or a state. Each observation contains one data value for each variable. In a database product table, an observation is analogous to a row. Unlike rows in a database product table or file, observations in a SAS data file have an inherent order.

Pass-Through Facility

a group of SQL procedure statements that send and receive data directly between a relational database management system and SAS. The Pass-Through Facility includes the CONNECT, DISCONNECT, and EXECUTE statements, and the CONNECTION TO component. SAS/ACCESS software is required in order to use the Pass-Through Facility.

PROC SQL view

a SAS data set (of type VIEW) that is created by the SQL procedure. A PROC SQL view contains no data. Instead, it stores information that enables it to read data values from other files, which can include SAS data files, SAS/ACCESS views, DATA step views, or other PROC SQL views. A PROC SQL view's output can be either a subset or a superset of one or more files.

query

a set of instructions that requests particular information from one or more data sources.

referential integrity

a set of rules that a DBMS uses to ensure that whenever a data value in one table is changed, the appropriate change is also made to any related values in other tables or in the same table. Referential integrity is also used to ensure that related data is not deleted or changed accidentally.

relational database management system

a database management system that organizes and accesses data according to relationships between data items. Oracle and DB2 are examples of relational database management systems.

rollback

in most databases, the process that restores the database to its state when changes were last committed, voiding any recent changes. The SQL ROLLBACK statement initiates the rollback processes. See also commit.

row

in relational database management systems, the horizontal component of a table. A row is analogous to a SAS observation.

SAS data file

a type of SAS data set that contains data values as well as descriptor information that is associated with the data. The descriptor information includes information such as the data types and lengths of the variables, as well as the name of the engine that was used to create the data. A PROC SQL table is a SAS data file. SAS data files are of member type DATA.

SAS data library

a collection of one or more SAS files that are recognized by SAS and that are referenced and stored as a unit. Each file is a member of the library.

SAS data set

a file whose contents are in one of the native SAS file formats. There are two types of SAS data sets: SAS data files and SAS data views. SAS data files contain data values in addition to descriptor information that is associated with the data. SAS data views contain only the descriptor information plus other information that is required for retrieving data values from other SAS data sets or from files whose contents are in other software vendors' file formats.

SAS data view

a file whose contents are in one of the native SAS file formats. There are two types of SAS data sets: SAS data files and SAS data views. SAS data files contain data values in addition to descriptor information that is associated with the data. SAS data views contain only the descriptor information plus other information that is required for retrieving data values from other SAS data sets or from files whose contents are in other software vendors' file formats.

SAS/ACCESS views

See view descriptor and SAS data view.

server

in a network, a computer that is reserved for servicing other computers in the network. Servers can provide several different types of services, such as file services and communication services. Servers can also enable users to access shared resources such as disks, data, and modems.

Structured Query Language (SQL)

the standardized, high-level query language that is used in relational database management systems to create and manipulate database management system objects. SAS implements SQL through the SQL procedure.

table

a two-dimensional representation of data, in which the data values are arranged in rows and columns.

trigger

a type of user-defined stored procedure that is executed whenever a user issues a data-modification command such as INSERT, DELETE, or UPDATE for a specified table or column. Triggers can be used to implement referential integrity or to maintain business constraints.

variable

a column in a SAS data set. A variable is a set of data values that describe a given characteristic across all observations.

view

a definition of a virtual data set. The definition is named and stored for later use. A view contains no data; it merely describes or defines data that is stored elsewhere. SAS data views can be created by the ACCESS and SQL procedures.

view descriptor

a file created by SAS/ACCESS software that defines part or all of the database management system (DBMS) data or PC file data that is described by an access descriptor. The access descriptor describes the data in a single DBMS table, DBMS view, or PC file.

wildcard

a file created by SAS/ACCESS software that defines part or all of the database management system (DBMS) data or PC file data that is described by an access descriptor. The access descriptor describes the data in a single DBMS table, DBMS view, or PC file.

Index

- A**
- abstract data
 - Informix data types 13
 - autopartitioning
 - Informix 8
- B**
- BYTE data type 13
- C**
- CHAR data type
 - Informix 13
 - character data
 - Informix data types 13
 - committed reads 11
 - cursor stability reads 12
- D**
- data set options
 - Informix 4
 - data types
 - Informix 12
 - date and time data
 - Informix data types 13
 - DATE data type
 - Informix 13
 - DATETIME data type 13
 - DBDATASRC environment variables 16
 - DBSLICE= option
 - autopartitioning 9
 - DBSLICEPARM= option
 - LIBNAME statement 9
 - DECIMAL data type
 - Informix 13
 - dirty reads 11
 - DOUBLE PRECISION data type
 - Informix 13
- F**
- FLOAT data type
 - Informix 13
- I**
- Informix, interface to 1
 - autopartitioning scheme 8
 - data set options 4
 - data types 12
 - Informix database servers, about 15
 - LIBNAME statement 1, 14
 - locking in 11
 - naming conventions 12
 - Pass-Through Facility 5, 15
 - passing joins to 10
 - passing SAS functions to 9
 - stored procedures, calling 5
 - INT8 data type 13
 - INTEGER data type
 - Informix 13
 - INTERVAL data type 13
- J**
- joins
 - passing to Informix 10
- L**
- LIBNAME statement
 - Informix specifics 1, 14
 - locking data, handling
 - Informix interface 11
- M**
- MONEY data type
 - Informix 13
- N**
- naming conventions
 - Informix, interface to 12
 - NCHAR data type 13
 - null values
 - Informix 14
 - numeric data
 - Informix data types 13
 - NUMERIC data type
 - Informix 13
 - NVARCHAR data type 13
- P**
- Pass-Through Facility
 - Informix specifics 5, 15, 16
 - PRESERVE_COL_NAMES= option
 - Informix 12
 - PRESERVE_TAB_NAMES= option, LIBNAME statement
 - Informix 12
- R**
- READ_ISOLATION_LEVEL= option
 - Informix interface 11
 - REAL data type
 - Informix 13
 - remote stored procedures
 - calling, Informix 5
 - repeatable reads 11
- S**
- SAS/ACCESS data set options
 - Informix specifics 4
 - SAS SQL functions
 - passing to Informix 9
 - SERIAL data type 13
 - SERIAL8 data type 13
 - SERVER= option
 - LIBNAME statement 3
 - SMALLFLOAT data type 13
 - SMALLINT data type
 - Informix 13

stored procedures
calling, Informix 5

T

TEXT data type
Informix 13

threaded reads
Informix, interface to 8

U

USER= option, LIBNAME statement
Informix 2
USING= option, LIBNAME statement 3

V

VARCHAR data type
Informix 13

Your Turn

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

