

CHAPTER

1

SAS ODBC Driver

<i>Overview: SAS ODBC Driver</i>	1
<i>What Is ODBC?</i>	1
<i>What Is the SAS ODBC Driver?</i>	2
<i>Types of Data Accessed with the SAS ODBC Driver</i>	3
<i>Understanding SAS</i>	4
<i>SAS Data Sets</i>	4
<i>Unicode UTF-8 Support</i>	5
<i>SAS Libraries</i>	5
<i>SAS Servers</i>	6
<i>SAS Scalable Performance Data (SPD) Server</i>	6
<i>SAS Terminology</i>	7
<i>What Software Do I Need?</i>	7
<i>Accessibility Features in the SAS ODBC Driver</i>	8

Overview: SAS ODBC Driver

This document describes the SAS ODBC driver for SAS 9.1 and SAS 9.2. If you are using the Version 8 driver, see *SAS ODBC Driver User's Guide and Programmer's Reference, Version 8*.

This document is intended for three audiences:

- users who use the SAS ODBC driver to access data that is stored on their own computers*
- system administrators who use the SAS ODBC driver to enable multiple users to access shared data on a remote machine
- application programmers and others who need information about how the SAS ODBC driver is implemented

This document assumes that you are familiar with your Windows operating environment and that you know how to use a mouse and keyboard to perform common Windows tasks.

What Is ODBC?

ODBC stands for Open Database Connectivity. ODBC is an interface standard that provides a common application programming interface (API) to access databases. Most

* For more information about what kinds of data you can access, see "Types of Data Accessed with the SAS ODBC Driver" on page 3 .

software that runs in the Windows operating environment follows this standard for data access so that you can access data created by a variety of software applications.

ODBC functionality is provided by three main components:

- the client application
- the ODBC Driver Manager
- the ODBC driver

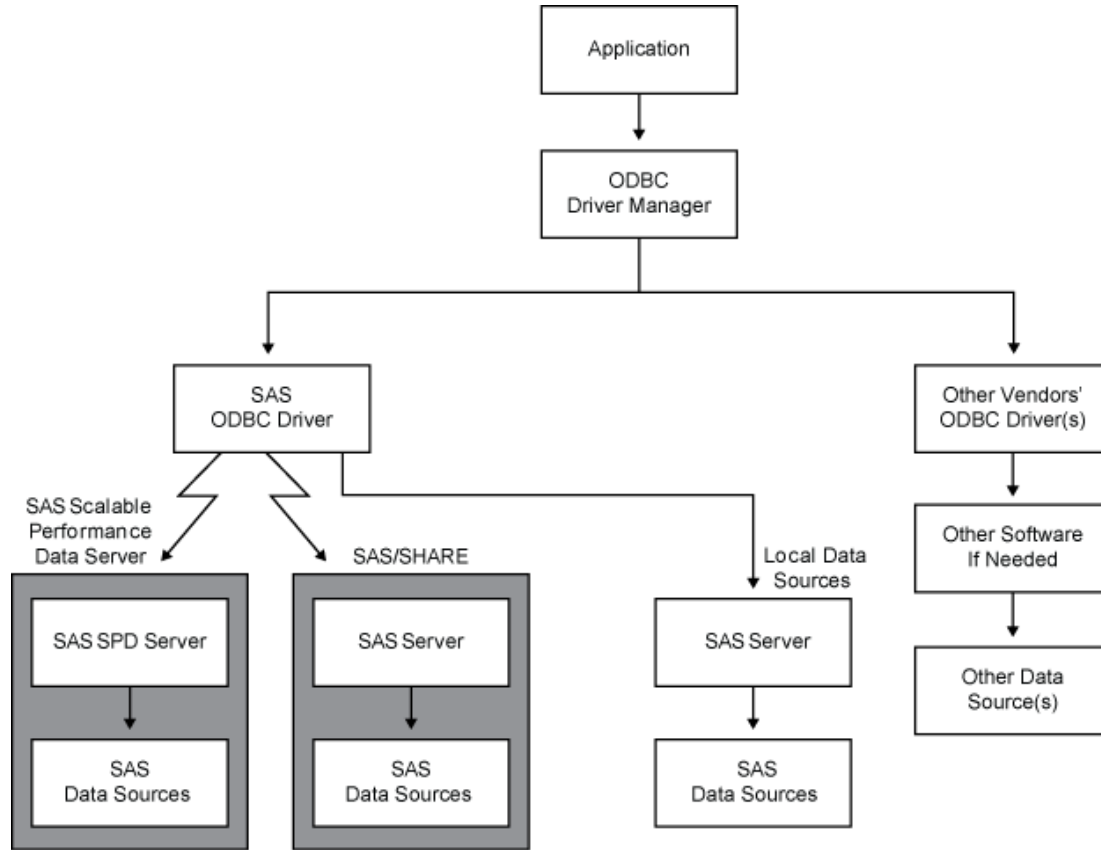
Figure 1.1 on page 3 displays components of ODBC functionality. The ODBC Driver Manager, which was developed by Microsoft, manages the interaction between a client application and one or more ODBC drivers.

What Is the SAS ODBC Driver?

The SAS ODBC driver is an implementation of the ODBC interface standard that enables you to access, manipulate, and update SAS data sources from applications that are ODBC compliant. As Figure 1.1 on page 3 shows, the SAS ODBC driver uses a SAS server to access data from your SAS data sources. The SAS server can be a local and minimized SAS session or a remote SAS/SHARE server. (For more information, see “SAS Servers” on page 6.) If you use other ODBC drivers (such as an ODBC driver for Oracle or SQL Server) to access other data sources, those ODBC drivers might require additional software components.

Note: To access ODBC data sources from SAS (the opposite of what the SAS ODBC driver enables you to accomplish), you must license SAS/ACCESS Interface to ODBC software. For more information, see *SAS/ACCESS for Relational Databases: Reference*. \triangle

Figure 1.1 Components of ODBC Functionality



Types of Data Accessed with the SAS ODBC Driver

In Figure 1.1 on page 3 and throughout this document, the term *SAS data sources* is used to describe data sources that you have defined in the SAS ODBC driver. These can include SAS data sets, flat files, and VSAM files, as well as data from many database management systems (DBMSs) through the use of SAS/ACCESS software. (For more information, see “SAS Data Sets” on page 4.)

If your PC is connected to a TCP/IP network, you can access both local data sources and remote data sources. (For more information about software requirements, see “What Software Do I Need?” on page 7.) *Local data* is data that you access through a SAS server on your local machine. The data can be stored either on your computer’s hard drive or on a network file system, such as a NetWare Windows NT file server, or a Windows shared folder, that makes the physical location of the data transparent to applications. *Remote data* is data that you access through a SAS/SHARE server that runs on another (remote) machine.

The ability to use the SAS ODBC driver with SAS/ACCESS software as a gateway to DBMS data is particularly useful in the following situations:

- When the DBMS vendor does not offer an ODBC driver. In this situation, you can use the SAS ODBC driver with SAS/ACCESS software on the SAS server to connect to the DBMS. SAS/ACCESS software to the DBMS must be licensed.
- You do not have a software license for the necessary software (either an ODBC driver for the DBMS or DBMS network access software) on your client machine.
- You want to merge DBMS data with other data.

Currently, SAS/ACCESS software is available for the following systems:

- ADABAS
- CA-IDMS
- DB2 under UNIX, Windows, and z/OS
- HP Neoview
- IMS
- Informix
- SQL Server
- MySQL
- Netezza
- ODBC
- OLE-DB
- Oracle
- PC Files
- Sybase
- System 2000
- Teradata

For more information about the individual SAS/ACCESS interfaces, see “Recommended Reading” on page 51.

Understanding SAS

To use the SAS ODBC driver, you should understand three components of SAS:

- SAS data sets
- SAS libraries
- SAS servers

You should also understand how ODBC terminology corresponds to SAS terminology. For more information, see “SAS Terminology” on page 7.

SAS Data Sets

A SAS data set is a file structured in a format that SAS can access. The physical object contains the following elements:

- data values that are stored in tabular form
- a descriptor portion that defines the types of data to SAS

The physical locations of the data values and the descriptor are not necessarily contiguous.

SAS data sets have two forms: *data files* and *data views*. SAS data files are relational tables with columns (or variables) and rows (or observations). The SAS data file structure can have many of the characteristics of a DBMS, including indexing, compression, and password protection.

SAS data views are definitions or descriptions of data that resides elsewhere. SAS data views enable you to use SAS to access many different data sources, including flat files, VSAM files, and DBMS structures, as well as native SAS data files. A data view eliminates the need to know anything about the structure of the data or the software

that created it. Data views need very little storage because they contain no data. They access the most current data from their defined sources because they collect the actual data values only when they are called.

You can use data views to define subsets of larger structures, or supersets of data that have been enhanced with calculated values. You can create SAS data views that combine views of dissimilar data sources. For example, you can combine a view of a relational DB2 table with a view of a SAS data file, a view of hierarchical IMS data, or even a view from a PC-based dBASE file.

You create SAS data views in three ways:

- with the DATA step (DATA step views)
- with the SQL procedure (PROC SQL views)
- with SAS/ACCESS software (SAS/ACCESS views)

You can use SAS/ACCESS software to work directly with DBMS tables, such as DB2 and Oracle, as if they were SAS data sets and data views by using the SAS/ACCESS LIBNAME statement. For more information, see *SAS/ACCESS for Relational Databases: Reference*.

These data views have some variation:

DATA step views

describe data from one or more sources, including flat files, VSAM files, and SAS data sets (either data files or other data views). You cannot use a DATA step view to update the view's underlying data because DATA step views only *read* other files.

For more information about how to create and use DATA step views, see "DATA Step Views," in *SAS Language Reference: Concepts*.

PROC SQL views

define either a subset or a superset of data from one or more SAS data sets. These data sets can be data files or data views, and can include data sets composed of DBMS data that are created with the SAS/ACCESS LIBNAME statement, and data views that are created with the SQL Pass-Through Facility to access DBMS data. You can also create PROC SQL views of DBMS data by using an embedded LIBNAME statement.

For example, an SQL procedure can combine data from PROC SQL views, DATA step views, and SAS/ACCESS views with data in a SAS data file. You cannot use a PROC SQL view to update the view's underlying files or tables. However, with some restrictions, you can use the UPDATE, DELETE, and INSERT statements in the SQL procedure to update data that is described by SAS/ACCESS views.

For more information about the SQL Pass-Through Facility, see "Overview: SQL Procedure," in *Base SAS Procedures Guide*.

Unicode UTF-8 Support

The SAS ODBC driver supports data sets with UTF-8 encoding.

SAS Libraries

SAS data sets are contained in libraries. Each SAS library has two names: a physical name and a logical name (libref). The physical name of the library fully identifies the directory or operating system data structure that contains the data sets. Therefore, the physical name must conform to the rules for naming files within your operating system.

You use the libref to identify a group of data sets (files or views) within SAS. The libref is a temporary name that you associate with the physical name of the SAS library

during each SAS job or session. After the libref is assigned, you can read, create, or update files in the library. A libref is valid only for the current SAS job or session and you can reference it repeatedly within that job or session.

You can use SAS/ACCESS software to associate a SAS libref with a DBMS database, schema, server, or group of tables and views, such as a DB2 database or group of Oracle tables and views. For more information about using SAS/ACCESS software, see *SAS/ACCESS for Relational Databases: Reference*.

For more information about SAS libraries, see “SAS Libraries,” in *SAS Language Reference: Concepts*.

SAS Servers

To access your SAS data sources, the SAS ODBC driver uses a SAS server. A SAS server is a SAS procedure (either PROC SERVER or PROC ODBCSEVER) that runs in its own SAS session. It accepts input and output requests from other SAS sessions and from the SAS ODBC driver on behalf of the applications that are ODBC compliant. While the server is running, the SAS session does not accept input from the keyboard.

The type of server that the driver uses depends on whether you are accessing local data or remote data:

local data

The SAS ODBC driver uses a SAS ODBC server to access your data. If you do not already have a SAS session running on your computer, the driver starts a SAS session and executes PROC ODBCSEVER, which automatically starts the SAS ODBC server when you connect to your local data source. For more information, see “Accessing Local Data Sources” on page 24. If you have a SAS session running on your computer (but not a SAS ODBC server), then you must either start the SAS ODBC server manually or end the SAS session before you connect to your local data sources. For more information, see “Starting a SAS ODBC Server” on page 24.

remote data

The SAS ODBC driver uses a SAS/SHARE server to access remote data. SAS/SHARE software must be licensed on the remote host. The driver requires TCP/IP software that is included with your operating system to communicate with the SAS/SHARE server. Your server administrator uses PROC SERVER to start the SAS/SHARE server on the remote host. For more information, see “Accessing Remote (SAS/SHARE) Data Sources” on page 25.

The SAS ODBC driver can communicate with a SAS/SHARE server or an SPD Server. You can interchange SAS data and SPD Server data by using the LIBNAME statement engine option in either SAS or SPD Server.

SAS Scalable Performance Data (SPD) Server

The SPD Server uses the latest parallel processing methods and data server capabilities to efficiently access large volumes of data and to serve large numbers of concurrent users. The SPD Server provides efficient data access for hundreds of users across multiple processors. The SAS ODBC driver can be configured for a direct connection to an SPD Server. ODBC connections to SPD SNET are not supported with SAS 9 and later versions of the SAS ODBC driver.

The SPD Server allows access to SAS data for intensive processing (queries and sorts) on the host server machine. It organizes and processes SAS data to take advantage of parallel processors on specific host servers.

You must have the SPD Server licensed on your client machine. Then, you can interchange SAS data and SPD Server data by using the LIBNAME statement engine option either in SAS or on the SPD Server. For more information, see *SAS Scalable Performance Data Server: User's Guide*. For information about configuring the SAS ODBC driver to connect to an SPD Server, see *SAS Scalable Performance Data Server: Administrator's Guide*.

SAS Terminology

Software products often include similar components or constructs that are known by different names. For the ODBC standard and SAS, the following correspondences exist:

<i>ODBC term</i>	<i>SAS term</i>
owner	library name (libref)
table	data set
qualifier	SAS data does not use a qualifier

Therefore, if your application that is ODBC compliant asks you to specify the owner for a SAS library, you should specify the libref. If the application asks for a table name, you should specify the name of the SAS data set. If a qualifier is requested, you can usually leave the field blank.

What Software Do I Need?

The SAS ODBC driver runs on 32-bit and 64-bit versions of Windows XP, Windows 2003, and Windows Vista. If the operating system is a server form such as Windows 2003 Server, then SAS/SHARE software must be licensed even if the data source is local data. Other software requirements depend on your hardware configuration and on the data sources that you want to access, as shown in Table 1.1 on page 7.

Note: The 16-bit version (shipped with Releases 6.10, 6.11, and 6.12 of SAS) can be downloaded from the SAS Web site and runs under Windows 3.1, 32s, and Windows for Workgroups. Δ

Table 1.1 Software Requirements

Data Source	Configuration	Requirements
local SAS data only	Stand-alone PC	Base SAS software*
local SAS data and local DBMS data	Stand-alone PC	Base SAS software* SAS/ACCESS interface for each DBMS

Data Source	Configuration	Requirements
remote SAS data	PC on a network	<i>On PC:</i> SAS ODBC driver** <i>On remote host:</i> Base SAS software SAS/SHARE software
remote SAS data and remote DBMS data	PC on a network	<i>On PC:</i> SAS ODBC driver** <i>On remote host:</i> Base SAS software SAS/SHARE software SAS/ACCESS interface for each DBMS

* For local data access, Base SAS software includes the SAS ODBC driver, but you install the driver separately. If the operating system is a server form, such as Windows 2003 Server, then SAS/SHARE software must be licensed for accessing local data.

** For remote data access, install the SAS ODBC driver on your client PC. The SAS license for the remote host must include SAS/SHARE software.

Accessibility Features in the SAS ODBC Driver

The SAS ODBC driver includes accessibility and compatibility features that improve usability of the product for users with disabilities. These features are related to accessibility standards for electronic information technology adopted by the U.S. Government under Section 508 of the U.S. Rehabilitation Act of 1973, as amended. If you have questions or concerns about the accessibility of SAS products, send e-mail to accessibility@sas.com.