

CHAPTER

2

What Is SAS/SHARE Software?

<i>Audience</i>	19
<i>SAS/SHARE Is a Multi-User Data Server</i>	19
<i>SAS/SHARE Offers a Path to Remote Data</i>	20
<i>SAS/SHARE Server Is the Hub between Data and Its Clients</i>	20
<i>Accessing SAS Files through the Operating Environment</i>	22
<i>Accessing SAS Files through a SAS/SHARE Server</i>	23
<i>Remote Library Services Offers Remote File Access</i>	24
<i>SAS/SHARE Software Components</i>	25
<i>SAS/SHARE Users</i>	25
<i>Using the SAS/SHARE and SAS/CONNECT Documentation Set</i>	26
<i>Version 6-Version 7 and Version 8 Compatibility</i>	26

Audience

All SAS/SHARE software users should read this chapter.

SAS/SHARE Is a Multi-User Data Server

SAS/SHARE software provides a multi-user data server that offers several advantages for local and remote SAS clients and other clients. The multi-user SAS/SHARE server enables two or more clients to write to the same SAS file (for example, a data set or a catalog) at the same time. This feature is traditionally known as *concurrent update access*.

The traditional scenario for using SAS/SHARE is an environment in which multiple client sessions all want to share (read and write) records in the same SAS data set.

The following list includes a sample of the types of operations that multiple clients can perform at the same time:

- While one user is creating a member in a data library, other users can create, read, delete, and update members in the same library.
- While one user attempts to use the SAS Explorer window on a data library, other users can open the same window to browse, delete, edit, or re-name members in the same library. You can also copy a member from the SAS Explorer window.
- While one user opens the CATALOG window on a catalog, other users can open the CATALOG window to browse, copy, delete, or rename entries in the same catalog.
- While one or more users are using the FSEDIT procedure, the FSVIEW procedure in edit mode, the UPDATE statement in the SQL procedure, or an SCL program to update a SAS data set, other users can

- update the SAS data set using the MODIFY, REMOVE, or REPLACE statement in a DATA step
- read the SAS data set as input data using the SET statement in a DATA step
- add observations to the SAS data set using the APPEND or SQL procedure or remove observations with the SQL procedure
- copy the SAS data set that is being updated (as well as other members of the library) into another library by using the COPY or the CATALOG procedure.

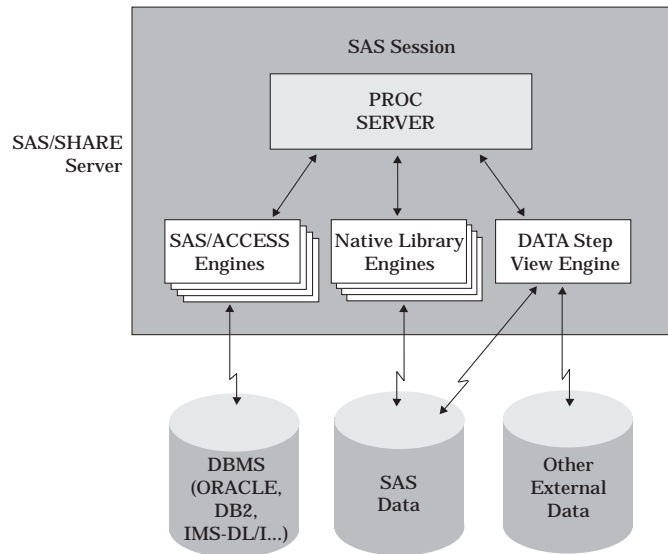
SAS/SHARE Offers a Path to Remote Data

The multi-user SAS/SHARE server also offers remote clients a path to shared data, even if they want only to read that data, without the overhead of a SAS/CONNECT sign on.

In this scenario, you might have a network of client machines that need read access to a data set that resides on a central server system machine. You could use SAS/CONNECT and have each of those clients create a remote session on the central server machine. However, if you need to read only a small-to-moderate amount of data, the overhead for each client that is signing on to the central host and starting a SAS session may be significant. Also, the additional load on the central machine that comes with each of these SAS sessions may be undesirable. You can avoid this overhead and additional load by having those client sessions access data through a SAS/SHARE server that is running on the central machine. Because the server is already running and it serves multiple users, connecting to the server and accessing the data takes very little time.

SAS/SHARE Server Is the Hub between Data and Its Clients

A SAS/SHARE server may be thought of as a hub that serves clients with data from many different sources. For example, a server must use a SAS/ACCESS engine to ORACLE in order to access data that is stored in an ORACLE DBMS. As another example, a server can get to SAS data through a Native Library engine. See your SAS/ACCESS documentation for details about how to specify an engine to access specific data. Figure 2.1 on page 21 shows a sample of the data sources that a SAS/SHARE server can provide to its clients.

Figure 2.1 Data Sources for a SAS/SHARE Server

With an identified DBMS, a SAS/SHARE server supplies the requesting client with data for its data processing needs. In Version 8, support extends to clients other than the classic SAS client.

With a license for SAS/SHARE*NET software, you may send requests to a SAS/SHARE server from a client that is not a SAS application. A SAS/SHARE*NET server is a SAS/SHARE server for which SAS/SHARE*NET software is licensed. SAS/SHARE*NET software includes the Data Services component of SAS/IntrNet software.

Examples of clients that are not SAS applications are:

htmSQL

runs a Web server and offers a gateway to your SAS data from a Web browser. It enables you to incorporate data into a Web page by using SQL queries.

Java applet or application

uses SAS/SHARE*NET Driver for JDBC, which enables you to write Java applets or applications that can view and update data through a direct connection to a SAS/SHARE*NET server.

C program

uses the SAS SQL Library for C, which is an API that enables you to create applications that use SQL queries and statements to access data in SAS data sets and in other database management systems.

Application that uses the ODBC driver, such as Microsoft Excel

uses the ODBC driver, which provides ODBC-compliant Windows applications with read and write access to local and remote SAS data sets.

OLE DB consumer or ADO application

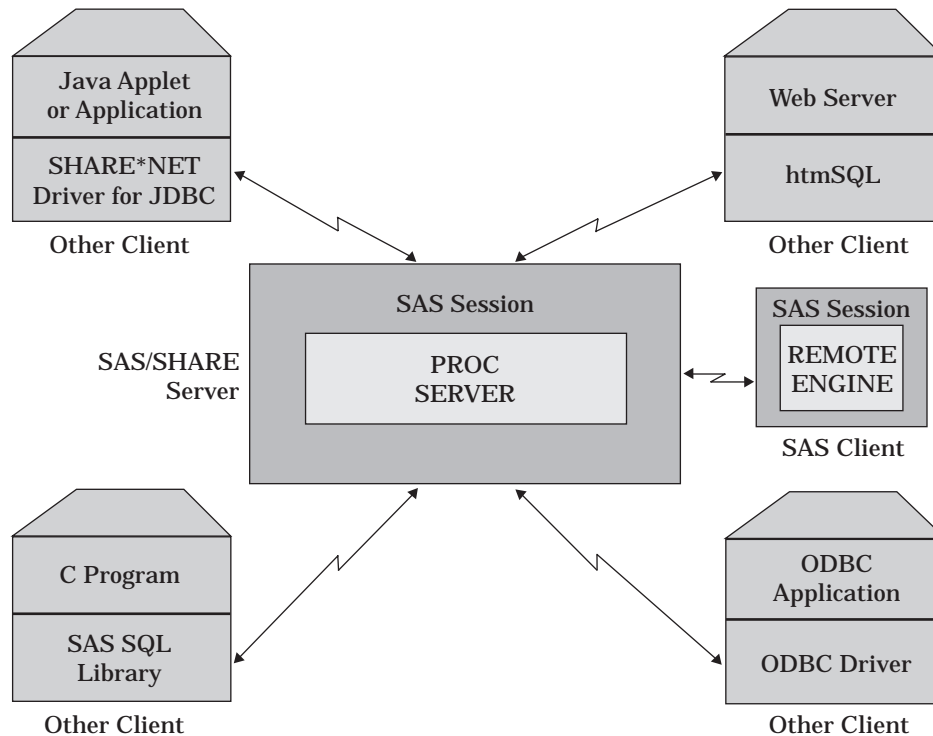
(new in Version 8) uses ShareProvider to view and update data through a direct connection to a SAS/SHARE*NET server. ShareProvider implements the Microsoft OLE DB specification and can be used by any OLE DB- compliant or ADO-enabled application.

Each of the preceding client interfaces or applications has its own documentation.

Figure 2.2 on page 22 shows a sample of the types of clients that a V8 server supports. The diagram shows a SAS/SHARE server that is running in a SAS session in a supported operating environment. A server administrator starts the SAS/SHARE

server session. SAS/SHARE clients can connect to the server from any machine on your network. SAS clients use the REMOTE engine to access data through a SAS/SHARE server.


Figure 2.2 SAS/SHARE Server Clients



The term *other client* refers to a client that is not a SAS application. For each of these clients, the appropriate client-side drivers and libraries must be invoked, as appropriate.

Accessing SAS Files through the Operating Environment

Figure 2.3 on page 23 shows two client sessions that access SAS files without a SAS/SHARE server.

Figure 2.3 Data Library Access Provided by the Operating Environment


In this model, your operating environment provides two types of access: read only and read-write. With read-only access, multiple users can simultaneously read the same member or different members in the same library.

SAS files that have read-write access are typically associated with only one user. Some hosts support only read access to libraries that are accessed directly by more than one user. Other hosts permit two users to write to different files in the same SAS library but do not permit them to write to the same file simultaneously. Still other hosts permit multiple users to write to the same file at the same time, even though doing so is not safe and invariably results in lost changes made by one or more users, at best. In all cases, the only safe way for more than one user to share a SAS file (and on some hosts an entire library) is for them to use a common third-party process to serialize access to the low-level parts of the file and coordinate updates to the data.

For more information about the SAS data library model, see *SAS Language Reference: Concepts*.

Accessing SAS Files through a SAS/SHARE Server

Figure 2.4 on page 24 introduces a SAS/SHARE server into the preceding diagram.

Figure 2.4 SAS Data Library Access Provided by a SAS/SHARE Server

The SAS/SHARE server enables multiple clients to effectively share the same SAS file at the same time. In this context, *share* means to allow access by multiple clients to a different unit (for example, an observation) in the same SAS file. In the diagram, Client 1 can read from and write to observation 2, and Client 2 can read from and write to observation 3 in the same SAS file. Both clients can read the same observation at the same time. However, only one client at a time can write to an observation.

The server enables client access to the lowest unit of the SAS data library hierarchy through its powerful lock manager facility. Locks are applied in either of two ways:

- A client may specify an explicit lock with a LOCK command or a LOCK statement.
- A client operation may automatically submit a request for an implicit lock.

The SAS/SHARE server evaluates each incoming client request to access a particular SAS data library unit (for example, a data library, a data set, or an observation) against a complex set of locking rules whose application seems transparent to clients. The server balances the client requests for access to data while ensuring the integrity of that data. The server grants permission to a qualifying client and denies a conflicting request. The denied client receives an informational message. For more information about server locking rules, see Chapter 5, “Locking SAS Data Objects,” on page 59. For information about the use of the LOCK command and the LOCK statement, see “Locking Objects Explicitly (LOCK Statement)” on page 63, and “Locking Explicitly in a SAS Window (LOCK Command)” on page 67.

Remote Library Services Offers Remote File Access

SAS/SHARE provides remote file access through its Remote Library Services (RLS). RLS provides transparent access to remote data libraries for moving data through the network as the local SAS session requests it.

RLS provides access to remote data through the REMOTE engine. Thus, SAS products can gain single-user or multi-user access to remote SAS data or third-party DBMS data, as applicable, by invoking a SAS/SHARE server and by assigning a library to the server through the REMOTE engine.

The LIBNAME statement associates a SAS library reference (libref) with a permanent SAS data library, which can be specified by using a host-specific full physical name. Typically, the SAS/SHARE server uses the BASE engine to access data; however, alternative engines can be assigned with the RENGINE= option, and attributes to the BASE or alternative engine can be supplied by the ROPTIONS= option. For details about the SAS/SHARE LIBNAME statement, see Chapter 9, “The LIBNAME Statement,” on page 103; for the SAS/CONNECT LIBNAME statement, see the *SAS/CONNECT User's Guide*. For information about the SAS/ACCESS LIBNAME statement, see *SAS/ACCESS Software for Relational Databases: Reference*.

SAS/SHARE Software Components

SAS/SHARE software consists of two procedures and an engine:

SERVER procedure

manages and performs input and output requests to SAS files on behalf of SAS client and clients that are not SAS applications.

OPERATE procedure

manages server, library, and client resources. For example, with PROC OPERATE, among the operations you can perform are: allocate a library to a server, free a library, stop a server, re-start a server, and display information about clients.

REMOTE engine

enables a client SAS session to access SAS data by means of a SAS/SHARE server. (The REMOTE engine is also licensed and distributed as part of SAS/CONNECT software.)

The REMOTE engine does not operate directly on files in a SAS data library; instead, it communicates with a SAS server. Each client runs a private copy of the REMOTE engine to communicate with a server.

A server uses one or more library engines or view engines to operate directly on files in SAS data libraries that clients access through the server. The server routes requests from the REMOTE engine to the appropriate engine for the SAS library or the file that the client accesses. A server's default library engine is specified by the ENGINE= system option. In Version 8, the V8 engine is used by default.

You may override the default and specify another engine when you start a server. Other engines are documented in the SAS documentation for your operating environment.

SAS/SHARE Users

SAS/SHARE software has various types of users, as described here:

end users

update or use concurrently accessed data in other ways. These users may not even be aware of SAS/SHARE software. They just need some basic information about how to update data sets and search for observations that they want to modify.

programmers or developers

write applications that use a server to access shared data. They need to know how to define a libref for a SAS data library that is accessed through a SAS/SHARE server and to be aware of how the server handles locking.

server administrators

start a server, manage its session, and evaluate its performance. They give a server access to SAS data libraries (taking into account library security) and supply clients with information to access these libraries. Server administrators then make a server accessible to clients who have permission to use it. At some sites, a server administrator might work with systems personnel to automate administrative functions and delegate certain functions to other personnel.

Using the SAS/SHARE and SAS/CONNECT Documentation Set

These books are used with SAS/SHARE and SAS/CONNECT software:

- ☐ *SAS/SHARE User's Guide*
- ☐ *SAS/CONNECT User's Guide*
- ☐ *Communications Access Methods for SAS/CONNECT and SAS/SHARE Software.*

For using SAS/SHARE or SAS/CONNECT software, you typically start with either the *SAS/SHARE User's Guide* or the *SAS/CONNECT User's Guide* and then refer to *Communications Access Methods for SAS/CONNECT and SAS/SHARE Software* for details about how to establish a connection between two hosts across a network. Furthermore, for information about using SAS/SHARE to access SAS files across hosts and SAS releases, see *Moving and Accessing SAS Files across Operating Environments*.

Version 6-Version 7 and Version 8 Compatibility

Accessing your data is a primary concern when migrating to a new version of the SAS software. Therefore, it is important to be aware of any restrictions in a cross-version environment.

Note: Version 7 and Version 8 features are virtually the same. Therefore, discussions of migration from Version 6 to Version 8 will also imply Version 7. Δ

Access to data depends on:

- ☐ the data object being accessed (data library, data file, view, or catalog).
- ☐ the version and release of SAS (Version 6 or Version 7/Version 8) being used
 - ☐ to create the data that is being accessed.
 - ☐ to develop the client application.
 - ☐ by the client session.
 - ☐ by the server session.

For information about Version 6-Version 7 and Version 8 compatibility issues for SAS/SHARE, see Appendix 6, "SAS/SHARE Version 6 to Version 7 and Version 8 Cross-Version Issues," on page 213.

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., *SAS/SHARE User's Guide, Version 8*, Cary, NC: SAS Institute Inc., 1999. pp. 247.

SAS/SHARE User's Guide, Version 8

Copyright © 1999 by SAS Institute Inc., Cary, NC, USA.

ISBN 1-58025-478-0

All rights reserved. Printed 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 the software by the government is subject to restrictions as 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, September 1999

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.

IBM®, AIX®, DB2®, OS/2®, OS/390®, RMT™, RS/6000®, System/370™, and System/390® are registered trademarks or trademarks of International Business Machines Corporation. ORACLE® is a registered trademark or trademark of Oracle Corporation. ® indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

The Institute is a private company devoted to the support and further development of its software and related services.