

CHAPTER

1

Introducing the SAS ODBC Driver

<i>Overview: The SAS ODBC Driver</i>	1
<i>What Is ODBC?</i>	2
<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>	5
<i>SAS Data Sets</i>	5
<i>SAS Data Libraries</i>	6
<i>SAS Servers</i>	6
<i>Scalable Performance Data Server</i>	7
<i>Converting DDE Servers to TCP Servers</i>	7
<i>SAS Terminology</i>	8
<i>What Software Do I Need?</i>	8

Overview: The SAS ODBC Driver

This book documents SAS 9.1 of the SAS ODBC driver. If you are using the Version 8 driver, you should see *SAS ODBC Driver User's Guide and Programmer's Reference, Version 8*.

This book 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
- applications programmers and others who need detailed information about how the SAS ODBC driver is implemented.

This book assumes that you are familiar with the Microsoft Windows NT, Windows 2000, or Windows XP operating environments and that you know how to use a mouse and keyboard to perform common Windows tasks.

* See "Types of Data Accessed with the SAS ODBC Driver" on page 3 for information about what kinds of data you can access.

What Is ODBC?

ODBC stands for Open Database Connectivity. It is an interface standard that provides a common application programming interface (API) to access databases. Many software products that run in the Windows operating environment adhere to this standard so that you can access data that other software has created.

ODBC functionality is provided by three main components:

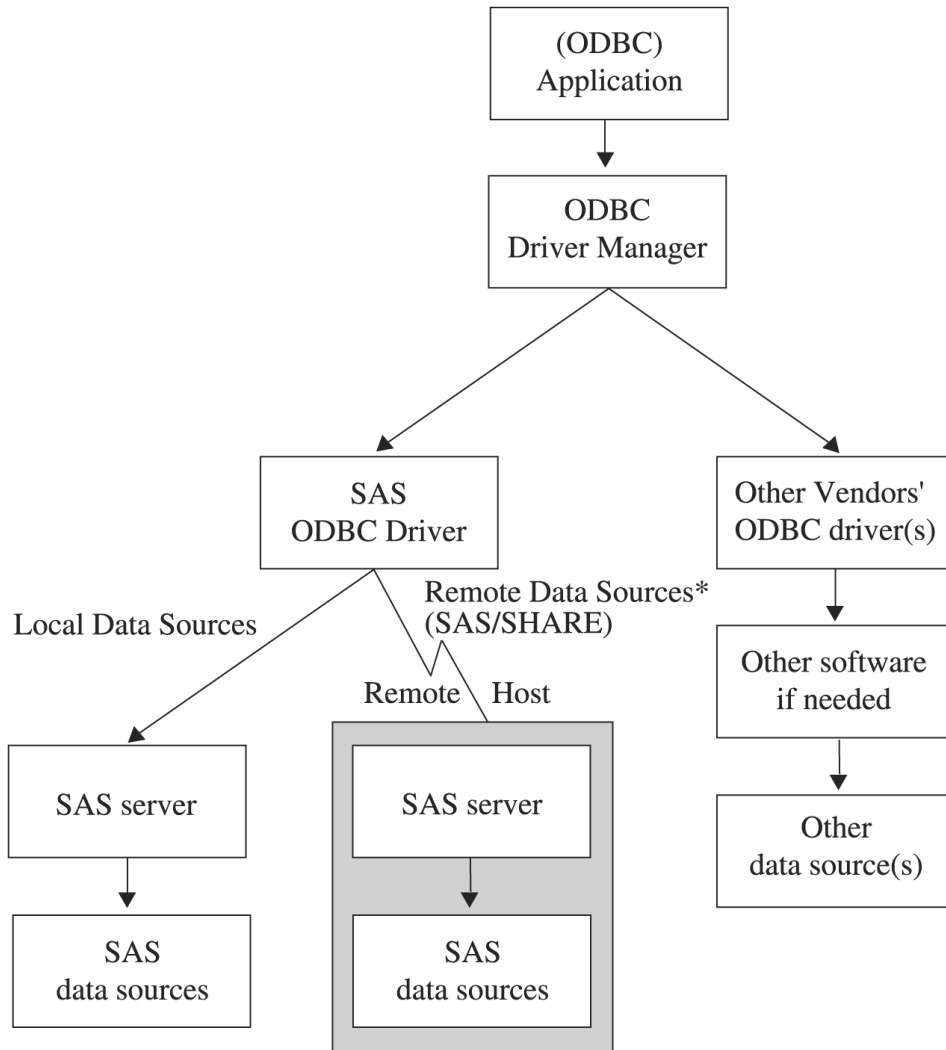
- the client application
- the ODBC driver manager
- the ODBC driver.

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 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 either a local, iconified SAS session or a remote SAS/SHARE server. (See “SAS Servers” on page 6 for more information.) If you use other ODBC drivers (such as those for ORACLE or SQL Server) to access other data sources, those drivers might require additional software components.

Note: To access ODBC data sources from within SAS (the opposite of what the SAS ODBC driver enables you to accomplish), you must license the SAS/ACCESS interface to ODBC. For additional information see *SAS/ACCESS for Relational Databases: Reference*. △

Figure 1.1 Components of ODBC Functionality

*To access remote data sources, the SAS ODBC driver requires TCP/IP communications software. You must also license SAS/SHARE and SAS/SHARE*NET on the remote host.

Types of Data Accessed with the SAS ODBC Driver

In Figure 1.1 on page 3 and elsewhere 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. (See “SAS Data Sets” on page 5 for details.)

If your personal computer is connected to a TCP/IP network, you can access both local data sources and remote data sources. (See “What Software Do I Need?” on page 8 for information about software requirements.) *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 or Windows NT

file server, 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 in conjunction with SAS/ACCESS software as a gateway to DBMS data is particularly useful under any of the following circumstances:

- There is no ODBC driver available for the DBMS. In this situation, you can use the SAS ODBC driver in conjunction with SAS/ACCESS on the SAS server to connect to the DBMS.
- You do not license the necessary software (either an ODBC driver for the DBMS or DBMS network access software) on your client machine.
- You want to join or merge DBMS data with other data.

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

ADABAS
 CA-Datcom/DB
 CA-IDMS
 CA-OpenIngres
 DB2 under OS/390,
 UNIX, and
 Windows
 IMS-DL/I
 Informix
 Microsoft SQL
 Server
 ODBC
 OLE DB Providers
 ORACLE
 Oracle Rdb
 PC File Formats*
 SYBASE
 SYSTEM 2000 Data
 Teradata

*This includes Lotus 1-2-3, Microsoft Excel (5, 7/95, 97, 2000, 2002), and Microsoft Access (97, 2000, 2002) file formats.

Note: You can use the SAS/ACCESS Interface to ODBC to access many sources that provide an ODBC driver, including AS/400, SQL Server, and Microsoft Access data. Δ

For information about the individual SAS/ACCESS interfaces, see “Recommended Reading” on page 49.

Understanding SAS

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

- SAS data sets
- SAS data libraries
- SAS servers.

You should also understand how ODBC terminology corresponds to SAS terminology. See “SAS Terminology” on page 8 for more details.

SAS Data Sets

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

- 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 essentially relational tables with columns (or variables) and rows (or observations). The SAS data file structure can have many of the characteristics of a database management system, including indexing, compression, and password protection.

SAS data views are definitions or descriptions of data that resides elsewhere. They 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 take up very little storage because they contain no data. They also access the most current data from their defined sources because they collect the actual data values only when 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 also 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-DL/I 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 the ACCESS procedure (SAS/ACCESS views).

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

Some variation occurs among these types of views:

DATA step views

describe data from one or more sources, including flat files, VSAM files, and SAS data sets (either files or other 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 the chapter “SAS Data 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 files or views, and can include data sets composed of DBMS data that are created with the SAS/ACCESS LIBNAME statement, and views that are created with the PROC SQL Pass-Through Facility to access DBMS data. You can also construct PROC SQL views of DBMS data by using an embedded LIBNAME statement.

For example, the 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 PROC SQL views to update the data in 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 information about the PROC SQL Pass-Through Facility, see the chapter “The SQL Procedure” in *Base SAS Procedures Guide*, or *SAS/ACCESS for Relational Databases: Reference*.

SAS/ACCESS views

describe views that are created with the ACCESS procedure in SAS/ACCESS software. They are usually called *view descriptors* to distinguish them from PROC SQL or DATA step views. You can use them to describe data from any DBMS for which you license a SAS/ACCESS interface (see the list in “Types of Data Accessed with the SAS ODBC Driver” on page 3). Each view descriptor describes all or some of the data in one DBMS table or in one DBMS view. Some SAS/ACCESS interfaces do not permit you to use a view descriptor to update the underlying DBMS data.

For more information about view descriptors, see *SAS/ACCESS for Relational Databases: Reference*.

SAS Data Libraries

SAS data sets are contained in data libraries. Each SAS data 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 data library during each SAS job or session. After the libref is assigned, you can read, create, or update files in the data 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 also 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. See *SAS/ACCESS for Relational Databases: Reference* for more information about using SAS/ACCESS software.

For more information about SAS data libraries, see the chapter “SAS Data 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.

Note: Beginning with Version 7 of SAS, the SAS ODBC driver uses the TCP/IP protocol to communicate with both local servers and remote SAS/SHARE servers. Support for the DDE protocol has been discontinued. Therefore, you should reconfigure your driver settings to change any DDE (local or network) servers to TCP/IP servers, as described in “Converting DDE Servers to TCP Servers” on page 7. △

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

local data

The 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 ODBCSEV, thereby automatically starting the server when you connect to your local data source. See “Accessing Local SAS Data Sources” on page 26 for more information. If you have a SAS session (but not a SAS ODBC server) running on your computer, then you must either start the SAS ODBC server manually or end the SAS session before you connect to your SAS data sources (except under Windows NT, which supports multiple concurrent SAS sessions). See “Starting a SAS ODBC Server” on page 26 for details.

remote data

The driver uses a SAS/SHARE server. This requires that you license the SAS/SHARE software on the remote host. The driver also requires TCP/IP software that is included with your operating environment to communicate with the server. Your server administrator uses PROC SERVER to start the server on your remote host. See “Accessing Remote (SAS/SHARE) Data Sources” on page 27 for more information.

The SAS ODBC driver can communicate with a SAS/SHARE server or a Scalable Performance Data (SPD) server. You can interchange SAS data and SPD server data by using the LIBNAME statement engine option in either SAS or SPD server.

Scalable Performance Data Server

The Scalable Performance Data (SPD) server utilizes the latest parallel processing methods and data server capabilities to efficiently access large volumes of data and serve large numbers of concurrent users. It provides efficient data access for hundreds of users across multiple processors.

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.

The SPD SNET server handles communication between SAS clients and an SPD server. 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 *Scalable Performance Data Server User's Guide*.

Converting DDE Servers to TCP Servers

If you still use SAS servers that are configured to use DDE, you must convert them to TCP servers to use them with this release of the SAS ODBC driver.

To convert a DDE (or Network DDE) server to a TCP server, perform these steps:

- 1 Install the new SAS ODBC driver.
- 2 Select the Servers tab from the SAS ODBC Driver Configuration dialog box.
- 3 Select the existing DDE server that you want to convert to TCP in the Servers list to the left of the Servers page.
- 4 Select to verify that the options are correct, and then select .

Note: If you are converting from a local DDE server, you need to remove the `-comamid dde` option from the SAS Startup Parameters text box. \triangle

- 5 Select and then to store your changes.

Note: For local TCP servers you will also have to verify that your TCP/IP Services file has an appropriate entry for the server that you are configuring. For more information, see “The TCP/IP Services File” on page 29. \triangle

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	not used

Therefore, if your application that is ODBC compliant asks you to specify the owner for a SAS data 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 under 32-bit Windows (Windows NT, Windows XP, and Windows 2000). Other software requirements depend on your hardware configuration and on the data source(s) that you want to access, as shown in Table 1.1 on page 9.

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

Table 1.1 Software Requirements

Data Source(s)	Configuration	Requirements
local SAS data only	Stand-alone PC	Base SAS*
local SAS data and local DBMS data	Stand-alone PC	Base SAS* SAS/ACCESS interface for each DBMS TCP/IP**
remote SAS data	PC on a network	<i>On PC:</i> SAS ODBC driver*** TCP/IP** <i>On remote host:</i> Base SAS SAS/SHARE TCP/IP**
remote SAS data and remote DBMS data	PC on a network	<i>On PC:</i> SAS ODBC driver*** TCP/IP** <i>On remote host:</i> Base SAS SAS/SHARE SAS/ACCESS interface for each DBMS TCP/IP**

* For local data access, base SAS software includes the SAS ODBC driver, but you install the driver separately.

** Requires a version of TCP/IP that supports the Microsoft Windows socket (“winsock”) API.

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

