

CHAPTER

1

Introducing the SAS ODBC Driver

<i>Using This Book</i>	1
<i>Additional Documentation</i>	1
<i>What Is ODBC?</i>	2
<i>What Is the SAS ODBC Driver?</i>	2
<i>What Kinds of Data Can I Access with the SAS ODBC Driver?</i>	3
<i>What Do I Need to Know about the SAS System?</i>	4
<i>SAS Data Sets</i>	5
<i>SAS Data Libraries</i>	6
<i>SAS Servers</i>	7
<i>Converting DDE Servers to TCP Servers</i>	7
<i>SAS Terminology</i>	8
<i>What Software Do I Need?</i>	8

Using This Book

This book documents Version 8 of the SAS ODBC driver. If you are using the Version 7 driver, you should refer to the *SAS ODBC Driver User's Guide and Programmer's Reference, Version 7-0*.

This book is intended for three audiences:

- users who are using the SAS ODBC driver to access data that is stored on their own computers*
- system administrators who are using the SAS ODBC driver to enable multiple users to access shared data
- applications programmers and others who need detailed information about how the SAS ODBC driver has been implemented.

In this book, we assume that all readers are familiar with the Microsoft Windows 95, Windows 98, or Windows NT operating environments and that they know how to use the mouse and keyboard to perform common Windows tasks.

Additional Documentation

The following list of documentation may be useful to you when you use the SAS ODBC driver. For more information, refer to the current *Publications Catalog*. See the SAS Publications Web page at www.sas.com/pubs.

* See "What Kinds of Data Can I Access with the SAS ODBC Driver?" on page 3 for information about what kinds of data you can access.

- *SAS/ACCESS Software for Relational Databases: Reference*. This book is bound in a three-ring binder and must be used in conjunction with a DBMS chapter. DBMS chapters are ordered separately and can be inserted into the binder.
- *SAS/ACCESS Software for PC File Formats: Reference*.
- *SAS/SHARE User's Guide*.

For other vendors' software publications, see the documentation shipped with your software.

What Is ODBC?

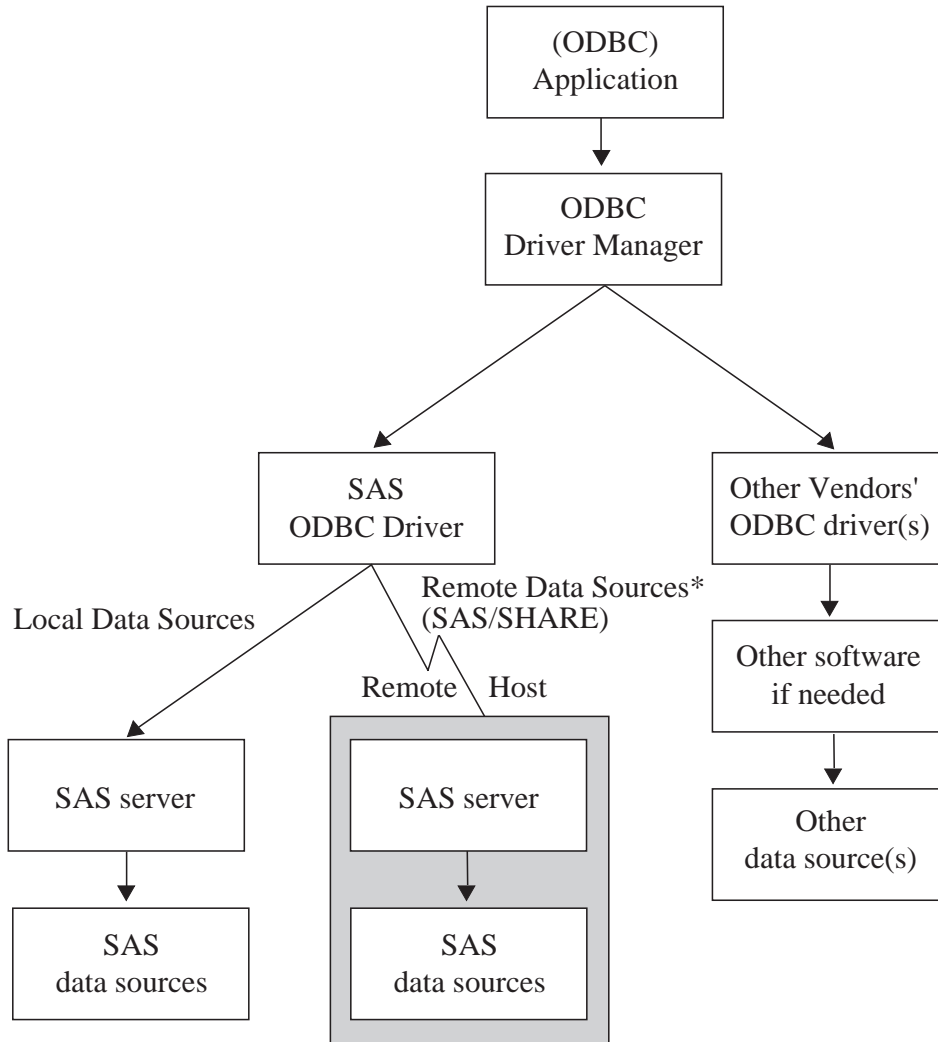
ODBC stands for Open Database Connectivity. It is an interface standard that provides a common application programming interface (API) for accessing databases. Many software products that run in the Windows operating environment adhere to this standard, giving users access to data that was created with other software.

ODBC functionality is provided by three main components: the client application, the ODBC driver manager, and the ODBC driver (see Figure 1.1 on page 3). 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 ODBC-compliant applications. 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 executes in an iconified, interactive SAS session. (See "SAS Servers" on page 7 for more information.) If you use other ODBC drivers (such as those for ORACLE or SQL Server) to access other data sources, those drivers may require additional software components.

Note: To access ODBC data sources from within the SAS System (the opposite of what the SAS ODBC driver enables you to accomplish), you must license the SAS/ACCESS interface to ODBC, which is documented in *SAS/ACCESS Software 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.

What Kinds of Data Can I Access with the SAS ODBC Driver?

In Figure 1.1 on page 3 and elsewhere throughout this report, we use the term *SAS data sources* to mean data sources that you have defined to the SAS ODBC driver. These can include not only SAS data sets, but also 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 PC is on a 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 may be stored either on your PC hard drive or on a PC 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 server that is running 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 for the remote DBMS. Therefore, you use the SAS ODBC driver in conjunction with SAS/ACCESS software in the remote SAS server to get ODBC connectivity.
- You don't have a license for 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 on a remote host.

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

ADABAS

CA-IDMS

CA-OpenIngres

DB2 under OS/390,
OS/2, UNIX, and
Windows

IMS-DL/I

Informix

ODBC

ORACLE

Oracle Rdb

PC File Formats*

SYBASE

SYSTEM 2000

*includes DBF, DIF, WK1, WK3, WK4, XLS4, XLS5, and XLS7 file formats.

Note: The SAS/ACCESS Interface to ODBC can be used 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 "Additional Documentation" on page 1.

What Do I Need to Know about the SAS System?

To use the SAS ODBC driver, you need to understand three components of the SAS System:

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

You also need to understand the correspondence between ODBC terminology and SAS terminology. These topics are explained in the following sections.

SAS Data Sets

A SAS data set is a file that the SAS System can access as if it were a physical object containing the following:

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

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 the SAS System to access many different data sources, including flat files, VSAM files, and DBMS structures, as well as native SAS data files. They also eliminate the need for users to know anything about the structure of the data or the software that was used to create it.

Because they contain no data, views take up very little storage. And because they collect the actual data values only when called, they always access the most current data from their defined sources.

Views can be used 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, a view of a relational DB2 table can be combined 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 can 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 Software for Relational Databases: Reference* for more information about using SAS/ACCESS software.

There is some variation among these types of views, as described in the following table.

Table 1.1 Types of SAS Data Views

Type of View	Description
DATA step views	<p>can describe data from one or more sources, including flat files, VSAM files, and SAS data sets (either files or other views). Because a DATA step view only <i>reads</i> other files, you cannot use it to update the view's underlying data.</p> <p>For more information about creating and using DATA step views see the chapter "SAS Data Views" in the <i>SAS Language Reference: Concepts</i>.</p>
PROC SQL views	<p>can 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. In Version 7, you can also construct PROC SQL views of DBMS data by using an embedded LIBNAME statement.</p> <p>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.</p> <p>For information about the PROC SQL Pass-Through Facility, see the chapter "The SQL Procedure" in <i>SAS Procedures Guide</i>, or the <i>SAS/ACCESS Software for Relational Databases: Reference</i>.</p>
SAS/ACCESS views	<p>are created with the ACCESS procedure in SAS/ACCESS software. They are generally called <i>view descriptors</i> 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 "What Kinds of Data Can I Access 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. When the SAS servers that are used by the SAS ODBC driver are involved, some SAS/ACCESS interfaces do not permit you to use a view descriptor to update the underlying DBMS data.</p> <p>For more information about view descriptors, see the <i>SAS/ACCESS Software for Relational Databases: Reference</i>.</p>

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, filetype or minidisk, or host system data structure that contains the data sets. It must therefore conform to the rules for naming files within your host system.

You use the libref to identify a group of data sets (files or views) to the SAS System. 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 can be referenced 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 Software 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 the *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 ODBCSEV) 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 ODBC-compliant application. While the server is running, the SAS session does not accept input from the keyboard.

Note: Beginning with Version 7 of the SAS System, the SAS ODBC driver uses the TCP/IP protocol to communicate both with local servers and with remote SAS/SHARE servers. Configuration of DDE to communicate with servers is no longer available and is not supported with the Version 7 (and later) driver. Therefore, it is recommended that you 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. Δ

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 PC, the driver starts a SAS session and executes PROC ODBCSEV, thereby automatically starting the server for you 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 PC, then you must either start the SAS ODBC server manually or end the SAS session before connecting to your SAS data sources.* 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 *NET facility and SAS/SHARE software on the remote host. The driver also requires TCP/IP software included with Windows 95 and Windows NT in order 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.

Converting DDE Servers to TCP Servers

If you still have SAS servers configured to use DDE, you must convert them to TCP servers to use them with this version of the SAS ODBC Driver.

To convert a DDE (or Network DDE) server to a TCP server, take the following steps:

- 1 Install the new SAS ODBC Driver.
- 2 From the SAS ODBC Driver Configuration dialog, select the Servers tab.
- 3 Select the existing DDE server that you want to convert to TCP in the list of servers to the left.

* except under Windows NT, which supports multiple concurrent SAS sessions.

- 4 Click **Configure** and verify that the options are correct.

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

Click **OK**.

- 5 Click **Update** and then **OK** 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 you are configuring. For more information, see “The TCP/IP Services File” on page 28. \triangle

SAS Terminology

Different software products often include similar components or constructs that are known by different names. For the ODBC standard and the SAS System, 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 ODBC-compliant application asks you to specify the owner for a SAS data library, you should specify the libref. If the application asks for a table name, supply the name of the SAS data set. If a qualifier is requested, you can generally leave the field blank.

What Software Do I Need?

The SAS ODBC driver included with Version 8 (and Version 7) runs under 32-bit Windows (Windows 95, Windows 98, and Windows NT).

Other software requirements depend on your hardware configuration and on the data source(s) that you want to access, as shown in Table 1.2 on page 8. The 16-bit version (shipped with Releases 6.10, 6.11 and 6.12 of the SAS System) is downloadable from the SAS Web site and runs under Windows 3.1, 32s, and Windows for Workgroups.

Table 1.2 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

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

* 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, the SAS ODBC driver is included with SAS/SHARE*NET. You install the driver on your client PCs, and you use the SAS SETINIT to enable SAS/SHARE*NET on the remote host.

