

Configuration Guide for SAS[®] 9.3 Foundation for Microsoft[®] Windows[®] for x64



Copyright Notice

The correct bibliographic citation for this manual is as follows: SAS Institute Inc., *Configuration Guide for SAS® 9.3 Foundation for Microsoft® Windows® for x64*, Cary, NC: SAS Institute Inc., 2012.

Configuration Guide for SAS® 9.3 Foundation for Microsoft® Windows® for x64

Copyright © 2012, SAS Institute Inc., Cary, NC, USA.

Some software included in SAS Foundation may display a release number other than 9.3.

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, by any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc. Limited permission is granted to store the copyrighted material in your system and display it on terminals, print only the number of copies required for use by those persons responsible for installing and supporting the SAS programming and licensed programs for which this material has been provided, and to modify the material to meet specific installation requirements. The SAS Institute copyright notice must appear on all printed versions of this material or extracts thereof and on the display medium when the material is displayed. Permission is not granted to reproduce or distribute the material except as stated above.

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.

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 trademarks of their respective companies.

Table of Contents

About this Guide	iv
Chapter 1: Post-Installation Setup for the Metabase Facility.....	1
Setting Up the System Repository Manager Files.....	1
Registering the SASHELP Repository in the Repository Manager	1
Converting Version 6 SAS/EIS Metabases to Version 8 Repositories	2
Chapter 2: Post-Installation Instructions for National Language Support (NLS)	3
Introduction	3
Selecting LOCALE During SASFoundation deployment	3
Chinese, Japanese, and Korean DBCS Support.....	4
Asian Font Catalogs	4
Installing Traditional Chinese Fonts	4
Specifying the Font Catalog in the Configuration File for Traditional Chinese Fonts	4
Specifying the Font Catalog in a SAS Session for Traditional Chinese Fonts	4
Chinese Localizations	4
Chapter 3: Post-Installation Setup for the SAS 9.3 Foundation.....	6
PATH Environment Variable in Windows XP	6
Other Windows Operating Systems.....	6
Chapter 4: Restricted Options	7
Global Restrictions	7
User Restrictions	7
Additional information.....	7
Chapter 5: Configuring Hadoop JAR Files	9
Updating Your Versioned Jar Repository	9
Install SAS-required Hadoop JAR Files on Your SAS Machines	9
Set the SAS_HADOOP_JAR_PATH Environment Variable	10
JAR Files for MapR-based Hadoop Systems	10
Supporting Multiple Hadoop Versions and Upgrading Hadoop Version	10
Chapter 6: Configuring SAS Analytics Accelerator for Teradata	12
Database Permission for Registering the UDFs.....	12
UDF Installation Step Requires LATIN1	12
Database Requirements and Configuration	12
Registering the UDFs.....	12
Alternative to PROC TSSQL.....	13
Re-enable the Default HTML ODS Destination.....	14
Documentation for Using the UDFs.....	14

Chapter 7: SAS/ACCESS Interfaces	15
Configuring SAS/ACCESS Interface to Aster nCluster	15
Installing and Configuring the ODBC Driver and Bulk Loader.....	15
Configuring SAS/ACCESS Interface to Greenplum	15
Bulkload.....	16
Registering SAS/ACCESS Interface to Greenplum.....	16
Configuring SAS/ACCESS Interface to Hadoop	17
Run the Hive Service	17
Data Integrity for Data Not in US-ASCII Format	17
Security Considerations.....	17
Read Access Security.....	17
Write Access Security	17
Default Hadoop HDFS Streaming and Hive ports	17
Successful SAS/ACCESS Connections.....	18
Unsuccessful SAS/ACCESS Connections	18
Starting with Hive	18
Proliferation of Hive Log Files in /tmp	19
Configuring SAS/ACCESS Interface to MySQL	19
Configuring SAS/ACCESS Interface to ODBC Software	20
Configuring SAS/ACCESS Interface to Oracle	21
Configuring SAS/ACCESS Interface to R/3	22
Configuring SAS/ACCESS Interface to Sybase	22
Configuring SAS/ACCESS Interface to Sybase IQ	22
Configuring SAS/ACCESS to Teradata.....	22
FastExporting	23
MultiLoad	23
Teradata Parallel Transporter.....	23
Configuring and Administering SAS In-Database Products.....	24
Chapter 8: Post-Installation Setup for SAS/ASSIST Software	25
Adding a Master Profile	25
Chapter 9: SAS/CONNECT Configuration.....	27
Storing and Locating SAS/CONNECT Script Files	27
TCP/IP	27
Configuring the SAS Windows Spawner Program	27
Chapter 10: Post-Installation Configuration for SAS/IntrNet Software	29
Overview	29
Installing and Configuring SAS/IntrNet Software	30
Install Your Web Server Software.....	30
Install Your SAS Software	30
CGI Tools Installation Dialogs.....	30
Installing CGI Tools and SAS Foundation on Machines with Different Operating Systems	32
Test the Web Server	33
Internet Information Services (IIS) 6.0+ Post-Install Configuration for Broker.....	33
Test the Application Broker.....	33
Configure a Default Application Dispatcher Service	34
Starting and Stopping the Default Service	35
Testing the Socket Service	35
Configure Additional Services.....	36

Chapter 11: Installing SAS/SECURE Software Client Components	37
SAS/SECURE for Windows Clients	37
SAS/SECURE for Java Clients	37
FIPS-Compliant Encryption.....	37
Chapter 12: Post-Installation Setup for SAS/SHARE Software.....	38
Using the TCP/IP Communications Access Method.....	38
System Configuration for the TCP/IP Access Method	38
Client-Side Components	39
SAS/SHARE Data Provider	39
SAS ODBC Driver	39
SAS/SHARE Driver for JDBC	39
SAS/SHARE SQL Library for C.....	39
NLS Information.....	39

About this Guide

This document provides configuration instructions for SAS 9.3 Foundation, which is made up of server-side Base SAS and a variety of server-side SAS products (the exact products vary by customer).

The server-side configuration instructions contained in this document are for the configuration of a generic SAS server. For information about configuring OLAP, workspace and stored process servers, please see *SAS 9.3 Intelligence Platform: Installation and Configuration Guide: Introduction to Installing the SAS Intelligence Platform*. This material can be found at <http://support.sas.com/documentation/cdl/en/biig/62611/HTML/default/p01intellplatform00installgd.htm>

Chapter 1: Post-Installation Setup for the Metabase Facility

In Version 7 of the SAS System, the SAS/EIS Metabase facility was converted to the V7 Common Metadata Repository. The Common Metadata Repository is a general-purpose metadata management facility that provides common metadata services to various metadata-driven applications.

Using the Common Metadata Repository requires a one-time setup of the system the repository manager. If the repository manager was set up in a previous release it may not need to be set up again. The steps in the following sections should be completed before you attempt to use the Metabase Facility. For Metabase Facility users who were using a release prior to Version 7, using the Common Metadata Repository requires a conversion. For more information, see the Version 8 OLAP Server Topic “Converting Legacy Metabases.”

Setting Up the System Repository Manager Files

Complete the following steps to set up the necessary system repository manager files. You must have write access to SASHELP in order to specify the system repository manager.

Note: *This process sets the default location for the repository manager for your site. Individual users may specify their own repository manager location by following the steps below and not selecting the Write values to system registry check box.*

1. Create a directory that will be dedicated exclusively to the storage of repository manager files, for example: !SASROOT\RPOSMGR.

This directory should not be used to store other SAS files.

2. At a SAS command line, type REPOSMGR and then select **Setup Repository Manager**.
3. In the Repository Manager Setup window, Library will default to RPOSMGR. For Path, specify the path from Step 1, above, and then select the **Write values to system registry** check box. Then select **OK**.
4. In the resulting dialog window, select **Yes** to generate the necessary repository manager files.

This completes the setup for the System Repository Manager. You can create additional repository managers (a user repository manager, for example) by repeating the steps above and by using a different path.

Registering the SASHELP Repository in the Repository Manager

The SASHELP repository is used in various samples, including the SAS/EIS Report Gallery templates. Before beginning the steps below, a repository manager must be created (see previous section). Complete the following steps to register the SASHELP repository in the Repository Manager:

1. At a SAS command line, type REPOSMGR and then select **Repository Registration**.
2. In the Repository Registration window, select **New**.
3. In the Register Repository (New) window, type SASHELP (in uppercase) in the **Repository** field, and then type the full directory path where the CORE catalog is located in the Path field, for example:

!SASROOT\CORE\SASHELP

4. In the **Description** field, you can type any character string (for example, SASHELP Repository). Select **OK** to close the Register Repository (New) window. Select **Close** to exit the Repository Registration window.

Note: *Repositories cannot span multiple directories because the path cannot contain concatenated directories. If you have existing metabases in concatenated directories, you should copy the metabases to a single path that will be referenced as a repository.*

Converting Version 6 SAS/EIS Metabases to Version 8 Repositories

For step-by-step instructions on converting Version 6 metabases to Version 8 repositories, please refer to the topic “Converting existing SAS/EIS metabases” in the SAS/EIS online Help.

Chapter 2: Post-Installation Instructions for National Language Support (NLS)

This chapter contains information on post-installation configuration for Asian and European language support.

Important: *To run a localized SAS image, your operating system must be configured to use a Windows for x64 regional setting appropriate to the localized image. You may experience unexpected results if the regional setting and the localized image are not consistent. If multiple SAS localizations have been installed on the system, you may need to change the regional setting for each localized SAS image before invocation. For information on how to change or to use regional settings, refer to your Microsoft Windows for x64 documentation.*

Introduction

When SAS 9.3 Foundation is installed, the installer creates a configuration file called `!SASROOT\sasv9.cfg`. When you open the file, you will find that the only option specified is the CONFIG option, which points to another configuration file. That configuration file contains the options used to establish your SAS session, including the LOCALE system option that establishes the NLS environment for SAS. Occasionally, other NLS options will also be included in the configuration file.

At least three NLS configuration files are created during the installation process:

- one in `!SASROOT\nls\en` for the default single-byte (SBCS) image that uses English messages
- one in `!SASROOT\nls\1d` for the default double-byte (DBCS) image that uses English messages
- one in `!SASROOT\nls\u8` for the SAS Unicode server

If you have chosen to install any of the localized SAS images, you will find a folder under `!SASROOT\nls` for each. The folders are named with a SAS specific 2-character language code. Each contains its own configuration file with a LOCALE option value set to match the localization. For example, the configuration file for French (fr) will have a LOCALE value with a French language code.

Selecting LOCALE During SASFoundation deployment

A new dialog was added to the SAS Deployment Wizard (SDW) in SAS 9.3 that allows the installer to select the locale. The locale that initially displays in this dialog is the user locale on the Windows machine where SASFoundation is being installed. If you prefer to use a different locale for SASFoundation, you can select a locale from dialog.

The selected locale is set as the value of the SAS LOCALE system option in the configuration file used to initialize SAS. If the selected locale matches a localization installed for the SASFoundation image, the `sasv9.cfg` file for that localization becomes the default configuration file for SAS. Otherwise, the default configuration file is the configuration file that uses English, which is one of these two files:

- `!SASROOT\nls\1d\sasv9.cfg` for any language that requires DBCS support
- `!SASROOT\nls\en\sasv9.cfg` for all other languages

For example, if the installer selects the French localization and the French Canadian locale, the LOCALE system option is set to `fr_CA` in both the French configuration file and the English configuration file. The French configuration file is used to initialize `SASFoundation`, by default.

Chinese, Japanese, and Korean DBCS Support

This section explains how to specify Asian font catalogs and how to determine the localization used for Chinese.

Asian Font Catalogs

The default configuration files for Asian language editions already contain font definitions. (However, the configuration file for DBCS extensions does not contain font definitions.) Asian fonts reside in the `SASHELP.FONT`s catalog. In SAS 9.3, **Simplified and Traditional Chinese fonts have been added.**

Asian font catalogs reside in subdirectories (by language) for easy installation. If you want to change the font catalog, you can specify it in either the configuration file or in your SAS session.

Installing Traditional Chinese Fonts

To use Traditional Chinese fonts, you must install the Traditional Chinese localization. You will also need to modify your configuration file, as described in the following section.

Specifying the Font Catalog in the Configuration File for Traditional Chinese Fonts

If you are not running a Traditional Chinese localization but you want to use the Traditional Chinese fonts, you need to modify the configuration file to set `GFONTx` to

```
-set gfontx !SASROOT/nls/zt/font-name
```

In this statement,

- `x` represents a value from 0-9.
- `font-name` represents the name of the font catalog you want to use.

Specifying the Font Catalog in a SAS Session for Traditional Chinese Fonts

To specify the font catalog in a SAS session, submit the following `LIBNAME` statement:

```
-libname gfontx !sasroot\nls\langcode\font-name
```

In this statement,

- `x` represents a value from 0-9.
- `font-name` represents the name of the font catalog you want to use.

Chinese Localizations

The installer has the option to install localizations for both Simplified Chinese and Traditional Chinese. Several Chinese-based locales are supported by SAS. In some cases, the localization selected for the locale may not be intuitive.

The following table shows which language SAS uses when you select one of the five Chinese locales. Note that the default language may be English.

Locale	Location of sasv9.cfg file	Language
Chinese (China) [zh_CN]	!SASROOT\nls\zh	Simplified Chinese
Chinese (Hong Kong) [zh_HK]	!SASROOT\nls\zt	Traditional Chinese
Chinese (Macau) [zh_MO]	!SASROOT\nls\1d	English
Chinese (Singapore) [zh_SG]	!SASROOT\nls\ \1d	English
Chinese (Taiwan) [zh_TW]	!SASROOT\nls\zt	Traditional Chinese

Chapter 3: Post-Installation Setup for the SAS 9.3 Foundation

If you add {SASHOME}\SASFoundation\9.3 to your system PATH, SAS 9.3 can run from a command console in any folder when you type SAS.

PATH Environment Variable in Windows XP

For Windows XP, The PATH Environment variable can be modified as follows:

1. Go to **Start -> Settings -> Control Panel**. In the Control Panel “classic view,” double-click the **System** icon.
2. Within the “System Properties” GUI, select the **Advanced** tab. Click the **Environment Variables** button (near the bottom of the GUI).
3. Under “System Variables,” scroll down and select **Path**.
4. Click **Edit**.
5. In the tiny “Variable Value:” window, scroll to the right and add the following to the end of the line:

```
;c:\program files\sas\sasfoundation\9.3\
```

Note: If you have installed your SAS software in a different location, use your actual sas.exe pathname instead.

Be sure to include the semicolon (;) to separate this path from all the other paths.

6. Click **OK, OK, and OK**.

Now that this folder is in your search PATH, you can easily execute any command in it by typing the name of that command (such as sas.exe) at the DOS command line prompt.

Other Windows Operating Systems

For information about setting the PATH Environment Variable for other versions and editions of the Windows Operating System, refer to Microsoft Windows Help.

Chapter 4: Restricted Options

SAS 9.3 Foundation options can be “restricted” by a site administrator so that once they are set by the administrator, they may not be changed by a user. An option can be restricted globally and by user. To restrict an option it must be added to the appropriate SAS 9.3 Foundation configuration file and this file must have the permissions set by the administrator so that it cannot be updated by users. The option files are processed by global and then by user. If an option is specified in multiple files then the last occurrence gets used.

Global Restrictions

Create the file `!SASROOT\rstropts\rsasv9.cfg` and add options to this file in the normal config file format.

User Restrictions

Create a file of the following format:

```
!SASROOT\rstropts\users\<user ID>_rsasv9.cfg
```

and add options to this file in the normal config file format.

Example:

For user smith, the file name is `smith_rsasv9.cfg`.

Additional information

To determine which options can be restricted, follow this example:

```
Proc Options LISTRESTRICT; run;
```

To verify that an option has been set correctly, follow this example:

1. Assume the option `-EMAILSYS=SMTP` was specified in one of the restricted configuration files.
2. Submit the following code:

```
proc options restrict; run;
```

The SAS log should then show a message similar to

```
Option Value Information For SAS Option EMAILSYS
```

```
Option Value: SMTP
```

```
Option Scope: SAS Session
```

```
How option value set: Site Administrator Restricted
```

The following describes the process when a user attempts to change the option value.

1. Assume the option `-NOTHEADS` was specified in one of the restricted configuration files.
2. Submit the following code:

```
options THREADS;
```

The SAS log should then show a message similar to the following:

```
options THREADS;
```

```
-----
```

```
36
```

```
WARNING 36-12: SAS option THREADS is restricted by your Site  
Administrator and cannot be updated.
```

Chapter 5: Configuring Hadoop JAR Files

Updating Your Versioned Jar Repository

If you add SAS/Access Interface to Hadoop to an existing installation, or if new Hadoop JAR files are added or updated, delete the existing `com.sas.app.launcher.cacheFile` file. The file is recreated when SAS is restarted and any new JAR files are discovered.

If default locations were chosen, the file can be found here:

```
C:\Program Files\SASHome\SASVersionedJarRepository\eclipse
```

This step clears the Versioned Jar Repository's cache, so updated versions of Hadoop JAR files (and other updated JAR files) in the VJR are used, rather than the old versions of those JAR files that have been cached.

Note: This instruction works if Hadoop JAR files have been updated in the Versioned Jar Repository, not in an alternate location.

Install SAS-required Hadoop JAR Files on Your SAS Machines

SAS components that access Hadoop require that Hadoop JAR files be copied from your Hadoop server onto the SAS machines in your organization that will access Hadoop. Create a directory on your SAS machine that is accessible to all SAS users. For older Hadoop releases (for example, Cloudera CDH3), copy the following Hadoop JAR files into that directory:

- `hive-exec`
- `hive-jdbc`
- `hive-metastore`
- `hive-service`
- `libfb303`
- `pig`
- `hadoop-core`

Newer Hadoop releases split `hadoop-core` into multiple JAR files. For newer Hadoop releases (for example, Cloudera CDH4), copy the following Hadoop JAR files into that directory:

- `hive-exec`
- `hive-jdbc`
- `hive-metastore`
- `hive-service`
- `libfb303`
- `pig`
- `guava`
- `hadoop-auth`
- `hadoop-common`
- `hadoop-hdfs`
- `protobuf-java`

Assistance from your Hadoop administrator may be required to locate the JAR files and network copy them to the SAS machine. Except for `libfb303`, these JAR files include version numbers. For example, on the Hadoop server, the pig JAR file might be `pig-0.8.0`, `pig-0.9.1`, or similar. Do not copy Thrift JAR files such as `libthrift` into the JAR directory.

Set the SAS_HADOOP_JAR_PATH Environment Variable

SAS must be able to find the JAR files. Create an operating environment variable named `SAS_HADOOP_JAR_PATH` as the directory path of the JAR files. For example, if the JAR files are copied to directory `C:\third_party\Hadoop\jars`, then the following DOS command sets the environment variable appropriately:

```
set SAS_HADOOP_JAR_PATH=C:\third_party\Hadoop\jars
```

Set `SAS_HADOOP_JAR_PATH` in a permanent manner for all SAS users who access Hadoop from this machine.

A `SAS_HADOOP_JAR_PATH` directory must not have multiple versions of a Hadoop `.jar`; this can cause unpredictable behavior in SAS.

Note: For SAS/ACCESS Interface to Hadoop to operate properly, your `SAS_HADOOP_JAR_PATH` directory must not contain any Thrift jars such as `libthrift*.jar`.

JAR Files for MapR-based Hadoop Systems

Along with the documented JAR files (`hive-*.jar`, etc.), you also need to point to the JAR files provided in the MapR client installation.

For example, set

```
SAS_HADOOP_JAR_PATH=C:\third_party\Hadoop\jars;C:\mapr\hadoop\hadoop-0.20.2\lib
```

where `C:\third_party\Hadoop\jars` is as described above, containing `hive-*.jars`, etc., and where `C:\mapr\hadoop\hadoop-0.20.2\lib` is the JAR directory laid down by the MapR client installation software.

In addition, SAS must be pointed to the MapR client installation directory that contains the MapRClient sharable library (for example, on Windows, the `MapRClient.dll`):

```
[SAS Invocation] -jreoptions (-Djava.library.path=C:\mapr\lib)
```

In most installations, the `-jreoptions` addition would be placed in the site-specific SAS configuration file.

Supporting Multiple Hadoop Versions and Upgrading Hadoop Version

The JAR files in the `SAS_HADOOP_JAR_PATH` directory must match the Hadoop server to which SAS connects. If you have multiple Hadoop servers running different Hadoop versions, then create and populate a separate directory with version-specific Hadoop JAR files on the SAS machine for each Hadoop version.

`SAS_HADOOP_JAR_PATH` then must be dynamically set depending on which Hadoop server a SAS job or SAS session will connect to. One means to dynamically set `SAS_HADOOP_JAR_PATH` is to create a wrapper script associated with each Hadoop version. SAS is invoked via a wrapper script

that sets `SAS_HADOOP_JAR_PATH` appropriately to pick up the JAR files that match the target Hadoop server.

Upgrading your Hadoop server version may involve having multiple Hadoop versions active. The same multi-version instructions apply.

Chapter 6: Configuring SAS Analytics Accelerator for Teradata

This chapter describes registering user-defined function (UDFs) that are included with your SAS Analytics Accelerator software in your Teradata database. The chapter assumes that you have successfully installed SAS Foundation, including SAS/STAT. To leverage Enterprise Miner and ETS UDFs, you also need to install SAS Enterprise Miner and SAS/ETS, respectively.

Database Permission for Registering the UDFs

Since the SAS Analytics Accelerator UDFs are going to be registered in database SYSLIB, the Teradata database user account you use to install the UDFs as described below must have the following privileges for the SYSLIB database.

- CREATE FUNCTION
- ALTER FUNCTION
- EXECUTE FUNCTION
- GLOP
- GLOP MEMBER

To obtain the appropriate permissions, contact your database administrator.

UDF Installation Step Requires LATIN1

SAS in-Database analytics procedures can be run with all supported encodings. However, the UDF installation step requires LATIN1 for the session encoding. If the system has a different encoding, it can be temporarily set to LATIN1 by changing the configuration file as described in *SAS 9.3 National Language Support (NLS): Reference Guide* at this location:
<http://support.sas.com/documentation/cdl/en/nlsref/63072/HTML/default/viewer.htm#titlepage.htm>.

Database Requirements and Configuration

To successfully install the UDFs and execute them, the Teradata database must be version 13.00 or higher. Additionally the database must have DBCEXTENSION installed to support operations on GLOP sets. Please contact your database administrator to ensure DBCEXTENSION has been installed prior to the next steps.

Registering the UDFs

To register the SAS Analytics Accelerator UDFs in your Teradata database, you should invoke three installation macros called `udftdstt.sas`, `udftdem.sas` and `udftdets.sas` that have been installed under `!SASROOT/stat/sasmacro/` as part of SAS Analytics Accelerator.

Note: During UDF installation, SAS creates temporary files under the folder referenced by the “work” library. If the absolute path of this folder is very long, UDF installation may fail due to the limit on external file names in Teradata Warehouse. If you encounter this situation, start the SAS session with the work library temporarily assigned to a directory with shorted path, e.g., “C:\” and proceed with the installation as described above. After installation is complete, the work library can be reassigned to

the original folder. For information about the options available when reassigning the work library, please refer to SAS documentation.

To run the macros, submit the following commands in the Program Editor from SAS:

```
ods html select none;
ods listing;

%let indconn = server=myserver user=myuserid password=XXXX
database=SYSLIB;

%udftdstt;
%udftdem;
%udftdets;

proc tssql nolibs noerrorstop noprompt="(&credentials.)";
  CREATE GLOP SET syslib.sas_vars;
  call DBCExtension.glop_add('syslib.sas_vars', 'SE', NULL,
'dmdb', 'N', 0, 'Y', 'M', 'E', 0, 256000, 1, '00'XB);
  CREATE GLOP SET syslib.sas_dmvars;
  call DBCExtension.glop_add('syslib.sas_dmvars', 'SE', NULL, 'dmine',
'N', 0, 'Y', 'M', 'E', 0, 256000, 1, '00'XB);
quit;
```

The first two statements in the program allow the registration macros to write and read temporary files without these files being redirected to the default HTML ODS destination of the SAS dms mode. The INDCONN macro variable provides credentials to connect to the Teradata machine. You must specify the server, user, password, and database to access the machine on which you have installed the Teradata data warehouse:

- *myserver* is the server on which the Teradata warehouse resides.
- *myuserid* is a valid user ID for that server that is granted the permissions described above.
- XXXX is the password for the user ID you are using.
- Because the SAS Analytics Accelerator UDFs must be registered in the SYSLIB database, database must be SYSLIB.

The statements executed by PROC TSSQL will create the GLOP sets and add the GLOPs used by SAS analytical procedures while executing the UDFs in Teradata database.

Alternative to PROC TSSQL

As an alternative to running PROC TSSQL, your database administrator may directly execute the following SQL commands on Teradata through the database client such as BTEQ.

```
CREATE GLOP SET syslib.sas_vars;
call DBCExtension.glop_add('syslib.sas_vars', 'SE', NULL,
'dmdb', 'N', 0, 'Y', 'M', 'E', 0, 256000, 1, '00'XB);
CREATE GLOP SET syslib.sas_dmvars;
call DBCExtension.glop_add('syslib.sas_dmvars', 'SE', NULL,
'dmine', 'N', 0, 'Y', 'M', 'E', 0, 256000, 1, '00'XB);
```

Re-enable the Default HTML ODS Destination

After the UDF registration macros have completed, you can re-enable the default HTML ODS destination by running the following command:

```
ods html select all;
```

Documentation for Using the UDFs

For information about how to use your newly registered UDFs, please see the *SAS Analytics Accelerator for Teradata: Guide* at

<http://support.sas.com/documentation/onlinedoc/analyticsaccel/index.html>.

Chapter 7: SAS/ACCESS Interfaces

Configuring SAS/ACCESS Interface to Aster nCluster

Installing and Configuring the ODBC Driver and Bulk Loader

Before you configure the ODBC driver, the bulk loader should be installed in `<SASHOME>\SASFoundation\9.3\` or somewhere that is in the PATH environment variable. After they have been installed, you can configure the ODBC data source. SAS recommends (but does not require) that you create a data source, using the instructions below. If you have more than one user who may use the same workstation to access the Aster nCluster database, you may choose to create it as a system DSN rather than a user DSN.

Complete the following steps to create a data source for Aster nCluster database access:

1. Run **ODBC Administrator** from the Control Panel or **Start** menu.
2. Select the **User DSN** or **System DSN** tab, depending on the number of users, as stated above.
3. Click **Add** and choose *Aster Data ODBC Driver for nCluster* from the pop-up window. Click **Finish**.
4. Enter any name in the **Data Source** field.
5. Enter the hostname or IP Address in the **Server** field.
6. Enter the port number in the **Port** field.
7. Enter the database name in the **Database** field.
8. Enter the user name in the **User Name** field.
9. Enter the password in the **Password** field.
10. Click **Test**. After a message appears that indicates that the test was successful, click **OK**.
11. Click **Save**.

Configuring SAS/ACCESS Interface to Greenplum

SAS/ACCESS Interface to Greenplum uses the DataDirect Technologies Greenplum Wire Protocol ODBC driver component, which should be downloaded from SAS. Please see the section entitled “SAS/ACCESS Interface to Greenplum Software” in the System Requirements documentation for information about the download. After you have received the ODBC driver, use the following steps to unzip it and put it in the appropriate location.

1. Use Windows Explorer to navigate to the `!SASROOT\access\sasmisc` directory where the `<platform>gplm60.zip` is located and where `<platform>` is specific to your operating system.
2. Unzip the `<platform>gplm60.zip` file.
3. Follow the instructions in the `Readme.txt` file contained in the zip file.

During the initial installation of SAS/ACCESS Interface to Greenplum, the SAS Deployment Wizard provides a dialog in which you can specify the location for the Microsoft SQL Server ODBC drivers that are required. If you want to update that location, use the SAS Deployment Manager. The procedure for updating the location is described in the “Configure SAS/ACCESS Interface to Greenplum” section of *SAS Deployment Wizard and SAS Deployment Manager 9.3: User’s Guide*, located at <http://support.sas.com/deploywizug93.html>.

After you download and install the software, configure a Data Source Name (DSN) to access a Greenplum database. You can use the ODBC administrator to add it. If you have more than one user who may use the same workstation to access the Greenplum database, you can create it as a system DSN instead of a user DSN.

Complete the following steps to create a data source for Greenplum database access:

1. Run **ODBC Administrator** from the Control Panel or Start menu.
2. Click the **User DSN** or **System DSN** tab, depending on the number of users, as stated above.
3. Click the **Add** button and choose **SAS ACCESS to Greenplum** from the pop-up window.
4. Enter any name in the **Data Source Name** field. This will be the name that you will use inside of SAS when referencing this database connection.
5. Enter a **Description** for this connection.
6. Enter the **Port Number** that the host is listening on for Greenplum connections. The default is 5432.
7. Enter the **Database Name** of the database on the Greenplum server.
8. Click **Test Connection** and enter your userid and password. A test connection will be initiated.
9. Click **OK** to close the ODBC Administrator and to save your new DSN entry.

Bulkload

SAS/ACCESS can interface with the Greenplum Client Loader interface for loading large volumes of data. To perform bulkloading, the Greenplum Client Loader Package must be present on the system where you installed SAS.

SAS recommends using the “gpfdist” protocol for bulkloading. For this protocol you must set the GPLOAD_HOME environment variable to point to the location where the gpfdist utility allocates the files to be loaded onto the Greenplum tables. Additional information about how to use the bulkloading feature can be found in the *SAS/ACCESS 9.3 for Relational Databases: Reference* documentation.

Registering SAS/ACCESS Interface to Greenplum

The following SAS procedure can be run so that the SAS/ACCESS Interface to Greenplum product is registered with the SAS system catalog:

```
PROC NICKNAME CAT=sashelp.core engine;  
  add nickname=greenplm module=sasiogpl desc="SAS/ACCESS to Greenplum"  
  preferred eng;  
quit;
```

Configuring SAS/ACCESS Interface to Hadoop

Note: Before you configure SAS/ACCESS Interface to Hadoop, please follow the instructions in “Configuring Hadoop JAR Files” on page 9 to set up the directory that contains Hadoop JAR files and the SAS_HADOOP_JAR_PATH variable.

Run the Hive Service

SAS/ACCESS reads Hadoop data via a JDBC connection to the Hive Service. The Hive Service typically is launched on the Hadoop Namenode. A good practice is to launch the Hive Service as a daemon that kicks off on system restarts, thereby assuring consistent service. For Hadoop administrators unfamiliar with Hive, a simple test before daemon setup is to start the Hive service at an operating system prompt. For example:

```
$ export HIVE_PORT=10000
$ HIVE_HOME/bin/hive --service hiveserver
```

Data Integrity for Data Not in US-ASCII Format

To avoid data integrity issues with data in UTF-8 format, install the HIVE-2137 patch for the hive-jdbc JAR file if it is not present. This patch should be installed in your SAS_HADOOP_JAR_PATH directory.

If your Hadoop data is not in US-ASCII or UTF-8 format, SAS/ACCESS cannot transparently convert the data for consumption by SAS. You will need to create a customized solution for this issue.

Security Considerations

Read Access Security

SAS/ACCESS supports standard USER= and PASSWORD= security options. SAS/ACCESS propagates the supplied username and password to the JDBC connection string passed to Hive. However, as of Hive 0.7.1, the Hive Service ignores JDBC user IDs and passwords. Permissions are instead those of the Linux user ID that started the Hive Service. The effect is that any SAS user can connect and have read access to all data accessible to the Hive Service. Hive JDBC security is being addressed by the Hadoop community.

Write Access Security

SAS/ACCESS leverages Hadoop HDFS Streaming to create a new Hive table or to append to an existing one. To write data, the SAS user name supplied must be a valid Linux user ID on the Hadoop server with write permission to both the Hadoop HDFS /tmp and the Hive warehouse directory. Hadoop HDFS Streaming currently ignores the password. SAS/ACCESS does not restrict the SAS user from specifying a user name that does not match that user's SAS machine login ID.

Default Hadoop HDFS Streaming and Hive ports

The default Hadoop HDFS Streaming port is 8020. The default Hive Service port is 10000. SAS/ACCESS honors these defaults. If you use the defaults, then SAS connection statements do not require additional override options, thus simplifying SAS code. If you override the defaults, you should communicate the appropriate port numbers to SAS users who will access Hadoop.

Successful SAS/ACCESS Connections

SAS code connects to Hive either with a LIBNAME or a PROC SQL CONNECT TO. The LIBNAME outputs information upon a successful connection whereas PROC SQL is silent on a successful connection. In these examples, Hive is listening on default port 10000 on Hadoop Namenode hadoop01.

Sample LIBNAME connection:

```
libname hdplib hadoop server=hadoop01 user=hadoop_usr  
password=hadoop_usr_pwd;
```

NOTE: Libref HDPLIB was successfully assigned as follows:

```
Engine:      HADOOP  
Physical Name: jdbc:hive://hadoop01:10000/default
```

Sample PROC SQL connection:

```
proc sql;  
connect to hadoop (server=hadoop01 user=hadoop_usr  
password=hadoop_usr_pwd);
```

Unsuccessful SAS/ACCESS Connections

SAS failure to connect can have different causes. Error messages will assist in diagnosing the issue.

In this sample failure, Hive is not active on port 10000 on Hadoop Namenode hadoop01.

```
libname hdplib hadoop server=hadoop01 port=10000 user=hadoop_usr  
password=hadoop_usr_pwd;
```

```
ERROR: java.sql.SQLException: Could not establish connecton to  
hadoop01:10000/default:
```

```
java.net.ConnectException: Connection refused: connect
```

```
ERROR: Unable to connect to server or to call the Java DriverManager.  
ERROR: Error trying to establish connection.  
ERROR: Error in the LIBNAME statement.
```

In this sample failure, the hive-metastore JAR file is missing from SAS_HADOOP_JAR_PATH.

```
libname hdplib hadoop server=hadoop01 port=10000 user=hadoop_usr  
password=hadoop_usr_pwd;
```

```
ERROR: java.lang.NoClassDefFoundError:  
org/apache/hadoop/hive/metastore/api/MetaException
```

```
ERROR: Unable to connect to server or to call the Java DriverManager.  
ERROR: Error trying to establish connection.  
ERROR: Error in the LIBNAME statement.
```

Starting with Hive

If you do not currently run Hive on your Hadoop server, then your Hadoop data likely resides in HDFS files initially invisible to Hive. To make HDFS files (or other formats) visible to Hive, a Hive CREATE TABLE is issued. The following simple scenario demonstrates how to access HDFS files from Hive using the Hive CLI. For more information, perform a Web search for "Hive CLI" and locate the appropriate Apache documentation.

1. Assume there are HDFS files `weblog1.txt` and `weblog2.txt` with data lines containing in order, a date field, a text integer field, and a string field. The fields are comma-delimited and lines `\n` terminated.

```
$ hadoop fs -ls /user/hadoop/web_data
Found 2 items

-rw-r--r--  3 hadoop [owner]      [size/date]
/user/hadoop/web_data/weblog1.txt

-rw-r--r--  3 hadoop [owner]      [size/date]
/user/hadoop/web_data/weblog2.txt
```

On the Hadoop Namenode, begin by terminating the Hive service if it is running. Next, at a Linux prompt, bring up the Hive CLI:

```
$ hive
```

2. At the Hive command prompt, make the weblogs visible to Hive:

```
hive> CREATE EXTERNAL TABLE weblogs (extract_date STRING,
extract_type INT, webdata STRING) ROW FORMAT DELIMITED FIELDS
TERMINATED BY ',' STORED AS TEXTFILE LOCATION
'/user/hadoop/web_data';
```

3. At the Hive command prompt, test that `weblog1.txt` is now accessible to Hive:

```
hive> SELECT * FROM weblogs LIMIT 1;
```

4. If the `SELECT` works, quit the Hive CLI and start the Hive Service on default port 10000. For example, if your namenode is `hadoop_cluster`, a test access from SAS would be as follows:

```
libname hdplib hadoop server=hadoop_cluster user=hadoop_usr
password=hadoop_usr_pwd;

data work.weblogs;
set hdplib.weblogs(obs=1);
put _all_;
run;
```

This is a complete but intentionally simple scenario intended for new Hive users. It is likely not representative of a mature Hive environment since the default Hive schema is used implicitly and the Hive default Derby metadata store may be in use. Consult Hadoop and Hive documentation to begin to explore Hive in detail. SAS/ACCESS user documentation provides more information on how SAS/ACCESS interacts with Hive.

Proliferation of Hive Logs Files in /tmp

Data access through Hive can create log files in Hadoop HDFS `/tmp`. Over time many log files can accumulate. You may want to disable the logging or periodically run a process to delete logs.

Configuring SAS/ACCESS Interface to MySQL

During the initial installation of SAS/ACCESS Interface to MySQL, the SAS Deployment Wizard provides a dialog in which you can specify the version of MySQL that you are using. If you want to update that information, use the SAS Deployment Manager. The procedure for updating the version is described in the “Configure SAS/ACCESS Interface to MySQL” section of *SAS Deployment Wizard*

and SAS Deployment Manager 9.3: User's Guide, located at <http://support.sas.com/deploywizug93.html>.

Before you can use SAS/ACCESS Interface to MySQL software, the following products are required:

- Base SAS software
- SAS/ACCESS Interface to MySQL software

Before you use SAS/ACCESS interface to MySQL software, verify that the MySQL client library (`libmysql.dll`) has been appended to Path, a system variable. The MySQL client library is often located in `C:\mysql\bin` or `C:\program files\mysql\mysql server 5.1\bin`.

To locate the Path system variable, follow the instructions for your operating system.

- For Windows NT, right-click **My Computer** and choose **Properties**. Select the **Environment** tab and you will see Path in your list of system variables.
- For Windows 2000, right-click **My Computer** and choose **Properties**. Select the **Advanced** tab and select the **Environment Variables** button. You will see Path in your list of System Variables.
- For Windows XP, click **Start** then select **Control Panel**. Double-click **System**. Select the **Advanced** tab and select the **Environment Variables** button. You will see Path in your list of System Variables.

For more information about SAS/ACCESS Interface to MySQL, refer to the MySQL chapter in *SAS/ACCESS 9.3 for Relational Databases: Reference*.

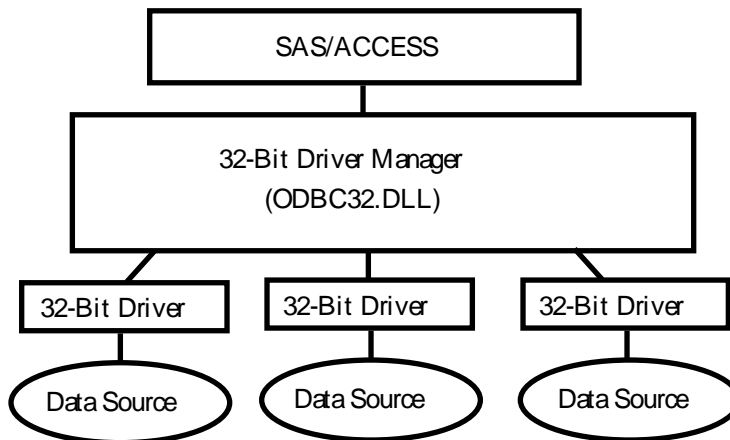
Configuring SAS/ACCESS Interface to ODBC Software

Before you can use SAS/ACCESS Interface to ODBC software, the following products are required:

- Base SAS Software
- SAS/ACCESS Interface to ODBC Software
- A 32-bit ODBC driver for the source from which you want to access data

The ODBC solution allows different technologies to be used by defining a standard interface. SAS/ACCESS Interface to ODBC software with ODBC drivers provides you with access to different databases. An ODBC driver processes ODBC function calls from--and returns results to--SAS/ACCESS software. You can get drivers from Microsoft, your database vendor, or other third party software vendors.

To run SAS/ACCESS software on 32-bit Windows, you need 32-bit drivers. The application calls the 32-bit Driver Manager (`ODBC32.DLL`), which in turn calls the 32-bit drivers. The following figure illustrates the architecture:



The ODBC Driver Manager and Administrator are Microsoft products that are included with all ODBC drivers. When you install an ODBC driver, the ODBC Driver Manager and Administrator are also installed. The ODBC Administrator should appear as an icon in the control panel. However, it can appear as an icon in the Start menu instead.

After you install the ODBC driver, you can use the ODBC Administrator to define and manage the data sources. A data source associates a particular ODBC driver with the data you want to access through that driver. It can consist of information about the data you want to access and its associated operating system, DBMS, and network platform (if any) used to access the DBMS. For information about how to configure your data sources, refer to the documentation provided with your ODBC drivers.

For more information about SAS/ACCESS to ODBC, refer to the ODBC chapter in *SAS/ACCESS 9.3 for Relational Databases: Reference*. For more information about ODBC, refer to *Microsoft ODBC 3.0 Programmer's Reference and SDK Guide*.

Configuring SAS/ACCESS Interface to Oracle

During the initial installation of SAS/ACCESS Interface to Oracle, the SAS Deployment Wizard provides a dialog in which you can specify the version of Oracle that you are using. If you want to update that information, use the SAS Deployment Manager. The procedure for updating the version is described in the "Configure SAS/ACCESS Interface to Oracle" section of *SAS Deployment Wizard and SAS Deployment Manager 9.3: User's Guide*, located at <http://support.sas.com/deploywizug93.html>.

Before SAS/ACCESS Interface to Oracle software can be used, the following products are required:

- Base SAS software
- SAS/ACCESS Interface to Oracle software
- Oracle client/server

After the Oracle client/server installation process is complete, ensure that the path to the Oracle client libraries is appended to the PATH system variable. The Oracle libraries are usually located in the directory ORACLE_HOME/bin. If the PATH variable does not have ORACLE_HOME/bin, you will see the following error:

ERROR: The SAS/ACCESS Interface to ORACLE cannot be loaded.

ERROR: Image SASORA found but not loadable.

For more information about SAS/ACCESS Interface to Oracle, refer to the Oracle chapter in *SAS/ACCESS 9.3 for Relational Databases: Reference*.

Configuring SAS/ACCESS Interface to R/3

SAS/ACCESS Interface to R/3 software requires extensive post-installation configuration before it can be used. For detailed information, refer to the *Configuration Instructions for SAS/ACCESS 4.4 Interface to R/3* on Install Center:

<http://support.sas.com/documentation/installcenter/en/ikr3cg/64225/PDF/default/config.pdf>

Configuring SAS/ACCESS Interface to Sybase

During the initial installation of SAS/ACCESS Interface to Sybase, the SAS Deployment Wizard provides a dialog in which you can specify the version of Sybase that you are using. If you want to update that information, use the SAS Deployment Manager. The procedure for updating the version is described in the “Configure SAS/ACCESS Interface to Sybase” section of *SAS Deployment Wizard and SAS Deployment Manager 9.3: User’s Guide*, located at <http://support.sas.com/deploywizug93.html>.

In SAS 9.3, the administrator or user must install two Sybase-stored procedures on the target server. Two files have been included in the !SASROOT\access\sasmisc directory to assist in the installation:

- sas-spcp.txt is a text file containing instructions on how to do the installation.
- sas-spdf.txt is the first of two actual stored procedure scripts for CTLIB 12.5x users
- sas-spdf_15.txt is the first of two actual stored procedure scripts for CTLIB 15 users
- sassp2df.txt is the second of two stored procedure scripts for CTLIB 12.5x users
- sassp2df_15.txt is the second of two stored procedure script for CTLIB 15 users.

The process utilizes two Sybase facilities, defncopy and isql.

To complete installation, you must add one of the following options to your sasv9.cfg file:

- For Sybase Open Client 15 users: -set sassyb sywx6150
- For Sybase Open Client 12.5 users: -set sassyb syw64125

Configuring SAS/ACCESS Interface to Sybase IQ

SAS/ACCESS Interface to Sybase IQ requires the installation of Sybase IQ client software, which is available from Sybase. On Windows for x64, there are no additional steps required to set up the environment and libraries.

Configuring SAS/ACCESS to Teradata

Before you use SAS/ACCESS Interface to Teradata software, verify connectivity by logging into your Teradata account with the Teradata BTEQ utility. If BTEQ is unavailable, establish connectivity as

documented in the *SAS/ACCESS to Teradata* white paper, available at <http://support.sas.com/resources/papers/teradata.pdf>.

If BTEQ fails to connect to the Teradata server, you may need an entry in the hosts file on your PC to provide the network address of the Teradata server. Typically, this means adding a dbccop1 entry to your hosts file. For more details about the entry, see your *Teradata Client for Windows Installation Guide*.

FastExporting

For optimal reads of large tables, SAS/ACCESS can perform *FastExporting*. To perform FastExporting, the Teradata FastExport Utility must be present on the system where you install SAS. You must also modify `Path`, a system variable. Append two directory paths to the end of `Path`:

1. The directory where `fexp.exe` (the FastExport Utility) is located. Often this is `C:\Program Files\Teradata\Client\13.0\bin`.
2. The directory where `sasaxsm.dll` is located (`sasaxsm.dll` is usually located in your SAS product tree, in the `!sasroot\access\sasexe` directory.)

The FastExport Utility is not required; SAS/ACCESS reads large tables quite efficiently without it. For further information, see the `DBSLICEPARM` option in your SAS/ACCESS to Teradata documentation. Contact Teradata if you want to obtain the Teradata FastExport Utility.

For more information about SAS/ACCESS Interface to Teradata software, please see the Teradata chapter in *SAS/ACCESS 9.3 for Relational Databases: Reference*, and the *SAS/ACCESS to Teradata* white paper, available on the SAS web site at <http://support.sas.com/resources/papers/teradata.pdf>.

MultiLoad

SAS/ACCESS can load large volumes of data to non-empty tables using MultiLoad. To perform MultiLoading, the Teradata MultiLoad Utility must be present on the system where you install SAS. You must also modify `Path`, a system variable. Append two directory paths to the end of `Path`:

1. The directory where `mload.exe` (the MultiLoad Utility) is located. Often this is `C:\Program Files\Teradata\Client\13.0\bin`.
2. The directory where `sasmlam.dll` and `sasmlne.dll` are located (`sasmlam.dll` and `sasmlne.dll` are usually located in your SAS product tree, in the `!sasroot\access\sasexe` directory.)

The MultiLoad Utility is not required; SAS/ACCESS provides other options for loading tables. For further information, see the `MULTISTMT` option in your *SAS/ACCESS to Teradata* documentation. Contact Teradata if you want to obtain the Teradata MultiLoad Utility.

Teradata Parallel Transporter

SAS/ACCESS supports the Teradata parallel transporter API for loading data using Multiload, Fastload and multi-statement inserts. The API also supports reading data using FastExport.

Note: *The Teradata Parallel Transporter API is not required; SAS/ACCESS provides other options for loading and reading data.*

If you plan to use the Teradata parallel transporter API, the following two requirements must be met:

1. The API must be installed on the system where SAS is installed.
2. The system variable TKPATHX8664 must be modified. If the variable does not exist already, it must be added.

To modify the variable, append the directory where `sasiotpt.dll` is located. (Usually `sasiotpt.dll` is located in your SAS product tree, in the `!sasroot\access\sasexe` directory.)

To locate the TKPATHX8664 system variable, follow the instructions for your operating system.

Configuring and Administering SAS In-Database Products

Deploying the SAS In-Database products requires detailed configuration and administration steps following the initial installation. Follow the steps in your Software Order E-mail and QuickStart Guide to perform the initial deployment. Then refer to the “Administrator's Guide” chapter of the *SAS In-Database Products: Administrator's Guide* for your particular database.

The “Administrator's Guide” chapter contains instructions on how to install and configure the in-database deployment package for your particular database. When you have completed the instructions described there, your software will be ready for use.

The SAS In-Database Products: Administrator's Guide is located at

<http://support.sas.com/documentation/onlinedoc/indbtech/index.html>.

Chapter 8: Post-Installation Setup for SAS/ASSIST Software

This chapter describes how to add an optional master profile to SAS/ASSIST software. You can use a master profile to override the default settings as sent by SAS. This allows you to provide a customized setup for SAS/ASSIST software. With the master profile, you can control the profile options of all SAS/ASSIST users from one central place. For information on the profile options, refer to *SAS/ASSIST Administrator's Guide*.

Adding a Master Profile

Complete the following steps to add a master profile to SAS/ASSIST software:

1. Specify the location of the master profile by creating a new directory that all users of SAS/ASSIST software will have read access to.

All users with write access to this directory will automatically also have write access to the master profile in SAS/ASSIST software. Select a name that conforms to the naming conventions at your installation. The name of this new directory must be stored in an entry in the SASHELP library. This requires that you have write access to the SASHELP library.

On line 1 of the Editor window of the SAS Display Manager System, type the physical pathname of the master profile directory. Execute the Save command to store this in the SASHELP.QASSIST catalog. For example:

```
SAVE SASHELP.QASSIST.PARMS.SOURCE
00001 S:\SAS\ASSIST\PARMS
00002
00003
```

The location of the master profile is now known by SAS/ASSIST software.

2. Create the master profile.

The first time SAS/ASSIST software is started, a master profile is created if SASHELP.QASSIST.PARMS.SOURCE contains the name of an existing physical pathname, and the person who starts SAS/ASSIST software has write access to this physical pathname.

3. Customize the master profile by starting SAS/ASSIST software and selecting **Setup > Profiles > Master/group...**

If you have write access to the SAS library containing the master profile, you can specify default values for your installation. New users will use these values as they start SAS/ASSIST software.

Note: If you restrict values by typing R in **Status**, users will not be allowed to change the values you define.

You can run SAS/ASSIST software in two different styles - workplace or block menu. The block menu can be new style or old style. You can control this using the profile options below.

Run workplace:

```
SAS/Assist style: Workplace
```

Run block menu new style:

SAS/Assist style:	Block Menu
Save selections on end:	Yes
Menu Style:	New

Run old style:

SAS/Assist style:	Block Menu
Save selections on end:	Yes
Menu Style:	Old

By setting the default values in the master profile, you can control if users should use the new or old style of SAS/ASSIST software. In addition, there are many other profile options. For more information on these options, refer to SAS/ASSIST Administrator's Guide.

4. Create group profiles.

From the master profile, it is possible to create group profiles to allow groups of users to have different setups. The master profile controls group profiles and user profiles when a user is not a member of any group. All users are indirectly controlled by the master profile when option values are set to a restricted (R) status.

From **Setup > Master/Group...**, select **Tools > Create Group Profile**. To add users to a group profile, select **Tools > Update User Group**. By default, the userid is found in the macro variable &SYSJOBID. This value is set in the option Userid in the master profile (option type System Administration). Change the value if your site uses another variable to keep the userid. If the value name starts with & , it is a macro variable; otherwise, it is an environment variable, which is set before the start of SAS 9.3.

Chapter 9: SAS/CONNECT Configuration

TCP/IP is the access method supported for SAS 9.3 on Windows for x64. For information about the access methods supported by other systems, refer to *Communications Access Methods for SAS/CONNECT 9.3 and SAS/SHARE 9.3*. This document can be found at <http://support.sas.com/documentation/onlinedoc/connect/index.html>

Note: The APPC access method is no longer supported for the Windows operating system.

Storing and Locating SAS/CONNECT Script Files

Several sample script files are shipped with SAS/CONNECT software. SAS/CONNECT software uses these script files to establish a connection to a remote SAS session.

The SASSCRIPT configuration option points to the location of the SAS/CONNECT script files. The SASSCRIPT option is used by SAS/ASSIST software and can be used by user-written SCL applications.

Under Windows, the script files are installed into the !SASROOT\CONNECT\SASLINK directory by default. The following line is added to the SASV9 .CFG file when SAS/CONNECT software is installed:

```
-SASSCRIPT !SASROOT\CONNECT\SASLINK
```

If you want to move the script files to another directory, you must edit the SASV9 .CFG file and update the SASSCRIPT option with the new directory location. This option can also be specified from the **Tools/Options/System/Communications/Networking and encryption** selection in DMSEXP mode.

TCP/IP

TCP/IP is the access method supported for SAS 9.3 on Windows for x64. SAS/CONNECT software supports Microsoft's TCP/IP System Driver, which is provided with Windows.

Configuring the SAS Windows Spawner Program

The SAS Windows Spawner is stored in the !SASROOT directory and can be executed manually from the !SASROOT directory at any time. You can run the SAS Windows Spawner as a service by executing SPAWNER.EXE with the -install option. By default, the SAS Windows Spawner will be installed to run with security. For complete documentation on the Windows Spawner and the supported options, see *Communications Access Methods for SAS/CONNECT 9.3 and SAS/SHARE 9.3*.

Note: If you are upgrading to SAS 9.3 and you run a spawner from a previous release as a service, you should stop and delete the existing spawner service. The SAS 9.3 CONNECT spawner can then be installed as a Windows service.

By default, when the SAS Windows Spawner is installed as a Windows service, it runs under the LocalSystem userid that has all required User Rights for running the SAS Windows Spawner.

If you do not install the SAS Windows Spawner as a Windows service (run it from your system prompt), the Windows userid that is used to start the SAS Windows Spawner must be the local Administrator of the machine and must have the following User Rights:

- act as part of the operating system
- bypass traverse checking (the default is everyone)
- increase quotas
- replace a process level token
- log on locally (the default is everyone)

The Windows user ID specified at sign on needs only the User Right “log on as a batch job.”

Chapter 10: Post-Installation Configuration for SAS/IntrNet Software

This chapter has information for your SAS/IntrNet installation. It will help you install, configure, and test your SAS/IntrNet components.

The procedures for installing SAS software using the SAS Deployment Wizard are described in other documentation and are not available in this chapter. In addition, the installation of a Web server is not described in SAS documentation because it is the administrator’s responsibility.

When SAS/IntrNet software has been installed, configured and tested using the procedures described in this chapter, review the latest version of the SAS/IntrNet product documentation online at <http://support.sas.com/documentation/onlinedoc/IntrNet/index.html>. The “What’s New” page at this Web site lists any recent changes to the product or documentation.

Overview

All SAS/IntrNet installations are made up of two components:

1. The SAS/IntrNet server (also referred to as *the Application Server*). This is where SAS Foundation is installed.
2. CGI Tools (also referred to as *the Broker*). This is where the `broker.cfg` file and its supporting files are installed.

When you install SAS/IntrNet, choose between these two installation configurations:

- Type A - The SAS/IntrNet server and CGI Tools components are both installed on the same system machine. The Web server must be installed before starting the SAS installation.
- Type B -The SAS/IntrNet server component is installed on one system machine and the CGI Tools component is installed on a different system machine. The Web server must be installed on the CGI Tools system machine prior to installing CGI Tools.

Type A and Type B require different installation steps:

Type A Installation Steps	Type B Installation Steps
Confirm that the Web server software (IIS, Apache etc.) is on the same system machine as your SAS/IntrNet software.	Confirm that the Web server software (IIS, Apache, etc.) is on the system machine where you will install CGI Tools.
Install your SAS products. Check “CGI Tools for the Web Server” in the “Select Products to Install” menu.	On your application server system machine, start your SAS installation. Uncheck “CGI Tools for the Web Server” in the “Select Products to Install” menu.
	On your Web server system machine, start your SAS installation. Uncheck all products except “CGI Tools for the Web Server” in the “Select Products to Install” menu. You can optionally check the IntrNet

Type A Installation Steps	Type B Installation Steps
	Monitor or Connect Drivers.
Test the Broker	
Configure a Socket Service	
Start the Socket Service	
Test the Socket Service	

Refer to your Web server’s documentation for its installation procedures.

Installing and Configuring SAS/IntrNet Software

Install Your Web Server Software

Refer to your Web server’s documentation for its installation procedures.

Install Your SAS Software

Refer to your *QuickStart Guide* for a description of how to start your SAS software installation.

If you are performing a Type A installation (as described in the “Overview” above), perform the following steps:

1. Confirm that your Web server software is installed before starting your SAS software installation.
2. Check “CGI Tools for the Web Server” in the **Select Products to Install** dialog.

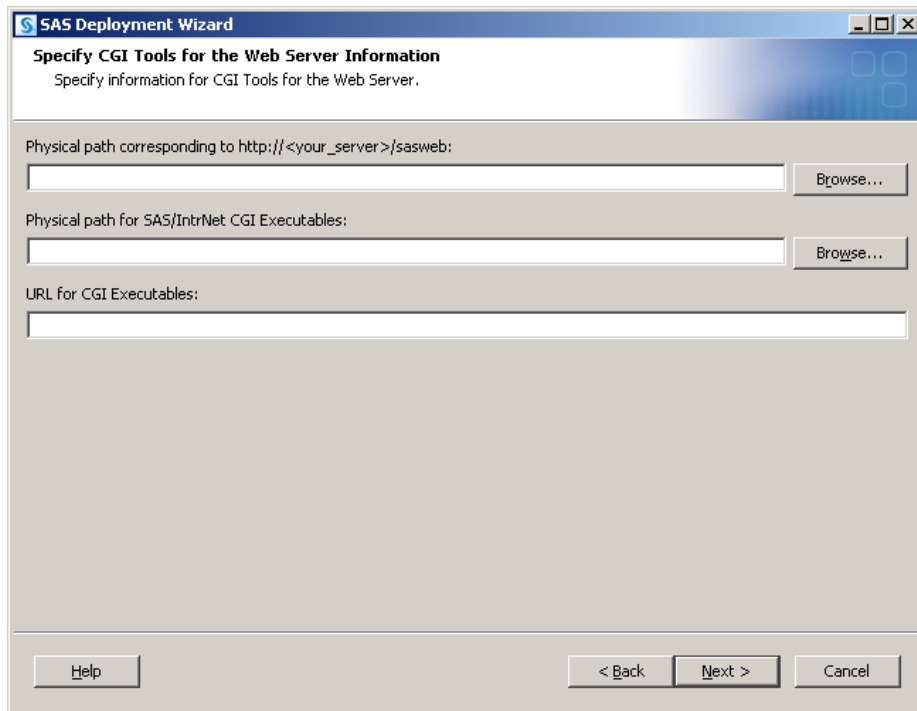
If you are performing a Type B installation (as described in the “Overview” above), do **both** of the following:

1. Install the SAS software on the SAS System side, unchecking “CGI Tools for the Web Server” in the **Select Products to Install** dialog.
2. Start the SAS software install on the Web server and check “CGI Tools for the Web Server” in the **Select Products to Install** dialog. SAS/IntrNet Monitor and SAS/CONNECT Driver for Java are optional selections. Uncheck everything else.

CGI Tools Installation Dialogs

The following windows appear for CGI Tools for the Web Server for all installations. For information about the fields, click **Help** on any dialog.

Customary entries are documented following each screen shown below. Customize the entries according to your environment.



The following are examples of common entries for popular Web servers. Customize your entries according to your own Web server environment. These fields will tell SAS where your Web server software is located.

Physical path corresponding to `http://<your_server>/sasweb`:

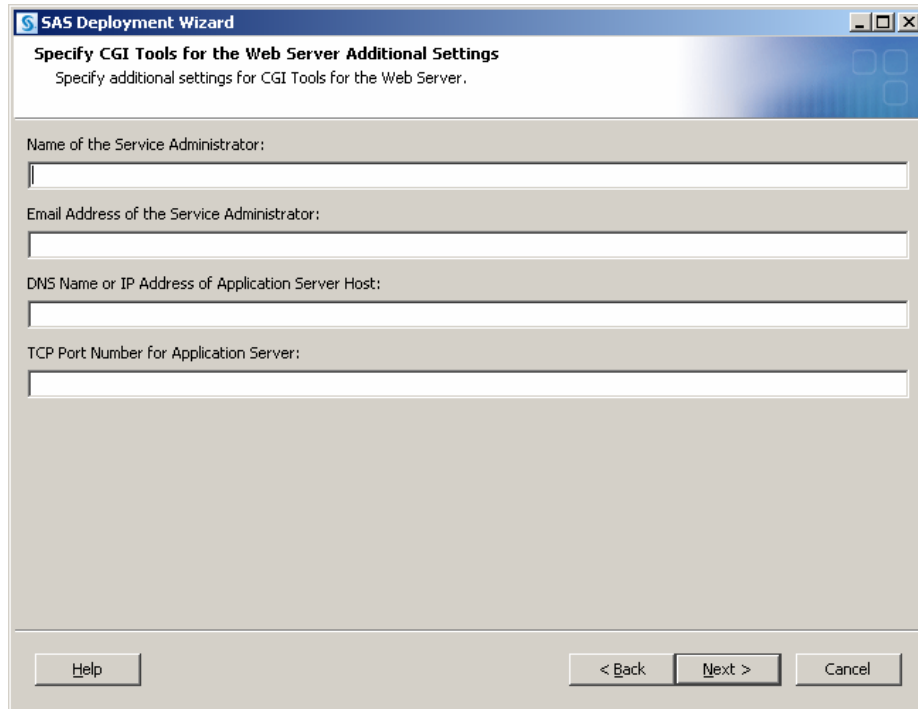
- IIS: `C:\Inetpub\wwwroot\sasweb`
- Apache (Windows): `C:\program files\Apache Software Foundation\Apache2.2\htdocs\sasweb`
- Apache (UNIX): `/usr/local/apache2/htdocs/sasweb`

Physical path for SAS/IntrNet CGI Executables:

- IIS: `C:\Inetpub\scripts`
- Apache (Windows): `C:\program files\Apache Software Foundation\Apache2.2\cgi-bin`
- Apache (UNIX): `/usr/local/apache2/cgi-bin`

URL for CGI Executables:

- IIS: `http://<web_servername>/scripts`
Example: `http://abcserver.comp.com/scripts`
- Apache (Windows): `http:// <web_servername>/cgi-bin`
Example: `http://abcserver.comp.com/cgi-bin`
- Apache (UNIX): `http://<web_servername>/cgi-bin`
Example: `http://abcserver.comp.com/cgi-bin`



Note that your entries for this dialog are added to the `broker.cfg` file. The `broker.cfg` file is a text file that can be edited after the installation is complete.

Name of the Service Administrator:

(optional) Enter the name of the administrator (for example, John Doe).

Email Address of the Service Administrator:

(optional) Enter the e-mail address of the administrator (for example, NetAdmin@comp.com).

DNS Name or IP Address of Application Server Host:

Enter the DNS name or IP address of the application server host where SAS Foundation is located.

TCP Port Number for Application Server:

The customary default port number is 5001, but you can use any valid available port on your system between 256 – 65535.

Installing CGI Tools and SAS Foundation on Machines with Different Operating Systems

Your SAS Foundation system’s operating system might be different than your CGI Tools system’s operating system. For example, your SAS Foundation might be installed on a Windows system and your CGI Tools might be installed on a UNIX system. The CGI Tools installation from the SAS Deployment Wizard will detect the destination operating system and install the appropriate operating system-specific software.

There are two methods to make the SAS Software Depot available to the installer on the destination CGI Tools system. Choose the method that is appropriate for your site’s available facilities:

- Launch the set-up from a SAS Software Depot that resides on a remote system. You might need to use a cross-platform file access method such as NFS or SAMBA to connect the two systems.
- Create media from an existing depot using the SAS Deployment Wizard and use that media on the host machine. This process is described more thoroughly in the *SAS Deployment Wizard and SAS Deployment Manager 9.3: User's Guide*, available from Install Center, <http://support.sas.com/documentation/installcenter/en/ikdeploywizug/64204/PDF/default/user.pdf>.

Note: SAS/IntrNet operation requires TCP/IP connectivity between the SAS Foundation system and the CGI Tools system, regardless of the operating systems on which these components are installed.

Test the Web Server

To determine if the Web server is running, launch the Web server's browser and enter `http://localhost`. This will return a Web page if the Web server is running.

If you do not receive a Web page, you must debug or reinstall your Web server before continuing.

Internet Information Services (IIS) 6.0+ Post-Install Configuration for Broker

In IIS 6.0 and later, you need to take manual steps in order for IIS to recognize the Broker as an application. For more information, please visit <http://support.sas.com/kb/10/496.html>.

Test the Application Broker

To verify that CGI Tools was installed correctly and can access the `broker.cfg` file, point your Web browser to the following URL:

Windows:

IIS: `http://<web_servername>/scripts/broker.exe`

Apache: `http://<web_servername>/cgi-bin/broker.exe`

Other hosts:

`http://your_webserver/cgi-bin/broker`

Replace *your_webserver* with the name of the Web server. The URL path might also need to be changed if you installed CGI Tools to a different directory. You should see a Web page similar to the following:

```
SAS/IntrNet Application Dispatcher
Application Broker Version 9.3 (Build 1495)
Application Dispatcher Administration
SAS/IntrNet Samples
SAS/IntrNet Documentation - requires Internet access
```

If you do not receive this page, you must debug your Web server installation before continuing. Verify that your Web server is enabled for CGI execution in the directory where you installed the Application Broker (`broker.exe` and `broker.cfg` files). This directory was determined by what

was entered for **Physical path for SAS/IntrNet CGI Executables** during the CGI Tools installation above.

Configure a Default Application Dispatcher Service

You must create a default Application Dispatcher service to run any of the sample programs supplied with SAS/IntrNet. You should create the default service with the SAS Deployment Manager which is accessible from the Start menu.

Perform the following steps to create and start the default service:

1. From the Start menu, select **Programs** ➔ **SAS** ➔ **SAS Deployment Manager**.
2. When the Choose Language window comes up, select the desired language and click **OK**.
3. The SAS Deployment Manager (SDM) will display the Select SAS Deployment Manager Task window. Under **SAS/IntrNet Service Tasks** select **Create Socket Service** and click **Next** to continue.
4. The Specify Service Name window displays. The default value for the Service Name field is `default`. Create this as your first service because this is what is used when you run the samples. Click **Next** to continue.
5. The Specify Service Directory window displays. The SDM selects a default service root directory based on the location that you chose for user files when you installed SAS Software. This default location is recommended for most users, although you can use the **Browse** button to select a different directory. Click **Next** to continue.
6. The Specify Service Ports window displays. Type the TCP/IP port number that you reserved for the default Application Dispatcher service. Click **Next** to continue.
7. The Specify Administrator Password window displays. A password is not necessary for the default service. You can add an administrator password later if you use this service for production applications. Click **Next** to continue.
8. The Summary window displays. It indicates `Stage 1: Create Socket Service`.
 - If the information you have entered is correct, click **Start**.
 - If you want to go back and change any of the values you entered, you can step backward through the windows by clicking the **Back** button.
9. The In Progress window displays while SDM creates the service.
10. The Deployment Complete screen displays when the task is finished.
 - If the service is created correctly, a green checkmark appears next to the **1. SAS/IntrNet** under **Stage 1: Create Socket Service**.
 - If there was a problem, a yellow exclamation or red X appears and you should check the log for a description of the problem. The logs reside in `<SASHOME>\SASFoundation\9.3\intrnet\sasmisc`.
11. From the Start menu, select **Programs** ➔ **SAS** (or other program group where SAS is installed) ➔ **IntrNet** ➔ **default Service** ➔ **Start Interactively**. Your default Application Server should now be running.

Starting and Stopping the Default Service

As stated above, from the Start menu, select Programs ➤ SAS (or other program group where SAS is installed) ➤ IntrNet ➤ default Service ➤ Start Interactively. Your default Application Server should now be running.

Testing the Socket Service

1. To make sure that the service was installed and started correctly, point your Web browser to this URL:

Windows:

IIS: `http://<web_servername>/scripts/broker.exe`

Apache: `http://<web_servername>/cgi-bin/broker.exe`

Other hosts:

`http://your_webserver/cgi-bin/broker`

Replace *your_webserver* with the name of the Web server. The URL path might also need to be changed if you installed the Application Broker to a different directory. You should see the following Web page: SAS/IntrNet Application Dispatcher

Application Broker Version 9.3 (Build 1495)

Application Dispatcher Administration

SAS/IntrNet Samples

SAS/IntrNet Documentation - requires Internet access

2. Click on the **Application Dispatcher Administration** link to see if the Application Broker can read the `broker.cfg` file. The Application Dispatcher Services Web page should open.
3. Verify connectivity between the Application Server and the Web server. Click on the **Application Dispatcher Administration** link and then click on the **ping** link under **SocketService default** heading. If the ping is successful, you should see:

Ping. The Application Server `<hostname>:<port_number>` is functioning properly.

4. To complete installation testing, type this URL in your browser address line:

Windows:

IIS:

`http://your_webserver/scripts/broker.exe?_service=default&_program=sample.webhello.sas`

Apache:

`http://your_webserver/cgi-bin/broker.exe?_service=default&_program=sample.webhello.sas`

Other hosts:

`http://your_webserver/cgi-bin/broker?_service=default&_program=sample.webhello.sas`

You should see the string “Hello World!” in large bold type in your browser. If you do not, add the debug option to create a log:

Windows:

IIS:

```
http://your_webserver/scripts/broker.exe?_service=default&_program=sample.webhello.sas&_debug=131
```

Apache:

```
http://your_webserver/cgi-bin/broker.exe?_service=default&_program=sample.webhello.sas&_debug=131
```

Other hosts:

```
http://yourWebserver/cgi-bin/broker?_service=default&_program=sample.webhello.sas&_debug=131
```

Save the log screen on the browser for SAS Technical Support.

Configure Additional Services

This chapter describes how to configure only a simple default Application Dispatcher service. There are many reasons you may want to configure additional services, including segregating applications by security or performance requirements and implementing more scalable servers. See the “Using Services” section of the SAS/IntrNet Application Dispatcher documentation at <http://support.sas.com/documentation/onlinedoc/intrnet/index.html> information on configuring additional services, using the Load Manager, and adding pool services.

Chapter 11: Installing SAS/SECURE Software Client Components

SAS/SECURE software includes client components that non-SAS System client applications can use to communicate with a SAS server in a secure environment. To use encryption between a non-SAS System client and a SAS server with SAS/SECURE software licensed, you must install the SAS/SECURE client components on the client machine. When SAS/SECURE software is included in a customer's order, these SAS/SECURE client components can be selected for installation via the SAS Deployment Wizard.

Note: This installation is not necessary if the SAS System is your client. The SAS System installs the components that it needs as part of the SAS System installation process.

SAS/SECURE for Windows Clients

The SAS/SECURE components needed by Windows clients can be installed by running the SAS Deployment Wizard to access the SAS/SECURE Windows Client Components.

SAS/SECURE for Java Clients

The SAS/SECURE components for Java clients provide encryption support for Java applications. You can incorporate this support into applications that are written using the following components:

- SAS/SHARE driver for JDBC
- SAS/CONNECT driver for Java
- IOM Bridge for Java

If the SAS/SECURE Java Component was part of your software order, the SAS/SECURE components needed by Java clients will be installed automatically when you run the SAS Deployment Wizard.

FIPS-Compliant Encryption

FIPS stands for Federal Information Processing Systems, and its 140-2 standard defines the security requirements for cryptographic modules. The 140-2 standard is detailed in the following document: <http://csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf>.

The ENCRYPTFIPS option is an option added to SAS/CONNECT spawners, IOM spawners, and Base SAS so that the communication encryption will be compliant with FIPS 140-2. The ENCRYPTFIPS option is required for FIPS-compliant encryption.

If you use ENCRYPTFIPS, the NETENCALG option must be set to AES or SSL. In Windows XP and in later versions of Windows, you need to enable the System cryptography: Use FIPS-compliant algorithms for encryption, hashing, and signing setting under your Local Security Policy or as part of Group Policy. This setting informs applications that they should use only those cryptographic algorithms that are FIPS 140-2 compliant and in compliance with FIP- approved modes of operation.

For more information about FIPS and encryption in general, refer to Encryption in SAS 9.3, available at <http://support.sas.com/documentation/onlinedoc/base/index.html#base93>.

Chapter 12: Post-Installation Setup for SAS/SHARE Software

This chapter discusses the use of the TCP/IP access methods with SAS/SHARE software. TCP/IP is the access method supported for SAS 9.3 on Windows. For information about the access methods supported by other systems, please visit <http://support.sas.com/documentation/onlinedoc/> and refer to *Communications Access Methods for SAS/CONNECT 9.3 and SAS/SHARE 9.3*.

Using the TCP/IP Communications Access Method

Communication between a SAS/SHARE server and user is handled by the TCP/IP communications access method, which is a part of SAS 9.3 that uses underlying communications software to exchange messages and data. The TCP/IP access method requires Microsoft's TCP/IP Network Protocol communications software.

Note: The APPC access method is no longer supported for the Windows operating system.

To use the TCP/IP access method, you must have the supporting software on each workstation on which a SAS/SHARE server or user will execute.

System Configuration for the TCP/IP Access Method

SAS/SHARE software supports Microsoft's TCP/IP Network Protocol, which is provided with Windows for x64.

To define server names in the TCP/IP SERVICES file, complete the following steps:

1. Locate the SERVICES file.

This file is located under the \windows directory path depending on the specific Windows operating system and upgrade method. For example, under a machine configured with a Windows 7 operating system, the directory path is named

```
<drive letter>:\Windows\System32\drivers\etc
```

2. Specify the server names and port assignments.

Each SAS/SHARE server that runs on a network must be defined as a service in the SERVICES file. Each entry in this file associates a service name with the port number and protocol used by that service. An entry for a SAS/SHARE server has the form:

```
<server name> <port number>/<protocol> # <comments>
```

The server name must be 1-8 characters in length. The first character must be a letter or underscore. The remaining seven characters can include letters, digits, underscores, the dollar (\$) sign, or the at (@) sign. The port number must be above 1024, as any port number equal to or less than 1024 is reserved. The protocol must always be TCP.

An entry for a server with the name MKTSERV might look like this:

```
mkt serv      5000/tcp    # SAS server for Marketing and Sales
```

The server name is specified with the SERVER= option in the PROC SERVER statement in the server's SAS session and in the PROC OPERATE and LIBNAME statements in user and server

administrator programs. If the server name is not defined, you must specify “__<port#>”, two consecutive underscores followed by the port number (e.g., server=__5012).

For more information about the options used with PROC SERVER and PROC OPERATE procedures, please refer to the *SAS/SHARE 9.3 User's Guide*.

Client-Side Components

SAS/SHARE software includes client components that are used outside of your SAS installation. These components are available from the SAS 9.3 Software Download site and are described below:

SAS/SHARE Data Provider

The SAS/SHARE data provider enables you to access, update, and manipulate SAS data using OLE DB- and ADO-compliant applications on Windows platforms.

SAS ODBC Driver

The SAS ODBC driver enables you to access, update, and manipulate SAS data from ODBC-compliant applications on Windows platforms.

SAS/SHARE Driver for JDBC

The SAS/SHARE driver for JDBC enables you to write applets, applications, and servlets that access and update SAS data. The Java Tools package that includes the SAS/SHARE driver for JDBC also includes the SAS/CONNECT driver for Java. If you are writing Java programs using these interfaces, you may also want to use the tunnel feature. This optional feature can be used with the Java applets you write to solve some common configuration problems.

SAS/SHARE SQL Library for C

The SAS SQL Library for C provides an application programming interface (API) that enables your applications to send SQL queries and statements through a SAS/SHARE server to data on remote hosts.

NLS Information

Sites that develop or support international applications that use SAS/SHARE software should refer to the chapter Post-Installation Instructions for National Language Support (NLS) on page 3.



THE
POWER
TO KNOW.

SAS is the leader in business analytics software and services, and the largest independent vendor in the business intelligence market. Through innovative solutions delivered within an integrated framework, SAS helps customers at more than 50,000 sites improve performance and deliver value by making better decisions faster. Since 1976, SAS has been giving customers around the world THE POWER TO KNOW®.