



THE
POWER
TO KNOW.

SAS[®] 9.2 Companion for UNIX Environments



The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2009. *SAS® 9.2 Companion for UNIX Environments*. Cary, NC: SAS Institute Inc.

SAS® 9.2 Companion for UNIX Environments

Copyright © 2009, SAS Institute Inc., Cary, NC, USA

ISBN 978-1-59994-792-1

All rights reserved. Produced in the United States of America.

For a hard-copy book: 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.

For a Web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

U.S. Government Restricted Rights Notice. Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions 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 electronic book, February 2009

2nd electronic book, May 2010

1st printing, March 2009

SAS® Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at support.sas.com/publishing or call 1-800-727-3228.

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 registered trademarks or trademarks of their respective companies.

Contents

<i>What's New</i>	<i>ix</i>
Overview	ix
Direct File I/O Options	ix
File Locking	ix
SAS Remote Browser	x
Multiple Work Directories	x
Installation and Configuration Changes	x
64-bit Architecture	x
SAS Logging Facility	x
Shared Executable Libraries	xi
SAS Language Elements	xi
IPv6 Standard	xiii
Documentation Enhancements	xiii

PART 1 **Running SAS Software Under UNIX** **1**

Chapter 1 \triangle Getting Started with SAS in UNIX Environments	3
Starting SAS Sessions in UNIX Environments	4
Running SAS in a Foreground or Background Process	6
Selecting a Method of Running SAS in UNIX Environments	7
SAS Windowing Environment in UNIX Environments	7
Interactive Line Mode in UNIX Environments	9
Batch Mode in UNIX Environments	10
Running SAS on a Remote Host in UNIX Environments	11
X Command Line Options	13
Executing Operating System Commands from Your SAS Session	14
Customizing Your SAS Registry Files	17
Customizing Your SAS Session by Using System Options	17
Customizing Your SAS Session by Using Configuration and Autoexec Files	20
Defining Environment Variables in UNIX Environments	23
Determining the Completion Status of a SAS Job in UNIX Environments	24
Exiting or Interrupting Your SAS Session in UNIX Environments	24
Ending a Process That Is Running as a SAS Server	27
Ending a SAS Process on a Relational Database	27
Chapter 2 \triangle Using SAS Files	31
Introduction to SAS Files, Libraries, and Engines in UNIX Environments	33
Common Types of SAS Files in UNIX Environments	34
Filename Extensions and Member Types in UNIX Environments	36
Using Direct I/O	37
Holding a File in Memory: The SASFILE Statement	38

Sharing SAS Files in a UNIX Environment	38
Compatible Computer Types in UNIX Environments	42
Migrating 32-Bit SAS Files to 64-Bit in UNIX Environments	44
Creating a SAS File to Use with an Earlier Release	45
Reading SAS Files from Previous Releases or from Other Hosts	46
Referring to SAS Files by Using Librefs in UNIX Environments	47
Specifying Pathnames in UNIX Environments	50
Assigning a Libref to Several Directories (Concatenating Directories)	51
Using Multiple Engines for a Library in UNIX Environments	53
Using Environment Variables as Librefs in UNIX Environments	53
Librefs Assigned by SAS in UNIX Environments	54
Sasuser Library	55
Work Library	58
Multiple Work Directories	59
Using One-Level Names to Access Permanent Files (User Library)	59
Accessing Disk-Format Libraries in UNIX Environments	60
Accessing Sequential-Format Libraries in UNIX Environments	60
Accessing BMDP, OSIRIS, or SPSS Files in UNIX Environments	62
Support for Links in UNIX Environments	66
Chapter 3 \triangle Using External Files and Devices	67
Introduction to External Files and Devices in UNIX Environments	68
Accessing an External File or Device in UNIX Environments	69
Specifying Pathnames in UNIX Environments	70
Assigning Filerefs to External Files or Devices with the FILENAME Statement	73
Concatenating Filenames in UNIX Environments	75
Assigning a Fileref to a Directory (Using Aggregate Syntax)	75
Using Environment Variables to Assign Filerefs in UNIX Environments	76
Filerefs Assigned by SAS in UNIX Environments	77
Reserved Filerefs in UNIX Environments	78
Sharing External Files in a UNIX Environment	79
Reading from and Writing to UNIX Commands (PIPE)	79
Sending Electronic Mail Using the FILENAME Statement (EMAIL)	82
Processing Files on TAPE in UNIX Environments	87
Chapter 4 \triangle Printing and Routing Output	91
Overview of Printing Output in UNIX Environments	92
Previewing Output in UNIX Environments	92
The Default Routings for the SAS Log and Procedure Output in UNIX Environments	93
Changing the Default Routings in UNIX Environments	93
Routing SAS Logging Facility Messages to SYSLOGD	95
Using the Print Dialog Box in UNIX Environments	95
Using Commands to Print in UNIX Environments	97
Using the PRINTTO Procedure in UNIX Environments	100
Using SAS System Options to Route Output	102

Printing Large Files with the PIPE Device Type in UNIX Environments	102
Changing the Default Print Destination in UNIX Environments	103
Changing the Default Print Command in UNIX Environments	103
Controlling the Content and Appearance of Output in UNIX Environments	104

Chapter 5 △ **Accessing Shared Executable Libraries from SAS** 107

Overview of Shared Libraries in SAS	108
The SASCBTBL Attribute Table	109
Special Considerations When Using Shared Libraries	114
Examples of Accessing Shared Executable Libraries from SAS	126

Chapter 6 △ **Viewing Output and Help in the SAS Remote Browser** 133

What Is Remote Browsing?	133
Using Remote Browsing with ODS Output	134
Installing the Remote Browser Server	134
System Options for Remote Browsing	134
Setting Up the SAS Remote Browser	135
Remote Browsing and Firewalls	136

PART 2 **SAS Windowing Environment** 137

Chapter 7 △ **Working in the SAS Windowing Environment** 139

Definition of the SAS Windowing Environment	140
Description of SAS in the X Environment	141
The SAS Session Manager (motifxsassm) in UNIX	143
Displaying Function Key Definitions in UNIX Environments	146
The SAS ToolBox in UNIX Environments	147
Opening Files in UNIX Environments	150
Changing Your Working Directory in UNIX Environments	152
Selecting (Marking) Text in UNIX Environments	153
Copying or Cutting and Pasting Selected Text in UNIX Environments	155
Using Drag and Drop in UNIX Environments	156
Searching For and Replacing Text Strings in UNIX Environments	157
Sending Mail from within Your SAS Session in UNIX Environments	158
Configuring SAS for Host Editor Support in UNIX Environments	160
Getting Help in UNIX Environments	161

Chapter 8 △ **Customizing the SAS Windowing Environment** 163

Overview of Customizing SAS in X Environment	164
Overview of X Resources	165
Methods for Customizing X Resources	165
Modifying X Resources Through the Preferences Dialog Box	167
Setting X Resources with the Resource Helper	172
Customizing Toolboxes and Toolsets in UNIX Environments	177
Customizing Key Definitions in UNIX Environments	184
Customizing Fonts in UNIX Environments	192

Customizing Colors in UNIX Environments	196
Controlling Drop-Down Menus in UNIX Environments	202
Customizing Cut and Paste in UNIX Environments	202
Customizing Session Workspace, Session Gravity, and Window Sizes in UNIX Environments	204
Specifying User-Defined Icons in UNIX Environments	206
Miscellaneous Resources in UNIX Environments	207
Summary of X Resources for SAS in UNIX Environments	209

PART 3 Application Considerations 213

Chapter 9 △ Data Representation 215

Numeric Variable Length and Precision in UNIX Environments	215
Missing Values in UNIX Environments	216
Reading and Writing Binary Data in UNIX Environments	216

PART 4 Host-Specific Features of the SAS Language 217

Chapter 10 △ Commands under UNIX 219

SAS Commands under UNIX	220
-------------------------	-----

Chapter 11 △ Data Set Options under UNIX 241

SAS Data Set Options under UNIX	241
Summary of SAS Data Set Options in UNIX Environments	241

Chapter 12 △ Formats under UNIX 251

SAS Formats under UNIX	251
------------------------	-----

Chapter 13 △ Functions and CALL Routines under UNIX 257

SAS Functions and CALL Routines under UNIX	257
--	-----

Chapter 14 △ Informats under UNIX 279

SAS Informats under UNIX	279
--------------------------	-----

Chapter 15 △ Macro Facility under UNIX 285

About the Macro Facility under UNIX	285
Automatic Macro Variables in UNIX Environments	285
Macro Statements in UNIX Environments	287
Macro Functions in UNIX Environments	287
SAS System Options Used by the Macro Facility in UNIX Environments	288
Using Autocall Libraries in UNIX Environments	288

Chapter 16 △ Procedures under UNIX 291

SAS Procedures under UNIX	291
---------------------------	-----

Chapter 17 △ Statements under UNIX 315

SAS Statements under UNIX	315
---------------------------	-----

Chapter 18	△ System Options under UNIX	341
	SAS System Options under UNIX	342
	Determining How a System Option Was Set	343
	Summary of All SAS System Options in UNIX Environments	343

PART 5 **Appendixes** **419**

Appendix 1	△ The !SASROOT Directory	421
	Introduction to the !SASROOT Directory	421
	Contents of the !SASROOT Directory	421
Appendix 2	△ Tools for the System Administrator	423
	The Utilities Directory in UNIX Environments	423
	Installing Manual Pages	423
	The UNIX Authentication API	424
	Utilities in the /utilities/bin Directory	424
Appendix 3	△ Commands That Are Not Specific to UNIX	429
	Commands That Are Not Specific to UNIX	430
Appendix 4	△ Recommended Reading	475
	Recommended Reading	475
Glossary		477
Index		485

What's New

Overview

SAS 9.2 in UNIX environments has the following new and enhanced features:

- direct file I/O options
- file locking capabilities
- SAS Remote Browser
- multiple work directories
- installation and configuration changes
- 64-bit architecture
- SAS logging facility
- shared executable libraries enhancements
- data set option enhancements
- function changes
- procedure enhancements
- statement changes
- system option changes
- IPv6 Standard
- documentation enhancements

Direct File I/O Options

Changes to direct file I/O include the addition of the following new options:

- ENABLEDIRECTIO LIBNAME statement option
- USEDIRECTIO= data set option
- USEDIRECTIO= LIBNAME statement option

File Locking

File locking has been enhanced in the following ways:

- The FILELOCKWAITMAX system option is new.
- A new syntax format has been added to the FILELOCKS system option.

SAS Remote Browser

The SAS Remote Browser enables you to view SAS documentation, URLs that are specified in the WBROWSE command, and some ODS HTML output in the default Web browser on your local computer. By displaying this documentation locally, you have faster access to the documentation. In addition, you free resources on the SAS server that were used by Netscape.

You configure remote browsing by using the following system options:

- HELPHOST
- HELPPORT System Option in *SAS Language Reference: Dictionary*.

Multiple Work Directories

SAS can make the distribution of Work libraries dynamic by distributing work libraries across several directories. When the argument to WORK is a list of directories in a file, you can specify a method for choosing which directory to use for Work. METHOD=RANDOM and METHOD=SPACE are the available methods.

Installation and Configuration Changes

- The default directory path where SAS is installed has changed.

In 9.2, the default directory path and executable are:

```
!SASROOT      = /usr/local/SAS/SASFoundation/9.2/
SAS executable = /usr/local/SAS/SASFoundation/9.2/sas
```

In previous releases, the default directory path and executable were:

```
!SASROOT      = /usr/local/SAS/SAS_x.x/
SAS executable = /usr/local/SAS/SAS_x.x/sas
```

- The sasv9_local.cfg file, which contains user-specified options and overrides the options in the default configuration file, has been added to the !SASROOT directory and to the order of precedence for SAS configuration files.

64-bit Architecture

You can now run a 64-bit version of SAS on Linux for x64.

SAS Logging Facility

The SAS 9.2 logging facility enables the categorization and collection of log event messages, and then writes them to a variety of output devices. The logging facility supports problem diagnosis and resolution, performance and capacity management, and auditing and regulatory compliance.

Shared Executable Libraries

Shared libraries in SAS can store useful routines that might be needed by multiple applications. When an application needs a routine that resides in an external shared library, the application loads the shared library, invokes the routine using the ROUTINE statement, and unloads the shared library upon completion of the routine. In the ROUTINE statement, the following three return types for the RETURNS= argument have been added:

- PTR
- [U]INT32
- [U]INT64

SAS Language Elements

Data Set Options

The new USEDIRECTIO= data set option turns on direct file I/O for the library that contains the file to which the ENABLEDIRECTIO option in the LIBNAME statement has been applied.

Functions

The LIBNAME function was removed from this Companion because it is portable.

Procedures

The SORT procedure now supports the DETAILS statement option. The DETAILS statement option specifies that PROC SORT write messages to the SAS log about whether the sort was performed in memory.

Statements

The following statements have been enhanced:

- The FILENAME statement now supports the LOCKINTERNAL statement option, which specifies SAS system locking for files that are associated with a fileref in the FILENAME statement. LOCKINTERNAL has the following arguments:
 - AUTO
locks a file so that in a SAS session, if a user has Write access to the file, then no other users can have Read or Write access to the file.
 - SHARED
locks a file so that in a SAS session, two users do not have simultaneous Write access to the file. The file can be shared simultaneously by one user who has Write access and multiple users who have Read access.
- The FILENAME statement now supports the Secure File Transfer Protocol (SFTP) access method, which enables you to access remote files by using SFTP.

- The LIBNAME statement now supports the following new statement options:
 - ENABLEDIRECTIO**
specifies that direct file I/O can be available for all of the files that are opened in the library that is identified in the LIBNAME statement.
 - FILELOCKWAIT**
specifies the number of seconds that SAS will wait for a locked file to become available to another process.
 - TRANSFERSIZE**
specifies the size of a large block of data that is read from a file that is opened.
 - USEDIRECTIO=**
if used with the ENABLEDIRECTIO option in the LIBNAME statement, turns on direct file I/O for all of the files that are associated with the libref that is listed in the LIBNAME statement.

System Options

The following system options are new:

- FILELOCKWAITMAX, which sets an upper limit on the time that SAS will wait for a locked file.
- FMTSEARCH, which specifies the order in which format catalogs are searched.
- HELPHOST, which specifies the name of the local computer where the SAS Remote Browser displays HTML pages.
- PRIMARYPROVIDERDOMAIN, which specifies the domain name of the primary authentication provider.

The following system options have been enhanced:

- The APPEND system option now supports the FMTSEARCH option.
- The INSERT system option now supports the FMTSEARCH option.
- The FULLSTIMER system option now specifies whether to display all available system performance statistics, as well as the datetime stamp in the SAS log.

The following system options are obsolete:

- ASYNCHIO
- COMAUX2
- DOCLOC
- FSDEVICE
- GISMAPS
- HELPENCMD
- INGOPTS
- NETMAC
- SEQENGINE
- TAPECLOSE

The following system options are no longer specific to UNIX, and are documented in *SAS Language Reference: Dictionary*:

- S
- S2


IPv6 Standard

SAS supports the next generation of IPv6 (Internet Protocol version 6), which is the successor to IPv4. Rather than replacing IPv4 with IPv6, SAS 9.2 supports both standards. A primary reason for the new version is that the limited supply of 32-bit IPv4 address spaces is being depleted. IPv6 uses a 128-bit address scheme, which provides more IP addresses than IPv4.

Documentation Enhancements


The following documentation enhancements have been made:

- A new section called “Sharing External Files in a UNIX Environment” has been added to the documentation about file locking in external files.
- A new section called “Using Direct I/O” has been added.
- A table of tools in the !SASROOT/utilities/bin directory has been added to Appendix 2, “Tools for the System Administrator.”
- A section called “Case Sensitivity in Data Set Names” has been added.
- In the performance tuning section of PROC SORT, a section called “How SAS Determines the Amount of Memory to Use” has been added. The section provides a discussion of the MEMSIZE, REALMEMSIZE, and SORTSIZE system options.
- The documentation for the SSLCALISTLOC, SSLCERTLOC, SSLCLIENTAUTH, SSLCRLCHECK, SSLCRLLOC, SSLPVTKEYLOC, and SSLPVTKEYPASS system options has been relocated to *Encryption in SAS*. The “Using SSL in a UNIX Environment” section has been moved to *Encryption in SAS*.
- Sections called “SAS Invocation Scripts” and “SAS Configuration Files” have been added.
- The documentation for UMASK has been enhanced. It now includes a section about changing the file permissions for your SAS session. The discussion of UMASK in the WORKPERMS system option has also been enhanced.
- The table of kill signals and their functions has been added. The table includes options and descriptions of **kill** commands, such as SIGINT, SIGNULL, and SIGTERM.
- An appendix has been added that lists and describes command line commands that are not specific to UNIX, but that you can use for editing in your SAS session.



Running SAS Software Under UNIX

<i>Chapter 1</i>	Getting Started with SAS in UNIX Environments	<i>3</i>
<i>Chapter 2</i>	Using SAS Files	<i>31</i>
<i>Chapter 3</i>	Using External Files and Devices	<i>67</i>
<i>Chapter 4</i>	Printing and Routing Output	<i>91</i>
<i>Chapter 5</i>	Accessing Shared Executable Libraries from SAS	<i>107</i>
<i>Chapter 6</i>	Viewing Output and Help in the SAS Remote Browser	<i>133</i>



CHAPTER

1

Getting Started with SAS in UNIX Environments

<i>Starting SAS Sessions in UNIX Environments</i>	4
<i>Invoking SAS</i>	4
<i>SAS Invocation Scripts</i>	5
<i>SAS Configuration Files</i>	5
<i>Regenerating SAS Invocation Scripts</i>	5
<i>Syntax of the SAS Command</i>	5
<i>Example: Invoke an Interactive SAS Session</i>	6
<i>What If SAS Does Not Start?</i>	6
<i>Running SAS in a Foreground or Background Process</i>	6
<i>Selecting a Method of Running SAS in UNIX Environments</i>	7
<i>SAS Windowing Environment in UNIX Environments</i>	7
<i>Introduction to the SAS Windowing Environment</i>	7
<i>What Is the Explorer Window?</i>	7
<i>What Are the Program Editor, Output, and Log Windows?</i>	8
<i>Invoking SAS in the Windowing Environment</i>	8
<i>Exiting SAS in the Windowing Environment</i>	9
<i>Interactive Line Mode in UNIX Environments</i>	9
<i>Introduction to Interactive Line Mode</i>	9
<i>Invoking SAS in Interactive Line Mode</i>	9
<i>Exiting SAS in Interactive Line Mode</i>	9
<i>Batch Mode in UNIX Environments</i>	10
<i>Introduction to Running SAS in Batch Mode</i>	10
<i>Invoking SAS in Batch Mode</i>	10
<i>Submitting a Program to the Batch Queue</i>	10
<i>Writing Data from an External File Using UNIX Pipes</i>	11
<i>Running SAS on a Remote Host in UNIX Environments</i>	11
<i>Introduction to Running SAS on a Remote Host</i>	11
<i>Steps for Running SAS on a Remote Host</i>	11
<i>Preventing SAS from Attempting to Connect to the X Server</i>	12
<i>Troubleshooting Connection Problems</i>	12
<i>X Command Line Options</i>	13
<i>How to Specify X Window System Options</i>	13
<i>Supported X Command Line Options</i>	13
<i>Unsupported X Command Line Options</i>	14
<i>Executing Operating System Commands from Your SAS Session</i>	14
<i>Deciding Whether to Run an Asynchronous or Synchronous Task</i>	14
<i>Executing a Single UNIX Command</i>	15
<i>Example 1: Executing a UNIX Command by Using the X Statement</i>	15
<i>Example 2: Executing a UNIX Command by Using the CALL SYSTEM Routine</i>	15
<i>How SAS Processes a Single UNIX Command</i>	15
<i>Executing Several UNIX Commands</i>	16

<i>Example: Executing Several Commands Using the %SYSEXEC Macro</i>	16
<i>How SAS Processes Several UNIX Commands</i>	16
<i>Changing the File Permissions for Your SAS Session</i>	16
<i>Executing X Statements in Batch Mode</i>	17
<i>Customizing Your SAS Registry Files</i>	17
<i>Customizing Your SAS Session by Using System Options</i>	17
<i>Ways to Customize Your SAS Session</i>	17
<i>Ways to Specify a SAS System Option</i>	18
<i>Overriding the Default Value for a System Option</i>	18
<i>How SAS Processes System Options Set in One Place</i>	19
<i>How SAS Processes System Options Set in Multiple Places</i>	20
<i>Order of Precedence for Processing System Options</i>	20
<i>Customizing Your SAS Session by Using Configuration and Autoexec Files</i>	20
<i>Customizing Your SAS Session</i>	20
<i>Introduction to Configuration and Autoexec Files</i>	20
<i>Differences between Configuration and Autoexec Files</i>	21
<i>Creating a Configuration File</i>	21
<i>Order of Precedence for Processing SAS Configuration Files</i>	21
<i>Specifying a Configuration File for SAS to Use</i>	22
<i>Defining Environment Variables in UNIX Environments</i>	23
<i>What Is a UNIX Environment Variable?</i>	23
<i>How to Define an Environment Variable for Your Shell</i>	23
<i>Bourne and Korn Shells</i>	23
<i>C Shell</i>	23
<i>Displaying the Value of an Environment Variable</i>	23
<i>Determining the Completion Status of a SAS Job in UNIX Environments</i>	24
<i>Exiting or Interrupting Your SAS Session in UNIX Environments</i>	24
<i>Methods for Exiting SAS</i>	24
<i>Methods for Interrupting or Terminating SAS</i>	25
<i>Using Control Keys</i>	25
<i>Using the SAS Session Manager</i>	25
<i>Using the UNIX kill Command</i>	26
<i>Messages in the SAS Console Log</i>	26
<i>Ending a Process That Is Running as a SAS Server</i>	27
<i>Ending a SAS Process on a Relational Database</i>	27
<i>How to Interrupt a SAS Process</i>	27
<i>Example: Interrupt Menu for PROC SQL</i>	28
<i>How to Terminate a SAS Process</i>	29
<i>What Happens When You Interrupt a SAS Process and the Underlying DBMS Process</i>	29

Starting SAS Sessions in UNIX Environments

Invoking SAS

A SAS session is invoked using a link in the **!SASROOT** directory. Your UNIX administrator can add this link to the list of commands for your operating environment.

Ask your system administrator for the command that invokes SAS at your site. At many sites, the command to invoke SAS is **sas**, but a different command might have been defined during the SAS installation process at your site. This documentation assumes that SAS is invoked by the **sas** command.

Note: Before you start your SAS session, review the different techniques for interrupting and terminating your SAS session (see “Exiting or Interrupting Your SAS Session in UNIX Environments ” on page 24). Also, if you cannot stop your SAS session, contact your system administrator. △

SAS Invocation Scripts

SAS is invoked by scripts that are located in the **!SASROOT/bin** directory. A SAS invocation script is created for each language that is installed. The invocation scripts are named using the language codes of the installed language. For example, **sas_en** invokes the English version of SAS. All languages are installed in all locations.

For more information about setting up SAS, refer to the installation documentation for the UNIX environment.

SAS Configuration Files

SAS creates a separate configuration file for each language that is installed. The language-specific configuration files have the form **!SASROOT/nls/<language>/sasv9.cfg** for each language. An additional configuration file that is language independent is **!SASROOT/sasv9.cfg**. This master configuration file in **!SASROOT** is used by all languages in addition to the language-specific files in **!SASROOT/nls/<language>/**. You can modify these configuration files to meet your needs. For information about how to customize SAS configuration files, see “Customizing Your SAS Session by Using Configuration and Autoexec Files” on page 20.

Regenerating SAS Invocation Scripts

The SAS invocation scripts that exist in **!SASROOT/bin** should not be modified. SAS Setup enables you to regenerate the default SAS invocation scripts that are located in **!SASROOT/bin**. To regenerate the invocation scripts, perform the following steps:

- 1 Run **SAS Setup** from **!SASROOT/sassetup**. Make sure that you have the appropriate privilege to update files in **SASROOT**.
- 2 Select **Run Setup Utilities** from the **SAS Setup Primary Menu**.
- 3 Select **Perform SAS Software Configuration**.
- 4 Select **Recreate the SAS Invocation Scripts**.

Syntax of the SAS Command

The general form of the SAS command is as follows:

```
sas <-option1...-option-n> <filename>
```

You can use these arguments with the SAS command:

-option1 ... -option-n

specifies SAS system options to configure your session or X command line options. See Chapter 18, “System Options under UNIX,” on page 341 and “X Command Line Options” on page 13 for more information. If you omit any options (either on the command line or in the configuration file), the SAS (or site-specific) default options are in effect.

filename

specifies the name of the file containing the SAS program to be executed. Specifying a filename on the SAS command invokes a batch SAS session. Omit the filename to begin an interactive session.

If the file is not in the current directory, specify its full pathname.

Example: Invoke an Interactive SAS Session

To invoke an interactive SAS session, without specifying any SAS system options, enter

```
sas
```

The execution mode will depend on your default settings. For more information, see “Selecting a Method of Running SAS in UNIX Environments” on page 7.

To specify the NODATE and LINESIZE system options, you could enter

```
sas -nodate -linesize 80
```

What If SAS Does Not Start?

If SAS does not start, the SAS log might contain error messages that explain the failure. However, error messages that SAS issues before the SAS log is initialized are written to the SAS console log.

Under UNIX, the STDOUT fileref specifies the location of the console log.

Running SAS in a Foreground or Background Process

UNIX is a multitasking operating system, so you can run multiple processes at the same time. For example, you can have one process running in the foreground and three in the background.

A *foreground process* executes while you wait for the prompt; that is, you cannot execute additional commands while the current command is being executed. After you enter a command, the shell starts a process to execute the command. After the system executes the command, the shell displays the prompt and you can enter additional commands. The following is an example of SAS executing as a foreground process:

```
sas
```

A *background process* executes independently of the shell. After you enter a command, the shell starts a process to execute the command, and then issues the system prompt. You can enter other commands or start other background processes without waiting for your initial command to execute. The following is an example of the command that is used to execute a background process:

```
sas&
```

Note: Both the C shell and the Korn shell include commands that enable you to move jobs among three possible states: running in the foreground, running in the background, and suspended. \triangle

Selecting a Method of Running SAS in UNIX Environments

You can run SAS in the following modes:

- SAS windowing environment
- interactive line mode
- batch mode

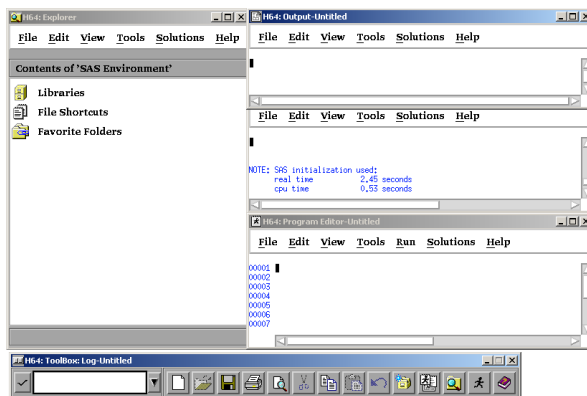
Ask your UNIX system administrator which interface or mode of operation is the default at your site.

SAS Windowing Environment in UNIX Environments

Introduction to the SAS Windowing Environment

You interact with SAS through windows using the keyboard, mouse, menus, and icons. The windowing environment includes, but is not limited to, the Explorer, Program Editor, Output, Log, and Results windows. The following display shows the Explorer, Output, Log, and Program Editor windows. The ToolBox window is also displayed.

Display 1.1 Windows in the SAS Windowing Environment



Your SAS session might default to the windowing environment interface. If you want to use the windowing environment, you can start your SAS session as a foreground process, or as a background process by adding an ampersand (&) to your SAS command line. See “Running SAS in a Foreground or Background Process” on page 6 for an example of these SAS commands.

For more information about using the windowing environment, see Chapter 7, “Working in the SAS Windowing Environment,” on page 139.

Note: If you are not using an X display, then you can invoke SAS in interactive line mode by using the NODMS system option. For more information, see “Interactive Line Mode in UNIX Environments” on page 9. \triangle

What Is the Explorer Window?

Explorer is a windowing environment for managing basic SAS software tasks such as viewing and managing data sets, libraries, members, applications, and output. The SAS Explorer is a central access point from which you can do the following:

- manipulate SAS data through a graphical interface
- access the Program Editor, Output, and Log windows (as well as other windows)
- view the results of SAS procedure output in the Results window
- import files into SAS

What Are the Program Editor, Output, and Log Windows?

The Program Editor, Output, and Log windows enable you to edit and execute SAS programs and display output. For more information about these windows, see the online SAS Help and Documentation.

Invoking SAS in the Windowing Environment

You can use the following commands to specify which windows open when a SAS session starts.

- You can open the Program Editor, Output, and Log windows by specifying the DMS system option:

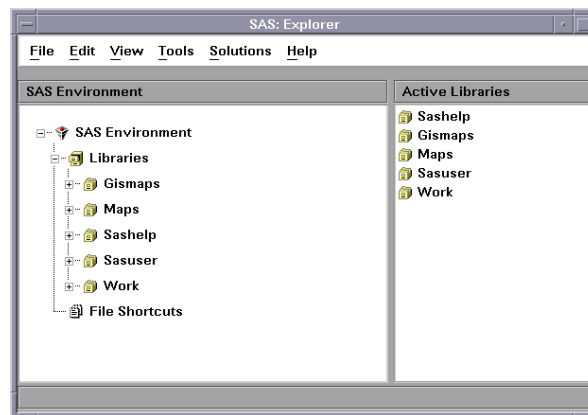
```
sas -dms
```

- You can open the Program Editor, Output, Log, and Results windows, as well as the Explorer window, by specifying the DMSEXP system option:

```
sas -dmsexp
```

- You can open only the Explorer window by specifying the EXPLORER system option:

```
sas -explorer
```



The default specification for invoking SAS is `sas -dms -dmsexp`. This command displays the Program Editor, Output, Log, and Results windows as well as the Explorer window. If you invoke SAS without the `-dmsexp` option, the Explorer window does not display.

SAS also opens a toolbox from which you can open additional SAS windows. For more information about the toolbox, see Chapter 7, “Working in the SAS Windowing Environment,” on page 139.

Exiting SAS in the Windowing Environment

To end your SAS session, enter the `BYE` or `ENDSAS` command in the command window, or select **File ► Exit** from the drop-down menu of the SAS session that you want to end.

Interactive Line Mode in UNIX Environments

Introduction to Interactive Line Mode

If you are not using an X display, you can invoke SAS in interactive line mode by using the `NODMS` system option.

You enter SAS statements line by line in response to prompts issued by SAS. SAS reads the source statements from the terminal as you enter them. `DATA` and `PROC` steps execute when one of the following occurs:

- a `RUN`, `QUIT`, or `DATALINES` statement is entered
- another `DATA` or `PROC` statement is entered
- the `ENDSAS` statement is entered

To use interactive line mode, you must run SAS in the foreground.

Invoking SAS in Interactive Line Mode

To start an interactive line mode session, invoke SAS with the `NODMS` or `NODMSEXP` system option:

```
sas -nodms
sas -nodmsexp
```

By default, SAS log and procedure output (if any) appear on your display as each step executes.

You can also invoke SAS in interactive line mode and pass parameters to it:

```
sas -sysparm 'A B C' progparm.sas
```

The value **A B C** is assigned to the `SYSPARM` macro variable, which can be read by the program `progparm.sas`.

After you invoke SAS, the `1?` prompt appears, and you can begin entering SAS statements. After you enter each statement, a line number prompt appears.

Exiting SAS in Interactive Line Mode

You can end the session by pressing the EOF key (usually `CTRL+D`; see “Using Control Keys” on page 25) or by issuing the `ENDSAS` statement:

```
endsas;
```

The session ends after all SAS statements have executed.

Batch Mode in UNIX Environments

Introduction to Running SAS in Batch Mode

To run SAS in batch mode, you specify your SAS program name in the SAS invocation command. You can run batch mode in the foreground, in the background by specifying an ampersand at the end of the SAS command, or submit your application to the batch queue by using the **batch**, **at**, **nohup**, or **cron** UNIX commands. (For more information, refer to the UNIX man pages for the **batch**, **at**, **nohup**, or **cron** commands.) If you start your application with one of these UNIX commands and you log off of your system, then your application will complete execution. If your application contains statements that start an interactive procedure such as FSEDIT, then you need to run your batch application in the foreground.

Invoking SAS in Batch Mode

To invoke SAS in batch mode, you must specify a filename in the SAS command. For example, if **weekly.sas** is the file that contains the SAS statements to be executed, and you want to specify the NODATE and LINESIZE system options, you would enter the following command:

```
sas weekly.sas -nodate -linesize 90
```

The command would run the program in the foreground. If you want to run the program in the background, add the ampersand to the end of the command:

```
sas weekly.sas -nodate -linesize 90 &
```

SAS creates a .log file and a .lst file in the current directory that contains the log and procedure output.

Submitting a Program to the Batch Queue

To submit your program to the batch queue, you can use the **batch**, **at**, **nohup**, or **cron** commands. For example, you could submit **weekly.sas** from your shell prompt as follows:

```
$ at 2am
sas weekly.sas
<control-D>
warning: commands will be executed using /usr/bin/sh
job 8400.a at Wed Jun 11 02:00:00 2008
$
```

If you create a file that contains the SAS command (for example, **cmdfile.sh**) that is necessary to run your program, then you can enter the following command at your shell prompt:

```
at 2am < cmdfile.sh
```

SAS sends the output to a file that has the same name as the program. The output file has an extension of .lst. The log file writes to a file with an extension of .log. Both of these files are written to your current directory. Refer to the UNIX man pages for these commands for more information on submitting jobs to the batch queue. For more

information about routing output, see Chapter 4, “Printing and Routing Output,” on page 91.

If you submit a file in batch mode, then a line that is greater than 256 bytes will be truncated. An explicit message about this truncation is written to the SAS log.

Note: If your program contains statements that start an interactive procedure such as the FSEDIT procedure, you will need to run your program as a foreground process. Δ

Writing Data from an External File Using UNIX Pipes

You can use a UNIX pipe to write data from an external file to a SAS program. For example, suppose that your data resides in the external file `mydata` and your SAS program `myprog.sas` includes this statement:

```
infile stdin;
```

Issue this command to have `myprog.sas` read data from the external file `mydata`:

```
cat mydata | sas myprog.sas
```

For information about using external files, see Chapter 3, “Using External Files and Devices,” on page 67. See “File Descriptors in the Bourne and Korn Shells” on page 78 for another way to have a SAS program read data from an external file.

Running SAS on a Remote Host in UNIX Environments

Introduction to Running SAS on a Remote Host

When you invoke SAS in an interactive mode, you can run SAS on your local host, or you can run SAS on a remote host and interact with the session through an X server running on your workstation. The server provides the display services that are needed for the X Window System.

Most of the time, the server name is derived from the computer’s name. For example, if your computer is named `green`, the name of the server is `green:0.0`. In most cases, the X server will already be running when you log in. If you need to start your server manually, consult the documentation that is provided with your X Window System software.

To run SAS on a remote host, you must tell SAS which display to use by either setting the DISPLAY environment variable or specifying the `-display` X command line option.

Steps for Running SAS on a Remote Host

To run SAS on a remote host, follow these steps:

- 1 Make sure that the clients running on the remote host have permission to connect to your server. With most X servers, authorization is controlled by using an `.Xauthority` file that is located in the user’s home directory. Additionally, the `xhost` command can be used to circumvent authority. To use the `xhost` client to permit all remote hosts to connect to your server, enter the following command at the system prompt on the system that is running your X server:

```
xhost +
```

If your system does not control access with the `xhost` client, consult your system documentation for information on allowing remote access.

For information about editing and displaying authorization information, see the UNIX man page for `xauth`.

- 2 Log in to the remote system, or use a remote shell.
- 3 Identify your server as the target display for X clients that are run on the remote host. You can identify your server in one of two ways:
 - a Set the `DISPLAY` environment variable. In the Bourne and Korn shells, you can set the `DISPLAY` variable as follows:

```
DISPLAY=green:0.0
export DISPLAY
```

In the Korn shell, you can combine these two commands:

```
export DISPLAY=green:0.0
```

In the C shell, you must use the UNIX `setenv` command:

```
setenv DISPLAY green:0.0
```

The `DISPLAY` variable will be used by all X clients on the system.

Note: To determine the shell for your current system, type `ps` at the UNIX command prompt or check the value of the `SHELL` environment variable. Δ

- b Use the `DISPLAY` system option. For example:

```
sas -display green:0.0
```

If you have trouble establishing a connection, you can try using an IP address instead of a display name, for example:

```
-display 10.22.1.1:0.0
```

Note: This option is a command line option for the X Window System, not for SAS. Specifying this option in a SAS configuration file or in the `SASV9_OPTIONS` environment variable might cause problems when you are running other interfaces. Δ

Preventing SAS from Attempting to Connect to the X Server

To prevent SAS from attempting to connect to the X server, unset the `DISPLAY` environment variable and use the `-noterminal` SAS option on the command line. The `-noterminal` option specifies that you do not want to display the SAS session. You must specify this option to generate a graph in batch mode. You must also specify this option when you use `PROC IMPORT` and `PROC EXPORT`. For more information, see "Running SAS/GRAPH Programs" in *SAS/GRAPH: Reference*.

Troubleshooting Connection Problems

If SAS cannot establish a connection to your display, it prints a message that indicates the nature of the problem and then terminates. An example of a message that you might receive is the following:

```
ERROR: The connection to the X display server could not be made.
Verify that the X display name is correct, and that you have
access authorization. See the online Help for more information
```

about connecting to an X display server.

Make sure that you have brought up the SAS session correctly. You might need to use the **xhost** client (enter **xhost +**) or some other method to change display permissions. You can also specify the NODMS system option when you invoke SAS to bring your session up in line mode.

If you are unable to invoke SAS, try running another application such as **xclock**. If you cannot run the application, you should contact your UNIX system administrator for assistance.

X Command Line Options

How to Specify X Window System Options

When you invoke some X clients, such as SAS, you can use command line options that are passed to the X Window System. In general, you should specify X Window System options after SAS options on the command line.

Supported X Command Line Options

The following list describes the X command line options that are available when you invoke a SAS session from the command prompt.

-display *host:server.screen*

specifies the name or IP address of the terminal on which you want to display the SAS session. For example, if your display node is **wizard** whose IP address is 10.22.1.1:0.0, you might enter

```
-display wizard:0.0
```

or

```
-display 10.22.1.1:0.0
```

-name *instance-name*

reads the resources in your SAS resource file that begin with *instance-name*. For example, **-name MYSAS** reads the resources that begin with **MYSAS**, such as

```
MYSAS.dmsfont: Cour14
MYSAS.defaultToolbox: True
```

-title *string*

specifies a title for your SAS session window. Titles can contain up to 64 characters. Window titles are displayed in the case in which they are entered, which can be lower case, mixed case, or upper case. To use multiple words in the title, enclose the words in single or double quotation marks. For example,

-title MYSAS produces **MYSAS:Explorer** in the title bar of the Explorer window.

-xrm *string*

specifies a resource to override any defaults. For example, the following resource turns off the Confirm dialog box when you exit SAS:

```
-xrm 'SAS.confirmSASExit: False'
```

Unsupported X Command Line Options

SAS does not support the following X command line options because their functionality is not applicable to SAS or is provided by SAS resources. Refer to “Overview of X Resources” on page 165 for more information on SAS resources.

-geometry

Window geometry is specified by the **SAS.windowHeight**, **SAS.windowWidth**, **SAS.maxWindowHeight**, and **SAS.maxWindowWidth** resources.

-background, -bg

These options are ignored.

-bordercolor, -bd

These options are ignored. Refer to “Defining Colors and Attributes for Window Elements (CPARMS)” on page 199 for a description of specifying the color of window borders.

-borderwidth, -bw

These options are ignored. The width of window borders is set by SAS.

-foreground, -fg

These options are ignored.

-font, -fn

SAS fonts are specified by the **SAS.DMSFont**, **SAS.DMSboldFont**, and **SAS.DMSfontPattern** resources.

-iconic

This option is ignored.

-reverse, -rv, +rv

These options are ignored. Refer to “Defining Colors and Attributes for Window Elements (CPARMS)” on page 199 for a description of specifying reverse video.

-selectionTimeout

Timeout length is specified by the **SAS.selectTimeout** resource.

-synchronous, +synchronous

The XSYNC command controls the X synchronization.

-xnlLanguage

This option is ignored.

Executing Operating System Commands from Your SAS Session

Deciding Whether to Run an Asynchronous or Synchronous Task

You can execute UNIX commands from your SAS session either asynchronously or synchronously. When you run a command as an *asynchronous* task, the command executes independently of all other tasks that are currently running. To run a command asynchronously, you must use the SYSTASK statement. See “SYSTASK Statement” on page 334 for information about executing commands asynchronously.

When you execute one or more UNIX commands *synchronously*, then you must wait for those commands to finish executing before you can continue working in your SAS session. You can use the CALL SYSTEM routine, %SYSEXEC macro program

statement, X statement, and X command to execute UNIX commands synchronously. The CALL SYSTEM routine can be executed with a DATA step. The %SYSEXEC macro statement can be used inside macro definitions, and the X statement can be used outside of DATA steps and macro definitions. You can enter the X command on any SAS command line. See “CALL SYSTEM Routine” on page 261 and “Macro Statements in UNIX Environments” on page 287 for more information.

Executing a Single UNIX Command

To execute only one UNIX command, you can enter the X command, X statement, CALL SYSTEM routine, or %SYSEXEC macro statement as follows:

```
X command
X command;
CALL SYSTEM ('command');
%SYSEXEC command;
```

Note: When you use the %SYSEXEC macro statement, if the UNIX command you specify includes a semicolon, you must enclose the UNIX command in a macro quoting function. Refer to *SAS Macro Language: Reference* for more information about quoting functions. △

Example 1: Executing a UNIX Command by Using the X Statement

You can use the X statement to execute the `ls` UNIX command (in a child shell) as follows:

```
x ls -l;
```

Example 2: Executing a UNIX Command by Using the CALL SYSTEM Routine

Inside a DATA step, you can use the CALL SYSTEM routine to execute a `cd` command, which will change the current directory of your SAS session:

```
data _null_;
  call system ('cd /users/smith/report');
run;
```

The search for any relative (partial) filenames during the SAS session will now begin in the `/users/smith/report` directory. When you end the session, your current directory will be the directory in which you started your SAS session.

For more information about the CALL SYSTEM routine, see “CALL SYSTEM Routine” on page 261.

How SAS Processes a Single UNIX Command

When you specify only one command, SAS checks to see whether the command is `cd`, `pwd`, `setenv`, or `umask` and, if so, executes the SAS equivalent of these commands. The SAS `cd` and `pwd` commands are equivalent to their Bourne shell counterparts. The SAS `setenv` command is equivalent to its C shell namesake. The SAS `umask` command is equivalent to the numeric mode of the `umask` command supported by the Bourne, Korn, and C shells. These four commands are built into SAS because they affect the environment of the current SAS session. When executed by SAS software, they affect only the SAS environment and the environment of any shell programs started by the

SAS session. They do not affect the environment of the shell program that began your SAS session.

If the command is not **cd**, **pwd**, or **setenv**, SAS starts a shell in which it executes the command that you specified. The shell that is used depends on the **SHELL** environment variable. If the command is **umask**, but you do not specify a *mask*, then SAS passes the command to the shell in which the current SAS session was started. For more information about the **umask** command, see “Changing the File Permissions for Your SAS Session” on page 16.

Executing Several UNIX Commands

You can also use the **X** command, **X** statement, **CALL SYSTEM** routine, and **%SYSEXEC** macro statement to execute several UNIX commands:

```
X 'command-1;...command-n'
X 'command-1;...command-n';
CALL SYSTEM ('command-1;...command-n' );
%SYSEXEC quoting-function(command-1;...command-n);
```

Separate each UNIX command with a semicolon (;).

Note: When you use the **%SYSEXEC** macro statement to execute several UNIX commands, because the list of commands uses semicolons as separators, you must enclose the string of UNIX commands in a macro quoting function. Refer to *SAS Macro Language: Reference* for more information on quoting functions. \triangle

Example: Executing Several Commands Using the %SYSEXEC Macro

The following code defines and executes a macro called **pwdls** that executes the **pwd** and **ls -l** UNIX commands:

```
%macro pwdls;
%sysexec %str(pwd;ls -l);
%mend pwdls;
%pwdls;
```

This example uses **%str** as the macro quoting function.

How SAS Processes Several UNIX Commands

When you specify more than one UNIX command (that is, a list of commands separated by semicolons), SAS passes the entire list to the shell and does not check for the **cd**, **pwd**, **setenv**, or **umask** commands, as it does when a command is specified by itself (without semicolons).

For more information about how SAS processes the **cd**, **pwd**, **setenv**, or **umask** commands, see “How SAS Processes a Single UNIX Command” on page 15.

Changing the File Permissions for Your SAS Session

At invocation, a SAS session inherits the file permissions from the parent shell. Any file that you create will inherit these permissions. If you want to change or remove file permissions from within SAS, issue the following command in the **X** statement: **umask**. The **umask** command applies a new “mask” to a file, that is, it sets new file permissions for any new file that you create. In this way, the **umask** command can provide file security by restricting access to new files and directories for the current process.

The default value for **umask** is 000, but you can set different values for **umask** in your `.kshrc` and `.cshrc` files. These values affect all child processes that are executed in the shell. Any subsequent file that you create during the current SAS session will inherit the permissions that you specified. The permissions of a file created under a given mask are calculated in octal representation.

Note: The value of a mask can be either numeric or symbolic. For more information about this command, see the UNIX man page for **umask**. △

Executing X Statements in Batch Mode

If you run your SAS program in batch mode and if your operating system supports job control, the program will be suspended when an X statement within the program needs input from the terminal.

If you run your SAS program from the batch queue by submitting it with the **at** or **batch** commands, SAS processes any X statements as follows:

- If the X statement does not specify a command, SAS ignores the statement.
- If any UNIX command in the X statement attempts to get input, it receives an end-of-file (standard input is set to `/dev/null`).
- If any UNIX command in the X statement writes to standard output or standard error, the output is mailed to you unless it was previously redirected.

Customizing Your SAS Registry Files

SAS registry files store information about the SAS session. The SAS registry is the central storage area for configuration data for SAS. The following list identifies some of the data that is stored in the registry:

- the libraries and file shortcuts that SAS assigns at start-up. These shortcuts could include secure information, such as your password.
- the printers that are defined for use and their print setup.
- configuration data for various SAS products.

The Sasuser registry file (called `regtry.sas7bitm`) contains your user defaults. These registry entries can be customized by using the SAS Registry Editor or by using PROC REGISTRY. For more information, see “The SAS Registry” in *SAS Language Reference: Concepts*.

CAUTION:

For experienced users only. Registry customization is generally performed by experienced SAS users and system administrators. △

Customizing Your SAS Session by Using System Options

Ways to Customize Your SAS Session

You can customize your SAS environment in several ways. One way is through the use of SAS system options. For information about other ways to customize a SAS session, see Chapter 8, “Customizing the SAS Windowing Environment,” on page 163.

Ways to Specify a SAS System Option

SAS options can be specified in one or more ways:

- in a configuration file
- in the SASV9_OPTIONS environment variable
- in the SAS command
- in an OPTIONS statement (either in a SAS program or an autoexec file)
- in the System Options window

Table 18.1 on page 344 shows where each SAS system option can be specified.

Any options that do not affect the initialization of SAS, such as CENTER and NOCENTER, can be specified and changed at any time.

Some options can be specified only in a configuration file, in the SASV9_OPTIONS variable, or in the SAS command. These options determine how SAS initializes its interfaces with the operating system and the hardware; they are often called configuration options. After you start a SAS session, these options cannot be changed. Usually, configuration files specify options that you would not change very often. In those cases when you need to change an option just for one job, specify the change in the SAS command.

Overriding the Default Value for a System Option

The default values for SAS system options will be appropriate for many of your SAS programs. However, you can override a default setting using one or more of the following methods:

configuration file

Modify your current configuration file (see “Order of Precedence for Processing SAS Configuration Files” on page 21) or create a new configuration file. Specify SAS system options in the file by preceding each with a hyphen. For ON/OFF options, just list the keyword corresponding to the appropriate setting. For options that accept values, list the keyword identifying the option followed by the option value. All SAS system options can appear in a configuration file.

For example, a configuration file might contain these option specifications:

```
-nocenter
-verbose
-linesize 64
```

SASV9_OPTIONS environment variable

Specify SAS system options in the SASV9_OPTIONS environment variable before you invoke SAS. See “Defining Environment Variables in UNIX Environments” on page 23.

Settings that you specify in the SASV9_OPTIONS environment variable affect SAS sessions that are started when the variable is defined.

For example, in the Korn shell, you would use:

```
export SASV9_OPTIONS='--fullstimer -nodate'
```

SAS command

Specify SAS system options in the SAS command. Precede each option with a hyphen:

```
sas -option1 -option2...
```


For ON/OFF options, list the keyword corresponding to the appropriate setting. For options that accept values, list the keyword that identifies the option, followed by the option value. For example,

```
sas -nodate -work mywork
```

Settings that you specify in the SAS command last for the duration of the SAS session or, for those options that can be changed within the session, until you change them. All options can be specified in the SAS command.

OPTIONS statement within a SAS session

Specify SAS system options in an OPTIONS statement at any point within a SAS session. The options are set for the duration of the SAS session or until you change them. When you specify an option in the OPTIONS statement, do not precede its name with a hyphen (-). If the option has an argument, use = after the option name.

For example,

```
options nodate linesize=72;
options editcmd='/usr/bin/xterm -e vi';
```

Refer to *SAS Language Reference: Dictionary* for more information about the OPTIONS statement. Not all options can be specified in the OPTIONS statement. To find out about a specific option, look up its name in Table 18.1 on page 344.

OPTIONS statement in an autoexec file

Specify SAS system options in an OPTIONS statement in an autoexec file. For example, your autoexec file could contain the following statements:

```
options nodate pagesize=80;
filename rpt '/users/myid/data/report';
```

System Options window

Change the SAS system options from within the System Options window.

In general, use quotation marks to enclose filenames and pathnames specified in the OPTIONS statement or the System Options window. Do not use quotation marks otherwise. Any exceptions are discussed under the individual option. You can use the abbreviations listed in Table 2.6 on page 51 to shorten the filenames and pathnames you specify.

How SAS Processes System Options Set in One Place

If the same option is set more than once within the SAS command, a configuration file, or the SASV9_OPTIONS environment variable, only the last setting is used; the others are ignored. For example, the DMS option is ignored in the following SAS command:

```
sas -dms -nodms
```

The DMS option is also ignored in the following configuration file:

```
-dms
-linesize 80
-nodms
```

By default, if you specify the HELPLOC, MAPS, MSG, SAMPLOC, SASAUTOS, or SASHELP system options more than one time, the last value that is specified is the value that SAS uses. If you want to add additional pathnames to the pathnames already specified by one of these options, you must use the APPEND or INSERT system

options. For more information, see “APPEND System Option” on page 356 and “INSERT System Option” on page 378.

How SAS Processes System Options Set in Multiple Places

When the same option is set in more than one place, the most recent specification is used. The following places are listed in order of precedence. For example, a setting made in the System Options window or OPTIONS statement will override any other setting, but if you set a system option using the SASV9_OPTIONS environment variable, then this option will override only the setting for the same system option in your configuration file.

Order of Precedence for Processing System Options

The precedence for processing system options is as follows:

- 1 System Options window or OPTIONS statement (from a SAS session or job).
- 2 autoexec file that contains an OPTIONS statement (after SAS initializes).
- 3 SAS command.
- 4 SASV9_OPTIONS environment variable.
- 5 configuration files (before SAS initializes). For more information, see “Order of Precedence for Processing SAS Configuration Files” on page 21.

For example, if a configuration file specifies NOSTIMER, you can override the setting in the SAS command by specifying `-FULLSTIMER`.

By default, if you specify the HELPLOC, MAPS, MSG, SAMPLOC, SASAUTOS, or SASHELP system option more than one time, the last value that is specified is the value that SAS uses. If you want to add additional pathnames to the pathnames already specified by one of these options, you must use the APPEND or INSERT system options to add the new pathname. See “APPEND System Option” on page 356 and “INSERT System Option” on page 378 for more information.

Customizing Your SAS Session by Using Configuration and Autoexec Files

Customizing Your SAS Session

You can customize your SAS environment in several ways. To customize your SAS environment at the point of invocation, you can use configuration and autoexec files. For information about how to customize a SAS session using the windowing environment, see Chapter 8, “Customizing the SAS Windowing Environment,” on page 163.

Introduction to Configuration and Autoexec Files

You can customize your SAS session by defining configuration and autoexec files. You can use these files to specify system options and to execute SAS statements automatically whenever you start a SAS session. (SAS system options control many aspects of your SAS session, including output destinations, the efficiency of program execution, and the attributes of SAS files and libraries. Refer to *SAS Language Reference: Dictionary* for a complete description of SAS system options.)

The configuration file (for SAS 9.2) is typically named `sasv9.cfg`, and the autoexec file is named `autoexec.sas`. These files typically reside in the directory where SAS was installed. By default, this directory is the **!SASROOT** directory.

You can have customized configuration and autoexec files in your user home directory. If you do, then SAS will use the customizations specified in these files when you start a SAS session. For more information about the order of precedence SAS uses when processing configuration files, see “Order of Precedence for Processing SAS Configuration Files” on page 21.

SAS system options can be restricted by a UNIX system administrator, so that once they are set by the administrator, they cannot be changed by a user. A system option can be restricted globally, by group, and by user. For more information, see the configuration guide for the UNIX environment on support.sas.com, and see in *SAS Language Reference: Dictionary*.

Differences between Configuration and Autoexec Files

The differences between configuration files and autoexec files are as follows:

- Configuration files can contain only SAS system option settings, while autoexec files can contain any valid SAS statement. For example, you might want to create an autoexec file that includes an `OPTIONS` statement to change the default values of various system options and `LIBNAME` and `FILENAME` statements for the SAS libraries and external files that you use most often.
- Configuration files are processed *before* SAS initializes, while autoexec files are processed immediately *after* SAS initializes but before it processes any source statements. An `OPTIONS` statement in an autoexec file is equivalent to submitting an `OPTIONS` statement as the first statement of your SAS session.

Creating a Configuration File

To create a configuration file, follow these steps:

- 1 Use a text editor to write the SAS system options into a UNIX file. Save the file as either `sasv9.cfg` or `.sasv9.cfg`. (See “Order of Precedence for Processing SAS Configuration Files” on page 21 for more information.)
- 2 Specify one or more system options on each line. Use the same syntax that you would use for specifying system options with the SAS command, but do not include the SAS command itself. For example, a configuration file might contain the following lines:

```
-nocenter
-verbose
-linesize 64
-work /users/myid/tmp
```

- 3 Save and close the configuration file.

Order of Precedence for Processing SAS Configuration Files

SAS is shipped with a default configuration file in the **!SASROOT** directory. Your on-site SAS personnel can edit this configuration file so that it contains whichever options are appropriate to your site.

You can also create one or more of your own configuration files. SAS reads option settings from each of these files in the following order:*

- 1 sasv9.cfg in the **!SASROOT** directory. (See Appendix 1, “The !SASROOT Directory,” on page 421.)
- 2 sasv9_local.cfg in the **!SASROOT** directory. (See Appendix 1, “The !SASROOT Directory,” on page 421.)
- 3 .sasv9.cfg in your home directory. (Notice the leading period.)
- 4 sasv9.cfg in your home directory.
- 5 sasv9.cfg in your current directory.
- 6 any restricted configuration files. Restricted configuration files contain system options that are set by the site administrator and cannot be changed by the user. Options can be restricted globally, by group, or by user. For more information about restricted configuration files, see the configuration guide for the UNIX environment.

For each system option, SAS uses the last setting it encounters; any other settings are ignored. For example, if the WORKPERMS system option is specified in sasv9.cfg in the **!SASROOT** directory and in sasv9.cfg in your current directory, SAS will use the value specified in sasv9.cfg in your current directory.

Specifying a Configuration File for SAS to Use

When you specify a configuration file for SAS to use, you bypass the search of the configuration files listed in “Order of Precedence for Processing SAS Configuration Files” on page 21.

Note: SAS still processes any restricted configuration files that exist. The settings in these files take precedence over the settings in the configuration file that you specify. \triangle

To specify a configuration file, complete one of the following steps:

- specify a configuration file with the CONFIG system option in the SAS command:

```
sas -config filename
```

- specify a configuration file in the SASV9_OPTIONS environment variable. See “Defining Environment Variables in UNIX Environments” on page 23. For example, in the Korn shell, you would use:

```
export SASV9_OPTIONS='-config filename'
```

- define the environment variable SASV9_CONFIG. See “Defining Environment Variables in UNIX Environments” on page 23. For example, in the Korn shell, you would use:

```
export SASV9_CONFIG=filename
```

filename is the name of a file that contains SAS system options.

If you have specified a configuration file in the SASV9_OPTIONS or SASV9_CONFIG environment variables, you can prevent SAS from using that file by specifying NOCONFIG in the SAS command.

If SAS cannot find SASV9_OPTIONS, the following message is written to the SAS log:

```
ERROR: Cannot open [/fullpath/filename]: No such
file or directory.
```

* For future releases of SAS, the names of these files will change accordingly.

Defining Environment Variables in UNIX Environments

What Is a UNIX Environment Variable?

UNIX *environment variables* are variables that apply to both the current shell and to any subshells it creates (for example, when you send a job to the background or execute a script). If you change the value of an environment variable, the change is passed forward to subsequent shells but not backward to the parent shell.

In a SAS session, you can use the SASV9_OPTIONS environment variable to specify system options and the SASV9_CONFIG environment variable to specify a configuration file. You can also use environment variables as filerefs and librefs in various statements and commands. Filerefs and librefs consist of uppercase letters, digits, and the underscore character in environment variable names. Other characters are not recognized by SAS.

Note: A SAS/ACCESS product initializes the environment variables it needs when loading. Any changes that you make to an environment variable after initialization will not be recognized. For more information, see the documentation for your SAS/ACCESS product. Δ

How to Define an Environment Variable for Your Shell

The way in which you define an environment variable depends on the shell that you are running. (To determine which shell you are running, type **ps** at the command prompt or **echo \$SHELL** to see the current value of the SHELL environment variable.)

Bourne and Korn Shells

In the Bourne shell and in the Korn shell, use the **export** command to export one or more variables to the environment. For example, these commands make the value of the variable **scname** available to all subsequent shell scripts:

```
$ scname=phonelist
$ export scname
```

In the Korn shell, you can combine these commands into one command:

```
$ export scname=phonelist
```

If you change the value of **scname**, then the new value affects both the shell variable and the environment variable. If you do not export a variable, only the shell script in which you define has access to its value.

C Shell

In the C shell (csh and tcsh), you set (define and export) environment variables with the **setenv** (set environment) command. For example, this command is equivalent to the commands shown previously:

```
% setenv scname phonelist
```

Displaying the Value of an Environment Variable

To display the values of individual environment variables, use the **echo** command and parameter substitution. An example is: **echo \$SHELL** which returns the current

value of the SHELL environment variable. Use the `env` (or `printenv`) command to display all environment variables and their current values.

Determining the Completion Status of a SAS Job in UNIX Environments

The exit status for the completion of a SAS job is returned in `$STATUS` for the C shell, and in `$?` for the Bourne and Korn shells. A value of 0 indicates normal termination. You can affect the exit status code by using the ABORT statement. The ABORT statement takes an optional integer argument, *n*, which can range from 0 to 255.

Note: Return codes of 0–6 and return codes greater than 977 are reserved for use by SAS. Δ

The following table summarizes the values of the exit status code.

Table 1.1 Exit Status Code Values

Condition	Exit Status Code
All steps terminated normally	0
SAS System issued warnings	1
SAS System issued errors	2
User issued ABORT statement	3
User issued ABORT RETURN statement	4
User issued ABORT ABEND statement	5
SAS could not initialize because of a severe error	6
User issued ABORT RETURN <i>n</i> statement	<i>n</i>
User issued ABORT ABEND <i>n</i> statement	<i>n</i>

If you specify the ERRORABEND SAS system option on the command line, and the job has errors, the exit status code is set to 5.

UNIX exit status codes are in the range 0-255. Numbers greater than 255 might not print what you expect because the code is interpreted as a signed byte.

Exiting or Interrupting Your SAS Session in UNIX Environments

Methods for Exiting SAS

Use one of the following methods to exit a SAS session:

- Select **File** ► **Exit** if you are using SAS in the windowing environment.
- Use `endsas;`.
- Enter **BYE** in the ToolBox if you are using SAS in the windowing environment.
- Use CTRL+D if this control key sequence is your EOF command and if you are using SAS in interactive line mode.

Methods for Interrupting or Terminating SAS

In addition to the methods for exiting SAS, SAS provides methods for interrupting or terminating a SAS session. SAS does not recommend that you use these methods until you have tried to exit SAS by one of the methods listed in “Methods for Exiting SAS” on page 24.

You can interrupt or terminate SAS in the following ways:

- Press the interrupt or quit control key.
- Use the SAS Session Manager.
- Enter the UNIX `kill` command. Use this command when all other methods of exiting SAS have failed.

Using the UNIX `kill` command on a SAS process that is running might corrupt data sets that are open for write or update access.

Using Control Keys

Control keys enable you to interrupt or terminate your session by pressing the interrupt or quit key sequence. However, control keys can be used only when your SAS program is running in interactive line mode or in batch mode in the foreground. You cannot use control keys to stop a background job.

Note: You cannot use control keys to stop a batch job that has been submitted with the `batch`, `at`, `nohup`, or `cron` command. △

Because control keys vary from system to system, issue the UNIX `stty` command to determine which key sends which signal. The `stty` command varies considerably among UNIX operating environments, so check the UNIX man page for `stty` before using the command. Usually, one of these forms of the command will print all of the current terminal settings:

```
stty
stty -a
stty everything
```

The output should contain lines similar to these:

```
intr = ^C; quit = ^\; erase = ^H;
kill = ^U; eof = ^D; eol = ^@
```

The caret (^) represents the CTRL key. In this example, CTRL+C is the interrupt key and CTRL+\ is the quit key.

Pressing the quit key is equivalent to specifying the `-SIGTERM` option on the `kill` command.

Using the SAS Session Manager

If you invoke SAS in the windowing environment, you can use the SAS Session Manager to interrupt or terminate your SAS session. The SAS Session Manager is automatically minimized when you start SAS. To interrupt or terminate your SAS session, open the SAS Session Manager window and click **Interrupt** or **Terminate**.

If asynchronous SAS/CONNECT tasks are running when you terminate a SAS session, these tasks are terminated and no warning message is displayed.

Note: Clicking **Interrupt** is equivalent to specifying the `-SIGINT` option on the `kill` command. Clicking **Terminate** is equivalent to specifying the `-SIGTERM` option on the `kill` command. △

For more information about the SAS Session Manager, see “The SAS Session Manager (motifxsassm) in UNIX” on page 143.

Using the UNIX kill Command

Note: Use the **kill** command only after you have tried all other methods to exit your SAS session. Δ

The **kill** command sends an interrupt or terminate signal to SAS, depending on which signal you specify. You can use the **kill** command to interrupt or terminate a SAS session running in any mode. The **kill** command cannot be issued from within a SAS session. You must issue it from another terminal or from another window (if your terminal permits it).

The format of the **kill** command is:

```
kill <-signal-name> pid
```

To send the interrupt signal, specify **-SIGINT**. To send the terminate signal, specify **-SIGTERM**. Use the **ps** command and its options to determine the process identification number (*pid*) of the SAS session that you want to interrupt or terminate.

The results of using the **ps** command differ in different operating environments. See the UNIX man page for your operating environment for specific information about the **ps** command and its options. Adding options helps to determine which process you want to kill if you have more than one SAS process running. Also, servers (metadata, OLAP, and so on) leave a process identification number in their start-up directories. You can use this number with the **kill** command.

The following table lists some of the important **kill** signals.

Table 1.2 Description of Important kill Signals

Signal	Option	Description
0	SIGNULL	Checks access to process identifier
1	SIGHUP	Causes SAS to terminate
2	SIGINT	Causes SAS to interrupt the session
3	SIGQUIT	Causes SAS to terminate and generates a core file
9	SIGKILL	Causes a forced termination of the SAS session
15	SIGTERM	Causes SAS to terminate

For more information, see the UNIX man pages for the **ps** and **kill** commands.

Messages in the SAS Console Log

If SAS encounters an error or warning condition when the SAS log is not available, then any messages that SAS issues are written to the SAS console log. Normally, the SAS log is unavailable only early in SAS initialization and late in SAS termination.

If you are using the **-STDIO** option, the log is displayed in **stderr**, and the listing is displayed in **STDOUT**.

Ending a Process That Is Running as a SAS Server

If you need to end a process running as a SAS server, use one of the following methods:

- If you are using the SAS Metadata Server, use the SAS Management Console to end a process.
- If you are using another SAS server, use the UNIX scripts that shipped with the servers to stop the process. You can also use these scripts to start (or restart) a server, as well as determine whether the server is already running. For more information about these scripts, contact your site administrator.

Note: If the server does not respond to the UNIX script, then you can use the **kill** command to end the server process. For more information, see “Using the UNIX kill Command” on page 26. △

Ending a SAS Process on a Relational Database

How to Interrupt a SAS Process

CAUTION:

When you interrupt a SAS process, you might terminate the current query. If you are using the current query to create a new data set, then the data set is still created even if the query is terminated. If you are using the current query to overwrite a data set, the data set is not overwritten if the query is terminated. In most cases you do not receive a warning that the query did not complete. △

The method that you use to interrupt a SAS process depends on how you invoke SAS.

- If you are running SAS in interactive line mode or in batch mode using a foreground process, then you can use either of the following methods to interrupt SAS:
 - Press the control key sequence that is set to interrupt in the shell that invoked SAS. In most cases, this control key sequence is CTRL+C. See the man page for the **stty** command to determine the appropriate key sequence for your environment.
 - Use the **-SIGINT** option in the **kill** command. For more information, see “Using the UNIX kill Command” on page 26.
- If you are running the SAS windowing environment in the foreground, then click **Interrupt** in the SAS Session Manager window.

Note: You can access the SAS Session Manager by invoking SAS with the **-dms** or **-dmsexp** option. Select SAS: Session Management from the menu. △

- If you are running SAS in batch mode, then you must click **Interrupt** in the SAS Session Manager window. You cannot use a control key sequence to interrupt the SAS process.

The interrupt signal is sent to the host supervisor. The supervisor determines which DATA steps or procedures are running, and gives you options to interrupt these DATA steps or procedures. The actions that you can take appear in the interrupt menu. After you select an action, the host supervisor performs the operation you selected. Because the options in the interrupt menu are dependent on what is currently executing, you might see a different interrupt menu for the following:

- each SAS procedure. The only option that procedures can act on is the Halt Datastep/Proc option.
- a DATA step. The options available when SAS is processing a DATA step are different from when SAS is processing a procedure.
- each SAS application. For example, SAS webAF has a different interrupt menu than the one for PROC SQL.

Note: Depending on the relational database, the interrupt signal might be handled differently. The interrupt signal is not handled until a safe point in the code is reached that allows the interrupt handler to be run safely. Δ

Example: Interrupt Menu for PROC SQL

The following is an example of the interrupt menu that you might see if you issue an interrupt signal while SAS is processing a PROC SQL statement:

```
Select:
  1. Cancel Submitted Statements
  2. Halt Datastep/Proc: SQL
  C. Cancel the dialog
  T. Terminate the SAS System
```

The following table explains each of these options:

Table 1.3 Description of Interrupt Menu Options for PROC SQL

Option	Description	What This Option Does
1	Cancel Submitted Statements	Selecting this option will end the current DATA step or procedure and the underlying DBMS process. Outstanding source code that is waiting to execute will be flushed from the system. In interactive mode, you will return to the command prompt.
2	Halt Datastep/Proc: SQL	<p>If you select this option and SAS is currently executing an SQL procedure, then the following menu appears:</p> <pre>Press: C to continue Q to cancel the current query S to cancel the submitted statements X to exit SQL procedure ?</pre> <ul style="list-style-type: none"> <input type="checkbox"/> If you select C, then the menu will disappear, but because the current query ended when you interrupted the SAS process, SAS will not return to the current query. Instead, SAS will begin processing the next line of code. <input type="checkbox"/> If you select Q, then SAS cancels the current query even if it is on a relational database. SAS continues processing the next statement. <input type="checkbox"/> If you select S, then all of the PROC SQL statements that you submitted are canceled. <input type="checkbox"/> If you select X, SAS exits the current SQL procedure and starts processing the next statement in the submit block.

Option	Description	What This Option Does
C	Cancel the dialog	Selecting this option returns you to normal processing; however, the current query might have been interrupted. If you are running a long query and the control is on the DBMS server, then selecting C will end the current query. If you are running a short query and SAS has the control, then selecting C will cause the interrupt menu to disappear and the current query will continue. To determine whether the query was interrupted while reading or writing out the DBMS data, use PROC PRINT to view the partially created DBMS table or SAS data set.
T	Terminate SAS	Selecting this option ends your SAS session as well as the current query.

How to Terminate a SAS Process

The method that you use to terminate a SAS process depends on how you invoke SAS.

- If you are running the SAS windowing environment in the foreground, then click **Terminate** in the SAS Session Manager window.
- If you are running an interactive SAS process in the background, then you must click **Terminate** in the SAS Session Manager window. You cannot use a control key sequence to terminate the SAS process.

If you click **Terminate** in the SAS Session Manager, then a dialog box appears confirming that you want to end the session. If you click **OK**, then both the SAS session and the current query are terminated. If you click **Cancel**, then you are returned to the SAS session.

What Happens When You Interrupt a SAS Process and the Underlying DBMS Process

CAUTION:


Interrupting a SAS process and the underlying DBMS process might kill all jobs that are running on your DBMS. Interrupting your SAS and DBMS processes should be an exception. Extensive care should be taken when you construct your queries.

If SAS sends SQL to an RDBMS, there is no way to interrupt the SQL statements because SAS no longer has control of them. The statements are running in the RDBMS. Δ

Note: In this section, SAS process refers to a series of events. It is not the process on the operating system. When you interrupt or terminate a SAS process, the process on the operating system might still be running. Δ

When you interrupt or terminate a query on a server, the following processes stop:

- processing of current extractions. For example, suppose you forgot to include a WHERE clause in your SQL query and are now extracting 1 billion rows into SAS. Issuing an interrupt stops the SAS process and the extract step in the DBMS.
- processing of queries that are in progress on the server. For example, you have a very complex extract query that runs for a long time before producing a result. Issuing an interrupt stops the SAS and DBMS processes. As a result, the complex query running on your DBMS server is interrupted and terminated.
- update, delete, and insert processing. For example, you are updating, deleting, or inserting many rows in your DBMS. An interrupt stops the SAS and DBMS processes.



CHAPTER

2

Using SAS Files

<i>Introduction to SAS Files, Libraries, and Engines in UNIX Environments</i>	33
SAS Files	33
What Is a SAS File?	33
Case Sensitivity in Data Set Names	33
SAS Libraries	34
What Are Libraries?	34
What Is a Libref?	34
Engines	34
What Is an Engine?	34
Additional Resources	34
<i>Common Types of SAS Files in UNIX Environments</i>	34
What Are Data Sets?	34
Descriptor Information and Data Values	34
SAS Data Files (Member Type DATA)	35
SAS Views (Member Type VIEW)	35
What Are Catalogs?	35
What Are Stored Program Files?	36
What Are Access Descriptor Files?	36
<i>Filename Extensions and Member Types in UNIX Environments</i>	36
<i>Using Direct I/O</i>	37
Introduction to Direct I/O	37
Turning On Direct I/O	38
<i>Holding a File in Memory: The SASFILE Statement</i>	38
<i>Sharing SAS Files in a UNIX Environment</i>	38
Sharing SAS Files	38
Options to Use for File Locking: SAS Files	39
File Locking for SAS Files: The FILELOCKS Statement Option	39
File Locking for SAS Files: The FILELOCKS System Option	39
Waiting to Use a Locked File	39
Conditions to Check When FILELOCKS=NONE	39
When FILELOCKS=CONTINUE	40
Sharing Files in a Network	40
Introduction to Sharing Files across Workstations	40
Accessing Files on Different Networks	41
Troubleshooting: Accessing Data over NFS Mounts	41
<i>Compatible Computer Types in UNIX Environments</i>	42
Characteristics of Compatible Computer Types	42
Compatible Computer Types for Release 6.12 through SAS 9.2	42
Determining Compatible Computer Types in SAS 9.2	43
<i>Migrating 32-Bit SAS Files to 64-Bit in UNIX Environments</i>	44
What Is File Migration?	45

<i>Benefits of Migrating SAS Files</i>	45
<i>How to Migrate a SAS Library</i>	45
<i>Additional Resources</i>	45
<i>Creating a SAS File to Use with an Earlier Release</i>	45
<i>Reading SAS Files from Previous Releases or from Other Hosts</i>	46
<i>Reading Version 6 Files</i>	46
<i>Reading Version 8 or Later Files from Compatible Computer Types</i>	46
<i>Reading Version 8 or Later Files from Incompatible Computer Types</i>	46
<i>Compatibility of Existing SAS Files with SAS 9.2</i>	46
<i>Accessing Version 8 or Later Files with CEDA</i>	47
<i>Referring to SAS Files by Using Librefs in UNIX Environments</i>	47
<i>Techniques for Referring to a SAS File</i>	47
<i>What Is a Libref?</i>	48
<i>Assigning Librefs</i>	48
<i>Methods for Assigning Librefs</i>	48
<i>Using the LIBNAME Statement</i>	48
<i>Using the LIBNAME Function</i>	48
<i>Use the DMLIBASSIGN Command</i>	48
<i>Use the LIBNAME Window</i>	49
<i>Use the SAS Explorer Window</i>	49
<i>Permanently Assigning a Libref</i>	49
<i>Accessing a Permanent SAS Library by Using a Libref</i>	50
<i>Specifying Pathnames in UNIX Environments</i>	50
<i>Rules for Specifying Directory and Pathnames</i>	50
<i>Example 1: Access a File That Is Not in the Current Directory</i>	50
<i>Example 2: Access a File in the Current Directory</i>	50
<i>Valid Character Substitutions in Pathnames</i>	51
<i>Assigning a Libref to Several Directories (Concatenating Directories)</i>	51
<i>Introduction to Concatenating Directories</i>	51
<i>How SAS Accesses Concatenated Libraries</i>	52
<i>Accessing Files for Input and Update</i>	52
<i>Accessing Files for Output</i>	52
<i>Accessing Data Sets with the Same Name</i>	52
<i>Using Multiple Engines for a Library in UNIX Environments</i>	53
<i>Using Environment Variables as Librefs in UNIX Environments</i>	53
<i>Librefs Assigned by SAS in UNIX Environments</i>	54
<i>Sasuser Library</i>	55
<i>What Is the Sasuser Library?</i>	55
<i>Contents of the Sasuser Library</i>	55
<i>Sasuser.Profile Catalog</i>	56
<i>Overview of the Sasuser.Profile Catalog</i>	56
<i>How SAS Accesses the Sasuser.Profile Catalog</i>	56
<i>When the Sasuser.Profile Catalog Does Not Exist</i>	56
<i>Checking for an Uncorrupted Sasuser.Profile Catalog</i>	56
<i>If Your Sasuser.Profile Catalog Is Locked or Corrupted</i>	57
<i>Sasuser.Registry Catalog</i>	57
<i>Overview of the Sasuser.Registry Catalog</i>	57
<i>How SAS Accesses the Sasuser.Registry Catalog</i>	57
<i>Sasuser.Prefs File</i>	58
<i>Work Library</i>	58
<i>Multiple Work Directories</i>	59
<i>Using One-Level Names to Access Permanent Files (User Library)</i>	59
<i>Introduction to One-Level Names</i>	59
<i>Techniques for Assigning the User Libref</i>	59

<i>Accessing Disk-Format Libraries in UNIX Environments</i>	60
<i>Accessing Sequential-Format Libraries in UNIX Environments</i>	60
<i>Benefits and Limitations of Sequential Engines</i>	60
<i>Reading and Writing SAS Files on Tape</i>	60
<i>Using a Staging Directory</i>	60
<i>Syntax of the LIBNAME Statement</i>	61
<i>Example: Assign a Libref to the Tape Device</i>	61
<i>Writing Sequential Data Sets to Named Pipes</i>	61
<i>Why Use Named Pipes?</i>	61
<i>Syntax of the LIBNAME Statement</i>	61
<i>Example: Creating a SAS Data Set Using a Named Pipe</i>	61
<i>Accessing BMDP, OSIRIS, or SPSS Files in UNIX Environments</i>	62
<i>Introduction to the BMDP, OSIRIS, and SPSS Files</i>	62
<i>The BMDP Engine</i>	62
<i>What Is the BMDP Engine?</i>	62
<i>Syntax for Accessing BMDP Save Files</i>	62
<i>Example: BMDP Engine</i>	63
<i>The OSIRIS Engine</i>	63
<i>What Is the OSIRIS Engine?</i>	63
<i>Notes About the OSIRIS Data Dictionary Files</i>	63
<i>Syntax for Accessing an OSIRIS File</i>	64
<i>Example: OSIRIS Engine</i>	64
<i>The SPSS Engine</i>	64
<i>What Is the SPSS Engine?</i>	64
<i>Syntax for Accessing an SPSS Export File</i>	64
<i>Reformatting SPSS Files</i>	65
<i>Example: SPSS Engine</i>	66
<i>Support for Links in UNIX Environments</i>	66

Introduction to SAS Files, Libraries, and Engines in UNIX Environments

SAS Files

What Is a SAS File?

Your data can reside in different types of files, including SAS files and files that are formatted by other software products, such as database management systems. Under UNIX, a SAS file is a specially structured UNIX file. Although the UNIX operating environment manages the file for SAS by storing it, the operating system cannot process it because of the structure built into the file by SAS. For example, you can list the filename with the `ls` command, but you cannot use the `vi` editor to edit the file. A SAS file can be permanent or temporary.

Case Sensitivity in Data Set Names

In UNIX operating environments, SAS data set names are written in all lowercase characters. Because of this requirement, SAS reads only data set names that are written in all lowercase characters.

If you use the UNIX utilities `mv` or `cp` to rename SAS data set names with uppercase or mixed-case characters, SAS can no longer read the data set names.

UNIX is case sensitive. Therefore, a data set named *xxx.sas7bdat* is not the same as a data set named *XXX.sas7bdat*. In fact, both of these data sets can reside in the same directory as completely different data sets.

SAS Libraries

What Are Libraries?

SAS files are stored in *SAS libraries*. A SAS library is a collection of SAS files within a UNIX directory. Any UNIX directory can be used as a SAS library. (The directory can also contain files called *external files* that are not managed by SAS. See Chapter 3, “Using External Files and Devices,” on page 67 for how to access external files.) SAS stores temporary SAS files in a Work library (see “Work Library” on page 58), which is automatically defined for you. You must specify a library for each permanent SAS file.

What Is a Libref?

SAS libraries can be identified with *librefs*. A libref is a name by which you reference the directory in your application. For more information about how to assign a libref, see “Referring to SAS Files by Using Librefs in UNIX Environments” on page 47.

Engines

What Is an Engine?

SAS files and SAS libraries are accessed through *engines*. An engine is a set of routines that SAS must use to access the files in the library. SAS can read from and, in some cases, write to the file by using the engine that is appropriate for that file type. For some file types, you need to tell SAS which engine to use. For others, SAS automatically chooses the appropriate engine. The engine that is used to create a SAS data set determines the format of the file.

Additional Resources

For more information about SAS files, libraries, and engines, see *SAS Language Reference: Concepts*.

Common Types of SAS Files in UNIX Environments

What Are Data Sets?

Descriptor Information and Data Values

Data sets consist of descriptor information and data values organized as a table of rows and columns that can be processed by one of the engines. The descriptor

information includes data set type, data set label, the names and labels of the columns in the data set, and so on. A SAS data set can also include *indexes* for one or more columns.

SAS data sets are implemented in two forms:

- If the data values and the data set's descriptor information are stored in one file, the SAS data set is called a *SAS data file*.
- If the file contains information about where to obtain a data set's data values and descriptor information, the SAS data set is called a *SAS view*.

The default engine processes the data set as if the data file or data view and the indexes were a single entity.

For more information, see “SAS Data Files (Member Type DATA)” on page 35 and “SAS Views (Member Type VIEW)” on page 35.

SAS Data Files (Member Type DATA)

The SAS data file is probably the most frequently used type of SAS file. These files have the extension **.sas7bdat**. SAS data files are created in the DATA step and by some SAS procedures. There are two types of data files:

- *Native data files* store data values and their descriptor information in files formatted by SAS. These data files are the traditional SAS data sets from previous releases of SAS.

Native SAS data files created by the default engine can be indexed. An *index* is an auxiliary file created in addition to the data file it indexes. The index provides fast access to observations within a SAS data file by a variable or key. Under UNIX, indexes are stored as separate files, but are treated as integral parts of the SAS data file by SAS.

CAUTION:

Do not remove index files using UNIX commands. Removing the index file can damage your SAS data set. Also, do not change its name or move it to a different directory. Use the DATASETS procedure to manage indexes. △

- *Interface data files* store data in files that have been formatted by other software and that SAS can read only. See “Accessing BMDP, OSIRIS, or SPSS Files in UNIX Environments” on page 62 for more information.

SAS Views (Member Type VIEW)

A SAS view contains only the information needed to derive the data values and the descriptor information. Depending on how the SAS view is created, the actual data can be located in other SAS data sets or in other vendors' files.

Views can be of two kinds:

- *Native SAS views* contain information about data in one or more SAS data files or SAS views. This type of view is created with the SQL procedure or DATA step.
- *Interface SAS views* contain information about data formatted by other software products such as a database management system. For example, the ACCESS procedure in SAS/ACCESS software creates an interface SAS view.

What Are Catalogs?

Catalogs are a special type of SAS file that can contain multiple entries. Many different types of entries can be kept in the same SAS catalog. For example, catalogs

can contain entries created by SAS/AF and SAS/FSP software, windowing applications, key definitions, SAS/GRAPH graphs, and so on.

Catalogs have the SAS member type of CATALOG.

What Are Stored Program Files?

Stored program files are compiled DATA steps generated by the Stored Program Facility. For information about the Stored Program Facility, see *SAS Language Reference: Dictionary*.

Stored program files have the SAS member type of PROGRAM.

What Are Access Descriptor Files?

Access descriptor files describe the data formatted by other software products such as the Oracle or the Sybase database management systems. Descriptor files created by the ACCESS procedure in SAS/ACCESS software have the SAS member type of ACCESS.

Filename Extensions and Member Types in UNIX Environments

Because SAS needs to distinguish between the different file types, it automatically assigns an extension to each file when it creates the file. Also, because each SAS file is a member of a library, SAS assigns each file a member type.

The following table lists the file extensions and their corresponding SAS member types.

CAUTION:

Do not change the file extensions of SAS files. File extensions determine how SAS accesses files; changing them can cause unpredictable results. \triangle

Table 2.1 File Extensions for SAS File Types

Version 6		Version 8, SAS 9		SAS Member Type	Description
Random Access Files	Sequential Access Files	Random Access Files	Sequential Access Files		
.sas	.sas	.sas	.sas	.sas	SAS program
.lst	.lst	.lst	.lst	.lst	Procedure output
.log	.log	.log	.log	.log	SAS log file
.ssdnn ¹	.sdqnn	.sas7bdat	.sas7sdat	DATA	SAS data file
.snxnn	.siqnn	.sas7bndx	.sas7sndx	INDEX	Data file index; not treated by the SAS system as a separate file
.sctnn	.scqnn	.sas7bcats	.sas7scats	CATALOG	SAS catalog
.sspnn	.ssqnn	.sas7bpgm	.sas7spgm	PROGRAM	Stored program (DATA step)
.ssvnn	.svqnn	.sas7bview	.sas7sview	VIEW	SAS view
.ssann	.saqnn	.sas7bacs	.sas7sacs	ACCESS	Access descriptor file
.sstnn	.stqnn	.sas7baud	.sas7saud	AUDIT	Audit file

Version 6		Version 8, SAS 9		SAS Member Type	Description
Random Access Files	Sequential Access Files	Random Access Files	Sequential Access Files		
.sfdnn	.sfqnn	.sas7bfdb	.sas7sfdb	FDB	Consolidation database
.ssmnn	.smqnn	.sas7bmdb	.sas7smdb	MDDB	Multidimensional database
.sdsnn	.soqnn	.sas7bods	.sas7sods	SASODS	Output delivery system file
.snmnn	.sqnnn	.sas7bdmd	.sas7sdmd	DMDB	Data mining database
.sitnn	.srqnn	.sas7bitm	.sas7sitm	ITEMSTOR	Item store file
.sutnn	.suqnn	.sas7butl	.sas7sutl	UTILITY	Utility file
.spunn	.spqnn	.sas7bput	.sas7sput	PUTILITY	Permanent utility file
.ssbnn	.sbqnn	.sas7bbak	.sas7sbak	BACKUP	Backup file

1 All Version 6 files end with a two-character code (*nn*) that identifies sets of compatible SAS files. See “Sharing SAS Files in a UNIX Environment” on page 38 for more information.

A UNIX directory can store a variety of files, but you might find it more practical to store files in separate directories according to their use. Also, you can keep libraries that are accessed by different engines in the same directory, but this is not recommended. For more information, see “Using Multiple Engines for a Library in UNIX Environments” on page 53.

Using Direct I/O

Introduction to Direct I/O

Direct I/O is a method for processing input and output files and is used in file handling. Direct I/O enables SAS to read files from and write files directly to storage devices without first going through the UNIX operating environment’s read and write caches. You can use direct I/O for SAS files. Using direct I/O might improve system performance, depending on the number and types of jobs that you are running.

SAS uses three related options that affect direct I/O:

- ENABLEDIRECTIO statement option
- USEDIRECTIO= statement option
- USEDIRECTIO= data set option

The ENABLEDIRECTIO option in the LIBNAME statement makes direct I/O processing available for data sets that are listed in the DATA statement. The libref that points to the data sets must have been defined in a LIBNAME statement that uses the ENABLEDIRECTIO option. Using ENABLEDIRECTIO itself does not turn on direct I/O.

A libref that is assigned to a directory with the ENABLEDIRECTIO option will not match another libref that is assigned to the same directory without the ENABLEDIRECTIO option. The two librefs will point to the same directory, but the files that are opened using the libref with ENABLEDIRECTIO will be read from and written to using direct I/O. Files that are opened using the other libref will be read from and written to using the regular disk I/O calls.

The `USEDIRECTIO=` data set option in the `DATA` statement or the `USEDIRECTIO=` statement option in the `LIBNAME` statement turns on direct I/O for data sets in which the `ENABLEDIRECTIO` statement option has been applied. Using `USEDIRECTIO=` without first applying the `ENABLEDIRECTIO` option has no effect on direct I/O in a data set.

Turning On Direct I/O

You can turn on direct I/O in two ways:

- Use both the `ENABLEDIRECTIO` and `USEDIRECTIO=` options in the `LIBNAME` statement.

This method opens for direct I/O all of the files that are referenced by the `libref` in the `LIBNAME` statement.

- Use the `ENABLEDIRECTIO` option in the `LIBNAME` statement to render direct I/O available, and use the `USEDIRECTIO=` data set option in a `DATA` statement to turn on direct I/O functionality.

This method opens for direct I/O only the data set that is referenced by the `libref` in the `LIBNAME` statement.

For more information about these options and how they are used, see:

`ENABLEDIRECTIO` in “Engine/Host Options” on page 331

“`USEDIRECTIO=` Data Set Option” on page 248

Holding a File in Memory: The SASFILE Statement

You can use the `SASFILE` statement to open a SAS data set. SAS attempts to allocate enough buffers to hold the entire data set in memory. If enough memory is available, then the entire data set is kept in memory until the data set is closed. If enough memory is not available, then SAS allocates as many buffers as it can. If your file is very large or if SAS is already using a large amount of memory, then using the `SASFILE` statement will not help.

When the `SASFILE` statement executes the first time, SAS opens the file.

Subsequent `DATA` and `PROC` steps use the file without having to open it again because the file remains in memory. The file remains open until a second `SASFILE` statement closes it, or until the program or session ends. For more information, see the `SASFILE` Statement in *SAS Language Reference: Dictionary*.

Sharing SAS Files in a UNIX Environment

Sharing SAS Files

If more than one SAS process has Write access to a SAS file (a data set, catalog, library, and so on) at the same time, you would get unpredictable results if the file was updated. SAS locks the file to prevent more than one user from having Write access to a file. When one SAS process opens a file with Write access, other processes are blocked

from Write access until the first process closes the file. SAS provides statement and system options to override this file protection. However, in almost all cases, you should leave file protection turned on.

Options to Use for File Locking: SAS Files

You can turn off file locking for SAS files in the following ways:

- Use the FILELOCKS option in the LIBNAME statement.
- Use the FILELOCKS system option.

File Locking for SAS Files: The FILELOCKS Statement Option

By default, SAS restricts Write access to one user. The FILELOCKS option in the LIBNAME statement overrides the default and allows multiple users to have Write access to a file. SAS files that are opened under the libref in the LIBNAME statement are the files that are locked. Multiple users have Read access to files.

The FILELOCKS statement option applies to most (but not all) of the SAS I/O files (for example, data sets and catalogs) that are opened under the libref in the LIBNAME statement.

For more information, see “LIBNAME Statement” on page 328.

File Locking for SAS Files: The FILELOCKS System Option

By default, SAS restricts Write access to one user. The FILELOCKS system option overrides this default for both SAS files and external files and allows multiple users to have Write access to a file. The FILELOCKS system option enables you to apply a behavior globally to individual files that are opened.

You can use the FILELOCKS system option at start-up, in the OPTIONS statement, or in the command line. You can specify multiple instances of the FILELOCKS system option. Each instance is added to an internal table of paths and settings. The FILELOCKS system option applies to most (but not all) of the SAS I/O files (for example, data sets and catalogs) that are opened under the libref in the LIBNAME statement. For more information, see “FILELOCKS System Option” on page 368 and the “LIBNAME Statement” on page 328.

Waiting to Use a Locked File

If you want to use a SAS file that is locked by another process, you can use the FILELOCKWAIT option in the LIBNAME statement to specify how long SAS will wait for the locked file to become available to another process. The FILELOCKWAIT statement option affects only those files that are opened under the libref in a LIBNAME statement. For more information, see “LIBNAME Statement” on page 328.

Conditions to Check When FILELOCKS=NONE

When file locking is turned off (that is, when the FILELOCKS system option is set to NONE), SAS attempts to open a file without checking for an existing lock on the file. These files are not protected from shared Update access.

CAUTION:

SAS recommends that you do not use the FILELOCKS=NONE option. If multiple users open the same file for Write access, then the file might become corrupted. The

FILELOCKS=NONE option is used primarily to determine whether a job failed because of a locked file. △

If the FILELOCKS system option is set to NONE, then you should perform one of the following tasks:

- Make sure that your **sasuser** directory is unique for each SAS session. Typically, the system administrator assigns this directory in the system configuration file. The specification in that file or in your personal configuration file helps ensure that the directory is unique as long as you run only one SAS session at a time.

If you run two or more SAS sessions simultaneously, you can guarantee unique user files by specifying different **sasuser** directories for each session. In the first session, you can use:

```
-sasuser ~/sasuser
```

In the *n* the session, you can use:

```
-sasuser ~/sasusern
```

For more information, see “Order of Precedence for Processing SAS Configuration Files” on page 21 and “SASUSER System Option” on page 401. The RSASUSER option can be used to control modifications to the Sasuser library when it is shared by several users (see “RSASUSER System Option” on page 396).

- If you run two or more SAS sessions simultaneously (either by using the X statement or by invoking SAS from two different windows) and you use the same Sasuser.Profile catalog, do not perform any actions (such as use the WSAVE command or change key assignments) within the SAS session to change the Sasuser.Profile catalog because both sessions can use the same catalog.
- Multiple users can read the same data sets at the same time. However, only one user at a time should write to or update a data set so that data is not overwritten or corrupted.

When FILELOCKS=CONTINUE

By default, SAS restricts Write access to one user. When you use the FILELOCKS=CONTINUE option, SAS fails to open a file if that file is locked by another user. However, if SAS returns a message that identifies some other error, then SAS disregards the lock on the file, opens the file, and continues to execute the job.

Sharing Files in a Network

Introduction to Sharing Files across Workstations

SAS can be licensed to run on one or more workstations in a network of similar computers. The license specifically lists the workstations on which SAS can run. Other workstations in the network might have access to the SAS executable files, but they might not be able to run SAS.

If the licensed workstations are connected through NFS mounts so that they share a file system, they can all share a single copy of the SAS executable files, although sharing of a single copy is not necessary. They can also share SAS files. However, if a SAS session attempts to update a data set or catalog, it must obtain an exclusive file lock on that file to prevent other SAS sessions from accessing that file.

If SAS is installed on different types of workstations that are connected through NFS, then each type of workstation must have its own copy of the SAS executable files. For information about how to move catalogs and data sets between hosts, see “Compatible Computer Types in UNIX Environments” on page 42.

Accessing Files on Different Networks

You can access a file on a different type of workstation if the two computers are connected to the same file system. You can access external files that were created under a different operating environment.

If you create a data set or catalog and save it to a directory, and another computer on a different network wants to access the file, then you have several alternatives for working with that file.

- You can log on to a remote computer and perform the work there. This alternative works best if a file is rarely used, or if the original file changes often.
- You can log on to a remote computer and perform the work there if the original file changes often. This alternative ensures that you are accessing the most current version of the file. If you use PROC CPORT to copy the file to your computer, the original file might have changed between the time it was copied and the time it was read.
- You can log on to a remote computer rather than transfer the file to your computer if you want to use the file once. It might not be efficient to use PROC CPORT, or you might not have enough disk space for PROC CPORT to run locally.
- You can use the File Transfer Protocol (FTP) or the Routing Control Processor (RCP) to transfer a file from a remote computer to your computer.
- You can do part of your work on your computer and the other part on a remote computer. One example of this alternative is to run a set of statements on a small test case on the local computer, and then submit the real work to be done on the remote computer. Similarly, you might want to subset a large data set on another computer, and then do local analysis on that subset. You can accomplish this task by using SAS/CONNECT software. For more information about Remote Library Services, see *SAS/CONNECT User's Guide*.

Troubleshooting: Accessing Data over NFS Mounts

SAS might hang when accessing data over NFS mounts if the FILELOCKS option is set to FAIL or CONTINUE. To alleviate the problem, make sure that all NFS file locking daemons are running on both computers (usually statd and lockd). Your UNIX system administrator can assist with starting statd and lockd.

Note: To test whether there is a problem with file locking, you can set the FILELOCKS system option to NONE temporarily. If setting FILELOCKS to NONE resolves the problem, then you know that there probably is a problem with the statd and lockd daemons. It is recommended that you do not set FILELOCKS to NONE permanently because it might cause data corruption or unpredictable results. △

Compatible Computer Types in UNIX Environments

Characteristics of Compatible Computer Types

Big endian, little endian, and bi-endian refer to the way an operating system orders integer values. These values are ordered in bytes, and the way in which bytes are ordered is called byte ordering. The ordering of bytes depends on the operating environment on which the integers were produced.

On big endian platforms, the value 1 is stored in binary and is represented here in hexadecimal notation. One byte is stored as 01, two bytes as 00 01, and four bytes as 00 00 00 01. On little endian platforms, the value 1 is stored in one byte as 01 (the same as big endian), in two bytes as 01 00, and in four bytes as 01 00 00 00.

If an integer is negative, the "two's complement" representation is used. The high-order bit of the most significant byte of the integer will be set on. For example, -2 would be represented in one, two, and four bytes on big endian platforms as FE, FF FE, and FF FF FF FE, respectively. On little endian platforms, the representation would be FE, FE FE, and FE FF FF FF. These representations result from the output of the integer binary value -2 expressed in hexadecimal notation.

Different computers store numeric binary data in different forms. Hewlett-Packard and IBM store data in big-endian format. Linux and Tru64 UNIX store data in little-endian format. Solaris SPARC stores data in big-endian format, and Solaris x64 stores data in little-endian format.

Some architectures, including Solaris SPARC V9 and PA-RISC, enable you to compute or pass data from big endian to little endian, or from little endian to big endian. This feature is called bi-endian and refers primarily to how a processor treats data access. A bi-endian platform can improve performance or simplify the logic of networking devices and software. You can use software to switch many bi-endian architectures so that they default to either big-endian or little-endian. For other architectures, the default of big-endian or little-endian is selected by the hardware and cannot be modified by using software.

SAS files can be transported between compatible computer types using various methods including NFS, FTP, and CD. For two computer types to be compatible, they must have the following characteristics in common:

- Same word length. Word lengths can be either 32-bit or 64-bit.
- Same ordering of bytes in memory. Computer types differ in whether the most significant byte (MSB) or the least significant byte (LSB) is loaded at the lower memory address. This byte order is often referred to as "big endian" or "little endian."

Compatible Computer Types for Release 6.12 through SAS 9.2

The tables in this section show the compatible computer types for Release 6.12 through SAS 9.2. After each table, a brief explanation is provided.

Table 2.2 Compatible Computer Types for Sharing Release 6.12 Files

Bits	Compatible Computer Types
32	Intel ABI, Linux, HP-UX, Solaris, AIX, IRIX
64	Tru64 UNIX

You can move a Release 6.12 SAS data set that was created on a 32-bit HP-UX host to a 32-bit AIX host using NFS, FTP, or CD. Because HP-UX and AIX are compatible computer types, you can use the V6 or V6TAPE engine to read the HP-UX data set on the AIX host.

The same 32-bit HP-UX data set can be moved to a 32-bit Intel ABI host. However, because these computer types are incompatible, you cannot use the V6 or V6TAPE engine to read the HP-UX data set.

For information about reading Release 6 data sets, see “Reading Version 6 Files” on page 46.

Table 2.3 Compatible Computer Types for Sharing Release 7 through SAS 9.2 Files

Bits	Platform	Compatible Computer Types
32	big endian	HP-UX for HP 9000 servers, Solaris for SPARC, AIX, and IRIX
32	little endian	Intel ABI, Linux
64	big endian	HP-UX for HP 9000 servers, Solaris for SPARC, AIX, and HP-UX for HP Integrity Servers
64	little endian	Linux Itanium, Tru64 UNIX, Solaris for x64, Linux x64

Note: In Release 8.2, both 32-bit and 64-bit SAS were available for the AIX, HP-UX, and Solaris operating environments. In SAS 9, SAS is 64-bit for these environments. Δ

You can move a Version 8 data set that was created on a 32-bit Solaris host to a 32-bit HP-UX host using methods such as NFS, FTP, or CD. Because this data set was created on compatible computer types, you will be able to read this data set in SAS.

The same 32-bit Solaris data set can be moved to a 64-bit HP-UX host. Because these computer types are incompatible, SAS will use Cross-Environment Data Access (CEDA) to read this data set. For more information, see “Reading Version 8 or Later Files from Incompatible Computer Types” on page 46.

Determining Compatible Computer Types in SAS 9.2

In SAS 9.2, the **Data Representation** field of the PROC CONTENTS output shows the compatible computer types for a SAS file. The following is a portion of the PROC CONTENTS output.

Output 2.1 Portion of PROC CONTENTS Output Using the V9 Engine

The CONTENTS Procedure			
Data Set Name	CLASSES.MAJORS	Observations	5
Member Type	DATA	Variables	6
Engine	V9	Indexes	0
Created	Monday, February 24, 2008 14:30:19	Observation Length	48
Last Modified	Monday, February 24, 2008 14:30:19	Deleted Observations	0
Protection		Compressed	NO
Data Set Type		Sorted	NO
Label	Second Data Set		
Data Representation	HP_UX_64, RS_6000_AIX_64, SOLARIS_64, HP_IA64		
Encoding	latin1 Western (ISO)		

In this example, the **Data Representation** field shows that the compatible computer types for this data set are: HP-UX 64 bit (**HP_UX_64**), AIX 64-bit (**RS_6000_AIX_64**), Solaris 64-bit (**SOLARIS_64**), and HP-UX Itanium 64-bit (**HP_IA64**). You can use NFS, FTP, or CD to move any data set to any computer and SAS will load the data set. If one data set has the same data representation, then SAS can read the data set natively. Otherwise, SAS must use CEDA. For more information about CEDA, see “Reading Version 8 or Later Files from Incompatible Computer Types” on page 46.

In SAS 9.2, you can use the Solaris x64 platform. In this case, the **Data Representation** field shows the following compatible computer types: **SOLARIS_X86_64**, **LINUX_X86_64**, **ALPHA_TRU64**, and **LINUX_IA64**.

The following table lists the possible values for the **Data Representation** field for SAS 9.2 and the corresponding computer types.

Table 2.4 Data Representation Value for Each Computer Type in SAS 9.2

Data Representation Value	Corresponding Computer Type
HP_UX_64	HP-UX PA-RISC 64-bit
HP_IA64	HP-UX Itanium processor family
LINUX	Linux on Intel 32-bit hardware
LINUX_X86_X64	Linux on 64-bit hardware
RS_6000_AIX_64	AIX 64-bit
SOLARIS_X64	Solaris for x64_x86
SOLARIS_64	Solaris 64-bit on SPARC

Note: The **Encoding** value might affect your ability to move SAS files between compatible computer types. It is important to note this value when you are transporting SAS files between hosts. For more information about encoding, see *SAS National Language Support (NLS): Reference Guide*. Δ

What Is File Migration?

File migration moves libraries forward to a new release of SAS. In many cases, SAS files from previous releases or from other hosts are compatible with SAS 9.2. If the files are not compatible, you can use PROC MIGRATE, a utility procedure, to migrate your files and libraries. Information to consider when you migrate your files includes the release of SAS in which your data currently resides, what member types exist in your libraries, and whether you must move members from 32-bit libraries to 64-bit libraries. PROC MIGRATE streamlines the process of moving files forward.

For information about using PROC MIGRATE and the Compatibility Calculator, see Migration at support.sas.com/migration.

Benefits of Migrating SAS Files

Migrating SAS files enables you to do the following:

- have Update access to unsupported data files
- have access to indexes, integrity constraints, and other features
- use long names for formats and informats
- use more than 32,767 variables
- use suppressed transcoding of a specified variable
- avoid the overhead of reading or writing to 32-bit files in a 64-bit SAS session

How to Migrate a SAS Library

To migrate a SAS library, use the MIGRATE procedure. If you are migrating from a 32-bit to a 64-bit environment, and catalogs are present in the library, you must have access to a 32-bit Release 8 SAS/CONNECT or SAS/SHARE server.

Note: If Release 8.2 files were created on a 64-bit UNIX computer, then these files are native on UNIX in SAS 9 for computers with the same data representation as those files that were created in Release 8.2. You do not need to migrate these files. △

For information about the MIGRATE procedure and how to use the PROC MIGRATE Calculator, see Migration at support.sas.com/migration.

Additional Resources

- For more information about the MIGRATE procedure and file compatibility, see Migration at support.sas.com/rnd/migration.
- For more information about reading Version 6 data sets, see “Reading Version 6 Files” on page 46.
- For more information about CEDA, see *SAS Language Reference: Concepts*.

Creating a SAS File to Use with an Earlier Release

The V9 and V9TAPE engines differ slightly from previous SAS engines. These engines support longer format and informat names than previous SAS engines. For a discussion about how to ensure compatibility between releases, see *SAS Language Reference: Concepts*. You can also go to Migration at support.sas.com/migration for information about cross-release compatibility.

Reading SAS Files from Previous Releases or from Other Hosts

Reading Version 6 Files

Using the V6 and V6TAPE read-only engines, SAS can read Release 6 data sets that were created by compatible computer types. In most cases, SAS invokes the V6 engine automatically and you do not have to specify it. The following examples demonstrate how you can use the V6 engine.

- If you are running SAS 9.2 on Linux, you can use the V6 engine to read Version 6 data sets that were created with any Intel ABI release of SAS, such as SCO UNIX.
- If you are running SAS 9.2 on HP-UX, you can use the V6 engine to read Release 6 data sets that were created on HP-UX, Solaris, AIX, or IRIX.

For a list of the compatible computer types for V6, see “Compatible Computer Types for Release 6.12 through SAS 9.2” on page 42. For more information about the compatibility of Version 6 files with SAS 9.2, see *SAS Language Reference: Concepts*.

Reading Version 8 or Later Files from Compatible Computer Types

If your files were created in 64-bit SAS, they are compatible with SAS 9.2. You do not need to use CEDA to read your files. To view tables that show compatible computer types for Release 6.12 through SAS 9.2, see “Compatible Computer Types for Release 6.12 through SAS 9.2” on page 42.

Reading Version 8 or Later Files from Incompatible Computer Types

Compatibility of Existing SAS Files with SAS 9.2

In Release 8.2, both 32-bit and 64-bit SAS were available for the AIX, HP-UX, and Solaris operating environments. In SAS 9, SAS for these environments is 64-bit only. Some SAS files that were created in 32-bit releases of SAS cannot be read by the V9 engine.

SAS automatically tries to use CEDA to read data sets. If you use CEDA to read a data set, and include `msglevel=i` in your code, then SAS writes a note to the log.

Release 8.2 files that are 64-bit and that were created on a UNIX computer are native on UNIX for SAS 9. You do not need to migrate these files.

The following table lists the supported processing for each SAS file under CEDA.

Table 2.5 Supported Processing for Release 8 32-Bit Files in SAS 9

File Type	Support
SAS files	input processing, output processing ¹
MDDb file	input processing
PROC SQL view	input processing
SAS/ACCESS view for Oracle or Sybase	input processing

File Type	Support
SAS/ACCESS view other than for Oracle or Sybase	no support
SAS catalog	no support
stored compiled DATA step program	no support
DATA step view	no support
item store	no support

1 In SAS 9, if you create a new data file from the 32-bit file, the new file is generally 64-bit. For more information about CEDA, see *SAS Language Reference: Concepts* or Migration at support.sas.com/migration.

Accessing Version 8 or Later Files with CEDA

CEDA enables a SAS data set that was created in Version 8 or later in any directory-based operating environment (such as UNIX, OpenVMS, and Windows) to be read by a SAS session that is running in another directory-based environment. In SAS 9.2, if you try to access a data set that was created in a previous release, then SAS automatically uses CEDA to process the file. For example, if you are running SAS 9.2 on Linux, SAS will use CEDA to process a data set that was created in Release 8 on a 64-bit Solaris host. With CEDA, you have Read and Write access to these files; however, you will not be able to update the file. For information about compatibility, see Migration at support.sas.com/migration.

If you need to access 32-bit SAS data sets, SAS/ACCESS views from Oracle or Sybase, SQL views, or MDDB files from a 64-bit SAS session, then you can access these files using CEDA. CEDA provides Read and Write access to these files. However, CEDA does not support Update processing. CEDA consumes additional resources each time that you read or write to these files.

Catalogs and other SAS files (not including SAS data sets) contain data structures that are known only to the application that created them. These files might contain data objects other than character or numeric objects and, therefore, cannot be shared between 64-bit SAS and earlier 32-bit releases of SAS.

Referring to SAS Files by Using Librefs in UNIX Environments

Techniques for Referring to a SAS File

If you want to read or write to a permanent SAS file, you can refer to the SAS file in one of two ways:

- Refer to the data file directly by using its pathname in the appropriate statements (such as DATA, SET, MERGE, UPDATE, OUTPUT, and PROC).
- Assign a libref to the SAS library (directory) that contains the data file and use the libref as the first level of a two-level filename.

What Is a Libref?

A libref is an alias that you can use to refer to the library during a SAS session or job. You will probably want to use a libref when one of the following is true:

- The data file pathname is long and must be specified several times within a program.
- The pathname might change. If the pathname changes, you need to change only the statement assigning the libref, not every reference to the file.
- Your application will be used on other platforms. Using librefs makes it easier to port an application to other operating environments.
- You need to concatenate libraries. See “Assigning a Libref to Several Directories (Concatenating Directories)” on page 51 for more information.

Librefs can be stored in the SAS registry. See “Customizing Your SAS Registry Files” on page 17 for more information.

Assigning Librefs

Methods for Assigning Librefs

You can use any of the following methods to assign a SAS libref:

- the LIBNAME statement
- the LIBNAME function
- the DMLIBASSIGN command
- the LIBNAME window
- the SAS Explorer window

A libref assignment remains in effect for the duration of the SAS job, session, or process unless you clear the libref or use the same libref in another LIBNAME statement or LIBNAME function.

If you assign a libref from a SAS process, that libref is valid only within that SAS process. If you clear a libref from within a SAS process, that libref is not cleared from other SAS processes.

Using the LIBNAME Statement

The LIBNAME statement identifies a SAS library to SAS, associates an engine with the library, enables you to specify options for the library, and assigns a libref to it. For information about LIBNAME statement syntax, see “LIBNAME Statement” on page 328.

Using the LIBNAME Function

The LIBNAME function takes the same arguments and options as the LIBNAME statement. For more information about the LIBNAME function, see *SAS Language Reference: Dictionary*.

Use the DMLIBASSIGN Command

Perform the following steps to assign a libref using the DMLIBASSIGN command:

- 1 Issue the DMLIBASSIGN command in the command window.
The New Library dialog box opens.
- 2 Specify the libref in the **Name** field.
- 3 Specify an engine for the libref in the **Engine** field by selecting the default engine or another engine from the menu. Depending on the engine that you select, the fields in the **Library Information** area might change.
- 4 Click **Enable at startup** to assign this libref when you invoke SAS.
- 5 Specify the necessary information for the SAS library in the **Library Information** area. Depending on the engine that you select, there might not be a **Path** field available for input.
- 6 Specify LIBNAME options in the **Options** field. These options can be specific to your host or engine, including options that are specific to a SAS engine that accesses another software vendor's relational database system.
- 7 Click **OK**.

Use the LIBNAME Window

Perform the following steps to assign a libref from the LIBNAME window:

- 1 Issue the LIBNAME command in the command window.
The LIBNAME window opens.
- 2 From the **File** menu, select **New**.
The New Library dialog box opens.
- 3 Complete the fields in the New Library dialog box as described in “Use the DMLIBASSIGN Command” on page 48.
- 4 Click **OK**.

Use the SAS Explorer Window

Perform the following steps to assign a libref from the SAS Explorer window:

- 1 From the **File** menu, select **New** when the **Libraries** node in the tree structure is active.
The New dialog box opens.
- 2 Select **Library**, and then click **OK**.
The New Library dialog box opens.
- 3 Complete the fields in the New Library dialog box as described in “Use the DMLIBASSIGN Command” on page 48.
- 4 Click **OK**.

Permanently Assigning a Libref

You might want to save a libref so that it is valid between SAS sessions. You can assign a libref permanently by using one of the following methods:

- Specify the LIBNAME statement or LIBNAME function in an autoexec file. For more information, see the LIBNAME function in *SAS Language Reference: Dictionary* or the “LIBNAME Statement” on page 328.
- Select **Enable at startup** when you assign a libref using the DMLIBASSIGN command, LIBNAME window, or SAS Explorer Window. Selecting this option will save the libref in the SAS registry. For more information about these methods, see “Assigning Librefs” on page 48.

- Use environment variables as librefs. Include these environment variables in your start-up files so that these variables are set when SAS is invoked.

Accessing a Permanent SAS Library by Using a Libref

After you have defined a libref, you can use the libref in one of two ways to access a permanent SAS library:

- as the first level of a two-level SAS filename:

libref.member-name

where *libref* is the first-level name referring to the directory where the file is stored, and *member-name* is the name of the file being read or created.

- as the value of the USER= option. (See “Using One-Level Names to Access Permanent Files (User Library)” on page 59 for information.)

For example, these SAS statements access the data file Final.sas7bdat in the Sales library that is stored in the `/users/myid/mydir` directory:

```
libname sales '/users/myid/mydir';
data sales.final;
```

Specifying Pathnames in UNIX Environments

Rules for Specifying Directory and Pathnames

Whether you specify a data filename directly in the various SAS statements, or you specify the library name in a LIBNAME statement and then refer to the libref, the same rules apply for specifying UNIX directory and file pathnames.

Specify directory and file pathnames in quotation marks. The level of specification depends on your current directory.

Example 1: Access a File That Is Not in the Current Directory

If `/u/2007/budgets` is not your current directory, then to access the data file named May, you must specify the entire pathname:

```
data '/u/2007/budgets/may';
```

If you wanted to use a libref, you would specify:

```
libname budgets '/u/2007/budgets';
data budgets.may;
```

Example 2: Access a File in the Current Directory

If `/u/2007/budgets` is your current directory, you could specify only the filenames:

```
data 'quarter1';
merge 'jan' 'feb' 'mar';
run;
```


Note: If you omit the quotation marks, then SAS assumes that these data sets are stored in the **Saswork** directory. △

If you wanted to use a libref, you would specify:

```
libname budgets '.';
data budgets.quarter1;
merge budgets.jan budgets.feb budgets.mar;
run;
```

Valid Character Substitutions in Pathnames

You can use the character substitutions in the following table to specify pathnames.

Table 2.6 Character Substitutions in Pathnames

Characters	Meaning
~/	\$HOME/ Can be used only at the beginning of a pathname.
~ <i>name</i> /	<i>name</i> 's home directory (taken from file <code>/etc/passwd</code>). Can be used only at the beginning of a pathname.
!sasroot	name of sasroot directory (see Appendix 1, "The !SASROOT Directory," on page 421). Specified only at the beginning of a pathname.
.	current working directory.
..	parent of current working directory.
\$ <i>VARIABLE</i>	environment variable <i>VARIABLE</i> .

Assigning a Libref to Several Directories (Concatenating Directories)

Introduction to Concatenating Directories

You can use the LIBNAME statement to assign librefs and engines to one or more directories, including the **work** directory.

If you have SAS data sets located in multiple directories, you can treat these directories as a single SAS library by specifying a single libref and concatenating the directory locations, as in the following example:

```
libname income ('/u/2007/revenue', '/u/2007/costs');
```

This statement indicates that the two directories, `/u/2007/revenue` and `/u/2007/costs`, are to be treated as a single SAS library.

If you have already assigned librefs to your SAS libraries, you can use these librefs to indicate that you want to concatenate the libraries, as in this example:

```
libname income ('/u/2007/corpsale', '/u/2007/retail');
libname costs ('/u/2007/salaries', '/u/2007/expenses');
libname profits (income, costs, '/u/2007/capgain');
```

This statement indicates that the five directories, `/u/2007/corpsale`, `/u/2007/retail`, `/u/2007/salaries`, `/u/2007/expenses`, and `/u/2007/capgain`, are to be treated as a single SAS library.

How SAS Accesses Concatenated Libraries

When you concatenate SAS libraries, SAS uses a protocol for accessing the libraries, which depends on whether you are accessing the libraries for read, write, or update. (A *protocol* is a set of rules.)

SAS uses the protocol in the following sections to determine which directory is accessed. (The protocol illustrated by these examples applies to all SAS statements and procedures that access SAS files, such as the DATA, UPDATE, and MODIFY statements in the DATA step, and the SQL and APPEND procedures.)

Accessing Files for Input and Update

When a SAS data set is accessed for input or update, the first SAS data set that is found by that name is the one that is accessed. For example, if you submit the following statements and the data set `old.species` exists in both directories, the one in the `mysasdir` directory is the one that is printed:

```
libname old ('mysasdir','saslib');
proc print data=old.species;
run;
```

The same would be true if you opened `old.species` for update with the FSEDIT procedure.

Accessing Files for Output

If the data set is accessed for output, it is always written to the first directory, provided that the directory exists. If the directory does not exist, an error message is displayed. For example, if you submit the following statements, SAS writes the `old.species` data set to the first directory (`mysasdir`), and replaces any existing data set with the same name:

```
libname old ('mysasdir','saslib');
data old.species;
x=1;
y=2;
run;
```

If a copy of the `old.species` data set exists in the second directory, it is not replaced.

Accessing Data Sets with the Same Name

If you use the DATA and SET statements to access data sets with the same name, the DATA statement uses the output rules and the SET statement uses the input rules. When you execute the following statements, assume that `test.species` originally exists only in the second directory, `mysasdir`. Execute the following statements:

```
libname test ('sas','mysasdir');
data test.species;
set test.species;
```

```

if value1='y' then
  value2=3;
run;

```

The DATA statement opens **test.species** for output according to the output rules; that is, SAS opens a data set in the first of the concatenated libraries (**sas**). The SET statement opens the existing **test.species** data set in the second directory (**mysasdir**), according to the input rules. Therefore, the original **test.species** data set is not updated. After the DATA step executes, two **test.species** data sets exist, one in each directory.

Using Multiple Engines for a Library in UNIX Environments

You can assign multiple librefs to a single directory, and specify a different engine with each libref. For example, after the following statements are executed, data sets that are referenced by **one** are created and accessed using the default engine, while data sets that are referenced by **two** are created and accessed using the sequential engine:

```

libname one v9 '/users/myid/educ';
libname two v8 '/users/myid/educ';

```

Note: Keeping different types of libraries in one directory is not recommended because you must remember the appropriate engine for accessing each library. SAS cannot determine the right engine for accessing libraries in a directory that contains libraries of different types. See “Omitting Engine Names from the LIBNAME Statement” on page 331 for more information. Δ

Using Environment Variables as Librefs in UNIX Environments

An environment variable can be used as a libref. The variable name must be in all uppercase characters, and the variable value must be the full pathname of the directory; that is, the name of the directory must begin with a slash.

Note: SAS on UNIX does not support the assignment of the User libref using the USER environment variable. Δ

Suppose you want to use the library in **/users/mydir/educ**, and you want to refer to it with the EDUC environment variable. You can define the variable at two times:

- Before you invoke SAS. See “Defining Environment Variables in UNIX Environments” on page 23. For example, in the Korn shell, you would use:

```
export EDUC=/users/mydir/educ
```

- After you invoke SAS, you can use the X statement (see “Executing Operating System Commands from Your SAS Session” on page 14) and the SAS **setenv** command:

```
x setenv EDUC /users/mydir/educ;
```

You cannot specify an engine when you define a libref as an environment variable, so SAS determines which engine to use as described in “Omitting Engine Names from the LIBNAME Statement” on page 331.

After the libref is defined, you can use it to access data sets stored in the library:

```
proc print data=educ.class;
run;
```

Note: If a variable and a libref have the same name, but refer to different libraries, SAS uses the libref. \triangle

Librefs Assigned by SAS in UNIX Environments

SAS automatically defines three librefs:

Sashelp

contains a group of catalogs that contain information that is used to control various aspects of your SAS session. The Sashelp library is in the **!SASROOT** directory. See Appendix 1, “The !SASROOT Directory,” on page 421.

Sasuser

contains SAS catalogs that enable you to tailor features of SAS (such as window size, font settings, and printer entries) for your needs. If the defaults in the Sashelp library are not suitable for your applications, you can modify them and store your personalized defaults in your Sasuser library.

Work

is the temporary, or scratch, library automatically defined by SAS at the beginning of each SAS session or job. The Work library stores two types of temporary files: those files you create, and those files that are created internally by SAS as part of normal processing.

These librefs and the library libref are reserved librefs. If your site also has SAS/GRAPH software, the maps libref might be automatically defined. All of these libraries are described in *SAS Language Reference: Dictionary*. Sasuser and Work have operating system dependencies.

Sasuser Library

What Is the Sasuser Library?

The Sasuser library contains the customizations (such as window size and positioning, colors, fonts, and printer entries) that you specified for your SAS session. When you invoke SAS, it looks for the **Sasuser** directory to find these customizations. If this directory does not exist, SAS uses the SASUSER system option to create it. The default directory is set in the system configuration file (sasv9.cfg) and is usually similar to the following:

```
-sasuser ~/sasuser.v91
```

This specification tells SAS to create a directory for the Sasuser libref in your home directory. To determine the value of this directory for your system, use PROC OPTIONS or **libname sasuser LIST**.

You can permit read-only access to the Sasuser library by using the RSASUSER system option. See Chapter 18, “System Options under UNIX,” on page 341 for information about the SASUSER and RSASUSER system options.

After the Sasuser library has been created, SAS automatically assigns the same Sasuser libref to it each time you start a SAS session. It cannot be cleared or reassigned during a SAS session. If you delete the library, SAS re-creates it the next time you start a session. Because SAS assigns the libref for you, you do not need to use a LIBNAME statement before referencing this library.

Contents of the Sasuser Library

Your customizations are stored in one of the following locations in the Sasuser library:

- Sasuser.Profile catalog
- Sasuser.Registry catalog
- Sasuser.Prefs file

Sasuser.Profile Catalog

Overview of the Sasuser.Profile Catalog

The Sasuser.Profile catalog is the profile.sas7bcat file in your Sasuser library. This catalog enables you to customize the way you work with SAS. SAS uses this catalog to store function key definitions, fonts for graphics applications, window attributes, and other information from interactive windowing procedures. SAS saves changes that you make to function key definitions, window attributes (such as size, color, and position), PMENU settings, and so on, in the Sasuser.Profile catalog. The information in the Sasuser.Profile catalog is accessed automatically by SAS when you need it for processing.

How SAS Accesses the Sasuser.Profile Catalog

SAS creates the Sasuser.Profile catalog the first time it tries to find it and it does not exist. If you are using an interactive windowing environment, then creating the Sasuser.Profile catalog occurs during system initialization in your first SAS session. If you are using one of the other modes of execution, the Sasuser.Profile catalog is created the first time you execute a SAS procedure that requires it.

When the Sasuser.Profile Catalog Does Not Exist

If the Sasuser.Profile catalog does not exist, then, at invocation, SAS checks for the Sashelp.Profile catalog. (This catalog will exist only if you have copied your Sasuser.Profile catalog to the Sashelp library.) If the Sashelp.Profile catalog exists, then SAS copies it to the Sasuser library, and this catalog becomes your new Sasuser.Profile catalog. If the Sashelp.Profile catalog does not exist, then SAS creates Sasuser.Profile using the default settings for a SAS session. The default settings for your SAS session are stored in several catalogs in the Sashelp library. If you make changes to key settings or other options, then the new information is stored in your Sasuser.Profile catalog. To restore the original default settings to the Sasuser.Profile catalog, use the CATALOG procedure or the CATALOG window to delete entries from your Sasuser.Profile catalog. By default, SAS then uses the corresponding entry from the Sashelp library.

Checking for an Uncorrupted Sasuser.Profile Catalog

When you invoke SAS, SAS checks for an existing, uncorrupted Sasuser.Profile catalog. If the catalog is found, SAS copies the Sasuser.Profile catalog to Sasuser.Profback. This backup catalog is used if Sasuser.Profile becomes corrupted.

If you invoke SAS and determine that your customizations have been lost, then your Sasuser.Profile catalog is either corrupted or locked by another SAS session that was started with the same user ID. If either of these conditions are true, then SAS uses Sashelp.Profile or Sasuser.Profback to replace the locked or corrupted Sasuser.Profile catalog.

If Your Sasuser.Profile Catalog Is Locked or Corrupted

If your Sasuser.Profile catalog is locked, then SAS checks for Sashelp.Profile. If Sashelp.Profile exists, SAS copies it to Work.Profile, and then saves the customizations to the Work.Profile catalog instead of the Sasuser.Profile catalog. This Work.Profile catalog is used for the duration of the SAS session. Because the contents of the **work** directory are temporary, any customizations that you save to the Work.Profile catalog will be lost at the end of the SAS session.

If your Sasuser.Profile catalog is corrupted, SAS copies the corrupted catalog to Sasuser.Badpro.SAS, and then checks for Sasuser.Profback. If Sasuser.Profback exists, then SAS copies it to Sasuser.Profile. Any changes that you made to the Sasuser.Profile catalog during the previous session will be lost. If your Sasuser.Profile catalog is being used by multiple SAS sessions, then you can specify the RSASUSER system option to permit read-only access to the Sasuser library. Because this permission is read-only, you will not be able to save any customizations to your Sasuser.Profile catalog during that SAS session.

For more information about the Sasuser.Profile catalog and its related catalogs, as well as information about recovering locked or corrupted profile catalogs, see *SAS Language Reference: Concepts*.

Sasuser.Registry Catalog

Overview of the Sasuser.Registry Catalog

The Sasuser.Registry catalog is the registry.sas7bitm file in your Sasuser library. If you change any Universal Printing entries or libref assignments during a SAS session, then SAS saves the changes in the Sasuser.Registry catalog.

How SAS Accesses the Sasuser.Registry Catalog

At invocation, SAS looks in the **sasuser** directory to see whether it can write to the Sasuser.Registry catalog. If SAS cannot write to this catalog, then the following warning appears in the SAS log:

```
WARNING: Unable to open SASUSER.REGISTRY. WORK.REGISTRY will be used instead.
NOTE: All registry changes will be lost at the end of the session.
```

If SAS can read the Sasuser.Registry catalog, then SAS copies the Sasuser.Registry catalog to create a Work.Registry catalog (in the Work library). This Work.Registry catalog will be used for the duration of the SAS session. Because the contents of the Work library are temporary, then any customizations that you save to the Work.Registry catalog will be lost at the end of the SAS session. However, the customizations saved in the Sasuser.Registry catalog will still exist.

If SAS cannot read the Sasuser.Registry catalog, then SAS creates the Work.Registry catalog using the default settings for a SAS session. In this case, SAS issues an additional warning to the SAS log:

```
WARNING: Unable to copy SASUSER.REGISTRY to WORK.REGISTRY.
```

Sasuser.Prefs File

The settings that you specify in the Preferences dialog box (with the exception of resources on the **General** tab) are saved in the Sasuser.Prefs file. For more information about these resources, see “Modifying X Resources Through the Preferences Dialog Box” on page 167.

Work Library

The Work library is the temporary library that is automatically defined by SAS at the beginning of each SAS session or job. The Work library stores temporary SAS files that you create, as well as files created internally by SAS.

To access files in the Work library, specify a one-level name for the file. The libref **Work** is automatically assigned to these files unless you have assigned the User libref.

When you invoke SAS, it assigns the Work libref to a subdirectory of the directory specified in the **WORK** system option described in Chapter 18, “System Options under UNIX,” on page 341. This subdirectory is usually named **SAS_workcode_nodename**, where:

workcode

is a 12-character code. The first four characters are randomly generated numbers. The next eight characters are based on the hexadecimal process identification number of the SAS session.

nodename

is the name of the UNIX computer where the SAS process is running.

This libref cannot be cleared or reassigned during a SAS session.

The **WORKINIT** and **WORKTERM** system options control the creation and deletion of the Work library. See *SAS Language Reference: Dictionary* for information.

Note: If a SAS session is terminated improperly (for example, using the **kill -9** command), SAS will not delete the **SAS_workcode_nodename** directory. You might want to use the **cleanwork** command to delete stragglng directories (see Appendix 2, “Tools for the System Administrator,” on page 423). \triangle

Multiple Work Directories

SAS can make the distribution of Work libraries dynamic by distributing Work libraries across several directories. This functionality eliminates the potential problem of filling up a single volume with all of the Work directories.

The WORK system option contains the PATHNAME argument, which can be a directory, or a file that contains a list of directories, that SAS can use for allocating Work libraries. Individual Work libraries will still reside in a single directory. You can use the WORK system option in a configuration file or in the command line.

When the argument to WORK is a list of directories in a file, you can specify a method for choosing which directory to use for WORK. If you specify METHOD=RANDOM, then SAS chooses at random a directory from the list of available directories. If you choose METHOD=SPACE, then SAS chooses the directory that has the most available space.

For more information, see “WORK System Option” on page 415 .

Using One-Level Names to Access Permanent Files (User Library)

Introduction to One-Level Names

SAS data sets are referenced with a one- or two-level name. The two-level name has the form *libref.member-name*, where *libref* refers to the SAS library in which the data set resides, and *member-name* refers to the particular *member* within that library. The one-level name has the form *member-name* (without a *libref*). In this case, SAS stores the files in the temporary Work library. To override this action and have files with one-level names stored in a permanent library, you must first assign the User libref to an existing directory. To refer to temporary SAS files while User is assigned, use a two-level name with Work as the libref.

Techniques for Assigning the User Libref

You have three ways to assign the User libref:

- Assign the User libref directly using the LIBNAME statement:

```
libname user '/users/myid/mydir';
```

- Specify the USER= system option before you start the SAS session. For example, you can assign the User libref when you invoke SAS:

```
sas -user /users/myid/mydir
```

- Specify the USER= system option after you start the SAS session. First, assign a libref to the permanent library. Then, use the USER= system option in an OPTIONS statement to equate that libref to User. For example, these statements assign the libref User to the directory with libref mine:

```
libname mine '/users/myid/mydir';
options user=mine;
```

See Chapter 18, “System Options under UNIX,” on page 341 for information about the USER= system option.

Note: SAS on UNIX does not support the assignment of the User libref using the USER environment variable. \triangle

Accessing Disk-Format Libraries in UNIX Environments

You will probably create and access libraries on disk more than any other type of library. The default engine and the compatibility engines allow Read, Write, and Update access to SAS files on disk. They also support indexing and compression of observations.

In the following example, the In libref is assigned to a directory that contains the Stats1 data set:

```
libname in '/users/myid/myappl';
proc print data=in.stats1;
run;
```

Remember, a *SAS-data-library* must exist before SAS can read from or write to this directory. For example, if you want to create the SAS data set Orders in a directory, use the X statement to issue the **mkdir** UNIX command. Then, you can use the LIBNAME statement to associate the libref with the directory:

```
x mkdir /users/publish/books;
libname books '/users/publish/books';
data books.orders;
... more SAS statements ...
run;
```

By default, the LIBNAME statement associates the V9 engine with the directory.

Accessing Sequential-Format Libraries in UNIX Environments

Benefits and Limitations of Sequential Engines

The sequential engines enable you to access libraries in sequential format on tape or disk. The sequential engines do not support indexing and compression of observations.

Note: Before using sequential engines, read the information about sequential libraries in *SAS Language Reference: Concepts*. \triangle

Reading and Writing SAS Files on Tape

Using a Staging Directory

Use a staging directory so that files can be processed directly from disk. You can use the UNIX **tar** command to move SAS data sets between the staging directory and tape. (Do not use the UNIX **cp** command.)

Syntax of the LIBNAME Statement

To access SAS 9.2 files on tape, you can specify the V9TAPE or TAPE engine in the LIBNAME statement:

```
LIBNAME libref V9TAPE 'tape-device-pathname';
```

The *tape-device-pathname* must be a pathname for a tape device; it should be the name of the special file that is associated with the tape device. (Check with your UNIX system administrator for information.) The name must be enclosed in quotation marks. You cannot specify remote tape devices in the LIBNAME statement.

Example: Assign a Libref to the Tape Device

The following LIBNAME statement assigns the libref Seq2 to the `/dev/tape2` tape device. Because the tape device is specified, the engine does not have to be specified.

```
libname seq2 '/dev/tape2';
```

Writing Sequential Data Sets to Named Pipes

Why Use Named Pipes?

You can send output to and read input from the operating environment by using named pipes. For example, you might want to compress a data set or send it to a tape management system without creating intermediate files.

Syntax of the LIBNAME Statement

You can read from and write to named pipes from within your SAS session by specifying the pipe name in the LIBNAME statement:

```
LIBNAME libref <TAPE> 'pipename';
```

Because you cannot position a pipe file, SAS uses the TAPE engine to ensure sequential access. You do not have to specify the engine name; TAPE is assumed.

Example: Creating a SAS Data Set Using a Named Pipe

To create a SAS data set and compress the data set without creating an intermediate, uncompressed data set, create a named pipe (such as **mypipe**) and enter the **compress** command:

```
mknod mypipe p compress <mypipe >sasds.z
```

In your SAS session, assign a libref to the pipe and begin writing to the data set:

```
libname x 'mypipe';
data x.a;
  ...more SAS statements...
output;
run;
```

The data is sent to **mypipe** and then compressed and written to the data set. When SAS closes the data set, compression finishes, and you have a compressed, sequential data set in **sasds.z**.

If you begin writing to a named pipe before the task on the other end (in this case, the **compress** command) begins reading, your SAS session will be suspended until the task begins to read.

Accessing BMDP, OSIRIS, or SPSS Files in UNIX Environments

Introduction to the BMDP, OSIRIS, and SPSS Files

SAS includes three interface library engines, BMDP, OSIRIS, and SPSS, that enable you to access external data directly from a SAS program. All of these engines are read-only.

Because they are sequential, these engines cannot be used with the POINT= option in the SET statement, or with the FSBROWSE, FSEDIT, or FSVIEW procedures. You can use PROC COPY, PROC DATASETS, or a DATA step to copy a BMDP or OSIRIS system file or an SPSS export file to a SAS data set, and then perform these functions on the SAS data set. Also, some procedures (such as PROC PRINT) give a warning message about the engine being sequential.

With these engines, the physical filename that is associated with a libref is an actual filename, not a directory. This association is an exception to the rules concerning librefs.

You can use the CONVERT procedure to convert BMDP, OSIRIS, and SPSS files to SAS files. See “CONVERT Procedure” on page 295 for more information.

The BMDP Engine

What Is the BMDP Engine?

The BMDP interface library engine enables you to read BMDP files from the BMDP statistical software application directly from a SAS program. The BMDP engine is a read-only engine. The following discussion assumes that you are familiar with the BMDP save file terminology.*

Note: This engine is available for AIX, HP-UX, and Solaris. \triangle

Syntax for Accessing BMDP Save Files

To read a BMDP save file, issue a LIBNAME statement that explicitly specifies the BMDP engine. In this case, the LIBNAME statement has the following form:

```
LIBNAME libref BMDP 'filename';
```

where:

libref specifies a SAS libref.

filename specifies a BMDP physical filename.

Note: If the libref appears previously as a fileref, omit *filename* because SAS uses the physical filename that is associated with the fileref. \triangle

This engine can read save files that are created only on UNIX.

* See the documentation provided by BMDP Statistical Solutions on the Web site for more information.

Because a single physical file can contain multiple save files, you reference the CODE= value as the member name of the data set within the SAS language. For example, if the save file contains CODE=ABC and CODE=DEF, and the libref is MyLib, you reference the files as MyLib.ABC and MyLib.DEF. All CONTENT types are treated the same. Even if member DEF has the value CONTENT=CORR, it is treated as if the value was CONTENT=DATA.

If you know that you want to access the first save file in the physical file or if there is only one save file, refer to the member name as `_FIRST_`. This reference is convenient if you do not know the CODE= value.

Example: BMDP Engine

Assume that the physical file `mybmdp.dat` contains the save file ABC. The following SAS code associates the libref `mylib` with the BMDP physical file and executes the CONTENTS and PRINT procedures on the save file:

```
libname mylib bmdp 'mybmdp.dat';
proc contents data=mylib.abc;
run;
proc print data=mylib.abc;
run;
```

The following example uses the LIBNAME statement to associate the libref `mylib2` with the BMDP physical file. Then, it writes the data for the first save file in the physical file:

```
libname mylib2 bmdp 'mybmdp.dat';
proc print data=mylib2._first_;
run;
```

The OSIRIS Engine

What Is the OSIRIS Engine?

The Inter-University Consortium for Political and Social Research (ICPSR) uses the OSIRIS file format for distribution of its data files. SAS provides the OSIRIS interface library engine to support the many users of the ICPSR data and to be compatible with PROC CONVERT.

With the OSIRIS engine, you can read OSIRIS data and dictionary files directly from a SAS program. The following discussion assumes that you are familiar with the OSIRIS file terminology and structure. If you are not, refer to the documentation provided by ICPSR.

Notes About the OSIRIS Data Dictionary Files

Because OSIRIS software does not run outside the z/OS environment, the layout of an OSIRIS data dictionary is consistent across operating environments. However, the OSIRIS engine is designed to accept a data dictionary from any other operating environment on which SAS runs. It is important that the dictionary and data files not be converted from EBCDIC to ASCII; the engine expects EBCDIC data.

The dictionary file should consist of fixed-length records of length 80. The data file should contain records that are large enough to hold the data that is described in the dictionary.

Syntax for Accessing an OSIRIS File

To read an OSIRIS file, issue a LIBNAME statement that explicitly specifies the OSIRIS engine. In this case, the syntax of the LIBNAME statement has the following form:

```
LIBNAME libref OSIRIS 'data-filename' DICT='dictionary-filename';
```

where:

libref

specifies a SAS libref.

'data-filename'

specifies the physical filename of the data file.

If the libref appears also as a fileref, omit *data-filename*.

DICT='*dictionary-filename*'

specifies the physical filename of the dictionary file. If *dictionary-filename* is an environment variable or a fileref, do not enclose it in quotation marks. The DICT= option is required.

OSIRIS data files do not have member names. Therefore, use whatever member name you want.

To use the same dictionary file with different data files, use a separate LIBNAME statement for each one.

Example: OSIRIS Engine

In the following example, the data file is `/users/myid/osr/dat`, and the dictionary file is `/users/myid/osr/dic`. The example associates the libref `mylib` with the OSIRIS files, and executes the CONTENTS and the PRINT procedures:

```
libname mylib osiris '/users/myid/osr/dat'
      dict='/users/myid/osr/dic';
proc contents data=mylib._first_;
run;
proc print data=mylib._first_;
run;
```

The SPSS Engine

What Is the SPSS Engine?

The SPSS engine is a read-only engine. With the SPSS interface library engine, you can read only SPSS export files. This engine does not read SPSS-X native files.

Syntax for Accessing an SPSS Export File

To read an SPSS export file, issue a LIBNAME statement that explicitly specifies the SPSS engine. In this case, the syntax of the LIBNAME statement has the following form:

```
LIBNAME libref SPSS 'filename';
```

where:

libref

specifies a SAS libref.

'*filename*' specifies the physical filename.

Note: If the libref appears also as a fileref, omit *filename* because SAS uses the physical filename that is associated with the fileref. △

Export files must be created by the SPSS EXPORT command and can originate from any operating environment. Export files must be transported to and from your operating environment in ASCII format. If they are transported in binary format, other operating environments will not be able to read them.

Because SPSS-X files do not have internal names, refer to them by any member name you like. A common extension for export files is .por, but this extension is not required.

SPSS can have system-missing and user-defined missing data. When you use the SPSS engine or PROC CONVERT, the missing values (user-defined or system-missing) are converted to system-missing values. User-defined missing values have to be recoded as valid values. When the data set is converted, you can use PROC FORMAT to make the translation. For example, -1 to .A and -2 to .B.

Reformatting SPSS Files

SAS cannot use an SPSS file that contains a variable with a numeric format that has a larger number of decimal places than the width of the entire variable. For example, if an SPSS file has a variable with a width of 17 and has 35 decimal places, SAS will return errors when you try to run a DATA step on the file or view it with the table viewer. To use the SPSS file with SAS, you have to reformat the variable.

You can reformat the variable by reducing the number of decimal spaces to a value that fits within the width of the variable. In the following example the statement **revision=cat(format,format1, '.2');** converts the number of decimal spaces to 2. This value reduces the number of decimal spaces so that the number is not greater than the width of the variable.

```
libname abc spss 'FILENAME.POR';
proc contents data=abc._all_ out=new;
run;

filename sascode temp;
data _null_;
  set new;
  file sascode
  if formatd > formatl then do;
    revision=cat(format,format1, '.2');
    put 'format' +1 name +1 revision ';' ;
  end;
run;

data temp;
  set abc._all_;
  %inc sascode/source2;
run;
```

Note: The OPTIONS NOFMterr statement does not allow SAS to use a data set with a DATA step or the table viewer. You have to reformat numeric variables that have a larger decimal value than their width before you can use a DATA step or the table viewer. △

Example: SPSS Engine

The following example associates the libref `mylib` with the physical file `/users/myid/mydir/myspssx.por` to execute the `CONTENTS` and `PRINT` procedures on the export file:

```
libname mylib spss '/users/myid/mydir/myspssx.por';
proc contents data=mylib._first_;
proc print data=mylib._first_;
run;
```

In the next example, the `FILENAME` statement associates the fileref `mylib2` with the `/users/myid/mydir/aspssx.por` SPSS physical file, and the `LIBNAME` statement associates the libref with the SPSS engine. The `PRINT` procedure writes the data from the portable file.

```
filename mylib2 '/users/myid/mydir/aspssx.por';
libname mylib2 spss;
proc print data=mylib2._first_;
run;
```

Support for Links in UNIX Environments

SAS provides limited support for hard links and symbolic links in UNIX environments. You can create links that point to a SAS data set or SAS catalog. If you reference the link in a SAS program, SAS will follow the link to find the data set or catalog.

For example, you can create a symbolic link in the `/tmp` directory to the `/home/user/mydata.sas7bdat` data set by typing the following command at the UNIX prompt:

```
ln -s /home/user/mydata.sas7bdat /tmp/mydata.sas7bdat
```


The following SAS code uses the symbolic link in the `/tmp` directory to find the `mydata.sas7bdat` data set. This code does not change the symbolic link, but it does sort the data in the data set.

```
libname tmp '/tmp';

proc sort data=tmp.mydata;
  by myvariable;
run;
```

If you are running in the SAS windowing environment, you can use the SAS Explorer window to view the symbolic links that are stored within a specific directory. Any symbolic link that points to a nonexistent SAS file will have a file size of `0.0KB` and a modified date of `31DEC59:19:00:00`.

Note: SAS does not support links for a version data set or for a data set that has an index. \triangle



CHAPTER

3

Using External Files and Devices

<i>Introduction to External Files and Devices in UNIX Environments</i>	68
<i>Accessing an External File or Device in UNIX Environments</i>	69
<i>Specifying a Pathname or a Fileref</i>	69
<i>What Is a Fileref?</i>	69
<i>Specifying Pathnames in UNIX Environments</i>	70
<i>Rules for Specifying Pathnames</i>	70
<i>Omitting Quotation Marks in a Filename</i>	70
<i>Working with Mixed Case or Uppercase Filenames</i>	71
<i>Interpreting the Messages in the SAS Log</i>	71
<i>Using Wildcards in Pathnames (Input Only)</i>	71
<i>Descriptions of the Valid Wildcards</i>	71
<i>Example 1: Selecting Files by Including a Wildcard in a String</i>	72
<i>Example 2: Reading Each File in the Current Directory</i>	72
<i>Example 3: Wildcards in Filenames When Using Aggregate Syntax</i>	72
<i>Example 4: Associating a Fileref with Multiple Files</i>	72
<i>Assigning Filerefs to External Files or Devices with the FILENAME Statement</i>	73
<i>Introduction to the FILENAME Statement</i>	73
<i>Accessing DISK Files</i>	73
<i>Debugging Code with DUMMY Devices</i>	73
<i>Sending Output to PRINTER Devices</i>	73
<i>Using Temporary Files (TEMP Device Type)</i>	74
<i>Accessing TERMINAL Devices Directly</i>	74
<i>Assigning Filerefs to Files on Other Systems (FTP, SFTP, and SOCKET Access Types)</i>	74
<i>Concatenating Filenames in UNIX Environments</i>	75
<i>Assigning a Fileref to a Directory (Using Aggregate Syntax)</i>	75
<i>Introduction to Aggregate Syntax</i>	75
<i>Example 1: Referring to a File Using Aggregate Syntax</i>	75
<i>Example 2: Using Aggregate Syntax with Filerefs Defined by Environment Variables</i>	76
<i>Assigning a Fileref to Several Directories</i>	76
<i>Using Environment Variables to Assign Filerefs in UNIX Environments</i>	76
<i>Requirements for Variable Names</i>	76
<i>Reading a Data File</i>	77
<i>Writing to an External File</i>	77
<i>Filerefs Assigned by SAS in UNIX Environments</i>	77
<i>Filerefs for Standard Input, Standard Output, and Standard Error</i>	77
<i>File Descriptors</i>	78
<i>What Is a File Descriptor?</i>	78
<i>File Descriptors in the Bourne and Korn Shells</i>	78
<i>Reserved Filerefs in UNIX Environments</i>	78
<i>Sharing External Files in a UNIX Environment</i>	79
<i>Sharing External Files</i>	79

Options to Use for File Locking: External Files	79
File Locking for External Files: The LOCKINTERNAL Statement Option	79
File Locking for External Files: The FILELOCKS System Option	79
Reading from and Writing to UNIX Commands (PIPE)	79
What Are Pipes?	79
Syntax of the FILENAME Statement to Assign a Fileref to a Pipe	80
Using the Fileref for Reading	80
Specifying a Fileref for Reading	80
Example 1: Sending the Output of the Process Command to a SAS DATA Step	80
Example 2: Using the Stdin Fileref to Read Input	81
Using the Fileref for Writing	81
Specifying a Fileref for Writing	81
Example 1: Sending Mail Using Pipes	81
Example 2: Starting a Remote Shell and Printing Output	81
Sending Electronic Mail Using the FILENAME Statement (EMAIL)	82
Advantages of Sending Electronic Mail from within SAS	82
Initializing Electronic Mail	82
Components of the DATA Step or SCL Code Used to Send E-mail	82
Syntax of the FILENAME Statement for Electronic Mail	82
Specifying E-mail Options in the FILE Statement	84
Defining the Body of the Message	84
Specifying E-mail Directives in the PUT Statement	84
Example: Sending E-mail from the DATA Step	85
Example: Sending E-mail Using SCL Code	86
Processing Files on TAPE in UNIX Environments	87
Using the PIPE Device Type	87
Working with External Files Created on the Mainframe	88

Introduction to External Files and Devices in UNIX Environments

At times during a SAS session, you might want to use *external files*, that is, files that contain data or text, or files in which you want to store data or text. These files are created and maintained by the operating system, not by SAS.

You can use external files in a SAS session to perform the following functions:

- hold raw data to be read with the INPUT statements
- store printed reports created by a SAS procedure
- submit a file containing SAS statements for processing
- store data written with PUT statements

For SAS, external files and devices can serve both as sources of input and as receivers of output. The input can be either raw data to be read in a DATA step or SAS statements to be processed by SAS. The output can be one of the following:

- the SAS log, which contains notes and messages produced by the program
- the formatted output of SAS procedures
- data written with PUT statements in a DATA step

You might also want to use peripheral devices such as a printer, plotter, or your own terminal. UNIX treats these I/O devices as if they were files. Each device is associated with a file, called a *special file*, which is treated as an ordinary disk file. When you write to a special file, the associated device is automatically activated. All special files reside in the `dev` directory or its subdirectories. Although there are some differences in how you use the various devices, the basic concept is the same for them all.

UNIX also enables you to use pipes to send data to and from operating system commands as if they were I/O devices.

If you need to access an external file containing a transport data library, refer to *Moving and Accessing SAS Files*.

Accessing an External File or Device in UNIX Environments

Specifying a Pathname or a Fileref

To access an external file or device, you need to specify its pathname or fileref in the appropriate SAS statements:

FILE

specifies the current output file for PUT statements.

%INCLUDE

includes a file that contains SAS source statements that are executed when you submit a program from the Program Editor.

Tip: If you use %INCLUDE, the line limit is 6000 bytes.

INFILE

identifies an external file that you want to read with an INPUT statement.

In the SAS statement, refer to the file or device in one of two ways:

- Specify the pathnames for the external files. For more information, see “Specifying Pathnames in UNIX Environments” on page 70.
- Assign a fileref to a device, one or more files, or a directory, and use the fileref when you want to refer to the file, directory, or device.

In most cases, you will want to use a fileref.

What Is a Fileref?

A fileref is nickname that you assign to a file or device. You assign the fileref once, and then use it as needed. Filerefs are especially useful under the following conditions:

- The pathname is long and has to be specified several times within a program.
- The pathname might change. If the pathname changes, you need to change only the statement that assigns the fileref, not every reference to the file.

You can assign filerefs in the File Shortcuts window of the Explorer, with the FILENAME statement, with the FILENAME function,* or by defining the fileref as an environment variable.

* For a complete description of the FILENAME statement and the FILENAME function, see *SAS Language Reference: Dictionary*.

Specifying Pathnames in UNIX Environments

Rules for Specifying Pathnames

You can reference an external file directly by specifying its pathname in the `FILE`, `INFILE`, or `%INCLUDE` statements, or you can reference the file indirectly by specifying a fileref and a pathname in the `FILENAME` statement and then using the fileref in the `FILE`, `INFILE`, or `%INCLUDE` statements.

Whether you reference a file directly or indirectly, you need to specify its pathname in the appropriate statement. In most cases, you must enclose the name in quotation marks. For example, the following `INFILE` statement refers to the file `/users/pat/cars`:

```
infile '/users/pat/cars';
```

The following `FILE` statement directs output to the specified computer:

```
file '/dev/ttypl';
```

Note: If a filename has leading blanks, then the blanks will be trimmed. Δ

The level of specification depends on your current directory. You can use the character substitutions shown in Table 2.6 on page 51 to specify the pathname. You can also use wildcards as described in “Using Wildcards in Pathnames (Input Only)” on page 71.

Omitting Quotation Marks in a Filename

You can omit the quotation marks in a filename if one of the following is true:

- There is not already a fileref defined with that filename.
- The file has the filename extension expected by the statement that you are using to refer to the file. If you do not enclose a filename in quotation marks, the `FILE` and `INFILE` statements assume a file extension of `.dat`, and the `%INCLUDE` statement assumes a file extension of `.sas`.
- The file is located in the current directory.
- The filename is written with all lowercase characters.

For example, if the current directory is `/users/mkt/report` and it includes file `qtr.sas`, you can reference `qtr.sas` in any of the following statements:

```
%include '/users/mkt/report/qtr.sas';
%include 'qtr.sas';
file 'qtr.sas';
```

If there is no `qtr` fileref already defined, you can omit the quotation marks and the filename extension in the `%INCLUDE` statement:

```
%include qtr;
```

Working with Mixed Case or Uppercase Filenames

Filenames in the UNIX operating environment are case sensitive. This means that a file named **PROGRAM** is not the same as a file named **program**. When you reference the name of a file that is written in mixed case or uppercase, and that filename is not enclosed in quotation marks, SAS converts the filename to lowercase. If the filename does not have a file extension, SAS adds the missing file extension.

For example, if you specify `%include code(PROGRAM);` in your program, SAS converts the filename PROGRAM to lowercase, and adds an extension of .sas to the filename. PROGRAM becomes **program.sas**.

Interpreting the Messages in the SAS Log

When you execute the following program, SAS converts **TEMP** to **temp**, and adds an extension of .sas to the filename:

```
filename inc_code 'your-directory';
%include inc_code(TEMP);
```

SAS writes the following messages to the SAS log:

```
WARNING: Physical file does not exist, A.../your-directory/TEMP.sas.
ERROR: Cannot %INCLUDE member TEMP in the aggregate INC_CODE.
```

The warning message shows only the original filename (TEMP.sas), and not the lowercase conversion (temp.sas). This situation might cause confusion if a file named TEMP.sas does exist.

To avoid this confusion, include the file extension with the filename if the filename contains an extension, or enclose the mixedcase or uppercase filename in quotation marks if the filename does not have an extension. For example:

```
%include code(TEMP.sas);
%include code("TEMP");
```

In both of these cases, SAS does not convert TEMP to lowercase.

Using Wildcards in Pathnames (Input Only)

Descriptions of the Valid Wildcards

You can use the *, ?, and [] wildcards to specify pathnames in the FILENAME (only if the fileref is to be used for input), INFILE, and %INCLUDE statements and the INCLUDE command.

- * matches one or more characters, except for the period at the beginning of filenames.
- ? matches any single character.
- [] matches any single character from the set of characters defined within the brackets. You can specify a range of characters by specifying the starting character and ending character separated by a hyphen.

Wildcards are supported for input only. You cannot use wildcards in the FILE statement.

Example 1: Selecting Files by Including a Wildcard in a String

The following example reads input from every file in the current directory that begins with the string `wild` and ends with `.dat`:

```
filename wild 'wild*.dat';
data;
  infile wild;
  input;
run;
```

Example 2: Reading Each File in the Current Directory

The following example reads input from every file in every subdirectory of the current working directory:

```
filename subfiles '*/**';
data;
  infile subfiles;
  input;
run;
```

If new files are added to any of the subdirectories, they can be accessed with the Subfiles fileref without changing the FILENAME statement.

Example 3: Wildcards in Filenames When Using Aggregate Syntax

You can also use wildcards in filenames, but not in directory names, when you use aggregate syntax:

```
filename curdir ".";
data;
  infile curdir('wild*');
  input;
run;
```

In the example above, the period in the FILENAME statement refers to the current directory. See Table 2.6 on page 51 for information about other character substitutions available on UNIX.

Example 4: Associating a Fileref with Multiple Files

The following statement associates the fileref `MyRef` with all files that begin with alphabetic characters. Files beginning with numbers or other characters such as the period or tilde are excluded.

```
filename myref '[a-zA-Z]*.dat';
```

The following statement associates `MyRef` with any file beginning with Sales (in either uppercase, lowercase, or mixed case) and a year between 1990 and 1999:

```
filename myref '[Ss][Aa][Ll][Ee][Ss]199[0-9].dat';
```

Assigning Filerefs to External Files or Devices with the FILENAME Statement

Introduction to the FILENAME Statement

The most common way to assign a fileref to an external file or device is with the FILENAME statement. There are several forms of the FILENAME statement, depending on the type of device you want to access. For more information, see “FILENAME Statement” on page 318.

Accessing DISK Files

The most common use of the FILENAME statement is to access DISK files. The FILENAME syntax for a DISK file is the following:

```
FILENAME fileref <DISK> 'pathname' <options>;
```

The following FILENAME statement associates the fileref **myfile** with the external file **/users/mydir/myfile**, which is stored on a disk device:

```
filename myfile disk '/users/mydir/myfile';
```

The following FILENAME statement assigns a fileref of **prices** to the file **/users/pat/cars**. The FILE statement then refers to the file using the fileref:

```
filename prices '/users/pat/cars';
data current.list;
  file prices;
  ...PUT statements...
run;
```

See “Concatenating Filenames in UNIX Environments” on page 75 for more information about using DISK files.

Note: If a filename has leading blanks, then they will be trimmed. Δ

Debugging Code with DUMMY Devices

You can substitute the DUMMY device type for any of the other device types. This device type serves as a tool for debugging your SAS code without actually reading or writing to the device. After debugging is complete, replace the DUMMY device name with the proper device type, and your program will access the specified device type.

Here is the FILENAME syntax for a DUMMY file:

```
FILENAME fileref DUMMY 'pathname' <options>;
```

Output to DUMMY devices is discarded.

Sending Output to PRINTER Devices

The PRINTER device type enables you to send output directly to a printer. Here is the FILENAME syntax to direct a file to a PRINTER:

```
FILENAME fileref PRINTER '<printer> <printer-options>' <options>;
```

For example, this SAS program sends the output file to the BLDG3 printer:

```
filename myfile printer 'bldg3';

data test;
  file myfile;
  put 'This will appear in bldg3 .';
run;
```

See “Printing the Contents of a Window” on page 98 and “Using the PRINTTO Procedure in UNIX Environments” on page 100 for more information.

Using Temporary Files (TEMP Device Type)

The TEMP device type associates a fileref with a temporary file stored in the same directory as the Work library. (See “Work Library” on page 58.) Using the TEMP device type enables you to create a file that lasts only as long as the SAS session.

Here is the FILENAME syntax for a TEMP file:

```
FILENAME fileref TEMP <options>;
```

For example, this FILENAME statement associates Tmp1 with a temporary file:

```
filename tmp1 temp;
```

Accessing TERMINAL Devices Directly

To access a terminal directly, use the TERMINAL device type. Here is the FILENAME syntax to associate a file with a terminal:

```
FILENAME fileref TERMINAL <'terminal-pathname'> <options>;
```

The *terminal-pathname* must be a pathname of the special file associated with the terminal. Check with your UNIX system administrator for information. Enclose the name in quotation marks. If you omit the terminal pathname, the fileref is assigned to your terminal.

For example, this FILENAME statement associates the fileref **here** with your terminal:

```
filename here terminal;
```

The following FILENAME statement associates the fileref **thatfile** with another terminal:

```
filename thatfile terminal '/dev/tty3';
```

Assigning Filerefs to Files on Other Systems (FTP, SFTP, and SOCKET Access Types)

You can access files on other systems in your network by using the FTP, SFTP, and SOCKET access methods. Here are the forms of the FILENAME statement:

```
FILENAME fileref FTP 'external-file' <ftp-options>;
```

```
FILENAME fileref SFTP 'external-file' <sftp-options>;
```

```
FILENAME fileref SOCKET 'external-file' <tcpip-options>;
```

```
FILENAME fileref SOCKET ':portno' SERVER <tcpip-options>;
```


These access methods are documented in *SAS Language Reference: Dictionary*. Under UNIX, the FTP access method supports an additional option:

MACH=*'machine'*

identifies which entry in the **.netrc** file should be used to get the user name and password. The **.netrc** file resides on the host on which the SAS program is running. See the UNIX man page for more information about the **.netrc** file. You cannot specify the MACH option together with the HOST option in the FILENAME statement.

If you are transferring a file to UNIX from the z/OS operating environment and you want to use either the S370V or S370VB format to access that file, then the file must be of type RECFM=U and BLKSIZE=32760 before you transfer it.

CAUTION:

When you use the FTP access method to create a remote file, the UNIX permissions for that file are set to *-rw-rw-rw-*, which makes the file world-readable and world-writable.

See the UNIX man page for **chmod** for information about changing file permissions. Δ

Concatenating Filenames in UNIX Environments

You can concatenate filenames in the FILENAME, %INCLUDE, and INFILE statements. Concatenating filenames enables you to read those files sequentially.

FILENAME *fileref* ("*pathname-1*" ... "*pathname-n*");

%INCLUDE '*filename-1*' ... '*filename-n*';

%INCLUDE ("*filename-1*" ... '*filename-n*');

INFILE '*filename-1*' ... "*filename-n*";

INFILE ("*filename-1*" ... '*filename-n*');

You can enclose the pathnames in single or double quotation marks and separate them with commas or blank spaces. You can use the characters shown in Table 2.6 on page 51 and the wildcards described in “Using Wildcards in Pathnames (Input Only)” on page 71 to specify the pathnames.

Assigning a Fileref to a Directory (Using Aggregate Syntax)

Introduction to Aggregate Syntax

Aggregate syntax enables you to assign a fileref to a directory and then work with any file in that directory by specifying its filename in parentheses after the fileref.

FILENAME *fileref* *directory-name*;

Aggregate syntax is especially useful when you have to refer to several files in one directory.

Example 1: Referring to a File Using Aggregate Syntax

To refer to a file in the directory, specify the fileref followed by the individual filename in parentheses. For example, you can refer to the file cars.dat in the directory **/users/pat** as shown in this example:

```
filename prices '/users/pat';
data current.list;
  file prices(cars);
  ...other SAS statements...
run;
```

Example 2: Using Aggregate Syntax with Filerefs Defined by Environment Variables

You can also use aggregate syntax with filerefs that have been defined using environment variables. (See “Using Environment Variables to Assign Filerefs in UNIX Environments” on page 76.) For example:

```
x setenv PRICES /users/pat;
data current.list;
  file prices(cars);
  ...other SAS statements...
run;
```

Assigning a Fileref to Several Directories

In the FILENAME statement, you can concatenate directory names and use the fileref to refer to any file within those directories:

```
FILENAME fileref ("directory-1" ... "directory-n");
```

When you concatenate directory names, you can use aggregate syntax to refer to a file in one of the directories. For example, assume that the Report.sas file resides in the directory associated with the MYPROGS environment variable. When SAS executes the following code, it searches for Report.sas in the pathnames that are specified in the FILENAME statement and it executes the program.

```
filename progs ("%MYPROGS" "/users/mkt/progs");
%inc progs(report);
```

SAS searches the pathnames in the order specified in the FILENAME statement until

- it finds the first file with the specified name. Even if you use wildcards (see “Using Wildcards in Pathnames (Input Only)” on page 71) in the filename, SAS matches only one file.
- it encounters a filename in the list of pathnames that you specified in the FILENAME statement.

Using Environment Variables to Assign Filerefs in UNIX Environments

Requirements for Variable Names

An environment variable can also be used as a fileref to refer to DISK files. The variable name must be in all uppercase characters, and the variable value must be the full pathname of the external file; that is, the filename must begin with a slash.

Note: If a variable and a fileref have the same name but refer to different files, SAS uses the fileref. For example, the %INCLUDE statement below refers to file `/users/myid/this_one`. \triangle

```
filename ABC '/users/myid/this_one';
x setenv ABC /users/myid/that_one;
%include ABC;
```

Reading a Data File

If you want to read the data file `/users/myid/educ.dat`, but you want to refer to it with the INED environment variable, you can define the variable at two times:

- Before you invoke SAS. See “Defining Environment Variables in UNIX Environments” on page 23. For example, in the Korn shell, you use the following:

```
export INED=/users/myid/educ.dat
```

- After you invoke SAS by using the X statement (see “Executing Operating System Commands from Your SAS Session” on page 14) and the SAS `setenv` command:

```
x setenv INED /users/myid/educ.dat;
```

After INED is associated with the file `/users/myid/educ.dat`, you can use `ined` as a fileref to refer to the file in the INFILE statement:

```
infile ined;
```

Writing to an External File

The same method applies if you want to write to an external file. For example, you can define OUTFILE before you invoke SAS:

```
OUTFILE=/users/myid/scores.dat
export OUTFILE
```

Then, use the environment variable name as a fileref to refer to the file:

```
file OUTFILE;
```

Filerefs Assigned by SAS in UNIX Environments

Filerefs for Standard Input, Standard Output, and Standard Error

Often a command’s arguments or options tell the command what to use for input and output, but in case they do not, the shell supplies you with three standard files: one for input (*standard input*), one for output (*standard output*), and one for error messages (*standard error*). By default, these files are all associated with your terminal: standard input with your keyboard, and both standard output and standard error with your terminal’s display. When you invoke SAS, it assigns a fileref to each file that it opens, including the three standard files. SAS assigns the filerefs Stdin, Stdout, and Stderr to standard input, standard output, and standard error, respectively.

File Descriptors

What Is a File Descriptor?

Each file has an internal *file descriptor* assigned to it. By default, 0 is the file descriptor for standard input, 1 is the file descriptor for standard output, and 2 is the file descriptor for standard error. As other files are opened, they get other file descriptors. In the Bourne shell and in the Korn shell, you can specify that data be written to or be read from a file using the file descriptor as described in “File Descriptors in the Bourne and Korn Shells” on page 78.

File Descriptors in the Bourne and Korn Shells

If you are using the Bourne shell or the Korn shell, SAS assigns filerefs of the following form to files that have a file descriptor (see “Filerefs Assigned by SAS in UNIX Environments” on page 77) larger than 2.

`FILDESnumber`

number is a two-digit representation of the file descriptor. You can use these filerefs in your SAS applications.

For example, if you invoke SAS with the following command, then the operating environment opens the file `sales_data` and assigns file descriptor 4 to it:

```
sas salespgm 4< sales_data
```

SAS assigns the fileref `FILDES04` to the file and executes the application `salespgm`. When the application reads input from `FILDES04`, it reads the file `sales_data`. Using file descriptors as filerefs enables you to use the same application to process data from different files without changing the application to refer to each file. In the command that you use to invoke the application, you assign the appropriate file descriptor to the file to be processed.

Reserved Filerefs in UNIX Environments

The following filerefs are reserved.

DATALINES fileref in the **INFILE** statement

specifies that input data immediately follow a **DATALINES** statement. You need to use **INFILE DATALINES** only when you want to specify options in the **INFILE** statement to read instream data.

LOG fileref in the **FILE** statement

specifies that output lines produced by **PUT** statements be written to the SAS log. **LOG** is the default destination for output lines.

PRINT fileref in the **FILE** statement

specifies that output lines produced by **PUT** statements be written to the same print file as output produced by SAS procedures.

Sharing External Files in a UNIX Environment

Sharing External Files

If more than one user has simultaneous Write access to an external file, or if a single user has Write access to the same file from different SAS sessions, the results of sharing the file can be unpredictable. To remedy this situation, you can use a statement or a system option to restrict Write access to one user, while allowing multiple users Read access. (For information about sharing SAS files, see “Sharing SAS Files” on page 38.)

Options to Use for File Locking: External Files

File locking applies to all files that are opened. You can turn off file locking for external files in the following ways:

- Use the LOCKINTERNAL option in the FILENAME statement.
- Use the FILELOCKS system option.

File Locking for External Files: The LOCKINTERNAL Statement Option

You can control file locking for external files by using the LOCKINTERNAL option in the FILENAME statement. The AUTO option value locks a file exclusively for Write access, or non-exclusively for Read access. For example, if a file is opened for update or output, then all other access from internal processes will be blocked. If a file is opened for input, then other users can also open the file for input. In this case, opening the file for update and output will be blocked. The SHARED option value allows for all of the behavior of the AUTO option, except that the file can be shared by one writer and multiple readers. The external file that is associated with the fileref is the file that is locked. By default, multiple users can simultaneously read an external file. For more information, see “FILENAME Statement” on page 318.

File Locking for External Files: The FILELOCKS System Option

You can control file locking for external files (as well as for SAS files) by using the FILELOCKS system option. This option enables you to apply a behavior globally to individual files or directories. Using FILELOCKS restricts writer access to one user. With file locking turned on, multiple SAS sessions are able to simultaneously read the same file. You can use FILELOCKS at start-up, in the OPTIONS statement, or in the command line. You can specify multiple instances of the FILELOCKS option. Each instance is added to an internal table of paths and settings. For more information, see “FILELOCKS System Option” on page 368.

Reading from and Writing to UNIX Commands (PIPE)

What Are Pipes?

Pipes enable your SAS application to receive input from any UNIX command that writes to standard output and to route output to any UNIX command that reads from

standard input. In UNIX commands, the pipe is represented by a vertical bar (|). For example, to find the number of files in your directory, you could redirect the output of the **ls** command through a pipe to the **wc** (word count) command:

```
ls | wc -w
```

Syntax of the FILENAME Statement to Assign a Fileref to a Pipe

Under UNIX, you can use the FILENAME statement to assign filerefs not only to external files and I/O devices, but also to a pipe. Here is the syntax of the FILENAME statement:

```
FILENAME fileref PIPE 'UNIX-command' <options>;
```

fileref

is the name by which you reference the pipe from SAS.

PIPE

identifies the device-type as a UNIX pipe.

'*UNIX-command*'

is the name of a UNIX command, executable program, or shell script to which you want to route output or from which you want to read input. The commands must be enclosed in either double or single quotation marks.

options

control how the external file is processed. See "FILENAME Statement" on page 318 for an explanation of these options.

Whether you are using the command as input or output depends on whether you use the *fileref* in a reading or writing operation. For example, if the fileref is used in an INFILE statement, then SAS assumes that the input comes from a UNIX command; if the fileref is used in a FILE statement, then SAS assumes that the output goes to a UNIX command.

Using the Fileref for Reading

Specifying a Fileref for Reading

When the fileref is used for reading, the specified UNIX command executes, and any output sent to its standard output or standard error is read through the fileref. In this case, the standard input of the command is connected to `/dev/null`.

Example 1: Sending the Output of the Process Command to a SAS DATA Step

The following SAS program uses the PIPE device-type keyword to send the output of the **ps** (process) command to a SAS DATA step. The resulting SAS data set contains data about every process currently running SAS:

```
filename ps_list pipe "ps -e|grep 'sas'";
data sasjobs;
  infile ps_list;
  length process $ 80;
  input process $ char80.;
run;
```

```
proc print data=sasjobs;
run;
```

The **ps -e** command produces a listing of all active processes in the system, including the name of the command that started the task. In BSD-based UNIX systems, you use the **ps -ax** command.

The operating environment uses pipes to send the output from **ps** to the **grep** command, which searches for every occurrence of the string **'sas'**. The FILENAME statement connects the output of the **grep** command to the fileref **ps_list**. The DATA step then creates a data set named **sasjobs** from the INFILE statement that points to the input source. The INPUT statement reads the first 80 characters on each input line.

Example 2: Using the Stdin Fileref to Read Input

In the next example, the Stdin fileref is used to read input through a pipe into the SAS command, which, in turn, executes the SAS program. By placing the piping operation outside the SAS program, the program becomes more general. The program in the previous example has been changed and stored in file **ps.sas**:

```
data sasjobs;
  infile stdin;
  length process $ 80;
  input process $ char80.;
run;
proc print data=sasjobs;
run;
```

To run the program, use pipes to send the output of **ps** to **grep** and from **grep** into the SAS command:

```
ps -e|grep 'sas'|sas ps.sas &
```

The output will be stored in **ps.lst**, and the log will be stored in **ps.log**, as described in “The Default Routings for the SAS Log and Procedure Output in UNIX Environments” on page 93.

Using the Fileref for Writing

Specifying a Fileref for Writing

When the fileref is used for writing, the output from SAS is read in by the specified UNIX command, which then executes.

Example 1: Sending Mail Using Pipes

In this example, any data sent to the **mail** fileref are piped to the **mail** command and sent to user **PAT**:

```
filename mail pipe 'mail pat';
```

Example 2: Starting a Remote Shell and Printing Output

Consider this FILENAME statement:

```
filename letterq pipe 'remsh alpha lp -dbldga3';
```

Any data sent to the **letterq** fileref is passed to the UNIX command, which starts a remote shell on the computer named Alpha. Note that the form of the command that

starts a remote shell varies among the various UNIX operating environments. The shell then prints the **letterq** output on the printer identified by the destination BLDGA3. Any messages produced by the **lp** command are sent to the SAS log.

Sending Electronic Mail Using the FILENAME Statement (EMAIL)

Advantages of Sending Electronic Mail from within SAS

SAS lets you send electronic mail using SAS functions in a DATA step or in SCL. Sending e-mail from within SAS enables you to do the following:

- Use the logic of the DATA step or SCL to subset e-mail distribution based on a large data set of e-mail addresses.
- Send e-mail automatically upon completion of a SAS program that you submitted for batch processing.
- Direct output through e-mail based on the results of processing.
- Send e-mail messages from within a SAS/AF frame application, customizing the user interface.

Initializing Electronic Mail

By default, SAS uses SMTP (Simple Mail Transfer Protocol) to send e-mail. SMTP, unlike some external scripts, supports attachments. This default is specified by the EMAILSYS system option. For information about how to change the e-mail protocol, see “EMAILSYS System Option” on page 366.

Before you can send e-mail from within SAS, your system administrator might need to set the EMAILHOST system option to point to the SMTP server. For more information about the EMAILHOST system option, see *SAS Language Reference: Dictionary*.

Components of the DATA Step or SCL Code Used to Send E-mail

In general, a DATA step or SCL code that sends electronic mail has the following components:

- a FILENAME statement with the EMAIL device-type keyword
- options specified in the FILENAME or FILE statements indicating the e-mail recipients, subject, and any attached files
- PUT statements that contain the body of the message
- PUT statements that contain special e-mail directives (of the form !EM_directive!) that can override the e-mail attributes (TO, CC, BCC, SUBJECT, ATTACH) or perform actions (such as SEND, ABORT, and start a NEWMSG)

Syntax of the FILENAME Statement for Electronic Mail

To send electronic mail from a DATA step or SCL, issue a FILENAME statement of the following form:

```
FILENAME fileref EMAIL 'address' <email-options>;
```

The FILENAME statement accepts the following *email-options*:

fileref

is a valid fileref.

'address'

is the destination e-mail address of the user to which you want to send e-mail. You must specify an address here, but you can override its value with the TO e-mail option.

email-options

can be any of the following:

TO=to-address

specifies the primary recipients of the electronic mail. If an address contains more than one word, enclose it in quotation marks. To specify more than one address, enclose the group of addresses in parentheses, enclose each address in quotation marks, and separate each address with a space. For example, **to='joe@someplace.org'** and **to=("joe@smplc.org" "jane@diffplc.org")** are valid TO values.

Note: You can send an e-mail without specifying a recipient in the TO= option as long as you specify a recipient in either the CC= or BCC= option. △

CC=cc-address

specifies the recipients you want to receive a copy of the electronic mail. If an address contains more than one word, enclose it in quotation marks. To specify more than one address, enclose the group of addresses in parentheses, enclose each address in quotation marks, and separate each address with a space. For example, **cc='joe@someplace.org'** and **cc=("joe@smplc.org" "jane@diffplc.org")** are valid CC values.

BCC=bcc-address

specifies the recipients you want to receive a blind copy of the electronic mail. Individuals listed in the **bcc** field will receive a copy of the e-mail. The BCC field does not appear in the e-mail header, so that these e-mail addresses cannot be viewed by other recipients.

If a BCC address contains more than one word, enclose it in quotation marks. To specify more than one address, enclose the group of addresses in parentheses, enclose each address in quotation marks, and separate each address with a space. For example, **bcc='joe@someplace.org'** and **bcc=("joe@smplc.org" "jane@diffplc.org")** are valid BCC values.

SUBJECT='subject'

specifies the subject of the message. If the subject text is longer than one word (that is, it contains at least one blank space), you must enclose it in quotation marks. You also must use quotation marks if the subject contains any special characters. For example, **subject=Sales** and **subject='June Report'** are valid subjects. Any subject not enclosed in quotation marks is converted to uppercase.

ATTACH='filename.ext' | ATTACH = ('filename.ext' <attachment-options>)

specifies the physical name of the files to be attached to the message and any options to modify attachment specifications. Enclose *filename.ext* in quotation marks. To attach more than one file, enclose the group of filenames in parentheses. For example, **attach='/u/userid/opinion.txt'** and **attach=("june98.txt" "july98.txt")** are valid file attachments.

By default, SMTP e-mail attachments are truncated at 256 characters. To send longer attachments, you can specify the LRECL= and RECFM= options from the FILENAME statement as the *attachment-options*. For more

information about the LRECL= and RECFM= options, see “FILENAME Statement” on page 318.

For more information about the options that are valid when you are using SMTP, see “FILENAME Statement, EMAIL (SMTP) Access Method” in *SAS Language Reference: Dictionary*.

Specifying E-mail Options in the FILE Statement

You can also specify the *email-options* in the FILE statement inside the DATA step. Options that you specify in the FILE statement override any corresponding options that you specified in the FILENAME statement.

Defining the Body of the Message

In your DATA step, after using the FILE statement to define your e-mail fileref as the output destination, use PUT statements to define the body of the message.

Specifying E-mail Directives in the PUT Statement

You can also use PUT statements to specify e-mail directives that change the attributes of your electronic message or perform actions with it. Specify only one directive in each PUT statement; each PUT statement can contain only the text associated with the directive it specifies.

The following are the directives that change the attributes of your message:

!EM_TO! *addresses*

Replace the current primary recipient addresses with *addresses*. In the PUT statement, specify *addresses* without single quotation marks.

!EM_CC! *addresses*

Replace the current copied recipient addresses with *addresses*. In the PUT statement, specify *addresses* without single quotation marks.

!EM_BCC! *addresses*

Replace the current blind copied recipient addresses with *addresses*. In the PUT statement, specify *addresses* without single quotation marks.

!EM_SUBJECT! *subject*

Replace the current subject of the message with *subject*.

!EM_ATTACH! *pathname*

Replace the names of any attached files with *pathname*.

The following are the directives that perform actions:

!EM_SEND!

Sends the message with the current attributes. By default, SAS sends a message when the fileref is closed. The fileref closes when the next FILE statement is encountered or the DATA step ends. If you use this directive, SAS sends the message when it encounters the directive, *and* again at the end of the DATA step.

!EM_ABORT!

Aborts the current message. You can use this directive to stop SAS from automatically sending the message at the end of the DATA step.

!EM_NEWMSG!

Clears all attributes of the current message, including TO, CC, SUBJECT, ATTACH, and the message body.

Example: Sending E-mail from the DATA Step

Suppose that you want to share a copy of your config.sas file with your coworker Jim, whose user ID is JBrown. If your e-mail program handles alias names and attachments, you could send it by submitting the following DATA step:

```
filename mymail email 'JBrown'
      subject='My CONFIG.SAS file'
      attach='config.sas';

data _null_;
  file mymail;
  put 'Jim,';
  put 'This is my CONFIG.SAS file.';
  put 'I think you might like the
      new options I added.';
run;
```

The following example sends a message and two attached files to multiple recipients. It specifies the e-mail options in the FILE statement instead of the FILENAME statement:

```
filename outbox email 'ron@acme.com';

data _null_;
  file outbox

      /* Overrides value in filename statement */
      to=('ron@acme.com' 'lisa@acme.com')
      cc=('margaret@yourcomp.com'
         'lenny@laverne.abc.com')
      subject='My SAS output'
      attach=('results.out' 'code.sas')
      ;
  put 'Folks,';
  put 'Attached is my output from the
      SAS program I ran last night.';
  put 'It worked great!';
run;
```

You can use conditional logic in the DATA step to send multiple messages and control which recipients get which message. For example, suppose you want to send customized reports to members of two different departments. If your e-mail program handles alias names and attachments, your DATA step might look like the following:

```
filename reports email 'Jim';

data _null_;
  file reports;
  infile cards eof=lastobs;
  length name dept $ 21;
  input name dept;
```

```

        /* Assign the TO attribute */
        put '!EM_TO!' name;

        /* Assign the SUBJECT attribute */
        put '!EM_SUBJECT! Report for ' dept;

        put name ',';
        put 'Here is the latest report for ' dept '.';

        /* ATTACH the appropriate report */
        if dept='marketing' then
            put '!EM_ATTACH! mktrept.txt';
        else

            put '!EM_ATTACH! devrept.txt';

        /* Send the message */
        put '!EM_SEND!';

        /* Clear the message attributes */
        put '!EM_NEWMSG!';

        return;

        /* Abort the message before the */
        /* RUN statement causes it to */
        /* be sent again. */
        lastobs: put '!EM_ABORT!';

        datalines;
        Susan      marketing
        Jim         marketing
        Rita       development
        Herb       development
        ;
        run;

```

The resulting e-mail message and its attachments are dependent on the department to which the recipient belongs.

Note: You must use the !EM_NEWMSG! directive to clear the message attributes between recipients. The !EM_ABORT! directive prevents the message from being automatically sent at the end of the DATA step. Δ

Example: Sending E-mail Using SCL Code

The following example is the SCL code behind a frame entry design for e-mail. The frame entry includes several text entry fields that let the user enter information:

<i>mailto</i>	the user ID to send mail to
<i>copyto</i>	the user ID to copy (CC) the mail to
<i>attach</i>	the name of a file to attach
<i>subject</i>	the subject of the mail
<i>line1</i>	the text of the message

The frame entry also contains a button named SEND that causes this SCL code (marked by the **send:** label) to execute.

```

send:

    /* set up a fileref */
rc = filename('mailit','userid','email');

    /* if the fileref was successfully set up
    open the file to write to */
if rc = 0 then do;
    fid = fopen('mailit','o');
    if fid > 0 then do;

        /* fput statements are used to
        implement writing the
        mail and the components such as
        subject, who to mail to, and so on. */
fputc1 = fput(fid,linel);
rc = fwrite(fid);

fputc2 = fput(fid,'!EM_TO! '||mailto);
rc = fwrite(fid);
fputc3 = fput(fid,'!EM_CC! '||copyto);
rc = fwrite(fid);

fputc4 = fput(fid,'!EM_ATTACH! '||attach);
rc = fwrite(fid);
fputc5 = fput(fid,'!EM_SUBJECT! '||subject);
rc = fwrite(fid);

        closerc = fclose(fid);
    end;
end;
return;

cancel:
    call execcmd('end');
return;

```

Processing Files on TAPE in UNIX Environments

Using the PIPE Device Type

You can use the PIPE device type together with the UNIX **dd** command to process a tape:

FILENAME *fileref* PIPE '*UNIX-commands*';

UNIX-commands are the commands needed to process the tape.

The PIPE device type and the **dd** command enables you to use remote tape drives. However, using UNIX commands in your application means that the application will have to be modified if it is ported to a non-UNIX environment.

For example, the following DATA step writes an external file to tape:

```
x 'mt -t /dev/rmt/0mn rewind';
filename outtape pipe 'dd of=/dev/rmt/0mn 2> /dev/null';
data _null_;
  file outtape;
  put '1 one';
  put '2 two';
  put '3 three';
  put '4 four';
  put '5 five';
run;
```

The following DATA step reads the file from tape:

```
x 'mt -t /dev/rmt/0mn rewind';
filename intape pipe 'dd if=/dev/rmt/0mn 2> /dev/null';
data numbers;
  infile intape pad;
  input digit word $8.;
run;
```

If the tape drive that you want to access is a remote tape drive, you can access the remote tape drive by adding **remsh computer-name** or the following:

```
rsh computer-name
```

The above code is added to the X and FILENAME statements. For example, if the remote computer name is **wizard**, then you could read and write tape files on **wizard** by modifying the X and FILENAME statements as follows:

```
x 'remsh wizard mt -t /dev/rmt/0mn rewind';
filename intape pipe 'remsh wizard \
  dd if=/dev/rmt/0mn 2> /dev/null';
```


Working with External Files Created on the Mainframe

There are three main points to remember when dealing with tapes on UNIX that were created on a mainframe:

- UNIX does not support IBM standard label tapes. IBM standard label tapes contain user data files and labels, which themselves are files on the tape. To process the user data files on these tapes, use a no-rewind device (such as **/dev/rmt/0mn**) and the **mt** command with the **fsf count** subcommand to position the tape to the desired user data file. Here is the formula for calculating **count**:

$$\text{count} = (3 \times \text{user_data_file_number}) - 2$$
- UNIX does not support multivolume tapes. To process multivolume tapes on UNIX, the contents of each tape must be copied to disk using the **dd** command. After all of the tapes have been unloaded, you can use the **cat** command to concatenate all of the pieces in the correct order. You can then use SAS to process the concatenated file on disk.
- You must know the DCB characteristics of the file. The records in files that are created on a mainframe are not delimited with end-of-line characters, so you must specify the original DCB parameters in the INFILE or FILENAME statement. In the INFILE statement, specify the record length, record format, and block size with the LRECL, RECFM, and BLKSIZE host options. In the FILENAME statement, if you use the PIPE device-type and the **dd** command, you must also

specify the block size with the **ibs** subcommand. For more information about host options in the INFILE statement, see “INFILE Statement” on page 326. For more information about the **ibs** subcommand, see to the man page for the **dd** command.



CHAPTER

4

Printing and Routing Output

<i>Overview of Printing Output in UNIX Environments</i>	92
<i>Previewing Output in UNIX Environments</i>	92
<i>Previewing Output Using Universal Printing</i>	92
<i>Previewing Output from within SAS/AF Applications</i>	92
<i>The Default Routings for the SAS Log and Procedure Output in UNIX Environments</i>	93
<i>Changing the Default Routings in UNIX Environments</i>	93
<i>Techniques for Routing Output</i>	93
<i>Determining Which Technique to Use When Changing the Routing</i>	93
<i>Routing SAS Logging Facility Messages to SYSLOGD</i>	95
<i>Using the Print Dialog Box in UNIX Environments</i>	95
<i>Printing from Text Windows</i>	95
<i>Open the Print Dialog Box from a Text Window</i>	95
<i>Default Printing Mode</i>	96
<i>Specifics for Forms Printing</i>	96
<i>Troubleshooting Print Server Errors</i>	96
<i>Printing from GRAPH Windows</i>	96
<i>Open the Print Dialog Box from the GRAPH Window</i>	96
<i>Specifics for SAS/GRAPH Drivers</i>	97
<i>Troubleshooting Print Server Errors</i>	97
<i>Using Commands to Print in UNIX Environments</i>	97
<i>Differences between the PRTFILE, PRINT, and FILE Commands</i>	97
<i>Sending Output to a UNIX Command</i>	98
<i>Specifying the Print File</i>	98
<i>Printing the Contents of a Window</i>	98
<i>Using PRTFILE and PRINT with a Fileref</i>	98
<i>Steps for Sending Output Directly to a Printer</i>	98
<i>Examples of FILENAME Statements Using PRINTER and PIPE</i>	99
<i>Using the FILE Command</i>	100
<i>Using the PRINTTO Procedure in UNIX Environments</i>	100
<i>Important Note about the PRINTTO Procedure</i>	100
<i>Using the LOG= and PRINT= Options</i>	100
<i>Routing Output to a Universal Printer</i>	101
<i>Routing Output to a Printer</i>	101
<i>Piping Output to a UNIX Command</i>	101
<i>Routing Output to a Terminal</i>	102
<i>Using SAS System Options to Route Output</i>	102
<i>Changing the Output Destination Using the LOG, PRINT, ALTLOG, and ALTPRINT System Options</i>	102
<i>Printing Large Files with the PIPE Device Type in UNIX Environments</i>	102
<i>Changing the Default Print Destination in UNIX Environments</i>	103
<i>Changing the Default Print Command in UNIX Environments</i>	103

<i>Controlling the Content and Appearance of Output in UNIX Environments</i>	104
<i>Overview of Controlling the Content and Appearance of Output</i>	104
<i>SAS Log Options</i>	104
<i>Procedure Output Options</i>	105

Overview of Printing Output in UNIX Environments

When you print text or graphics, SAS needs to know where the output should go, how it should be written, and how the output should look. Universal Printing is the default printing mechanism in UNIX. Universal Printing supports PostScript, PCL, GIF, PNG, SVG, and PDF files in all environments. For more information about Universal Printing, see *SAS Language Reference: Concepts*.

Forms printing is an older method of text printing available from SAS. It involves using the FORM subsystem, which consists of the Form window. For information, see FORMS printing in *SAS Language Reference: Dictionary*.

If you are printing graphics, the output is controlled by native SAS/GRAPH drivers. See the online Help for SAS/GRAPH for information about native SAS/GRAPH drivers.

Previewing Output in UNIX Environments

Previewing Output Using Universal Printing

With Universal Printing, you can preview your output before you send it to a printer, plotter, or external file. To preview your output, you first need to define a previewer for your system. For more information about Universal Printing, see *SAS Language Reference: Concepts*.

Previewing Output from within SAS/AF Applications

To preview output from within a SAS/AF application, use the DMPRTMODE and DMPRTPREVIEW commands to turn on preview mode, print the output, open the Print Preview dialog box, and then turn preview mode off. For example, the following code prints the GRAPH1 object using the host drivers and displays it in the Preview dialog box:

```
/* Turn on preview mode. */
CALL EXECCMDI ("DMPRTMODE PREVIEW");

/* Print the graph */
GRAPH1._PRINT_();

/* Open the Preview dialog box */
CALL EXECCMDI ("DMPRTPREVIEW");

/* Turn off preview mode */
CALL EXECCMDI ("DMPRTMODE NORMAL");
```

The Default Routings for the SAS Log and Procedure Output in UNIX Environments

For each SAS job or session, SAS automatically creates two types of output:

SAS log

contains information about the processing of SAS statements. As each program step executes, notes are written to the SAS log along with any applicable error or warning messages.

SAS output

is also called the *procedure output file* or *print file*. Whenever a SAS program executes a PROC step or a DATA step that produces printed output, SAS sends the output to the SAS output file.

Table 4.1 on page 93 shows the default routings of the SAS log and output files.

Table 4.1 Default Routings of the SAS Log and Output Files

Processing Mode	SAS Log File	SAS Output File
batch	<i>filename.log</i>	<i>filename.lst</i>
windowing environment	Log window	Output window
interactive line	terminal	terminal

By default, both the log file and the output file are written to your current directory. Your system administrator might have changed these default routings.

Changing the Default Routings in UNIX Environments

Techniques for Routing Output

There are four primary methods for routing your output:

- Using the Print dialog box. The Print dialog box is available when you are using the SAS windowing environment.
- Issuing windowing environment commands. The PRTFILE, PRINT, and FILE commands can be issued from any command line and can be used to send output to external files or to other devices defined with the FILENAME statement.
- Using the PRINTTO procedure. You can use the PRINTTO procedure in any mode. Using the FILENAME statement with the PRINTTO procedure is the most flexible way of routing your output.
- Using SAS system options, such as PRINT, LOG, ALTPRINT, or ALTLOG, to specify alternate destinations.

Determining Which Technique to Use When Changing the Routing

Use the following table to help you decide which method you should choose to change the routing.

Table 4.2 Decision Table: Changing the Default Destination

Output destination for your SAS log or procedure output	Processing mode	Method	See
a printer	any mode	FILENAME statement (UPRINTER or PRINTER device type) and PRINTTO procedure	“Using the PRINTTO Procedure in UNIX Environments” on page 100
	windowing environment	DMPRINT command	“Using the Print Dialog Box in UNIX Environments” on page 95
		Print dialog box	“Using the Print Dialog Box in UNIX Environments” on page 95
an external file	any mode	FILENAME statement and PRTFILE, PRINT, and FILE commands	“Printing the Contents of a Window” on page 98
		PRINTTO procedure and FILENAME statement	“Using the PRINTTO Procedure in UNIX Environments” on page 100
	batch	LOG and PRINT system options	“Using SAS System Options to Route Output” on page 102
a UNIX command (pipe)	any mode	FILENAME statement and PRINTTO procedure	“Using the PRINTTO Procedure in UNIX Environments” on page 100
	windowing environment	FILENAME statement and PRTFILE and PRINT commands	“Printing the Contents of a Window” on page 98
its usual location <i>and</i> to an external file	any mode	ALTLOG and ALTPRINT system options	“Using SAS System Options to Route Output” on page 102
	windowing environment	FILE command	“Using the FILE Command” on page 100
		Print dialog box	“Using the Print Dialog Box in UNIX Environments” on page 95
a terminal	batch	FILENAME statement and PRINTTO procedure	“Routing Output to a Terminal” on page 102

Routing SAS Logging Facility Messages to SYSLOGD

The SAS 9.2 logging facility enables the categorization and collection of log event messages, and then writes them to a variety of output devices. The logging facility supports problem diagnosis and resolution, performance and capacity management, and auditing and regulatory compliance. The following features are provided:

- Log events are categorized using a hierarchical naming system that enables you to configure logging at a broad or a fine-grained level.
- Log events can be directed to multiple output destinations, including files, operating system facilities, databases, and client applications. For each output destination, you can specify:
 - the categories and levels of log events to report
 - the message layout, including the types of data to be included, the order of the data, and the format of the data
 - filters based on criteria such as diagnostic levels and message content
- Logging diagnostic levels can be adjusted dynamically without starting and stopping processes.
- Performance-related log events can be generated for processing by an Application Response Measurement (ARM) 4.0 server.

The logging facility is used by most SAS server processes. You can also use the logging facility within SAS programs.

In the UNIX operating environment, logging facility messages can be written to SYSLOGD.

For information about using the logging facility in the UNIX operating environment, see *SAS Logging: Configuration and Programming Reference*.

Using the Print Dialog Box in UNIX Environments

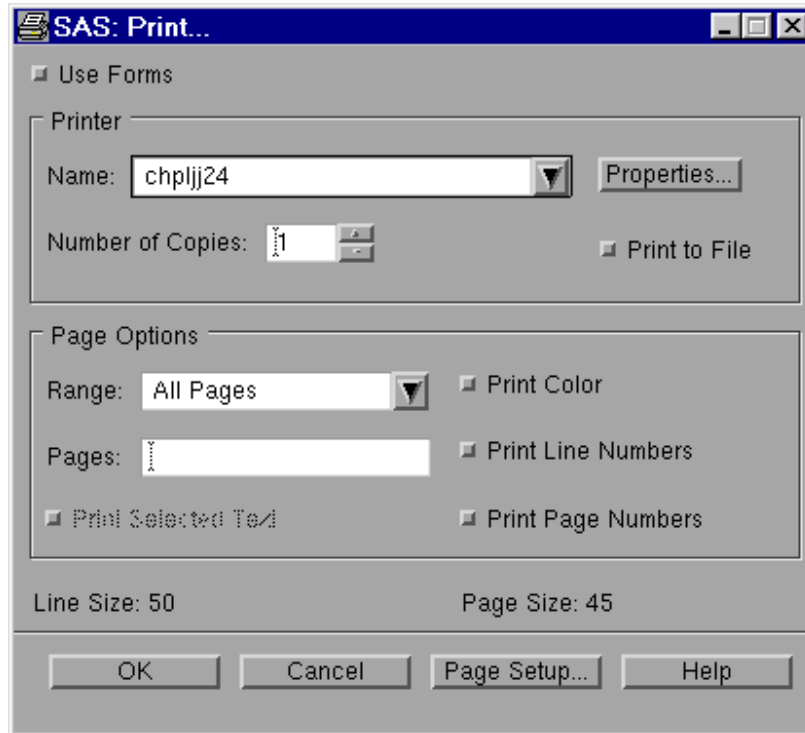
Printing from Text Windows

Open the Print Dialog Box from a Text Window

To print part or all of the contents of a window, complete the following steps:

- 1 Click in the window to make it the active window. If you want to mark and print only selected lines of text, mark the text before you open the Print dialog box.
- 2 Issue the DMPRINT command or select **File ► Print** to open the Print dialog box.

Display 4.1 Print Dialog Box



Default Printing Mode

In UNIX, the default printing mode is Universal Printing. For more information about how to use Universal Printing, click **help** on the Print dialog box.

Specifics for Forms Printing

To use forms for printing, select **Use forms**. SAS prompts you to enter a spool command and the name of your system printer. When you click **OK**, SAS prints the contents of the active window using the command and printer name that you specified and additional information from your default form. See *SAS Language Reference: Dictionary* for more information about forms printing.

Troubleshooting Print Server Errors

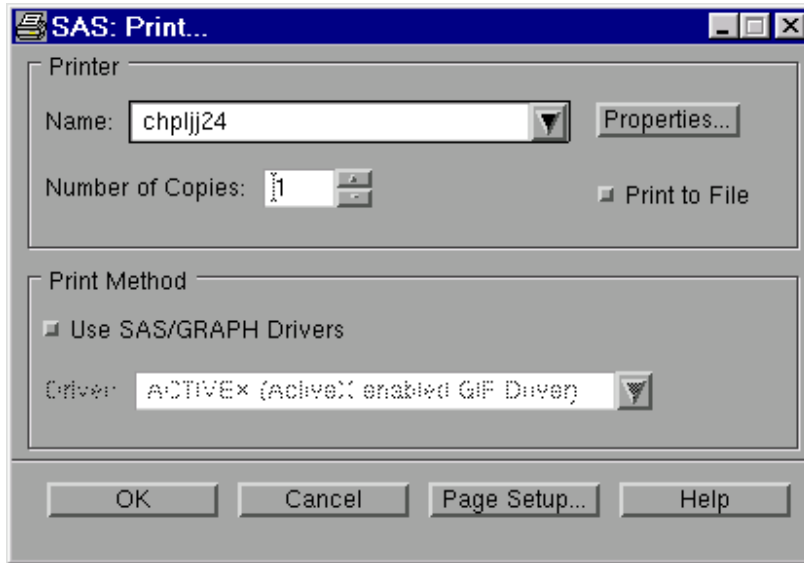
After clicking **OK**, if SAS displays a clock icon for a long time and you are sending output to a network printer, your printer server might be down. If so, you will eventually see a message in the shell where you invoked your SAS session that indicates that the server is down.

Printing from GRAPH Windows

Open the Print Dialog Box from the GRAPH Window

With Universal Printing, you can use the Print dialog box to print the contents of a GRAPH window. Click in the window to make it the active window, and then issue the DMPRINT command or select **File ► Print** to open the Print dialog box.

Display 4.2 Print Dialog Box for Graphs



Note: In most cases, fonts set through the Print dialog box have no effect when you print from GRAPH windows. However, some SAS/GRAPH drivers use Universal Printing and can be affected by the fonts set in the dialog box. Make sure that you specify the correct options on a GOPTIONS statement. Δ

Specifics for SAS/GRAPH Drivers

To print output using a SAS/GRAPH driver, select **Use SAS/GRAPH Drivers**. Select the down arrow beside the **Driver** field to display the available drivers. Make sure that your printer destination has been set inside the device using the GDEVICE procedure or the GOPTIONS statement. For complete information about printing from GRAPH windows, refer to *SAS/GRAPH: Reference* and the online Help for SAS/GRAPH.

Troubleshooting Print Server Errors

After clicking **OK**, if SAS displays a clock icon for a long time and you are sending output to a network printer, your printer server might be down. If so, you will eventually see a message in the shell where you invoked your SAS session that indicates that the server is down.

Using Commands to Print in UNIX Environments

Differences between the PRTFILE, PRINT, and FILE Commands

In the SAS windowing environment, you can use the PRTFILE, PRINT, and FILE commands to send the contents of the active window to an output device.

The following table lists the results of each of these commands.

Table 4.3 Routing Output Commands

Command	Action Performed
PRTFILE	specifies the filename or fileref for your output.
FILE	sends the contents of the active window to the filename or fileref you specify.
PRINT	sends the contents of the active window either: <ul style="list-style-type: none"> <input type="checkbox"/> to your default printer when issued from the command line of the window. <input type="checkbox"/> to the location specified with the PRTFILE command.

Sending Output to a UNIX Command

If you want to send your output to a UNIX command, you can use the FILENAME statement. The FILENAME statement enables you to create filerefs that point to printers, plotters, or external files or filerefs that pipe to a UNIX command. For more information, see “FILENAME Statement” on page 318.

Specifying the Print File

When you issue the PRINT command, SAS sends your output to your default printer unless you specify a print file. You can specify a print file in two ways:

- entering the PRTFILE command:

PRTFILE *file-spec* **CLEAR** | **APPEND** | **REPLACE**

The *file-spec* can be either a fileref or a filename.

- selecting **File** \blacktriangleright **Print Utilities** \blacktriangleright **Set Print File** and entering the name of the print file if you are using forms. This option is available only when Universal Printing is turned off.

Printing the Contents of a Window

Using PRTFILE and PRINT with a Fileref

You can use the PRTFILE command, followed by the PRINT command, to print the contents of windows. PRTFILE establishes the destination, and PRINT sends the contents of the window to that destination. If you do not specify a destination with the PRTFILE command, PRINT automatically sends the window contents to your default printer.

Steps for Sending Output Directly to a Printer

If you want to send output directly to a printer, you must first submit the FILENAME statement to assign a fileref to the PRINTER or PIPE device. For example, to print the contents of your Output window, complete the following steps:

Table 4.4 Printing the Contents of Your Output Window

Step	Action	Example
1	Submit a FILENAME statement or FILENAME function to associate a fileref with a system printer (PRINTER device type) or a UNIX command (PIPE device type). Enclose the printer name or UNIX command in either single or double quotation marks.	filename myrpt printer 'bldga2'; or filename ascout pipe 'lp -dmyljet'; For more information, see “Examples of FILENAME Statements Using PRINTER and PIPE” on page 99.
2	Issue the PRTFILE command as described in “Specifying the Print File” on page 98. Specify the fileref from your FILENAME statement or FILENAME function.	prtfile myrpt
3	Issue the PRINT command from the command line of the windows whose contents you want to print. If you are sending output to a system printer or if you are using forms-based printing, then you can print the contents of more than one window.	
4	Enter A in the dialog box that appears to warn you that the destination file already exists. The A value tells SAS to append the window contents to the destination file.	
5	Submit a FILENAME statement or FILENAME function to clear (deassign) the fileref.	filename myrpt clear;

To clear the print file setting, issue the PRTFILE CLEAR command.

Examples of FILENAME Statements Using PRINTER and PIPE

The following statement associates MyRpt with the system printer named BldgA2 and specifies two copies of every printout:

```
filename myrpt printer 'bldga2 -n2';
```

(See the documentation for your print command for information about other options that you can specify.)

The following statement enables you to print output using the **lp** command on the printer named myljet:

```
filename ascout pipe 'lp -dmyljet';
```

The following statement sends output to the **lp** command and redirects any error messages produced by this command to the LpError file in your home directory:

```
filename myrpt pipe 'lp 2>${HOME}/lperror';
```

Note: Redirecting standard error is allowed only in the Bourne and Korn shells. Δ

If you frequently use the same print command and destination, you can add the appropriate FILENAME statement to your autoexec file. See “Customizing Your SAS Session by Using System Options” on page 17 for more information.

Using the FILE Command

You can use the FILE command to copy the contents of many different windows to external files. Issue the FILE command on the command line of the window whose contents you want to copy. For example, to copy the contents of the Log window to `/u/myid/log/app1`, issue the following command on the command line of the Log window:

```
file '/u/myid/log/app1'
```

If the file does not exist, SAS creates it. If the file already exists, a dialog box asks you whether you want to replace it or to append data to the existing data.

If you have already associated a fileref with your external file, you can use the fileref instead of the filename:

```
file myref
```

Note: If you use the FILE command to save your output, carriage-control information is not saved (that is, page breaks are removed from the output). You might want to use the PRINT command with the FILE option instead:

```
PRINT FILE=fileref | 'pathname'
```

Δ

Using the PRINTTO Procedure in UNIX Environments

Important Note about the PRINTTO Procedure

Any time you use PROC PRINTTO to route output, you must close the output device before PROC PRINTTO will release the output or log and send it to the destination you have specified. To close the output device, issue PROC PRINTTO without any parameters:

```
proc printto;
run;
```

Issuing PROC PRINTTO without any parameters closes the output device, generates output, and reroutes the log and procedure output to their default destinations. See Table 4.1 on page 93 for a list of the default destinations.

For more information, see “PRINTTO Procedure” on page 304 and *Base SAS Procedures Guide*.

Using the LOG= and PRINT= Options

When you use the PRINTTO procedure with its LOG= and PRINT= options, you can route the SAS log or SAS procedure output to an external file or a fileref from any mode. Specify the external file or the fileref in the PROC PRINTTO statement. The following example routes procedure output to `/u/myid/output/prog1`:

```
proc printto print='/u/myid/output/prog1' new;
run;
```

The NEW option causes any existing information in the file to be cleared. If you omit the NEW option from the PROC PRINTTO statement, the SAS log or procedure output is appended to the existing file.

If you plan to specify the same destination several times in your SAS program, you can assign a fileref to the file using a FILENAME statement. (See “Assigning Filerefs to External Files or Devices with the FILENAME Statement” on page 73 for information and examples.)

Routing Output to a Universal Printer

You can direct output directly to your Universal Printer by using the UPRINTER device type:

```
filename myoutput uprinter;
proc printto print=myoutput;
run;
```

Output will be sent to your default Universal Printer. This output will be in PostScript or PCL format.

Routing Output to a Printer

You can direct output directly to your system printer by using the PRINTER device type:

```
filename myoutput printer;
proc printto print=myoutput;
run;
```

Output will be sent to your default system printer or, if you have specified the SYSPRINT system option, to the printer specified with that option. This method will produce output in ASCII format.

Piping Output to a UNIX Command

You can also use the PIPE device type to send output to a UNIX command. When you specify the print command, you might also want to specify a destination for (redirect) any error messages produced by the print command. Enclose the UNIX command in either single or double quotation marks. The following example associates the fileref MyOutput with the print command **lp**, which will send output to the printer named myljet:

```
filename myoutput pipe 'lp -dmyljet';
proc printto print=myoutput;
run;
```

You can send the SAS log to the same printer by using the LOG= option:

```
filename mylog pipe 'lp -dmyljet';
proc printto log=mylog;
run;
```

The log and procedure output continue to be routed to the designated external file until another PROC PRINTTO statement reroutes them.

Routing Output to a Terminal

In batch mode, you can direct output to a terminal by associating a fileref with a terminal and then using PROC PRINTTO to send output to that fileref. In the FILENAME statement, specify the TERMINAL device-type and the special file associated with the terminal. For example, the following statements send the SAS log to the terminal that is associated with the `/dev/tty3` special file:

```
filename term terminal '/dev/tty3';
proc printto log=term;
run;
```

Using SAS System Options to Route Output

Changing the Output Destination Using the LOG, PRINT, ALTLOG, and ALTPRINT System Options

You can use SAS system options to change the destination of the SAS log and procedure output. The options that you use depend on which task you want to accomplish:

- To route your SAS log or procedure output to an external file *instead of* to their default destinations, use the LOG and PRINT system options.
- To route the log or output to an external file *in addition* to their default destinations, use the ALTLOG and ALTPRINT system options. This method works in all modes of running SAS.

LOG and PRINT are normally used in batch and interactive line modes. These system options have no effect in the windowing environment. If you are running in the windowing environment, use the ALTLOG and ALTPRINT system options.

You can specify these options in following locations:

- the SAS command
- a configuration file
- the SASV9_OPTIONS environment variable

For example, you could specify these options in the SAS command as follows:

```
sas -log '/u/myid/log' -print '/u/myid/prt'
sas -altlog '/u/myid/log' -altprint '/u/myid/prt'
```

See “Ways to Specify a SAS System Option” on page 18 for more information.

Printing Large Files with the PIPE Device Type in UNIX Environments

When you print a file with the `lp` command, a symbolic link is created from the file to the `/usr/spool` directory. When you pipe output to the `lp` command, the output is copied under the `/usr/spool` directory.

If you experience problems printing large files using the PIPE device type, you can circumvent the problem in either of the following ways:

- save the print file to a disk file and then print it with the **lp** command. Issue the PRINT command from the output or Log window:

```
print file='bigfile'
```

Exit your SAS session and print the file, or use the SAS X command to print the file from within your SAS session:

```
x 'lp -dmylsrjt bigfile'
```

- create a fileref using the PIPE device type that can handle large files. For example, the following fileref saves the print file to disk, prints the saved file, and then removes the file:

```
filename myfile pipe 'cat >bigfile;lp -dmylsrjt bigfile;rm bigfile;';
```

Changing the Default Print Destination in UNIX Environments

When you print a file, SAS looks in the following locations to determine where to send output. The locations are listed in order of precedence:

- 1 The destination specified in Universal Printing or the form printer device that you are using. See Universal Printing or forms printing in *SAS Language Reference: Dictionary* for more information.
- 2 The value specified in the SYSPRINT system option. You can use the SYSPRINT option to set your default print destination. Use the SYSPRINT system option to specify the destination option that is used with your print command. For example, if your print command is **lp**, you can set the default destination to the printer named myljet by entering the following OPTIONS statement:

```
options sysprint='-dmyljet';
```

- 3 The value of the \$LPDEST environment variable. See “Defining Environment Variables in UNIX Environments” on page 23 for more information.

SAS uses the first destination that it finds. If you specify a destination in all three locations, SAS uses the destination specified by Universal Printing.

Changing the Default Print Command in UNIX Environments

UNIX uses **lp** as the default print command. You can use the PRINTCMD system option to specify a different print command. For example, you can change your default print command to **lpr** by entering the following at SAS invocation:

```
sas -printcmd 'lpr'
```

You can also customize your default print command in your SAS configuration file. If you use this method, then you will not have to change the default print command every time you invoke SAS. For more information, see “PRINTCMD System Option” on page 393.

Controlling the Content and Appearance of Output in UNIX Environments

Overview of Controlling the Content and Appearance of Output

Some of the attributes of the SAS log and procedure output depend on the destination to which they are being sent. For example, if the log and output are being sent to your display, the default line and page size are derived from your display. If one or both of these files are sent to the system printer or written to a file, the default line size and page size depend on your printer and page setup. The line size and page size for your current settings can be seen in the Print dialog box.

Some of the attributes of the SAS log and procedure output depend on the mode in which you are running. For example, if you are running in interactive line mode, SAS source statements are not echoed to the SAS log. If you are using the SAS windowing environment all source statements are written to the log as they are submitted. In batch mode, the log and procedure output are formatted for a standard system printer.

See “Customizing Your SAS Session by Using System Options” on page 17 for information about specifying system options.

SAS Log Options

Use the following options to control the contents of the log. See Chapter 18, “System Options under UNIX,” on page 341 for information about specifying options.

FULLSTIMER

NOFULLSTIMER

controls whether a list of resources (such as I/O performed, page faults, elapsed time, and CPU time) used for each PROC or DATA step is written to the log.

NOFULLSTIMER is the default.

LINESIZE=*width*

controls the line length used. *Width* can be any value from 64 to 256.

NEWS

NONEWS

controls whether messages are written to the SAS log. NEWS is the default.

NOTES

NONOTES

controls printing of NOTES on the log. NOTES is the default setting for all execution modes. Specify NOTES unless your SAS program is completely debugged.

PAGESIZE=*n*

controls the number of lines that are printed on each page. *N* can be any number from 15 to 32767.

SOURCE

NOSOURCE

controls whether SAS source statements are written to the log. NOSOURCE is the default setting in interactive line mode; otherwise, SOURCE is the default.

SOURCE2
NOSOURCE2

controls whether SAS statements that are included with %INCLUDE statements are written to the log. NOSOURCE2 is the default setting for all execution modes.

STIMER
NOSTIMER

controls whether user CPU time and elapsed time are written to the log. STIMER is the default.

Procedure Output Options

Use these system options to control the contents of the procedure output:

CENTER
NOCENTER

controls whether the printed results are centered or left-aligned on the procedure output page. CENTER is the default.

DATE
NODATE

controls whether the date is written at the top of each procedure output page. DATE is the default.

LINESIZE=*width*

controls the line length used. *Width* can be any value from 64 to 256.

NUMBER
NONUMBER


controls whether the output page number is written on each procedure output page. NUMBER is the default.

PAGENO=*n*

resets the current page number in the print file. The default page number at the beginning of the SAS session is 1. The pages are numbered sequentially throughout the SAS session unless the PAGENO option is specified in an OPTIONS statement during the session.

PAGESIZE=*n*

controls the number of lines that are printed on each page. *N* can be any number from 15 to 32,767.



CHAPTER

5

Accessing Shared Executable Libraries from SAS

<i>Overview of Shared Libraries in SAS</i>	108
<i>What Is a Shared Library?</i>	108
<i>Invoking Shared Libraries from within SAS</i>	108
<i>Steps for Accessing an External Shared Library</i>	108
<i>The SASCBTBL Attribute Table</i>	109
<i>Introduction to the SASCBTBL Attribute Table</i>	109
<i>What Is the SASCBTBL Attribute Table?</i>	109
<i>Syntax of the Attribute Table</i>	109
<i>ROUTINE Statement</i>	109
<i>ARG Statement</i>	112
<i>The Importance of the Attribute Table</i>	113
<i>Special Considerations When Using Shared Libraries</i>	114
<i>32-Bit and 64-Bit Considerations</i>	114
<i>Compatibility between Your Shared Libraries and SAS</i>	114
<i>Memory Storage Allocated by the Shared Library</i>	114
<i>Naming Considerations When Using Shared Libraries</i>	115
<i>Example of Creating a Symbolic Link</i>	115
<i>Using PEEKLONG Functions to Access Character String Arguments</i>	116
<i>Accessing Shared Libraries Efficiently</i>	117
<i>Grouping SAS Variables as Structure Arguments</i>	117
<i>Example: Grouping Your System Information as Structure Arguments</i>	118
<i>Using Constants and Expressions as Arguments to the MODULE Function</i>	119
<i>Specifying Formats and Informats to Use with MODULE Arguments</i>	120
<i>C Language Formats</i>	120
<i>FORTRAN Language Formats</i>	120
<i>PL/I Language Formats</i>	121
<i>COBOL Language Formats</i>	121
<i>%CSTRw. Format</i>	122
<i>%BYVALw. Format</i>	123
<i>Understanding MODULE Log Messages</i>	124
<i>Examples of Accessing Shared Executable Libraries from SAS</i>	126
<i>Example 1: Updating a Character String Argument</i>	126
<i>Example 2: Passing Arguments by Value</i>	127
<i>Example 3: Using PEEKCLONG to Access a Returned Pointer</i>	128
<i>Example 4: Using Structures</i>	129
<i>Example 5: Invoking a Shared Library Routine from PROC IML</i>	131

Overview of Shared Libraries in SAS

What Is a Shared Library?

Shared libraries in UNIX are libraries that contain executable programs that are written in any of several programming languages. In UNIX, the names of these programs typically end with a .so or .sl extension. However, they are not constrained to this naming convention.

Shared libraries are a mechanism for storing useful routines that might be needed by multiple applications. When an application needs a routine that resides in an external shared library, it loads the shared library, invokes the routine, and unloads the shared library upon completion.

Invoking Shared Libraries from within SAS

SAS provides routines and functions that let you invoke these external routines from within SAS. You can access the shared library routines from the DATA step, the IML procedure, and SCL code. You use the MODULE family of SAS CALL routines and functions (including MODULE, MODULEN, and MODULEC), as well as the SAS/IML functions and CALL routines (including MODULEIC, MODULEIN, and MODULEI), to invoke a routine that resides in a shared library. This documentation refers to the MODULE family of CALL routines and functions generically as the MODULE function.

For information about MODULE, MODULEN, and MODULEC, see the *SAS Language Reference: Dictionary*. For information about MODULEIC, MODULEIN, and MODULEI, see the *SAS/IML User's Guide*.

Steps for Accessing an External Shared Library

Use the following steps to access an external shared library routine:

- 1 Create a text file that describes the shared library routine you want to access, including the arguments it expects and the values it returns (if any). This attribute file must be in a special format, as described in “The SASCBTBL Attribute Table” on page 109.
- 2 Use the FILENAME statement to assign the SASCBTBL fileref to the attribute file you created.
- 3 In a DATA step or SCL code, use a call routine or function (MODULE, MODULEN, or MODULEC) to invoke the shared library routine. The specific function you use depends on the type of expected return value (none, numeric, or character). (You can also use MODULEI, MODULEIN, or MODULEIC within a PROC IML step.) The MODULE functions are described in “CALL MODULE Routine” on page 258.

CAUTION:

Only experienced programmers should access external routines in shared libraries. By accessing a function in a shared library, you transfer processing control to the external function. If done improperly, or if the external function is not reliable, you might lose data, get unreliable results, or receive severe errors. \triangle

The SASCBTBL Attribute Table

Introduction to the SASCBTBL Attribute Table

Because the MODULE function invokes an external routine that SAS knows nothing about, you must supply information about the routine's arguments so that the MODULE function can validate them and convert them, if necessary. For example, suppose you want to invoke a routine that requires an integer as an argument. Because SAS uses floating-point values for all of its numeric arguments, the floating-point value must be converted to an integer before you invoke the external routine. The MODULE function looks for this attribute information in an attribute table that is referred to by the SASCBTBL fileref.

What Is the SASCBTBL Attribute Table?

The attribute table is a sequential text file that contains descriptions of the routines you can invoke with the MODULE function. The table defines how the MODULE function should interpret supplied arguments when it builds a parameter list to pass to the called routine.

The MODULE function locates the table by opening the file that is referenced by the SASCBTBL fileref. If you do not define this fileref, the MODULE function simply calls the requested shared library routine without altering the arguments.

CAUTION:

Using the MODULE function without defining an attribute table can cause SAS to crash, produce unexpected results, or result in severe errors. You need to use an attribute table for all external functions that you want to invoke. △

Syntax of the Attribute Table

The attribute table should contain the following items:

- a description in a ROUTINE statement for each shared library routine you intend to call
- descriptions in ARG statements for each argument that is associated with the routine you intend to call

At any point in the attribute table file, you can create a comment using an asterisk (*) as the first non-blank character of a line or after the end of a statement (following the semicolon). You must end the comment with a semicolon.

ROUTINE Statement

Here is the syntax of the ROUTINE statement:

```
ROUTINE name MINARG=minarg MAXARG=maxarg
  <CALLSEQ=BYVALUE | BYADDR>
  <TRANSPPOSE=YES | NO> <MODULE=shared-library-name>
  <RETURNS=DBLPTR | CHAR<n> | DOUBLE | LONG | PTR | SHORT |
  [U]INT32 | [U]INT64 | ULONG | USHORT>
```

The following are descriptions of the ROUTINE statement attributes:

ROUTINE *name*

starts the ROUTINE statement. You need a ROUTINE statement for every shared library function you intend to call. The value for *name* must match the routine name or ordinal you specified as part of the 'module' argument in the MODULE function, where *module* is the name of the shared library (if not specified by the MODULE attribute) and the routine name or ordinal. For example, in order to specify `libc,getcwd` in the MODULE function call, the ROUTINE *name* should be `getcwd`.

The *name* argument is case sensitive, and is required for the ROUTINE statement.

MINARG=*minarg*

specifies the minimum number of arguments to expect for the shared library routine. In most cases, this value will be the same as MAXARG; but some routines do allow a varying number of arguments. This attribute is required.

MAXARG=*maxarg*

specifies the maximum number of arguments to expect for the shared library routine. This attribute is required.

CALLSEQ=BYVALUE | BYADDR

indicates the calling sequence method used by the shared library routine. Specify BYVALUE for call-by-value and BYADDR for call-by-address. The default value is BYADDR.

FORTRAN and COBOL are call-by-address languages. C is usually call-by-value, although a specific routine might be implemented as call-by-address.

The MODULE function does not require that all arguments use the same calling method. You can identify any exceptions by using the BYVALUE and BYADDR options in the ARG statement.

TRANSPOSE=YES | NO

specifies whether SAS transposes matrices that have both more than one row and more than one column before it calls the shared library routine. This attribute applies only to routines called from within PROC IML with MODULEI, MODULEIC, and MODULEIN.

TRANSPOSE=YES is necessary when you are calling a routine that is written in a language that does not use row-major order to store matrices. (For example, FORTRAN uses column-major order.)

For example, consider this matrix with three columns and two rows:

```

                columns
                1  2  3
                -----
rows  1 | 10 11 12
      2 | 13 14 15

```

PROC IML stores this matrix in memory sequentially as 10, 11, 12, 13, 14, 15. However, FORTRAN routines will expect this matrix as 10, 13, 11, 14, 12, 15.

The default value is NO.

MODULE=*shared-library-name*

names the executable module (the shared library) in which the routine resides. You do not need to specify this attribute if the name of the shared library is the same name as the routine. If you specify the MODULE attribute here in the ROUTINE statement, then you do not need to include the module name in the *module* argument of the MODULE function (unless the shared library routine

name you are calling is not unique in the attribute table). The MODULE function is described in “CALL MODULE Routine” on page 258.

You can have multiple ROUTINE statements that use the same MODULE name. You can also have duplicate routine names that reside in different shared libraries.

The MODULE function searches the directories that are defined in each operating system’s library path environment variable when it attempts to load the shared library argument provided in the MODULE attribute. The following table lists this environment variable for each UNIX operating system that SAS supports.

Table 5.1 Shared Library Environment Variable Name

Operating Environment	Environment Variable Name
Solaris	\$LD_LIBRARY_PATH
AIX/R	\$LIBPATH
HP-UX	\$LD_LIBRARY_PATH or \$SHLIB_PATH
Linux	\$LD_LIBRARY_PATH

Note: For more information about these environment variables, see the man pages for your operating environment. Δ

You can also use the PATH system option to point to the directory that contains the shared library specified in the MODULE= option. Using the PATH system option overrides your system’s environment variable when you load the shared library. For more information, see “PATH System Option” on page 392.

RETURNS=DBLPTR | CHAR<n> | DOUBLE | LONG | PTR | SHORT | [U]INT32 | [U]INT64 | ULONG | USHORT

specifies the type of value that the shared library routine returns. This value will be converted as appropriate, depending on whether you use MODULEC (which returns a character) or MODULEN (which returns a number). The following are the possible return value types:

DBLPTR

pointer to a double-precision floating point number (instead of using a floating-point register). See the documentation for your shared library routine to determine how it handles double-precision floating-point values.

CHAR<n>

pointer to a character string up to *n* bytes long. The string is expected to be null-terminated and will be blank-padded or truncated as appropriate. If you do not specify *n*, the MODULE function uses the maximum length of the receiving SAS character variable.

DOUBLE

double-precision floating-point number.

LONG

long integer.

PTR

character string being returned.

SHORT

short integer.

[U]INT32

32-bit unsigned integer.

[U]INT64
64-bit unsigned integer.

ULONG
unsigned long integer.

USHORT
unsigned short integer.

If you do not specify the RETURNS attribute, you should invoke the routine with only the MODULE and MODULEI call routines. You will get unpredictable values if you omit the RETURNS attribute and invoke the routine using the MODULEN/MODULEIN or MODULEC/MODULEIC functions.

ARG Statement

The ROUTINE statement must be followed by as many ARG statements as you specified in the MAXARG= option. The ARG statements must appear in the order that the arguments will be specified within the MODULE function.

Here is the syntax for each ARG statement:

```
ARG argnum NUM|CHAR <INPUT|OUTPUT|UPDATE>
<NOTREQD|REQUIRED> <BYADDR|BYVALUE> <FDSTART>
<FORMAT=format>;
```

Here are the descriptions of the ARG statement attributes:

ARG *argnum*
defines the argument number. This a required attribute. Define the arguments in ascending order, starting with the first routine argument (ARG 1).

NUM | CHAR
defines the argument as numeric or character. This attribute is required.
If you specify NUM here but pass the routine a character argument, the argument is converted using the standard numeric informat. If you specify CHAR here but pass the routine a numeric argument, the argument is converted using the BEST12 informat.

INPUT | OUTPUT | UPDATE
indicates the argument is either input to the routine, an output argument, or both. If you specify INPUT, the argument is converted and passed to the shared library routine. If you specify OUTPUT, the argument is *not* converted, but is updated with an outgoing value from the shared library routine. If you specify UPDATE, the argument is converted, passed to the shared library routine, *and* updated with an outgoing value from the routine.

You can specify OUTPUT and UPDATE only with variable arguments (that is, no constants or expressions are allowed).

NOTREQD | REQUIRED
indicates whether the argument is required. If you specify NOTREQD, then the MODULE function can omit the argument. If other arguments follow the omitted argument, identify the omitted argument by including an extra comma as a placeholder. For example, to omit the second argument to routine XYZ, you would specify:

```
call module('XYZ',1,,3);
```

CAUTION:

Be careful when using NOTREQD; the shared library routine must not attempt to access the argument if it is not supplied in the call to MODULE. If the routine does attempt to access it, you might receive unexpected results or severe errors. Δ

The REQUIRED attribute indicates that the argument is required and cannot be omitted. REQUIRED is the default value.

BYADDR | BYVALUE

indicates whether the argument is passed by reference or by value.

BYADDR is the default value unless CALLSEQ=BYVALUE was specified in the ROUTINE statement, in that case, BYVALUE is the default. Specify BYADDR when you are using a call-by-value routine that also has arguments to be passed by address.

FDSTART

indicates that the argument begins a block of values that are grouped into a structure whose pointer is passed as a single argument. Note that all subsequent arguments are treated as part of that structure until the MODULE function encounters another FDSTART argument.

FORMAT=*format*

names the format that presents the argument to the shared library routine. Any formats supplied by SAS, PROC FORMAT style formats, or SAS/TOOLKIT formats are valid. Note that this format must have a corresponding valid informat if you specified the UPDATE or OUTPUT attribute for the argument.

The FORMAT= attribute is not required, but is recommended because format specification is the primary purpose of the ARG statements in the attribute table.

CAUTION:

Using an incorrect format can produce invalid results, cause SAS to crash, or result in serious errors. Δ

The Importance of the Attribute Table

The MODULE function relies heavily on the accuracy of the information in the attribute table. If this information is incorrect, unpredictable results can occur (including a system crash).

Consider an example routine **xyz** that expects two arguments: an integer and a pointer. The integer is a code indicating what action takes place. For example, action 1 means that a 20-byte character string is written into the area that is pointed to by the second argument, the pointer.

Now suppose you call **xyz** using the MODULE function, but you indicate in the attribute table that the receiving character argument is only 10 characters long:

```
routine xyz minarg=2 maxarg=2;
arg 1 input num byvalue format=ib4.;
arg 2 output char format=$char10.;
```

Regardless of the value given by the LENGTH statement for the second argument to MODULE, MODULE passes a pointer to a 10-byte area to the **xyz** routine. If **xyz** writes 20 bytes at that location, the 10 bytes of memory following the string provided by MODULE are overwritten, causing unpredictable results:

```
data _null_;
  length x $20;
  call module('xyz',1,x);
run;
```

The call might work fine, depending on which 10 bytes were overwritten. However, overwriting can cause you to lose data or cause your system to crash.

Also, note that the PEEKLONG and PEEKCLONG functions rely on the validity of the pointers you supply. If the pointers are invalid, it is possible that severe errors will result. For example, this code causes an error:

```
data _null_;
  length c $10;
  /* trying to copy from address 0!!!*/
  c = peekclong(0,10);
run;
```

Special Considerations When Using Shared Libraries

32-Bit and 64-Bit Considerations

Compatibility between Your Shared Libraries and SAS

Starting in SAS 9, SAS is a 64-bit application that runs on all supported UNIX environments that are 64-bit enabled. The only exception is the Linux version of SAS which is a 32-bit application. When you call external routines in shared libraries, the shared library needs to be compatible with SAS.

For example, if you are running a 64-bit version of SAS on Solaris, you need to call a routine that is located in `libc.so`. In order for this shared library to be compatible with SAS, it needs to be a 64-bit shared library.

To determine whether a vendor supplied library is 32-bit or 64-bit, you can use the `FILE` command. The following output shows the results of using this command on Solaris for a 32-bit and 64-bit library:

```
$ file libc.so
libc.so: ELF 32--bit MSB dynamic lib SPARC Version 1, dynamically linked,
not stripped

$ file ./libc.so
./libc.so: ELF 64--bit MSB dynamic lib SPARCV9 Version 1, dynamically linked,
not stripped
```

Memory Storage Allocated by the Shared Library

When specifying your SAS format and informat for each routine argument in the `FORMAT` attribute of the `ARG` statement, you need to consider the amount of memory the shared library allocates for the parameters that it receives and returns. To determine how much storage is being reserved for the input and return parameters of the routine in the external shared library, you can use the `sizeof()` C function.

The following table lists the typical memory allocations for C data types for 32-bit and 64-bit systems:

Table 5.2 Memory Allocations for C Data Types

Type	32-Bit System Size (Bytes)	32-Bit System Size (Bits)	64-Bit System Size (Bytes)	64-Bit System Size (Bits)
char	1	8	1	8
short	2	16	2	16
int	4	32	4	32
long	4	32	8	64
long long	8	64	8	64
float	4	32	4	32
double	8	64	8	64
pointer	4	32	8	64

For information about the SAS formats to use for your data types, see “Specifying Formats and Informats to Use with MODULE Arguments” on page 120.

Naming Considerations When Using Shared Libraries

SAS loads external shared libraries that meet the following naming constraints:

- The name is eight characters or less.
- The name does not contain a period.

If the name of your external shared library is greater than eight characters or contains a period, then you can create a symbolic link to point to the destination of the shared library. Once the link is created, you can add the name of the symbolic link to the MODULE statement in the SASCBTBL attribute table. When you are ready to execute your SAS program, use the PATH system option to point to the directory that contains the symbolic link.

Example of Creating a Symbolic Link

The Hewlett-Packard shared library `libc.sl` that is installed in the `/usr/lib/pa20_64` directory contains a period in the name. Before SAS will load this shared library, you need to create a symbolic link that meets the naming convention of eight characters or less and no period. The symbolic link shown in the following example points to the target location of `libc.sl`:

```
$ ln -s /usr/lib/pa20_64/libc.sl /tmp/libclnk
```

After the symbolic link is created, you can then update the `MODULE=` option in the SASCBTBL attribute table, as shown in the following code:

```
routine name minarg=2 maxarg=2 returns=short module=libclnk;
arg 1 char output byaddr fdstart format=%cstr9.;
arg 2 char output format=%cstr9.;
```

To load the shared library during your invocation of SAS, type the following command:

```
/usr/local/sasv91/sas -path /tmp module.sas
```

Using PEEKLONG Functions to Access Character String Arguments

Because the SAS language does not provide pointers as data types, you can use the SAS PEEKLONG functions to access the data stored at these address values.

For example, the following program demonstrates how the address of a pointer is supplied and how it can set the pointer to the address of a static table containing the contiguous integers 1, 2, and 3. It also calls the `useptr` routine in the `useptr` shared library on a 64-bit operating system.

```
static struct MYTABLE {
  int value1;
  int value2;
  int value3;
} mytable = {1,2,3};

useptr(toset)
char **toset;
{
  *toset = (char *)&mytable
}
```

The following is the SASCBTBL attribute table entry:

```
routine useptr minarg=1 maxarg=1;
arg 1 char update format=$char20.;
```

The following is the SAS code:

```
data _null_;
  length ptrval $20 thedata $12;
  call module('*i','useptr',ptrval);
  thedata=peekclong(ptrval,12);

  /* Converts hexadecimal data to character data */
  put thedata=$hex24.;

  /* Converts hexadecimal positive binary values to fixed or floating point value */
  ptrval=hex40.;
run;
```

SAS writes the following output to the log.

Output 5.1 Log Output for Accessing Character Strings with the PEEKLONG Function

```
thedata=000000010000000200000003 ptrval=800003FFFF0C
```

In this example, the PEEKCLONG function is given two arguments, a pointer via a numeric variable and a length in bytes. PEEKCLONG returns a character string of the specified length containing the characters at the pointer location.

For more information about the PEEKLONG functions, see “PEEKLONG Function” on page 275.

Accessing Shared Libraries Efficiently

The `MODULE` function reads the attribute table that is referenced by the `SASCBTBL` fileref once per step (DATA step, PROC IML step, or SCL step). It parses the table and stores the attribute information for future use during the step. When you use the `MODULE` function, SAS searches the stored attribute information for the matching routine and module names. The first time you access a shared library during a step, SAS loads the shared library and determines the address of the requested routine. Each shared library you invoke stays loaded for the duration of the step, and is not reloaded in subsequent calls. All modules and routines are unloaded at the end of the step.

In the following example, the attribute table has the following basic form:

```
* routines XYZ and BBB in FIRST.Shared Library;
routine XYZ minarg=1 maxarg=1 module=FIRST;
arg 1 num input;
routine BBB minarg=1 maxarg=1 module=FIRST;
arg 1 num input;
* routines ABC and DDD in SECOND.Shared Library;
routine ABC minarg=1 maxarg=1 module=SECOND;
arg 1 num input;
routine DDD minarg=1 maxarg=1 module=SECOND;
arg 1 num input;
```

The DATA step code looks like the following:

```
filename sascbtbl 'myattr.tbl';
data _null_;
  do i=1 to 50;
    /* FIRST.Shared Library is loaded only once */
    value = modulen('XYZ',i);
    /* SECOND.Shared Library is loaded only once */
    value2 = modulen('ABC',value);
    put i= value= value2=;
  end;
run;
```

In this example, `MODULEN` parses the attribute table during DATA step compilation. In the first loop iteration ($i=1$), `FIRST.Shared Library` is loaded and the `XYZ` routine is accessed when `MODULEN` calls for it. Next, `SECOND.Shared Library` is loaded and the `ABC` routine is accessed. For subsequent loop iterations (starting when $i=2$), `FIRST.Shared Library` and `SECOND.Shared Library` remain loaded, so the `MODULEN` function simply accesses the `XYZ` and `ABC` routines. SAS unloads both shared libraries at the end of the DATA step.

Note that the attribute table can contain any number of descriptions for routines that are not accessed for a given step. The presence of the attribute table does not cause any additional overhead (apart from a few bytes of internal memory to hold the attribute descriptions). In the above example, `BBB` and `DDD` are in the attribute table but are not accessed by the DATA step.

Grouping SAS Variables as Structure Arguments

A common need when calling external routines is to pass a pointer to a structure. Some parts of the structure might be used as input to the routine, while other parts might be replaced or filled in by the routine. Even though SAS does not have structures

in its language, you can indicate to the `MODULE` function that you want a particular set of arguments grouped into a single structure. You indicate this grouping by using the `FDSTART` option of the `ARG` statement to flag the argument that begins the structure in the attribute table. SAS gathers that argument and all the arguments that follow (until encountering another `FDSTART` option) into a single contiguous block, and passes a pointer to the block as an argument to the shared library routine.

Example: Grouping Your System Information as Structure Arguments

This example uses the `uname` routine, which is part of the `/usr/lib/pa20_64/libc.sl` shared library in the HP-UX operating environment. This routine returns the following information about your computer system:

- the nodename on which you are executing SAS.
- the version of the operating system.
- the vendor of the operating system.
- the computer identification number.
- model type of your computer.
- the unique identification number of your class of hardware. This value could be a serial number.

The following is the C prototype for this routine:

```
int uname(struct utsname *name);
```

In C, the `utsname` structure is defined with the following members:

```
#define UTSLEN 9
#define SNLEN 15

char sysname[UTSLEN];
char nodename[UTSLEN];
char release[UTSLEN];
char version[UTSLEN];
char machine[UTSLEN];
char idnumber[SNLEN];
```

Each of the above structure members are null-terminated strings.

To call this routine using the `MODULE` function, you use the following attribute table entries:

```
* attribute table entry;
routine uname minarg=6 maxarg=6 returns=short module=libc;
arg 1 char output byaddr fdstart format=%cstr9.;
arg 2 char output          format=%cstr9.;
arg 3 char output          format=%cstr9.;
arg 4 char output          format=%cstr9.;
arg 5 char output          format=%cstr9.;
arg 6 char output          format=%cstr15.;
```

The following example shows the SAS source code to call the `uname` routine from within the `DATA` step:

```
x 'if [ ! -L ./libc ]; then ln -s /usr/lib/pa20_64/libc.sl ./libc ; fi' ;
x 'setenv LD_LIBRARY_PATH ./usr/lib:/lib:/usr/lib/pa20_64'

data _null_;
    length sysname $9 nodename $9 release $9 version $9 machine $9 idnumber $15.
```

```

retain sysname nodename release version machine idnumber ' ' ' ';
rc=modulen('uname', sysname, nodename, release, version, machine, idnumber)
put rc = ;
put sysname = ;
put nodename = ;
put release = ;
put version = ;
put machine = ;
put idnumber = ;
run;

```

SAS writes the following output to the log:

Output 5.2 Grouping SAS Variables as a Structure

```

rc=0
sysname=HP-UX
nodename=garage
release=B.11.00
version=u
machine=9000/800
idnumber=103901537

```

Using Constants and Expressions as Arguments to the MODULE Function

You can pass any kind of expression as an argument to the MODULE function. The attribute table indicates whether the argument is for input, output, or update.

You can specify input arguments as constants and arithmetic expressions. However, because output and update arguments must be able to be modified and returned, you can pass only a variable for them. If you specify a constant or expression where a value that can be updated is expected, SAS issues a warning message pointing out the error. Processing continues, but the MODULE function cannot perform the update (meaning that the value of the argument you wanted to update will be lost).

Consider these examples. Here is the attribute table:

```

* attribute table entry for ABC;
routine abc minarg=2 maxarg=2;
arg 1 input format=ib4.;
arg 2 output format=ib4.;

```

Here is the DATA step with the MODULE calls:

```

data _null_;
  x=5;
  /* passing a variable as the      */
  /* second argument - OK          */
  call module('abc',1,x);

  /* passing a constant as the     */
  /* second argument - INVALID     */
  call module('abc',1,2);

  /* passing an expression as the  */
  /* second argument - INVALID     */

```

```

    call module('abc',1,x+1);
run;

```

In the above example, the first call to `MODULE` is correct because `x` is updated by the value that the `abc` routine returns for the second argument. The second call to `MODULE` is not correct because a constant is passed. `MODULE` issues a warning indicating you have passed a constant, and passes a temporary area instead. The third call to `MODULE` is not correct because an arithmetic expression is passed, which causes a temporary location from the `DATA` step to be used, and the returned value to be lost.

Specifying Formats and Informats to Use with MODULE Arguments

You specify the SAS format and informat for each shared library routine argument by specifying the `FORMAT` attribute in the `ARG` statement. The format indicates how numeric and character values should be passed to the shared library routine and how they should be read back upon completion of the routine.

Usually, the format you use corresponds to a variable type for a given programming language. The following sections describe the proper formats that correspond to different variable types in various programming languages.

C Language Formats

C Type	SAS Format/Informat for 32-Bit Systems	SAS Format/Informat for 64-Bit Systems
double	RB8.	RB8.
float	FLOAT4.	FLOAT4.
signed int	IB4.	IB4.
signed short	IB2.	IB2.
signed long	IB4.	IB8.
char *	IB4.	IB8.
unsigned int	PIB4.	PIB4.
unsigned short	PIB2.	PIB2.
unsigned long	PIB4.	PIB8.
char[w]	\$CHARw. or \$CSTRw. (see "\$CSTRw. Format" on page 122)	\$CHARw. or \$CSTRw. (see "\$CSTRw. Format" on page 122)

Note: For information about passing character data other than as pointers to character strings, see "\$BYVALw. Format" on page 123. Δ

FORTRAN Language Formats

FORTRAN Type	SAS Format/Informat
integer*2	IB2.
integer*4	IB4.
real*4	RB4.
real*8	RB8.
character*w	\$CHARw.

The MODULE function can support FORTRAN character arguments only if they are not expected to be passed by a descriptor.

PL/I Language Formats

PL/I Type	SAS Format/Informat
FIXED BIN(15)	IB2.
FIXED BIN(31)	IB4.
FLOAT BIN(21)	RB4.
FLOAT BIN(31)	RB8.
CHARACTER(w)	\$CHARw.

The PL/I descriptions are added here for completeness. These descriptions do not guarantee that you will be able to invoke PL/I routines.

COBOL Language Formats

COBOL Format	SAS Format/Informat	Description
PIC Sxxxx BINARY	IBw.	integer binary
COMP-2	RB8.	double-precision floating point
COMP-1	RB4.	single-precision floating point
PIC xxxx or Sxxxx	Fw.	printable numeric
PIC yyyy	\$CHARw.	character

The following COBOL specifications might not match properly with the formats supplied by SAS because zoned and packed decimal are not truly defined for systems based on Intel architecture.

COBOL Format	SAS Format/Informat	Description
PIC Sxxxx DISPLAY	ZDw.	zoned decimal
PIC Sxxxx PACKED-DECIMAL	PDw.	packed decimal

The following COBOL specifications do not have true native equivalents and are usable only in conjunction with the corresponding S370Fxxx informat and format, which enables IBM mainframe-style representations to be read and written in the UNIX environment.

COBOL Format	SAS Format/Informat	Description
PIC xxxxx DISPLAY	S370FZDUw.	zoned decimal unsigned
PIC Sxxxx DISPLAY SIGN LEADING	S370FZDLw.	zoned decimal leading sign
PIC Sxxxx DISPLAY SIGN LEADING SEPARATE	S370FZDSw.	zoned decimal leading sign separate
PIC Sxxxx DISPLAY SIGN TRAILING SEPARATE	S370FZDTw.	zoned decimal trailing sign separate
PIC xxxxx BINARY	S370FIBUw.	integer binary unsigned
PIC xxxxx PACKED-DECIMAL	S370FPDUw.	packed decimal unsigned

\$CSTRw. Format

If you pass a character argument as a null-terminated string, use the \$CSTRw. format. This format looks for the last non-blank character of your character argument and passes a copy of the string with a null terminator after the last non-blank character. For example, consider the following attribute table entry:

```
* attribute table entry;
routine abc minarg=1 maxarg=1;
arg 1 input char format=$cstr10.;
```

With this entry, you can use the following DATA step:

```
data _null_;
  rc = module('abc', 'my string');
run;
```

The \$CSTR format adds a null terminator to the character string **my string** before passing it to the **abc** routine. Adding a null terminator to the character string and then passing the string to the **abc** routine is equivalent to the following attribute entry:

```
* attribute table entry;
routine abc minarg=1 maxarg=1;
arg 1 input char format=$char10.;
```

The entry would have the following DATA step:

```
data _null_;
  rc = module('abc', 'my string' || '00'x);
run;
```


The first example is easier to understand and easier to use when using variable or expression arguments.

The \$CSTR informat converts a null-terminated string into a blank-padded string of the specified length. If the shared library routine is supposed to update a character argument, use the \$CSTR informat in the argument attribute.

\$BYVAL w . Format

When you use a MODULE function to pass a single character by value, the argument is automatically promoted to an integer. If you want to use a character expression in the MODULE call, you must use the special format/informat called \$BYVAL w . The \$BYVAL w . format/informat expects a single character and will produce a numeric value, the size of which depends on w . \$BYVAL2. produces a short, \$BYVAL4. produces a long, and \$BYVAL8. produces a double. Consider this example using the C language:

```
long xyz(a,b)
  long a; double b;
  {
  static char c = 'Y';
  if (a == 'X')
    return(1);
  else if (b == c)
    return(2);
  else return(3);
  }
```

In this example, the **xyz** routine expects two arguments, a long and a double. If the long is an **x**, the actual value of the long is 88 in decimal. This result happens because an ASCII **x** is stored as hexadecimal 58, and this value is promoted to a long, represented as 0x00000058 (or 88 decimal). If the value of **a** is **x**, or 88, then a 1 is returned. If the second argument, a double, is **y** (which is interpreted as 89), then 2 is returned.

If you want to pass characters as the arguments to **xyz** then in the C language, you would invoke them as follows:

```
x = xyz('X',(double)'Z');
y = xyz('Q',(double)'Y');
```

The characters are invoked in this way because the **x** and **Q** values are automatically promoted to integers (which are the same as longs for the sake of this example), and the integer values corresponding to **Z** and **Y** are cast to doubles.

To call **xyz** using the MODULEN function, your attribute table must reflect the fact that you want to pass characters:

```
routine xyz minarg=2 maxarg=2 returns=long;
arg 1 input char byvalue format=$byval4.;
arg 2 input char byvalue format=$byval8.;
```

Note that it is important that the BYVALUE option appears in the ARG statement as well. Otherwise, MODULEN assumes that you want to pass a pointer to the routine, instead of a value.

Here is the DATA step that invokes MODULEN and passes it characters:

```
data _null_;
  x = moduln('xyz','X','Z');
  put x= ' (should be 1)';
  y = moduln('xyz','Q','Y');
  put y= ' (should be 2)';
```

```
run;
```

Understanding MODULE Log Messages

If you specify **i** in the control string parameter to **MODULE**, SAS prints several informational messages to the log. You can use these messages to determine whether you have passed incorrect arguments or coded the attribute table incorrectly.

Consider this example that uses **MODULEIN** from within the **IML** procedure. It uses the **MODULEIN** function to invoke the **changi** routine (which is stored in theoretical **TRYMOD.so**). In the example, **MODULEIN** passes the constant 6 and the matrix **x2**, which is a 4x5 matrix to be converted to an integer matrix. The attribute table for **changi** is as follows:

```
routine changi module=trymod returns=long;
arg 1 input num format=ib4. byvalue;
arg 2 update num format=ib4.;
```

The following **IML** step invokes **MODULEIN**:

```
proc iml;
  x1 = J(4,5,0);
  do i=1 to 4;
    do j=1 to 5;
      x1[i,j] = i*10+j+3;
    end;
  end;
  y1= x1;
  x2 = x1;
  y2 = y1;
  rc = modulein('*i','changi',6,x2);
  ....
```

The **'*i'** control string causes the lines shown in the following output to be printed in the log.

Output 5.3 MODULEIN Log Output

```
---PARAM LIST FOR MODULEIN ROUTINE--- CHR PARM 1 885E0AA8 2A69 (*i)
CHR PARM 2 885E0AD0 6368616E6769 (changi)
NUM PARM 3 885E0AE0 0000000000001840
NUM PARM 4 885E07F0
0000000000002C400000000000002E40000000000000304000000000000031400000000000003240
000000000000384000000000000039400000000000003A400000000000003B400000000000003C40
000000000000414000000000000040414000000000
---ROUTINE changi LOADED AT ADDRESS 886119B8 (PARMLIST AT 886033A0)--- PARM 1 06000000 <CALL-BY-VALUE>
PARM 2 88604720
0E0000000F000000010000000110000001200000018000000190000001A0000001B0000001C000000
22000000230000002400000025000000260000002C0000002D0000002E0000002F00000030000000
---VALUES UPON RETURN FROM changi ROUTINE--- PARM 1 06000000 <CALL-BY-VALUE>
PARM 2 88604720
140000001F0000002A0000003500000040000000820000008D00000098000000A3000000AE000000
F0000000FB00000006010000110100001C0100005E01000069010000740100007F0100008A010000
---VALUES UPON RETURN FROM MODULEIN ROUTINE--- NUM PARM 3 885E0AE0 0000000000001840
NUM PARM 4 885E07F0
00000000000034400000000000003F40000000000045400000000000804A400000000000005040
00000000004060400000000000A061400000000000063400000000006064400000000000C06540
0000000000006E400000000000606F4000000000
```

The output is divided into four sections.

- The first section describes the arguments passed to MODULEIN.

The CHR PARM n portion indicates that character parameter n was passed. In the example, 885E0AA8 is the actual address of the first character parameter to MODULEIN. The value at the address is hexadecimal 2A69, and the ASCII representation of that value ('*i') is in parentheses after the hexadecimal value. The second parameter is printed similarly. Only these first two arguments have their ASCII equivalents printed because other arguments might contain unreadable binary data.

The remaining parameters appear with only hexadecimal representations of their values (NUM PARM 3 and NUM PARM 4 in the example).

The third parameter to MODULEIN is numeric, and it is at address 885E0AE0. The hexadecimal representation of the floating-point number 6 is shown. The fourth parameter is at address 885E07F0, which points to an area containing all the values for the 4x5 matrix. The *i option prints the entire argument. Be careful if you use this option with large matrices, because the log might become quite large.

- The second section of the log lists the arguments that are to be passed to the requested routine and, in this case, changed. This section is important for determining whether the arguments are being passed to the routine correctly. The first line of this section contains the name of the routine and its address in memory. It also contains the address of the location of the parameter block that MODULEIN created.

The log contains the status of each argument as it is passed. For example, the first parameter in the example is call-by-value (as indicated in the log). The second parameter is the address of the matrix. The log shows the address, along with the data to which it points.

Note that all the values in the first parameter and in the matrix are long integers because the attribute table states that the format is IB4.

- In the third section, the log contains the argument values upon return from **changi**. The call-by-value argument is unchanged, but the other argument (the matrix) contains different values.
- The last section of the log output contains the values of the arguments as they are returned to the MODULEIN calling routine.

Examples of Accessing Shared Executable Libraries from SAS

Example 1: Updating a Character String Argument

This example uses the `tmpnam` routine in the shared library supplied by Solaris, `libc.so`, which is installed in the `/usr/lib/sparcv9` directory. The `tmpnam` routine generates a unique filename that can be used safely as a temporary filename. The temporary filename is typically placed in the `/var/tmp` directory.

Here is the C prototype for this routine:

```
char * tmpnam(char *s);
```

The attribute table for this prototype would be the following:

```
routine tmpnam minarg=1 maxarg=1 returns=char255. module=libc;
arg 1 char output byaddr format=%cstr255;
```

The SAS source code would be the following:

```
x 'if [ ! -L ./libc ] ; then ln -s /usr/lib/sparcv9/libc.so.1 ./libc ; fi' ;
x 'setenv LD_LIBRARY_PATH ./usr/lib/sparcv9:/usr/lib:/lib';
```

```
data _null_;
  length tempname $255 tname $255;
  retain tempname tname ' ' ;
  tname = modulec ('tmpnam', tempname);
  put tempname = ;
  put tname = ;
run;
```

The SAS log output would be the following:

Output 5.4 Updating a Character String Argument

```
tempname=/var/tmp/aaaKraydG
tname=/var/tmp/aaaKraydG
```

The POSIX standard for the maximum number of characters in a pathname is defined in `/usr/include/limits.h` to be 255 characters, so this example uses 254 as the length of the generated filename (`tempname`) with one character reserved for the null terminator. The `$CSTR255.` informat ensures that the null terminator and all subsequent characters are replaced by trailing blanks when control returns to the DATA step.

Example 2: Passing Arguments by Value

This example calls the `access` routine that is supplied by most UNIX vendors. This particular `access` routine is in the Hewlett-Packard shared library, `libc.sl`, which is installed under the `/usr/lib/pa20_64` directory.

Here is the C prototype for this routine:

```
int access(char *path, int amode);
```

The `access` routine checks the file that is referenced by the accessibility path according to the bit pattern contained in `amode`. You can use the following integer values for `amode` that correspond to the type of permission for which you are testing:

```
4  Read access
2  Write access
1  Execute (search) access
0  Check existence of file
```

A return value of 0 indicates a successful completion and the requested access is permitted. A return value of -1 indicates a failure and the requested access is not permitted.

Because the `amode` argument is a pass-by-value, this example includes the BYVALUE specification for `arg 2` in the attribute table. If both arguments were a pass-by-value, one could use the `CALLSEQ=BYVALUE` attribute in the ROUTINE statement and it would not be necessary to specify the BYVALUE option in `arg 2`.

The attribute table would be the following:

```
routine access minarg=2 maxarg=2 returns=short module=libc;
arg 1 char input byaddr format=$cstr200.;
arg 2 num input byvalue format=ib4.;
```

The SAS source code would be the following:

```
x 'if [ ! -L ./libc ] ; then ln -s /usr/lib/pa20_64/libc.so ; fi' ;
x 'setenv LD_LIBRARY_PATH ./usr/lib/pa20_64:/usr/lib:/lib' ;

data _null_;
  length path $200.;
  path='/dev';

  /* A non-root user is testing for write permission in the /dev directory */
  rc = modulen("*ie", 'access', path, 2);
  put rc = ;
run;
```

The SAS log output would be the following:

Output 5.5 Log Output If Request Access Is Permitted

```
rc=-1
```

If you changed the SAS source code to check for a Write permission in the user's \$HOME directory, the output would be different:

```
data _null_;
  length homedir $200.;
  homedir=sysget('HOME');

  /* A user is testing for write permissions in their $HOME directory */
  rc = modulen('*ie','access',homedir,2);
  put rc = ;
run;
```

In this case, the SAS log output would be the following:

Output 5.6 Log Output for Successful Completion (Access Permitted)

```
rc=0
```

Example 3: Using PEEKCLONG to Access a Returned Pointer

This example uses the **strcat** routine, which is part of the Red Hat Linux shared library **libc-2.2.3.so**. This library is typically installed in the **/lib/i686** directory. This routine concatenates two strings and returns a pointer to the newly concatenated string.

Here is the C prototype for this routine:

```
char *strcat(char, *dest, const char *src);
```

The proper SASCBTBL attribute table would be the following:

```
routine strcat minarg=2 maxarg=2 returns=ulong module=libc;
arg 1 char input format=$cstr200.;
arg 2 char input format=$cstr200.;
```

The following example shows the SAS code:

```
filenamesascbtbl './sascbtbl.txt';

data _null_;
  file sascbtbl;
  put 'routine strcat minarg=2 maxarg=2 returns=ulong module=libc;';
  put 'arg 1 char input format=$cstr200.';
  put 'arg 2 char input format=$cstr200.';
run;

data _null_;
  length string1 string2 newstring $200;
```

```

length chptr $20;
string1='This is string one and';
string2=' this is string two.';
chptr=modulec('strcat', string1, string2);
newstring=peekclong(chptr,200);
put newstring=;
run;

```

SAS writes the following output to the log:

Output 5.7 Log Output for Using PEEKCLONG to Access a Returned Pointer

```

newstring=This is string one and this is string two.

```

The PEEKCLONG function was used here because the Red Hat Linux shared library `/lib/i686/libc-2.2.3.so` is a 32-bit library. The following output demonstrates this:

```

$pwd
/lib/i686

$file ./libc-2.2.3.so
libc-2.2.3.so: ELF 32-bit LSB shared object, Intel 80386, version 1, not stripped

```

For more information about the PEEKLONG functions, see “PEEKLONG Function” on page 275.

Example 4: Using Structures

“Grouping SAS Variables as Structure Arguments” on page 117 describes how to use the FDSTART attribute to pass several arguments as one structure argument to a shared library routine. The passing of several arguments as one structure is another example of using structures with another routine in an external shared library.

The `statvfs` routine that is available under most UNIX operating systems retrieves file system information. This example uses the `statvfs` routine that is in the Solaris `libc.so.1` shared library and typically installed in the `/usr/lib/sparcv9` directory.

Here is the C prototype for this routine:

```
int statvfs(const char *path, struct statvfs *buf);
```

The `statvfs` routine will return a 0 if the routine completes successfully and `-1` if there is a failure.

The `statvfs` structure is defined with the following members:

```

unsigned long f_ysize;           /* preferred file system block size */
unsigned long f_frsize;         /* fundamental file system block */
unsigned long f_blocks;        /* total number of blocks on file system in units */
unsigned long f_bfree;         /* total number of free blocks */
unsigned long f_bavail;        /* number of free blocks available to non-superuser */
unsigned long f_files;         /* total number of file nodes (inodes) */
unsigned long f_ffree;         /* total number of free file nodes */
unsigned long f_favail;        /* number of inodes available to non-superuser */
unsigned long f_fsid;          /* file system id (dev for now) */
char          f_basetype[16];  /* target fs type name, null-terminated */
unsigned long f_flag;          /* bit mask of flags */

```

```

unsigned long g f_namemax;      /* maximum filename length */
char          f_fstr[32];      /* file system specific string */

```

The SASCBTBL attribute table would be the following:

```

routine statvfs
  minarg=14
  maxarg=14
  returns=short
  module=libc;
arg 1 char input  byaddr      format=$char256.;
arg 2 num  output byaddr fdstart format=piB8.;
arg 3 num  output                    format=piB8.;
arg 4 num  output                    format=piB8.;
arg 5 num  output                    format=piB8.;
arg 6 num  output                    format=piB8.;
arg 7 num  output                    format=piB8.;
arg 8 num  output                    format=piB8.;
arg 9 num  output                    format=piB8.;
arg 10 num output                    format=piB8.;
arg 11 char output                   format=$cstr16.;
arg 12 num output                    format=piB8.;
arg 13 num output                    format=piB8.;
arg 14 char output                   format=$cstr32.;

```

The SAS source code to call the **statvfs** routine from within the DATA step would be the following:

```

x 'if [ ! -L ./libc ]; then ln -s /usr/lib/sparcv9/libc.so.1 ./libc ; fi' ;
x 'setenv LD_LIBRARY_PATH ./usr/lib/sparcv9:/usr/lib:/lib';

data _null_;
  length f_basetype $16. f_fstr $32.;
  retain f_bsize f_frsize f_blocks f_bfree f_bavail f_files f_ffree f_favail
         f_fsid f_flag f_namemax 0;
  retain f_basetype f_fstr ' ';
  rc=modulen ('statvfs' , '/tmp', f_bsize, f_frsize, f_blocks, f_bfree, f_bavail,
            f_files, f_ffree, f_favail, f_fsid, f_basetype, f_flag,
            f_namemax, f_fstr);

  put rc = ;
  put f_bsize = ;
  put f_frsize = ;
  put f_blocks = ;
  put f_bfree = ;
  put f_bavail = ;
  put f_files = ;
  put f_ffree = ;
  put f_favail = ;
  put f_fsid = ;
  put f_basetype = ;
  put f_flag = ;
  put f_namemax = ;
  /* Determining the total bytes available in the file system and then dividing the
     total number of bytes by the number of bytes in a gigabyte */
  gigsfree = ((f_bavail * f_bsize)/1073741824);
  put 'The total amount of space available in /tmp is 'gigsfree 4.2' Gigabytes.';

```



```
run;
```

The SAS log output would be the following:

Output 5.8 Log Output for Using Structures

```
rc=0
f_bsize=8192
f_frsize=8192
f_blocks=196608
f_bfree=173020
f_bavail=173020
f_files=884732
f_ffree=877184
f_favail=877184
f_fsid=2
f_basetype=tmpfs
f_flag=4
f_namemax=255

The total amount of space available in /tmp is 1.32 Gigabytes.
```

Example 5: Invoking a Shared Library Routine from PROC IML

This example shows how to pass a matrix as an argument within PROC IML. The example creates a 4x5 matrix. Each cell is set to $10x+y+3$, where x is the row number and y is the column number. For example, the cell at row 1 column 2 is set to $(10*1)+2+3$, or 15.

The example invokes several routines from the theoretical TRYMOD shared library. It uses the **changd** routine to add $100x+10y$ to each element, where x is the C row number (0 through 3) and y is the C column number (0 through 4). The first argument to **changd** specifies the extra amount to sum. The **changdx** routine works just like **changd**, except that it expects a transposed matrix. The **changi** routine works like **changd** except that it expects a matrix of integers. The **changix** routine works like **changdx** except that integers are expected.

Note: A maximum of three arguments can be sent when invoking a shared library routine from PROC IML. Δ

In this example, all four matrices x_1 , x_2 , y_1 , and y_2 should become set to the same values after their respective MODULEIN calls. Here are the attribute table entries:

```
routine changd module=trymod returns=long;
arg 1 input num format=rb8. byvalue;
arg 2 update num format=rb8.;
routine changdx module=trymod returns=long
transpose=yes;
arg 1 input num format=rb8. byvalue;
arg 2 update num format=rb8.;
routine changi module=trymod returns=long;
arg 1 input num format=ib4. byvalue;
arg 2 update num format=ib4.;
routine changix module=trymod returns=long
transpose=yes;
arg 1 input num format=ib4. byvalue;
arg 2 update num format=ib4.;
```

Here is the PROC IML step:

```


proc iml;
  x1 = J(4,5,0);
  do i=1 to 4;
    do j=1 to 5;
      x1[i,j] = i*10+j+3;
    end;
  end;
  y1= x1; x2 = x1; y2 = y1;
  rc = modulein('changd',6,x1);
  rc = modulein('changdx',6,x2);
  rc = modulein('changi',6,y1);
  rc = modulein('changix',6,y2);
  print x1 x2 y1 y2;
run;

```

The following are the results of the PRINT statement:

Output 5.9 Invoking a Shared Library Routine from PROC IML

X1				
20	31	42	53	64
130	141	152	163	174
240	251	262	273	284
350	361	372	383	394
X2				
20	31	42	53	64
130	141	152	163	174
240	251	262	273	284
350	361	372	383	394
Y1				
20	31	42	53	64
130	141	152	163	174
240	251	262	273	284
350	361	372	383	394
Y2				
20	31	42	53	64
130	141	152	163	174
240	251	262	273	284
350	361	372	383	394



CHAPTER

6

Viewing Output and Help in the SAS Remote Browser

<i>What Is Remote Browsing?</i>	133
<i>Using Remote Browsing with ODS Output</i>	134
<i>Installing the Remote Browser Server</i>	134
<i>System Options for Remote Browsing</i>	134
<i>HELPHOST System Option</i>	135
<i>Setting Up the SAS Remote Browser</i>	135
<i>Setting Up the SAS Remote Browser at SAS Invocation</i>	135
<i>Setting Up the SAS Remote Browser During a SAS Session</i>	135
<i>Remote Browsing and Firewalls</i>	136
<i>For General Users</i>	136
<i>For System Administrators</i>	136

What Is Remote Browsing?

Remote browsing enables you to view SAS documentation, URLs that are specified in the WBROWSE command, and ODS output in the Web browser on your local computer. In the past, all Web documentation was displayed by executing a Netscape browser on the SAS server. By displaying this documentation locally, you have faster access to the documentation and you free up resources on the SAS server that were used by Netscape.

A small software agent called the remote browser server runs on your local computer. When SAS needs to display HTML content, it connects to the remote browser server and sends the URL that references the content. The remote browser server then passes the URL to a browser for display. If the remote browser server is not running on your computer, SAS displays a dialog box that contains the URL that you need to use to download the remote browser server.

Two system options are provided to configure remote browsing: HELPHOST and HELPPORT. These options specify the host name and port number of the computer where HTML content is to be displayed. In most cases, these options do not need to be set. HELPHOST defaults to the host name that is specified in the X11 DISPLAY environment variable, and HELPPORT defaults to the standard port for the remote browser server.

Using Remote Browsing with ODS Output

The SAS Output Delivery System (ODS) can be used to generate graphical reports of your SAS data. Remote browsing enables you to view your output directly from a SAS session as the output is generated or on demand from the Results window.

Remote browsing displays ODS output in many formats. If your browser does not have the appropriate plug-in for output that is not HTML, the browser displays a dialog box rather than the output. This dialog box enables you to download your output to your computer and view it using a local program such as Excel for an XSL file.

The automatic display of ODS output (HTML, PDF, and RTF only) is turned off by default. You can turn on the automatic display of ODS output by issuing the AUTONAVIGATE command in the Results window or by selecting **View results as they are generated** from the **Results** tab of the **Preferences** dialog box.

Installing the Remote Browser Server

You can install the remote browser server directly from your SAS session. If SAS is unable to make a connection for remote browsing, SAS displays a dialog box that contains the URL that you need to download the installer. Use this URL to download and install the remote browser server. Do not exit SAS. To install the remote browser server, follow these steps:

- 1 Type the URL that appears in the dialog box into your browser and press ENTER, or use the **Copy URL** button in the dialog box to copy the URL, and then paste it into your browser.
- 2 After the download page is displayed, download the installer that is appropriate for your computer.
- 3 Run the installer.
 - In the Windows environment, the remote browser server is added to your start-up items, so that the server will start whenever you log in. An icon is displayed in your system tray to indicate that the remote browser server is running.
 - In the Linux environment, manually add the command **rbrowser** to the start-up script for your windowing environment. The remote browser server will initially run minimized.

System Options for Remote Browsing

After the remote browser server is running on your computer, you can run the remote browsing system by specifying the HELPHOST and HELPPORT system options.

- The HELPHOST system option specifies the name of your host computer. If you do not specify this option, then the host name specified in the X display name is used. For more information, see “HELPHOST System Option” on page 374.
- The HELPPORT system option specifies the port number for the remote browser server that is installed on your computer. Under UNIX, you can use the default value for this option. For more information, see the HELPPORT System Option in *SAS Language Reference: Dictionary*.

You can set these options in your configuration file, at SAS invocation, or during your SAS session in the OPTIONS statement or in the SAS System Options window.

HELPHOST System Option

Specifies the name of the host computer where the remote browsing system is to be displayed.

Category: Environment control: Help

See: "HELPHOST System Option" on page 374

Setting Up the SAS Remote Browser

Setting Up the SAS Remote Browser at SAS Invocation

The following syntax is specific for UNIX operating environments and shows how you might set up the SAS Remote Browser if your remote browser server is using network port 12000:

```
sas92 -helpport 12000
```

Because you did not specify the HELPHOST system option, SAS uses the host name that is specified in the X display name.

Operating Environment Information: For the syntax for OpenVMS environments, see the documentation for each Help system option. Δ

Setting Up the SAS Remote Browser During a SAS Session

The syntax in this example applies to both UNIX and OpenVMS environments.

You can set up the remote browsing system during a SAS session by using the OPTIONS statement or the SAS System Options window. The following example uses the OPTIONS statement to change the value of the HELPPORT system option:

```
options helpport=12000;
```

Because you did not specify the HELPHOST system option, its value remains unchanged.

Remote Browsing and Firewalls

For General Users

If your network has a firewall between desktop computers and the computer that is hosting SAS, Web browsers cannot display Web pages from your SAS session. Usually, this problem is indicated by a timeout or connection error from the Web browser. If you receive a timeout or connection error, contact your system administrator.

For System Administrators


To enable the display of Web pages when a firewall exists between desktop computers and the computer that is hosting SAS, a firewall rule must be added that allows a Web browser to connect to SAS. The firewall rule specifies a range of network ports for which SAS remote browsing connections are allowed. Contact the appropriate system administrator who can select and configure a range of network ports for remote browsing. The range depends on the number of simultaneous SAS users. A value of approximately three times the number of simultaneous SAS users should reserve a sufficient number of network ports.

After the firewall rule is added, SAS must be configured to listen for network connections in the network port range. Normally, SAS selects any free network port, but the `HTTPSERVERPORTMIN` and the `HTTPSERVERPORTMAX` system options limit the network ports that SAS can select. Add these system options to your SAS configuration file. Set `HTTPSERVERPORTMIN` to the lowest port in the network range. Set `HTTPSERVERPORTMAX` to the highest port in the network range. For example, if the system administrator defined a network port range of 8000 to 8200, the system options would be the following:

```
httpserverportmin=8000
httpserverportmax=8200
```

After these system options are set, desktop computers can display Web pages. If there is an insufficient number of network ports, or the system options are specified incorrectly, a message appears in the SAS log.


For more information about these system options, see `HTTPSERVERPORTMIN= System Option` and `HTTPSERVERPORTMAX= System Option` in *SAS Language Reference: Dictionary*.



SAS Windowing Environment

Chapter 7.....**Working in the SAS Windowing Environment** 139

Chapter 8.....**Customizing the SAS Windowing Environment** 163



CHAPTER

7

Working in the SAS Windowing Environment

<i>Definition of the SAS Windowing Environment</i>	140
<i>Description of SAS in the X Environment</i>	141
<i>Definition of X Window System</i>	141
<i>X Window Managers</i>	141
<i>SAS Window Session ID</i>	141
<i>Workspace and Gravity in a SAS Session</i>	141
<i>Window Types</i>	142
<i>Top-Level Windows</i>	142
<i>Interior Windows</i>	142
<i>The SAS Session Manager (motifxsassm) in UNIX</i>	143
<i>What Is the SAS Session Manager?</i>	143
<i>Features of the SAS Session Manager</i>	143
<i>Using the Host Editor from within Your SAS Session</i>	145
<i>Closing the SAS Session Manager</i>	145
<i>Disabling the SAS Session Manager</i>	145
<i>Displaying Function Key Definitions in UNIX Environments</i>	146
<i>Benefits of Assigning Function Key Definitions</i>	146
<i>How to Display Function Key Definitions</i>	146
<i>The SAS ToolBox in UNIX Environments</i>	147
<i>Introduction to the SAS ToolBox</i>	147
<i>Customizing the Default SAS ToolBox</i>	147
<i>Default Configuration for the Command Window and the Toolbar</i>	147
<i>Opening and Closing the Command Window and the Toolbar</i>	148
<i>Executing Commands</i>	149
<i>Opening Files in UNIX Environments</i>	150
<i>Open the Open Dialog Box</i>	150
<i>Description of the Open Dialog Box Options</i>	151
<i>Specifying the Initial Filter and Directory Using SAS Resources</i>	151
<i>Using Regular Expressions in Filenames</i>	151
<i>Changing Your Working Directory in UNIX Environments</i>	152
<i>What Is Your Working Directory?</i>	152
<i>Changing Your Working Directory</i>	152
<i>The Change Working Directory Dialog Box</i>	152
<i>Selecting (Marking) Text in UNIX Environments</i>	153
<i>Difference between Marking Character Strings and Blocks</i>	153
<i>Techniques for Selecting Text</i>	154
<i>Select Text with the Mouse</i>	154
<i>Select Text with the MARK Command</i>	155
<i>Select Text by Using the Edit Menu</i>	155
<i>Copying or Cutting and Pasting Selected Text in UNIX Environments</i>	155
<i>Techniques for Copying or Cutting and Pasting Selected Text</i>	155

<i>How SAS Uses the Automatic Paste Buffer</i>	155
<i>Disabling the Automatic Paste Buffer</i>	156
<i>Copying and Pasting Text between SAS and Other X Clients</i>	156
<i>Using Drag and Drop in UNIX Environments</i>	156
<i>Difference between Default and Non-Default Drag and Drop</i>	156
<i>Limitations of Drag and Drop in UNIX</i>	156
<i>How to Drag and Drop Text</i>	156
<i>Searching For and Replacing Text Strings in UNIX Environments</i>	157
<i>What Are the Find and Replace Dialog Boxes?</i>	157
<i>Open the Find Dialog Box</i>	157
<i>Description of the Find Dialog Box Options</i>	157
<i>Open the Replace Dialog Box</i>	157
<i>Description of the Replace Dialog Box Options</i>	157
<i>Sending Mail from within Your SAS Session in UNIX Environments</i>	158
<i>Default E-mail Protocol in SAS</i>	158
<i>What Is the Send Mail Dialog Box?</i>	158
<i>Send E-mail by Using the Send Mail Dialog Box</i>	159
<i>Sending the Contents of a Text Window</i>	159
<i>Sending the Contents of a Non-Text Window</i>	160
<i>Change the Default File Type</i>	160
<i>Configuring SAS for Host Editor Support in UNIX Environments</i>	160
<i>Requirements for Using a Host Editor</i>	160
<i>Invoking and Using Your Host Editor</i>	160
<i>How to Open and Use the Host Editor</i>	160
<i>Example 1: Invoking SAS to Use xedit with the HOSTEDIT Command</i>	161
<i>Example 2: Invoking SAS to Use vi</i>	161
<i>Troubleshooting the Transfer of Text Attributes</i>	161
<i>Getting Help in UNIX Environments</i>	161

Definition of the SAS Windowing Environment

The SAS windowing environment refers to the windows that open when you invoke SAS. These windows include: the Program Editor, Log, Output, Explorer, and Results. These windows appear when you start SAS from your X workstation or through an X emulator. For more information about these windows, see the online SAS Help and Documentation.

The SAS windowing environment supports the use of X-based graphical user interfaces (GUIs). In UNIX environments, SAS provides an X Window System interface that is based on the Motif style.

Many features of the SAS windowing environment are controlled by X resources. For example, colors, window sizes, the appearance of the SAS ToolBox, and key definitions are all controlled through X resources. Chapter 8, “Customizing the SAS Windowing Environment,” on page 163 provides general information about resources, such as how to specify resources, and describes all of the resources that you can use to customize the interface.

Description of SAS in the X Environment

Definition of X Window System

The X Window System is a networked windowing system. If several computers are on a network, you can run an X server that, in turn, serves X applications (as clients) from all the other computers in the network.

X Window Managers

In UNIX environments, SAS features an X Window System interface that is based on Motif. This interface uses the window manager on your system to manage the windows on your display. Any window manager that is compliant with the *Inter-Client Communication Conventions Manual* (ICCCM) can be used with the Motif interface to SAS. Vendors provide at least one window manager with the X Window System environment. A common window manager is the Common Desktop Environment (CDE). If you are using another window manager, such as GNOME, you should read the documentation that is supplied by the vendor for that window manager.

All window managers perform the same basic functions, but they differ in their style and in their advanced functions. The appearance and function of the interface to SAS depends to some extent on your X window manager. Most window managers provide some kind of frame around a window. The window manager also governs the placement, sizing, stacking, and appearance of windows, as well as their interaction with the keyboard. The basics of interacting with SAS are the same for all window managers: opening menus, moving windows, responding to dialog boxes, dragging text, and so on.

SAS Window Session ID

When you run SAS on an X workstation, SAS shares the display with other X applications, including other SAS sessions. To enable you to distinguish between different applications and SAS sessions, SAS generates a SAS window session ID for each session by appending a number to the application name, which, by default, is **SAS**. This session ID appears in the window title bar for each SAS window and in the window icon title. The SAS sessions are assigned sequentially. Your first SAS session is not assigned a number, so the session ID is **SAS**; your second SAS session is assigned the session ID **SAS2**, and so on. Although the default application name is **SAS**, you can use the **-name** X option or the **-title** X option to change the instance name. The instance name can be up to 64 characters long and is displayed in the case in which it was entered, which can be lowercase, mixed case, or uppercase.

Workspace and Gravity in a SAS Session

When you use SAS on an X workstation, the display might be shared by many concurrent applications. When SAS windows from different sessions and windows from other applications appear on the display, the display can become cluttered. To help alleviate this problem, the windows for a SAS session first appear within an *application workspace* (AWS). The AWS defines a rectangular region that represents a virtual display in which SAS windows are initially created. SAS attempts to position the AWS in relation to the upper-left corner of your display. In other words, the workspace

gravitates toward a certain direction (*session gravity*) on the display. Some window manager configurations might override the placement that SAS has chosen for a window.

If you issue windowing commands or execute SAS procedures that create new SAS windows, the same rules of initial position and size apply to these windows: they are initially placed in the SAS AWS. You can use the WSAVE command to save the current window positions (or geometry). See “Customizing Session Workspace, Session Gravity, and Window Sizes in UNIX Environments” on page 204 for information.

Window Types

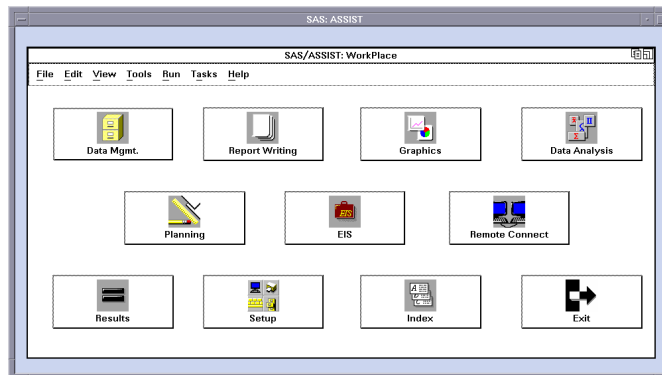
Top-Level Windows

SAS uses primary and interior windows. Some SAS applications consist of one or more primary windows controlled by the X window manager, in addition to the interior windows controlled by SAS. The SAS windowing environment primary windows, as well as most SAS application windows, initially appear as *top-level windows*. Top-level windows interact directly with the X window manager. They have a full title bar along with other window manager decorations. You can manipulate them individually after they appear on the display.

Interior Windows

Interior windows behave differently from primary windows. SAS/ASSIST software is an application with interior windows. Interior windows are contained within *container windows*, which might not be primary windows. The following display shows an interior window in SAS/ASSIST software.

Display 7.1 Sample Interior Window



SAS provides some degree of window management for interior windows. Specifically, interior windows have the following sizing and movement capabilities:

- You can move interior windows by clicking the interior window title bar and dragging the window to the desired location. If the destination of the interior window is outside the bounds of the container window, the container window changes according to the value of the **SAS.awsResizePolicy** resource. (The space within the container window is the application workspace, which is described in “Workspace and Gravity in a SAS Session” on page 141.) See “Overview of X Resources” on page 165 for more information.

- Interior windows cannot be minimized individually. Clicking on the container window icon button minimizes the container window and its interior windows.
- A **push-to-back** button (the small overlapping squares in the upper-right corner) is available with interior windows. However, you cannot push an active window behind an inactive window.

The SAS Session Manager (motifxsassm) in UNIX

What Is the SAS Session Manager?

The SAS Session Manager for X (**motifxsassm**) is an X client that is run by SAS when you use the SAS windowing environment. The Session Manager is automatically minimized when you start SAS. The SAS: Session Management dialog box for the SAS Session Manager appears as shown in the following display:

Display 7.2 SAS: Session Management Dialog Box



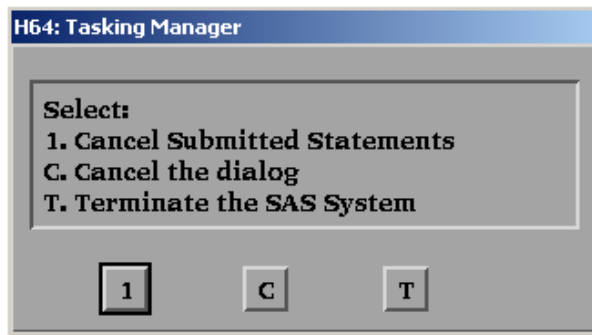
The SAS: Session Management dialog box lists the following information:

- which SAS session it controls
- the host computer from which the SAS session was invoked
- the UNIX process identifier of the SAS session

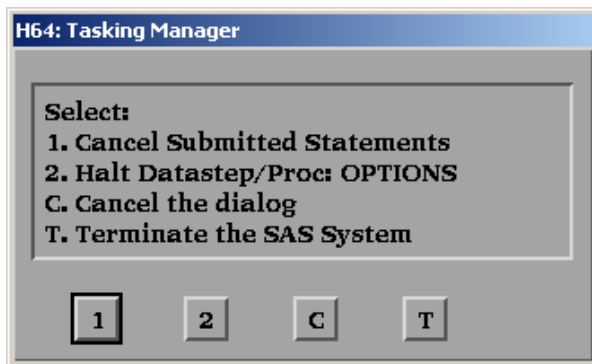
Features of the SAS Session Manager

The buttons in the SAS Session Manager enable you to do the following tasks:

Minimize	maps and minimizes all windows of the SAS session. This function is performed with standard X library calls and will work with most X window managers.
Restore	restores all of the windows that are open in the SAS session that is controlled by that SAS Session Manager. This function is performed with standard X library calls and will work with most X window managers.
Interrupt	sends a UNIX signal to SAS. When SAS receives the signal, it displays the following dialog box if no PROC or DATA steps are executing:



If a PROC or DATA step is executing, SAS displays the following dialog box:




You make your selection by clicking one of the buttons:

- | | |
|---|---|
| 1 | Causes the current PROC or DATA step statements to be deleted. |
| 2 | Causes the current PROC or DATA step to receive a request to interrupt processing. You are prompted to confirm this action. |
| C | Closes the dialog box without affecting SAS processing. |
| T | Forces SAS to terminate the SAS session. You are prompted to confirm this action. |

Terminate

displays a dialog box that asks you to confirm whether you want to terminate the SAS session:



If you click **OK**, the SAS Session Manager sends a UNIX signal to the SAS session that forces the session to terminate.

CAUTION:

Terminating your SAS session might result in data loss or data corruption. Before terminating your session, you should attempt to end SAS using one of the methods described in “Methods for Exiting SAS” on page 24. △

Help provides Help for the SAS: Session Management dialog box.

Using the Host Editor from within Your SAS Session

When you issue the HOSTEDIT command, SAS passes the request to the SAS Session Manager, which then invokes your host editor, so the SAS Session Manager must be running for the HOSTEDIT command to take effect. When you issue the HOSTEDIT command, SAS creates a temporary file that contains the data from the active SAS window and passes this file to your host editor. (These temporary files are stored in the directory specified by the SAS WORK option.) When you save your file in the host editor, the file is copied back into the SAS window if the window is writable, and the temporary files are deleted when the SAS session ends. See “Configuring SAS for Host Editor Support in UNIX Environments” on page 160 for more information.

Closing the SAS Session Manager

If you close the SAS: Session Management dialog box, you cannot retrieve the SAS Session Manager. To display the SAS Session Manager again, you can reinvoked **!SASROOT/utilities/bin/motifxsassm** with the *-pid* or the *-sessionid* arguments. Execute these commands at the UNIX prompt or use them with the X statement:

```
!SASROOT/utilities/bin/motifxsassm -pid pid
```

```
!SASROOT/utilities/bin/motifxsassm -sessionid integer
```

Disabling the SAS Session Manager

You can disable the SAS Session Manager by performing one of the following steps:

- Select **Tools ► Options ► Preferences**.

On the **General** tab, deselect the **Start Session manager** check box.

- Specify the following X resource, in lowercase, on the SAS command line at invocation:

```
sas -xrm 'SAS.startSessionManager: False'
```

Specifying the **sas.startSessionManager** X resource will deselect the **Start Session manager** check box in the Preferences dialog box.

Note: SAS saves the settings in the Preferences dialog box when it exits. If you have disabled the SAS Session Manager during your session, then the next time you invoke SAS, the SAS Session Manager will not run. To start the SAS Session Manager, select the **Start Session manager** check box in the Preferences dialog box or specify the following command, in lowercase, on the SAS command line at invocation:

```
sas -xrm 'SAS.startSessionManager: True'
```

△

Displaying Function Key Definitions in UNIX Environments

Benefits of Assigning Function Key Definitions

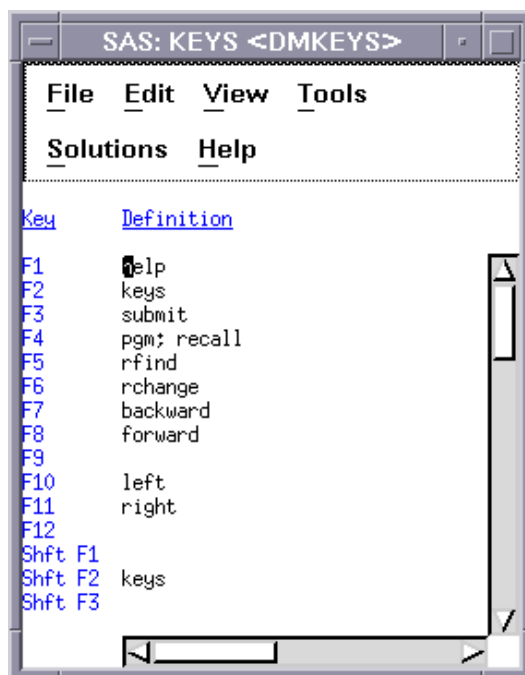
Function keys provide quick access to commands. They enable you to issue commands, insert text strings, and insert commands in programs. Function key definitions can be different on different terminals. These definitions are fully customizable.

How to Display Function Key Definitions

You can open the KEYS (DMKEYS) window to display all of your function key definitions in one of the following ways:

- Press **F2**.
- Issue the KEYS command.
- Select **Tools** \blacktriangleright **Options** \blacktriangleright **Keys**.

Display 7.3 KEYS (DMKEYS) Window



To view a single key definition without bringing up the KEYS window, use the KEYDEF command and specify the key definition that you want to view. For example, the following command displays the definition for key F4:

```
keydef f4
```

For information about customizing key definitions, see “Customizing Key Definitions in UNIX Environments” on page 184. See *Base SAS Help* for more information about the Keys window and the KEYDEF command.

The SAS ToolBox in UNIX Environments

Introduction to the SAS ToolBox

The SAS ToolBox has two parts as illustrated in the following display:

Display 7.4 The SAS ToolBox



- A *command window* that enables you to quickly enter any command in the active SAS window. For information about commands that are available under UNIX, see Chapter 10, “Commands under UNIX,” on page 219 and the SAS commands section in the Base SAS section in the online SAS Help and Documentation.
- A *toolbar* that contains several tool icons. When you select a tool icon, SAS immediately executes the command that is associated with that icon. The toolbar and the tool icons are completely customizable. For more information, see “Using the Tool Editor” on page 179.

The name of the active window is displayed in the title bar of the SAS ToolBox. For example, if the Log window were active, the title bar would say SAS ToolBox: Log instead of SAS ToolBox: Program Editor.

Under UNIX, the default SAS ToolBox automatically appears at the bottom of the SAS windows stack by default. To control its configuration, you use the Preferences dialog box. (See “Modifying the SAS ToolBox Settings” on page 170.)

Customizing the Default SAS ToolBox

The default SAS ToolBox is automatically copied to your SASUSER.PROFILE.DMS.TOOLBOX regardless of whether you customize the ToolBox. If you invoke an application that does not have an associated PMENU entry, the default toolbox is displayed for that application. If you then customize the toolbox for that application, the customized toolbox is stored in SASUSER.PROFILE.DEFAULT.TOOLBOX, where DEFAULT is the same entry name as the PMENU entry for the window or application.

You can customize the default SAS ToolBox, create multiple toolboxes and switch between them, and create application-specific toolboxes (such as with SAS/AF applications) that are automatically loaded when the application is loaded. Only one toolbox is displayed at a time, and the tools in the toolbox change as you move between applications. For more information about customizing your toolboxes, see “Customizing Toolboxes and Toolsets in UNIX Environments” on page 177.

Default Configuration for the Command Window and the Toolbar

By default, the toolbar and the command window are joined and are automatically displayed when SAS initializes unless one of the following conditions applies:

- You executed your SAS job in a non-windowing environment mode.

- The `SAS.defaultToolBox` or `SAS.defaultCommandWindow` resource is set to **False**. The default value is **True**. For more information about the resources that control the toolbox, see “X Resources That Control Toolbox Behavior” on page 178.
- You deselect **Display tools window**, **Display command window**, or **Combine windows** from the **ToolBox** tab in the Preferences dialog box.

The following display shows the command window and the toolbar in their default configuration.

Display 7.5 Default Configuration for Command Window and Toolbar



Opening and Closing the Command Window and the Toolbar

The following table lists the steps that you can use to open and close the command window and toolbar.

Table 7.1 Steps for Opening and Closing the Command Window and the Toolbar

Window	How to Open	How to Close
Command Window and Toolbar	<p>To open both windows, complete any of the following steps:</p> <ul style="list-style-type: none"> □ Issue the <code>COMMAND WINDOW</code> command. □ Issue the <code>TOOLLOAD</code> command. See “<code>TOOLLOAD</code> Command” on page 235 for more information. □ Select Tools ▶ Options ▶ Toolbox 	<p>To close these windows, complete any of the following steps:</p> <ul style="list-style-type: none"> □ Select Close from the ToolBox window menu. □ Enter the <code>TOOLCLOSE</code> command as described in “<code>TOOLCLOSE</code> Command” on page 234. □ Select Tools ▶ Options ▶ Toolbox so that ToolBox is deselected.

Window	How to Open	How to Close
Command Window	<p>To open only the command window, complete the following steps:</p> <ol style="list-style-type: none"> 1 Deselect Combine windows on the ToolBox tab of the Preferences dialog box. 2 Complete any of the following steps: <ul style="list-style-type: none"> □ Select Display command window in the ToolBox tab of the Preferences dialog box. □ Issue the COMMAND WINDOW command. 	<p>To close only the command window, complete the following steps:</p> <ol style="list-style-type: none"> 1 Deselect Display command window on the ToolBox tab of the Preferences dialog box. 2 Select Close from the window menu.
Toolbar	<p>To open only the toolbar, complete the following steps:</p> <ol style="list-style-type: none"> 1 Deselect Combine windows on the ToolBox tab of the Preferences dialog box. 2 Complete any of the following steps: <ul style="list-style-type: none"> □ Select Display tools window on the ToolBox tab of the Preferences dialog box. □ Issue the TOOLLOAD command. See “TOOLLOAD Command” on page 235 for more information. □ Select Tools ► Options ► Toolbox. 	<p>To close only the toolbar, complete the following steps:</p> <ol style="list-style-type: none"> 1 Deselect Combine windows on the ToolBox tab of the Preferences dialog box. 2 Complete any of the following steps: <ul style="list-style-type: none"> □ Deselect Display tools window on the ToolBox tab of the Preferences dialog box. □ Issue the TOOLCLOSE command as described in “TOOLCLOSE Command” on page 234. □ Select Tools ► Options ► Toolbox so that ToolBox is deselected.

Executing Commands

You can execute commands from either the command window or the toolbar. The following table gives more details about how to execute commands.

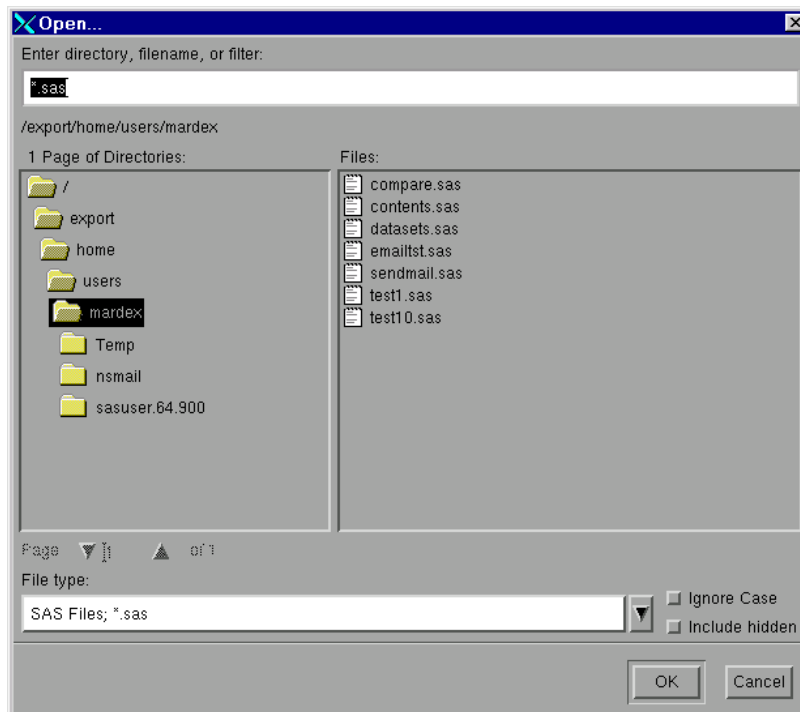
Table 7.2 Executing Commands in the Command Window and the Toolbar

Location	Execution
Command Window	<p>To execute a command, complete the following steps:</p> <ol style="list-style-type: none"> 1 Click in the command window. 2 Type in the command. 3 Press ENTER or click the check mark. <p>The command is executed in the active SAS window. You can use the up and down arrow keys to scroll through previously entered commands, or you can select a previous command from the drop-down list. Use the left mouse button to select a command from the drop-down list. Use the right mouse button to select and execute a command from the list.</p>
Toolbar	<p>To execute a command, click a tool icon in the toolbar to execute the command or commands that are associated with that icon. If you place the cursor over an icon for the amount of time specified by the SAS.toolBoxTipDelay resource, a pop-up window displays text that describes the command for that icon.</p>

Opening Files in UNIX Environments

Open the Open Dialog Box

The Open dialog box enables you to select files from the host file system. To open this dialog box, select **File** \blacktriangleright **Open**.

Display 7.6 Open Dialog Box

Description of the Open Dialog Box Options

The following table describes the options found in the Open dialog box.

Table 7.3 Options in the Open Dialog Box

Option	Description
Enter directory, filename, or filter	<p>is where you can type in the name of the directory, file, or file filter (file type) that you want to open.</p> <p>The directory shown in the Filter field is the currently selected directory. You can change this directory either by selecting a name from the Page of Directories list or by typing the new name directly into the field. The dialog box displays non-readable directories with a different icon.</p> <p>To display a list of all the files in a directory, enter the asterisk (*) wildcard in the Filter field or select All Files; * as the file type.</p>
Page of Directories	contains the names of the directories specified in the Filter and Page fields.
Files	contains the files in the selected directory that match the filter specified.
Page	enables you to change the directories that are listed in the Page of Directories list. A new page is defined when the number of entries in the Page of Directories list exceeds twice the screen height. To change pages, use the right or left arrows next to the Page field.
File type	enables you to select the type or types of files to be shown in the Files list. You can display a list of possible file filters by selecting the down arrow next to the field. Click on a file filter to select it.
Ignore Case	specifies that both uppercase and lowercase names be included in the display. (If you select All Files ; * as the filter, both uppercase and lowercase names are displayed whether you select Ignore Case .)
Include hidden	includes or excludes hidden files and directories from the graphical display.

Specifying the Initial Filter and Directory Using SAS Resources

You can specify the initial filter in the **File type** field by assigning a value to the **SAS.pattern** resource. However, the Open dialog box retains its filter between invocations, so the **SAS.pattern** resource applies only to the first invocation of the Open dialog box. You can also use the **SAS.directory** resource to specify the directory that you want when you first invoke the Open dialog box.

For more information about specifying SAS resources, see “Overview of X Resources” on page 165.

Using Regular Expressions in Filenames

Everything that you enter into the Open dialog box is treated as a regular expression. When you are opening or saving a file and you want to use a regular

expression special character as part of the filename, precede the character with a backslash (\). For example, to write to a file named \$Jan, enter \Jan as the filename.

For more information on regular expressions, refer to UNIX man page 5 for regexp:

```
man 5 regexp
```

Changing Your Working Directory in UNIX Environments

What Is Your Working Directory?

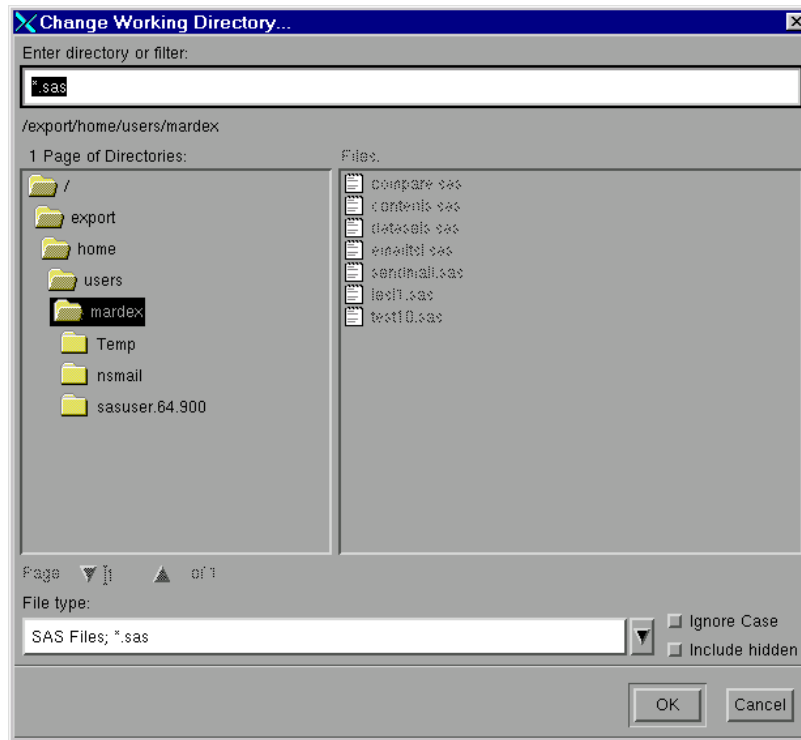
The working directory is the operating system directory to which many SAS commands and actions apply. By default, SAS uses the current directory as the working directory when you begin your SAS session.

Changing Your Working Directory

You can change the working directory during your SAS session. You can use the Change Working Directory dialog box to select a new directory, or you can use the X command, the X statement, the CALL SYSTEM routine, or the %SYSEXEC macro statement to issue the change directory (**cd**) command. For information on the X command and statement, the CALL SYSTEM routine, and the %SYSEXEC macro statement, see “Executing Operating System Commands from Your SAS Session” on page 14.

The Change Working Directory Dialog Box

To open the Change Working Directory dialog box, issue the DLGCDIR command or select **Tools ▶ Options ▶ Change Directory**.

Display 7.7 The Change Working Directory Dialog Box

The Change Working Directory dialog box works exactly the same as the Open dialog box, except that you cannot select a file from the list. For an explanation of the options in the Change Working Directory dialog box, see “Description of the Open Dialog Box Options” on page 151.

Selecting (Marking) Text in UNIX Environments

Difference between Marking Character Strings and Blocks

When you select text in a SAS window, you can select character strings or blocks. Character strings include the text in successive columns of one or more rows, as shown in the following display. Blocks are rectangular blocks that include the same columns from successive rows, as shown in Display 7.9 on page 154.

Display 7.8 Strings That Are Marked

```

SAS: Program Editor-compare.sas
File Edit View Tools Run Solutions Help

00001 libname students '/u/myid/students';
00002
00003 data students.one(label='First Data Set');
00004   input student year state $ grade1 grade2;
00005   label year='Year of Birth';
00006   format grade1 4.1;
00007   datalines;
00008 1000 1998 NC 85 87
00009 1042 1998 MD 92 92
00010 1095 1997 PA 78 72
00011 1187 1997 MA 87 94
00012 ;
00013 run;
00014
00015 data students.two(label='Second Data Set');
00016   input student $ year state $ grade1 grade2 major $;
00017   label state='Home State';
00018   format grade1 5.2;

```

Display 7.9 Blocks That Are Marked

```

SAS: Program Editor-compare.sas
File Edit View Tools Run Solutions Help

00001 libname students '/u/myid/students';
00002
00003 data students.one(label='First Data Set');
00004   input student year state $ grade1 grade2;
00005   label year='Year of Birth';
00006   format grade1 4.1;
00007   datalines;
00008 1000 1998 NC 85 87
00009 1042 1998 MD 92 92
00010 1095 1997 PA 78 72
00011 1187 1997 MA 87 94
00012 ;
00013 run;
00014
00015 data students.two(label='Second Data Set');
00016   input student $ year state $ grade1 grade2 major $;
00017   label state='Home State';
00018   format grade1 5.2;

```

Techniques for Selecting Text

Select Text with the Mouse

To select your text, complete the following steps:

- 1 Position the cursor at the beginning of the text that you want to mark.
- 2 Press and hold the left mouse button. If you want to select a block instead of a string, press and hold the CTRL key before you press the left mouse button.
- 3 Drag the mouse pointer over the text that you want to mark.
- 4 Press and hold down the ALT key (or EXTEND char key or META key, depending on your keyboard) while you release the mouse button. The marks that are generated by the mouse are called *drag marks*.

To extend an area of marked text, press and hold the SHIFT key, and use the left mouse button and the ALT key (and the CTRL key, if you are marking a block) to mark the new ending position. To unmark the selected text, press the mouse button anywhere in the window.

Select Text with the MARK Command

You can issue the MARK command from the command line, or you can assign it to a function key. With the MARK command, you can select more than one area of text in the same window at the same time. For more information about the MARK command, see the online SAS Help and Documentation.

To select your text, complete the following steps:

- 1 Position the cursor at the beginning of the text that you want to mark.
- 2 Issue the MARK command. If you want to select a block instead of a string, add the BLOCK argument to the MARK command.
- 3 Move the cursor to the end of the text that you want to mark.
- 4 Issue the MARK command a second time.

To unmark the selected text, issue the UNMARK command.

Select Text by Using the Edit Menu

To select your text using the **Edit** menu, complete the following steps:

- 1 Position the cursor at the beginning of the text that you want to mark.
- 2 Select **Edit ▶ Select**.
- 3 Position the cursor at the end of the text that you want to mark.
- 4 Press the left mouse button.

To unmark the selected text, select **Edit ▶ Deselect**.

Copying or Cutting and Pasting Selected Text in UNIX Environments

Techniques for Copying or Cutting and Pasting Selected Text

After you have marked text, you can copy or cut the text and paste it in another location.

- To copy text, select the Copy icon from the toolbox, issue the STORE or WCOPY command, or select **Edit ▶ Copy**.
- To cut text, select the Cut icon from the toolbox, issue the CUT or WCUT command, or select **Edit ▶ Cut**.
- To paste the cut or copied text, select the Paste icon from the toolbox, issue the PASTE or WPASTE command, or select **Edit ▶ Paste**.

For more information about the CUT, PASTE, and STORE commands, refer to online SAS Help and Documentation.

How SAS Uses the Automatic Paste Buffer

When you end a drag mark by releasing the mouse button without holding down the ALT key, SAS performs an end-of-mark action that might automatically generate a

STORE command to save the contents of the mark into a SAS paste buffer. If the STORE command is generated automatically, you do not have to explicitly copy the text before you paste it.

Disabling the Automatic Paste Buffer

You can disable the automatic paste buffer in the following ways:

- Set the **SAS.markPasteBuffer** resource.
- Deselect **Automatically store selection** on the **Editing** tab in the Preferences dialog box: **Tools** \blacktriangleright **Options** \blacktriangleright **Preferences**.

For more information, see “Customizing Cut and Paste in UNIX Environments” on page 202.

Copying and Pasting Text between SAS and Other X Clients

You can cut or copy and paste text between X clients if you associate the default SAS paste buffer with a paste buffer specific to X. For example, if you associate the default SAS paste buffer with the paste buffer, you can copy and paste text between xterm windows and SAS windows. To associate the SAS buffer with an X buffer, specify the **SAS.defaultPasteBuffer** resource:

```
SAS.defaultPasteBuffer: XTERM
```

For more information about using paste buffers, see “Customizing Cut and Paste in UNIX Environments” on page 202.

Using Drag and Drop in UNIX Environments

Difference between Default and Non-Default Drag and Drop

The SAS windowing environment on UNIX offers two types of drag and drop: default and non-default. Default drag and drop enables you to move text from one place to another. Non-default drag and drop enables you to choose whether to move or copy the text, submit the text if you are dragging SAS code, or cancel the drag and drop operation. With default drag and drop, you can drag text between SAS windows in different SAS sessions and between SAS windows and other Motif applications, such as Netscape, that support drag and drop. Non-default drag and drop is available only between windows in the same SAS session.

Limitations of Drag and Drop in UNIX

Under UNIX, you cannot drag and drop files or RTF (Rich Text Format) text.

How to Drag and Drop Text

To drag and drop text, first mark the text in one of the ways described in “Selecting (Marking) Text in UNIX Environments” on page 153. To use default drag and drop, use

the middle mouse button to drag the text where you want it. To use non-default drag and drop, press and hold the ALT (or EXTEND CHAR) key before you release the mouse button.

Searching For and Replacing Text Strings in UNIX Environments

What Are the Find and Replace Dialog Boxes?

The Find and Replace dialog boxes enable you to search for and replace strings in SAS text editor windows such as the Program Editor, the SCL editor, or Notepad.

Open the Find Dialog Box

To search for a string, open the Find dialog box by issuing the DLGFIND command or by selecting **Edit** \blacktriangleright **Find**.

Description of the Find Dialog Box Options


The Find dialog box works like the Replace dialog box, except it does not have the **Replace** field or the **Replace** and **Replace All** buttons.

For a description of the options in the Find dialog box, see “Description of the Replace Dialog Box Options” on page 157.

Open the Replace Dialog Box

To replace one text string with another, open the Replace dialog box by issuing the DLGREPLACE command or by selecting **Edit** \blacktriangleright **Replace**.

Display 7.10 Replace Dialog Box



Description of the Replace Dialog Box Options

To find a character string, type the string in the **Find** field, and click **Find**. To change a character string, type the string in the **Find** field, type its replacement in the

Replace field, and click **Replace**. To change every occurrence of the string to its replacement string, click **Replace All**.

You can tailor your find or replace operation using the following buttons:

Match Case

tells the search to match the uppercase and lowercase characters exactly as you entered them.

Match Word

searches for the specified string delimited by space, end-of-line, or end-of-file characters.

Previous

searches from the current cursor position toward the beginning of the file.

Next

searches from the current cursor position toward the end of the file.

Sending Mail from within Your SAS Session in UNIX Environments

Default E-mail Protocol in SAS

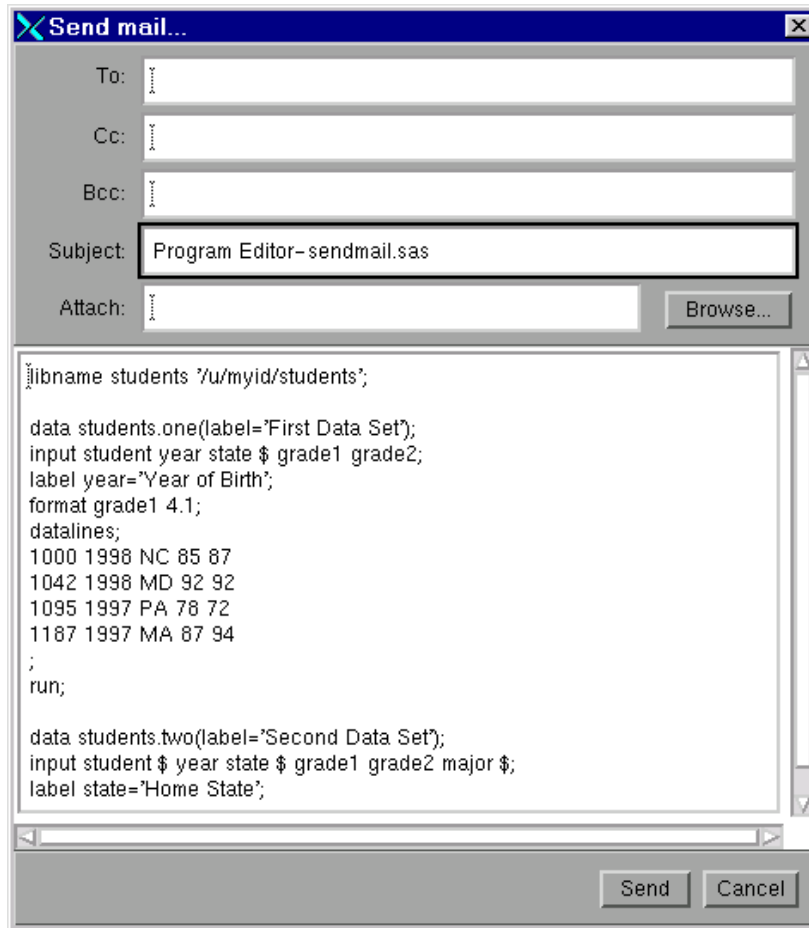
By default, SAS uses SMTP (Simple Mail Transfer Protocol) to send e-mail from within your SAS session. You can use the EMAILSYS system option to specify which script or protocol you want to use for sending electronic mail. See “EMAILSYS System Option” on page 366 for more information.

For more information about the SMTP e-mail interface, see *SAS Language Reference: Concepts*.

What Is the Send Mail Dialog Box?

The Send Mail dialog box enables you to send e-mail without leaving your current SAS session. To invoke the dialog box, issue the DLGSMail command or select **File ► Send Mail**.

Display 7.11 Send Mail Dialog Box



Send E-mail by Using the Send Mail Dialog Box

To send e-mail, complete the following steps as needed:

- 1 Enter the IDs of the e-mail recipients in the **To**, **Cc**, and **Bcc** fields. Separate multiple addresses with either spaces or commas.
- 2 Edit the entry in the **Subject** field as needed.
- 3 Enter the name of the file that you want to send in the **Attach** field. Separate multiple filenames with spaces. You can also use **Browse** to select a file.

Note: Some external scripts do not support sending e-mail attachments. △

- 4 Type your message in the message area or edit the contents grabbed from the active SAS text window.
- 5 Click **Send**.

To cancel a message, click **Cancel**.

Sending the Contents of a Text Window

You can e-mail the contents of an active SAS text window (such as the Program Editor or the Log) by using the Send mail dialog box. To open the Send mail dialog box,

select **File** \blacktriangleright **Send Mail**. SAS automatically copies the contents of the active SAS window and includes the text in the body of your e-mail. You can change or add to the e-mail message in the Send mail dialog box.

If you do not want to include the contents of the active SAS window in your message, select **Edit** \blacktriangleright **Clear All** before invoking the Send mail dialog box.

Sending the Contents of a Non-Text Window

To send the contents of a non-text window (such as a graph generated by SAS/GRAPH or an image from your PROC REPORT output), select **File** \blacktriangleright **Send Mail** from the active SAS window. SAS automatically copies the image data to a temporary file and enters that filename into the **Attach** field of the Send mail dialog box. To change the default file type for this temporary file, see “Change the Default File Type” on page 160.

SAS only copies the portion of the image that is visible in the active window, along with the window frame and title. This behavior is similar to using the DLGSCRDUMP command. For more information, see “DLGSCRDUMP Command” on page 226.

If you do not want to attach this image to your e-mail, clear the contents of the **Attach** field.

Note: Some external scripts do not support sending e-mail attachments. \triangle

Change the Default File Type

You can change the default file type for the temporary file that SAS creates by using the Preferences dialog box. To open the Preferences dialog box, do the following:

- 1 Select **Tools** \blacktriangleright **Options** \blacktriangleright **Preferences**.
- 2 On the **DMS** tab in the **Image type for Email attachments** box, select one of the following file types:
 - Portable Network Graphics (.png)
 - Graphics Interchange Format (.gif)
 - Tagged Image File Format (.tif)

Configuring SAS for Host Editor Support in UNIX Environments

Requirements for Using a Host Editor

SAS supports the use of a host text editor with the Motif interface, so you can use an editor such as vi or Emacs with your SAS session. There is no host editor set as the default host editor, so you must specify one to use this feature. Host editor support requires the use of the motifxsassm client. (See “The SAS Session Manager (motifxsassm) in UNIX” on page 143 for more information.)

Invoking and Using Your Host Editor

How to Open and Use the Host Editor

To use your host text editor with SAS, complete the following steps:

- 1 Specify the command required to invoke your editor with the EDITCMD system option.
- 2 Invoke the editor as needed with the HOSTEDIT command.

The HOSTEDIT command passes data from a SAS window to the host editor. When you save data in the host editor, the data is copied back into the SAS window if the window can be written to.

After you return to the SAS text editor window, you can issue the UNDO command to undo all of the changes that you made with your host editor. You must issue the UNDO command a second time to return to the state of the window before the HOSTEDIT command was issued. If you issue the HOSTEDIT command in a read-only window, you can save your editing changes to an external file, but the SAS text editor window remains unchanged.

See “EDITCMD System Option” on page 365 and “HOSTEDIT Command” on page 231 for more information.

Example 1: Invoking SAS to Use xedit with the HOSTEDIT Command

Some systems have an X-based editor installed that is called xedit. If you want to use XEDIT with the HOSTEDIT command, you can invoke SAS with the following command:

```
sas -editcmd '/usr/local/bin/xedit'
```

Example 2: Invoking SAS to Use vi

The vi editor is a terminal-based editor that requires a terminal window. The xterm client's `-e` option runs a program when the xterm client is invoked. To use the EDITCMD option to display an xterm client in conjunction with vi, invoke SAS as follows:

```
sas -editcmd '/usr/bin/X11/xterm -e /usr/bin/vi'
```

Troubleshooting the Transfer of Text Attributes

Text attributes, such as color and highlighting, are not transferred between a host editor window and a SAS text editor window. Issue the HEATTR ON command to display a dialog box that will warn you if you are editing text with highlighting and color attributes that will be removed by the host editor. This dialog box prompts you to continue or abort the HOSTEDIT command. Specify HEATTR OFF to suppress this dialog box.

Getting Help in UNIX Environments

The **Help** menu is always available within your SAS session. Here are descriptions of the Help topics that are available from the **Help** menu:

Using This Window

provides help information that is relevant to the active window. You can access the same information by clicking the **Help** button or pressing the F1 key.

SAS Help and Documentation

provides tutorials and sample programs to help you learn how to use SAS, comprehensive documentation for all products installed at your site, and information about contacting SAS for additional support.

Note: If you set the block unrequested pop-up windows option in your browser's Preferences dialog box, then the online SAS Help and Documentation might not display. Δ

Getting Started with SAS Software


opens a tutorial that will help you get started with SAS.

SAS on the Web

provides links to useful areas on the SAS Web site, including the customer support center, frequently asked questions, sending feedback to SAS, and the SAS home page.

About SAS 9

opens the About SAS 9 dialog box, which provides information about SAS software, your operating environment, and Motif.



CHAPTER

8

Customizing the SAS Windowing Environment

<i>Overview of Customizing SAS in X Environment</i>	164
<i>Overview of X Resources</i>	165
<i>Introduction to X Resources</i>	165
<i>Syntax for Specifying X Resources</i>	165
<i>Methods for Customizing X Resources</i>	165
<i>Modifying X Resources Through the Preferences Dialog Box</i>	167
<i>What Is the Preferences Dialog Box?</i>	167
<i>Open the Preferences Dialog Box</i>	167
<i>Description of the Options in the Preferences Dialog Box</i>	168
<i>Modifying the General Settings</i>	168
<i>Modifying the DMS Settings</i>	168
<i>Modifying the Editing Settings</i>	169
<i>Modifying the Results Settings</i>	170
<i>Modifying the SAS ToolBox Settings</i>	170
<i>Setting X Resources with the Resource Helper</i>	172
<i>Introduction to the Resource Helper</i>	172
<i>How to Start the Resource Helper</i>	172
<i>Start the Resource Helper from a SAS Session</i>	172
<i>Start the Resource Helper from a Shell Prompt</i>	172
<i>Defining Keys with the Resource Helper</i>	173
<i>How to Define a Key</i>	173
<i>Troubleshooting Incorrect Key Definitions</i>	174
<i>Modifying the Color of a SAS Window Using the Resource Helper</i>	174
<i>How to Use the Color Window</i>	174
<i>Example: Change the Color of a SAS Window</i>	175
<i>Return to the Default Settings</i>	176
<i>Permanently Save Your Color Settings</i>	176
<i>How the Resource Helper Searches for X Resources</i>	176
<i>Customizing Toolboxes and Toolsets in UNIX Environments</i>	177
<i>Techniques for Customizing Toolboxes</i>	177
<i>X Resources That Control Toolbox Behavior</i>	178
<i>Using the Tool Editor</i>	179
<i>What Is a Toolset?</i>	179
<i>Invoking the Tool Editor</i>	179
<i>After You Invoke the Tool Editor</i>	180
<i>Changing the Appearance of the Entire Toolbox</i>	180
<i>Changing the Attributes of an Existing Tool</i>	180
<i>Adding Tools to the Toolbox</i>	182
<i>Change the Order of the Tools in the Toolbox</i>	182
<i>Deleting Tools from the Toolbox</i>	182
<i>Returning to the Default Settings</i>	182

<i>Saving Changes to the Toolbox or Toolset</i>	182
<i>Creating a New Toolbox</i>	183
<i>Create or Customize an Application- or Window-Specific Toolbox</i>	183
<i>Create or Customize an Application- or Window-Specific Toolset</i>	183
<i>Customizing Key Definitions in UNIX Environments</i>	184
<i>Techniques for Customizing Your Key Definitions</i>	184
<i>Defining Key Translations</i>	185
<i>What Is a Key Translation?</i>	185
<i>What Is the SAS.keyboardTranslations Resource?</i>	185
<i>Steps for Creating a Key Definition</i>	185
<i>Determining Keysyms</i>	186
<i>Syntax of the SAS.keyboardTranslations Resource</i>	187
<i>Syntax of the SAS.keysWindowLabels Resource</i>	188
<i>SAS Keyboard Action Names</i>	188
<i>Examples: Defining Keys Using SAS Resources</i>	191
<i>Customizing Fonts in UNIX Environments</i>	192
<i>Difference between the System Font and Fonts That Are Used in the Windowing Environment</i>	192
<i>How SAS Determines Which Font to Use</i>	192
<i>Customizing Fonts by Using the Fonts Dialog Box</i>	193
<i>Introduction to the Fonts Dialog Box</i>	193
<i>How to Change the Default Font</i>	193
<i>Specifying Font Resources</i>	193
<i>Specifying Font Aliases</i>	195
<i>Example: Substitute the Lucida Font for Palatino</i>	195
<i>Customizing Colors in UNIX Environments</i>	196
<i>Methods for Customizing the Color Settings in Your SAS Session</i>	196
<i>Customizing Colors by Using the SASCOLOR Window</i>	196
<i>Syntax of the COLOR Command</i>	197
<i>Defining Color Resources</i>	197
<i>Types of Color Resources</i>	197
<i>Specifying RGB Values or Color Names for Foreground and Background Resources</i>	198
<i>Defining Colors and Attributes for Window Elements (CPARMS)</i>	199
<i>Controlling Contrast</i>	202
<i>Controlling Drop-Down Menus in UNIX Environments</i>	202
<i>Customizing Cut and Paste in UNIX Environments</i>	202
<i>Instructions for Cutting and Pasting Text</i>	202
<i>Types of Paste Buffers</i>	203
<i>Selecting a Paste Buffer</i>	203
<i>Manipulating Text Using a Paste Buffer</i>	203
<i>Notes about Preserving Text and Attribute Information</i>	204
<i>Customizing Session Workspace, Session Gravity, and Window Sizes in UNIX Environments</i>	204
<i>Specifying User-Defined Icons in UNIX Environments</i>	206
<i>Why Specify User-Defined Icons?</i>	206
<i>How SAS Locates a User-Defined Icon</i>	206
<i>X Resources for Specifying User-Defined Icons</i>	206
<i>Miscellaneous Resources in UNIX Environments</i>	207
<i>Summary of X Resources for SAS in UNIX Environments</i>	209

Overview of Customizing SAS in X Environment

The SAS windowing environment supports the use of X-based graphical user interfaces (GUIs). In UNIX environments, SAS provides an X Window System interface

that is based on the Motif style. For more information about SAS in the X environment, see “Description of SAS in the X Environment” on page 141.

You can customize your working environment by using X resources.

Overview of X Resources

Introduction to X Resources

X clients usually have characteristics that can be customized; these properties are known as *X resources*. Because SAS functions as an X Windows client, many aspects of the appearance and behavior of the SAS windowing environment are controlled by X resources. For example, X resources can be used to define a font, a background color, or a window size. The resources for an application, such as SAS, are placed in a *resource database*.

SAS functions correctly without any modifications to the resource database. However, you might want to change the default behavior or appearance of the interface. There are several ways to specify your customizations. Some methods modify all SAS sessions displayed on a particular X server. Some methods affect all SAS sessions run on a particular host. Other methods affect only a single SAS session.

If you need more information about X Window System clients and X resources, see the documentation provided by your vendor.

Syntax for Specifying X Resources

A resource specification has the following format:

resource-string: value

The *resource string* usually contains two identifiers and a separator. The first identifier is the client or application name (**SAS**), the separator is a period (.) or asterisk (*) character, and the second identifier is the name of the specific resource. The *value* given can be a Boolean value (**True** or **False**), a number, or a character string, depending on the resource type.

The application name and resource name can both specify an *instance value* or a *class value*. A specification for a class applies to a larger scope than a single instance.

The following are sample resource specifications:

```
SAS.startSessionManager: True
SAS.maxWindowHeight: 100
SAS.awsResizePolicy: grow
```

See your X Window System documentation for more information about resource specifications.

Methods for Customizing X Resources

The following list describes the methods that you can use to customize X resources.

- Use the Font dialog box, the Preferences dialog box, or the Resource Helper to customize your SAS session. All of these tools write X resource definitions out to a

location that SAS will read the next time you start a SAS session. See “Modifying X Resources Through the Preferences Dialog Box” on page 167, “Setting X Resources with the Resource Helper” on page 172, and “Customizing Fonts in UNIX Environments” on page 192 for more information on these tools.

Note: The settings that you specify in the Preferences dialog box will override any command line settings. Δ

- \square Specify session-specific resources by using the `-xrm` option on the command line for each invocation of SAS. For example, the following command specifies that SAS will not display the Confirm dialog box when you exit your SAS session:

```
sas -xrm 'SAS.confirmSASExit: False'
```

You can specify the `-xrm` option as many times as needed. You must specify the `-xrm` option for each resource.

Note: If you normally invoke SAS with a shell script, you should protect the quotation marks from the shell with the backslash (`\`) character:

```
sasscript -xrm \'SAS.confirmSASExit: False\'
```

Δ

- \square Add resource definitions to a file in your home directory. If you place resources in a file that X Toolkit normally searches for when applications are invoked, these resources will be loaded when you invoke SAS. For information about where the X Toolkit searches for resources, see the documentation for the X Window System.

You can also add resources to the resource database after SAS has initialized by running the `xrdb` utility. For example, the following command merges the definitions in the `MyResources` file into the resource database:

```
xrdb -merge myresources
```

- \square Create a subdirectory for storing resource definitions. (This subdirectory is usually named `app-defaults`.) Set the `XUSERFILESEARCHPATH` environment variable to the pathname of this subdirectory. You can use `%N` to substitute an application class name for a file when specifying the `XUSERFILESEARCHPATH` environment variable. Specify the definition for this environment variable in the initialization file for your shell (for example, the `$HOME/.login`, `$HOME/.cshrc`, or `$HOME/.profile` files), to ensure that the `XUSERFILESEARCHPATH` variable is defined for each shell that is started.

Create a file called `SAS` in the subdirectory identified by `XUSERFILESEARCHPATH`. Include your resource definitions in this file.

Note: Alternatively, you could set the `XAPPLRESDIR` environment variable to the pathname of the subdirectory that stores your resource definitions. The `XAPPLRESDIR` and `XUSERFILESEARCHPATH` environment variables use a slightly different syntax to specify the location of your resource definitions. The location specified by the `XUSERFILESEARCH` environment variable takes precedence over the location specified by the `XAPPLRESDIR` variable. For more information, see the UNIX X man page. Δ

- \square If you want the customized resource definitions to be used for all users on a particular host, create a file called `SAS` to contain your resource definitions, and store this file in the system `app-defaults` directory.

For more information about X resources, see the X Window System documentation supplied by your vendor or to other documentation about the X Window System.

Modifying X Resources Through the Preferences Dialog Box

What Is the Preferences Dialog Box?

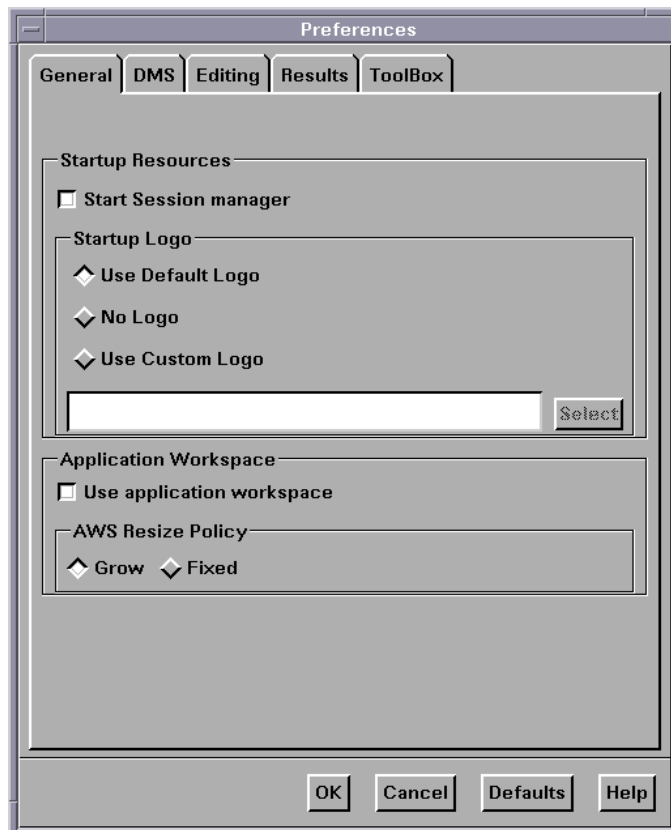
The Preferences dialog box enables you to control the settings of certain X resources. Changes made through the Preferences dialog box (with the exception of those resources on the **General** tab) become effective immediately, and the settings are saved in the SasuserPrefs file in your Sasuser directory.

Note: The settings that you specify in the Preferences dialog box will override any command line settings for the current session. △

Open the Preferences Dialog Box

You can invoke the Preferences dialog box by issuing the DLGPREF command or by selecting **Tools ► Options ► Preferences**.

Display 8.1 Preferences Dialog Box



Description of the Options in the Preferences Dialog Box

Modifying the General Settings

To modify the General settings, select the **General** tab in the Preferences dialog box, and then select from the items in the window:

Start Session manager

specifies whether you want the SAS Session Manager to be started automatically when you start your SAS session. If you want to use your host editor in your SAS session, the SAS Session Manager must be running. The SAS Session Manager enables you to interrupt or terminate your SAS session and minimize and restore all of the windows in a SAS session. See “The SAS Session Manager (motifxsassm) in UNIX” on page 143 and “Configuring SAS for Host Editor Support in UNIX Environments” on page 160 for more information. Clicking the **Start Session manager** box sets the **SAS.startSessionManager** resource.

Startup Logo

specifies whether you want SAS to display an XPM file while your SAS session is being initialized and, if so, which file.

If you select **Use Default Logo**, SAS uses the default file for your site. If you select **No Logo**, then no file is displayed. If you select **Use Custom Logo**, then you can either enter the XPM filename directly in the text field or click **Select** to open the File Selection dialog box. Selecting this box sets the **SAS.startupLogo** resource.

Use application workspace

confines all windows displayed by an application to a single application workspace. Selecting this box sets the **SAS.noAWS** resource. You must exit and reopen the windows for changes to this resource to take effect.

AWS Resize Policy

controls the policy for resizing AWS windows as interior windows are added and removed. (See “Workspace and Gravity in a SAS Session” on page 141 and “Window Types” on page 142.)

Grow

The AWS window attempts to grow any time an interior window is grown or moved (to make all of its interior windows visible), but it will not shrink to remove unused areas.

Fixed

The AWS window attempts to size itself to the size of the first interior window and will not attempt any further size changes. Selecting this box sets the **SAS.awsResizePolicy** resource.

Modifying the DMS Settings

To modify the DMS settings, select the **DMS** tab in the Preferences dialog box.

Use menu access keys

activates menu mnemonics. When mnemonics are turned on, you can select menu items by typing the single, underlined letter in the item. Selecting this box sets the **SAS.usePmenuMnemonics** resource.

Confirm exit

displays the Exit dialog box when you exit your SAS session. Selecting this box sets the **SAS.confirmSASExit** resource.

Save Settings on Exit

tells SAS to issue the WSAVE ALL command when you exit your SAS session. This command saves the global settings, such as window color and window position, that are in effect for all windows that are currently open. These settings are saved in your Sasuser.Profile catalog. Selecting this box sets the **SAS.wsaveAllExit** resource.

Note: For the WSAVE command to work, your window manager must support explicit window placement. See the documentation for your window manager to determine how to configure your window manager. For example, if you are running Exceed, open the Screen Definition Settings dialog box and deselect **Cascade Windows**. Δ

Backup Documents

enables you to specify whether you want SAS to automatically save (at the interval specified by the **SAS.autoSaveInterval** resource) the documents that you currently have open. Selecting this box sets the **SAS.autoSaveOn** resource.

Image type for Email attachments

specifies the default file type for the temporary file that SAS creates when sending the contents of a non-text window via e-mail. Examples of non-text windows include a graph generated by SAS/GRAPH or an image from your PROC REPORT output. For more information, see “Sending the Contents of a Non-Text Window” on page 160.

Modifying the Editing Settings

To modify the Editing settings, select the **Editing** tab in the Preferences dialog box.

Default paste buffer

defines an alias for the default SAS buffer. The following list describes the paste buffer alias names and the X buffer with which each name is associated.

XPRIMARY

X primary selection (**PRIMARY**)

XSCNDARY

X secondary selection (**SECONDARY**)

XCLIPBRD

X clipboard (**CLIPBOARD**)

XTERM

exchange protocol used by the xterm client

XCUTn

X cut buffer where n is between 0 and 7, inclusive

Selecting this box sets the **SAS.defaultPasteBuffer** resource. See “Controlling Drop-Down Menus in UNIX Environments” on page 202 for more information about cut-and-paste buffers.

Automatically store selection

generates a STORE command every time you mark a region of text with the mouse. Selecting this box sets the **SAS.markPasteBuffer** resource.

Cursor

controls the editing mode in SAS text editor windows. Selecting the **Insert** and **Overtyp**e boxes sets the **SAS.insertModeOn** resource to **True** and **False**, respectively.

Modifying the Results Settings

The items on the **Results** tab affect only output that is produced through the Output Delivery System (ODS). To modify these settings, select the **Results** tab in the Preferences dialog box.

For a complete description of ODS, see the *SAS Output Delivery System: User's Guide*.

Create Listing

opens the ODS Listing destination, which produces monospace output. Selecting this box is equivalent to entering the ODS LISTING SELECT ALL statement.

Create HTML

opens the ODS HTML destination, which produces output that is formatted in HyperText Markup Language (HTML). Selecting this box is equivalent to entering the ODS HTML SELECT ALL statement.

Folder

specifies a destination for HTML files. Specifying a directory in this field is equivalent to specifying a directory with the PATH option in the ODS HTML statement.

Use WORK Folder

tells the ODS to send all HTML files to your WORK directory. Selecting this box is equivalent to specifying the pathname of your WORK directory with the PATH option in the ODS HTML statement.

Style

specifies the style definition to use for HTML output. The style definition controls such aspects as color, font name, and font size. Specifying a style in this field is equivalent to specifying a style with the STYLE option in the ODS HTML statement. You can specify any style that is defined in the \ODS\REFERENCES\STYLES key in the SAS registry. You can open the SAS registry by issuing the REGEDIT command or by selecting **Solutions ► Accessories ► Registry Editor**.

View results as they are generated

tells SAS to automatically display results files when they are generated. If you select this box, make sure that **Password protect HTML file browsing** is deselected.

Password protect HTML file browsing

tells SAS to prompt you for your password before sending HTML files to your browser. If you select this box, make sure that **View results as they are generated** is deselected. Selecting this box sets the **SAS.htmlUsePassword** resource.

Modifying the SAS ToolBox Settings

The items on the **ToolBox** tab of the Preferences dialog box affect both the toolbox and the command window. To modify these settings, select the **ToolBox** tab in the Preferences dialog box.

Display tools window

determines whether to display the default toolbox. Selecting this check box sets the **SAS.defaultToolBox** resource.

Display command window

determines whether to display the command window. Selecting this check box sets the **SAS.defaultCommandWindow** resource.

Auto Complete Commands

specifies whether SAS automatically fills in the remaining letters of a command as you type a command in the command window that begins with the same letter as a command that you have entered previously. If both this box and **Save Commands** are selected, then SAS can automatically fill in commands that were entered in previous sessions. Selecting this check box sets the **SAS.autoComplete** resource.

Save Commands

specifies whether SAS saves the commands that you enter in the command window and how many commands are saved. You can specify a number from 0 to 50. If you specify 0, no commands will be saved. If you specify 1 or more, that number of commands is saved in the file **commands.hist** in your Sasuser directory. If this box is selected, then SAS will be able to automatically fill in (see **Auto Complete Commands**) commands that were entered in previous sessions. Selecting this field sets the **SAS.commandsSaved** resource.

Combine windows

combines the toolbox and command window into one window. The toolbox and command window are combined by default. Selecting this check box sets the **SAS.useCommandToolBoxCombo** resource.

Use arrow decorations

adds arrows to both ends of the combined toolbox/command window. Selecting this check box sets the **SAS.useShowHideDecorations** resource.

Always on top

keeps the toolbox or the combined toolbox/command window on top of the window stack. This check box is selected by default, which might cause problems with window managers and other applications that want to be on top of the window stack. If you have such a situation, turn off this feature. Selecting this check box sets the **SAS.toolboxAlwaysOnTop** resource.

Toolbox Persistent

specifies whether the toolbox that is associated with the Program Editor window stays open when you close the Program Editor. By default, the Program Editor toolbox stays open whenever you close the Program Editor window. If you deselect this box, then the toolbox will close if you close the Program Editor. Selecting this check box sets the **SAS.isToolBoxPersistent** resource.

The items in the Tools area affect the individual tools in the toolbox.

Use large tools

controls whether tool icons are displayed as 24x24 or 48x48 pixels. The default is 24x24. Selecting this check box sets the **SAS.useLargeToolBox** resource.

Use tip text

specifies whether the ToolTip text is displayed when you position your cursor over a tool in the toolbox. Some window managers might place the toolbox tip behind the toolbox. If the toolbox tip is placed behind the toolbox in your environment, deselect this box. Selecting this check box sets the **SAS.useToolBoxTips** resource.

delay

controls the delay in milliseconds before popping up the toolbox tip. Selecting this check box sets the `SAS.toolboxTipDelay` resource. You can enter a value directly into the field or use the arrows to the right of the field to change the value.

Setting X Resources with the Resource Helper

Introduction to the Resource Helper

With Resource Helper, you can customize the key definitions and the colors of the SAS interactive interface. Resource Helper creates SAS resource definitions and stores them in a location where the Resource Manager can find them. See “How the Resource Helper Searches for X Resources” on page 176 for a list of the locations that Resource Helper searches for resource definitions. Resource settings that are saved with Resource Helper will take effect the next time you start a SAS session.

You can start Resource Helper from within a SAS session or from your shell prompt.

How to Start the Resource Helper

Start the Resource Helper from a SAS Session

Start the SAS Resource Helper from a SAS window by entering the following command on the command line in the command window:

```
reshelper
```

Display 8.2 Main Window for Resource Helper



Start the Resource Helper from a Shell Prompt

Resource Helper is installed in the `/utilities/bin` subdirectory in the directory where SAS is installed (`!SASROOT`). The name of the executable module is `reshelper`. For example, if SAS is installed in `/usr/local/sas92`, you start Resource Helper by typing the following command:

```
/usr/local/sas92/utilities/bin/reshelper &
```

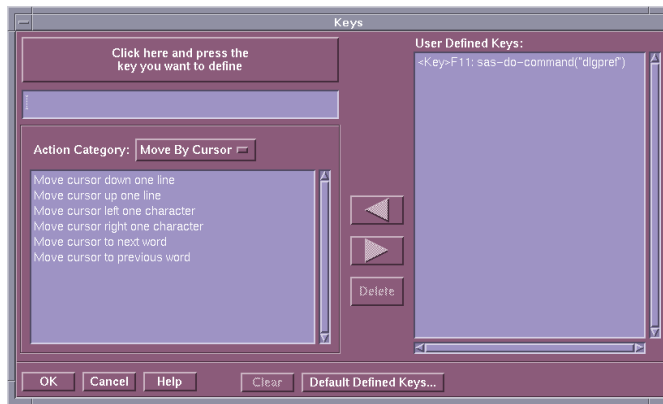
Defining Keys with the Resource Helper

How to Define a Key

To define a key, follow these steps:

- 1 Start the Resource Helper (see “How to Start the Resource Helper” on page 172) and select the Keys icon.

Display 8.3 Keys Window for Resource Helper



Key definitions are divided into several **Action Categories**:

- Move By Cursor**
- Move By Field**
- Edit**
- Miscellaneous**
- All Actions**

- 2 Select **Click here and press the keys you want to define**.
- 3 Press the key or combination of keys that you want to assign an action to. For example, press F12. If a default SAS translation has already been assigned to the key combination, Resource Helper displays the default translation.
- 4 Select the action category menu button to open a list of action categories. Select the action category that you want. For example, if you want to define a key to delete the current field, select **Edit** as your **Action Category**. Resource Helper will display a list of actions in that category.
- 5 Select an action from the list. For example, **Delete current field**. Resource Helper can assign only one action to a translation. If the action that you select requires an argument (such as **sas-action-routine**), Resource Helper prompts you for the argument.

Resource Helper displays the key combination and its new definition:

```
None<Key>F12: sas-delete()
```

Note: If you select the **sas-action-routinesas-function-key** action routine, then the key definition is automatically displayed in the Keys window. If you choose another action routine and if you want the definition to appear in the Keys

window, you will need to define a window label for the key. See “Syntax of the SAS.keysWindowLabels Resource” on page 188 for information about defining labels in the Keys window. Δ

- 6 Select the right arrow to add this key translation to the list of **User-Defined Keys**.
- 7 Click **OK** to exit the Keys window after you have finished defining key translations.
- 8 To save your translations permanently, from the Resource Helper drop-down menus, select **File** \blacktriangleright **Save Resources**.

To modify a key definition that is already in the **User-Defined Keys** list, select the definition, select the left arrow to remove the definition from the list, and edit the definition.

To delete a definition from **User-Defined Keys**, select it and click **Delete**.

Clear clears the key definition edit window.

Default Defined Keys displays the default key definitions for your system.

Troubleshooting Incorrect Key Definitions

In most cases, using Resource Helper is much easier and faster than defining the resources yourself. However, the X Window System searches for resources in several places, so it is possible for Resource Helper to pick up the wrong key symbol for the key you are trying to define. If you get unexpected results while using Resource Helper, you might need to define your key resources yourself. See “Defining Key Translations” on page 185 for more information.

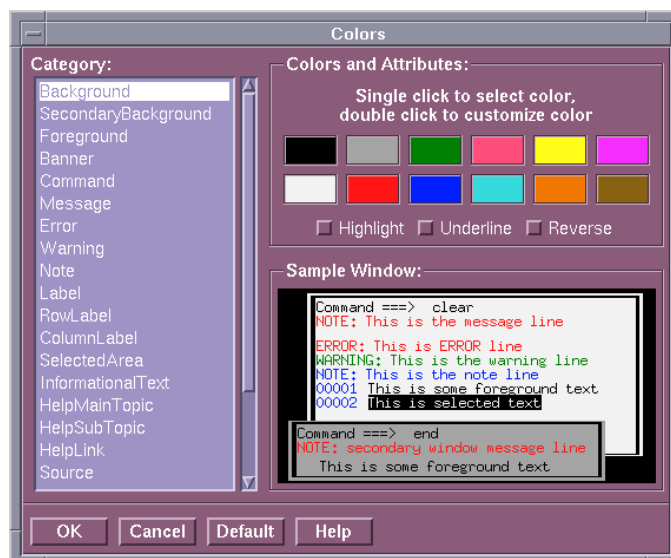
Modifying the Color of a SAS Window Using the Resource Helper

How to Use the Color Window

You can modify the color of part of a SAS window as follows:

- 1 Start Resource Helper and select the Colors icon.


Display 8.4 Colors Window for Resource Helper



- 2 Select a category from the **Category** area.
- 3 Click a color or attribute in the Colors and Attributes window, or double-click a color to open the **Customize Colors** area, shown in the following display.

You can also change the attributes of some categories of SAS windows. The attributes options are HIGHLIGHT, UNDERLINE , or REVERSE.

Display 8.5 Customize Colors Window for Resource Helper



You can customize a color by doing the following:

- selecting a new **Alias**
 - moving the **Red**, **Blue**, or **Green** sliders
 - selecting **Grab** and clicking a color anywhere on your screen
- 4 Click **OK** to exit the Colors window after you have finished defining your color settings.

The result is displayed in the **Sample Window**. The hexadecimal value of the color is displayed at the bottom of the window.

Example: Change the Color of a SAS Window

The following example shows how to change the color of a SAS window.

- 1 Double-click **Red** in the Colors window.

The **From:** display shows the red currently used by the SAS windowing environment.
- 2 Click **Aquamarine** under **Aliases** and observe the change in the **To:** display.
- 3 Move the **Red**, **Green**, and **Blue** sliders with your mouse and note the changes in the color of the **To:** display.

- 4 Click **Apply** and note the difference in the color displayed as **Red** in the Colors area.
- 5 Click **OK** to save your changes.

Return to the Default Settings

Click **Defaults** to restore your color settings to their default values.

Permanently Save Your Color Settings

To save your color settings permanently, from the Resource Helper drop-down menus, select **File ► Save Resources**.

How the Resource Helper Searches for X Resources

The following list describes the locations where the Resource Helper searches for resource definitions and the order in which it searches these locations.

- 1 Resource Helper loads the resources in the file pointed to by the `XENVIRONMENT` environment variable. If `XENVIRONMENT` is not set, Resource Helper loads the resources in the `~/.xdefaults-hostname` file, where *hostname* is the name of the computer on which Resource Helper is running.
- 2 Resource Helper loads the resources defined in the `RESOURCE_MANAGER` property. If the `RESOURCE_MANAGER` property is the first location in which Resource Helper finds resources, the `RESOURCE_MANAGER` property will override any resources that you generate with Resource Helper.

To determine whether any resources have been defined in your `RESOURCE_MANAGER` property, issue the following command:

```
xrdb -q | more
```

If no listing is returned, the `RESOURCE_MANAGER` property does not exist. In this case, Resource Helper loads the resources defined in the `~/.xdefaults` file.

- 3 Resource Helper loads the resources in the file pointed to by the `XUSERFILESEARCHPATH` environment variable.

You can use `%N` to substitute an application class name for a file when specifying the `XUSERFILESEARCHPATH` environment variable. For example, to point to `/usr/local/resources` as the location of all the resources for any application, issue the following command in the Bourne or Korn shells:

```
export XUSERFILESEARCHPATH=\
/usr/local/resources/%N
```

In the C shell, the command is the following:

```
setenv XUSERFILESEARCHPATH \
/usr/local/resources/%N
```

As a result, when SAS is invoked, the file pointed to by `XUSERFILESEARCHPATH` is the following:

```
/usr/local/resources/SAS
```

SAS is the application class name for SAS.

- 4 Resource Helper loads the resources in the file specified by the `XAPPLRESDIR` environment variable. The application's class name is appended to the `XAPPLRESDIR` environment variable and the resulting string is used to search

for resources. For example, you can issue the following command in the Bourne or Korn shells:

```
export XAPPLRESDIR=/usr/local/app-defaults
```

If you do this, then at the next invocation of SAS, the application's class name is appended to the path:

```
/usr/local/app-defaults/SAS
```

In the C shell, the command is the following:

```
setenv XAPPLRESDIR /usr/local/app-defaults
```

- 5 Resource Helper loads the resources in the file named **~/SAS**.
- 6 Resource Helper loads the resources in the file or substitution specified by the **XFILESEARCHPATH** environment variable.

Note: To determine whether an environment variable has been set, you can issue the following command:

```
env|grep <environment_variable>
```

Δ

- 7 Resource Helper loads the resources defined in **/usr/lib/x11/app-defaults**. Resource Helper does not need to have Write access to this file, but it must be able to read the file and add the SAS resources to a resource file that has Write access. Resource Helper does not generate a warning message if the file is not present or if it cannot read the file.
- 8 Resource Helper loads the fallback resources that are defined in the SAS code.

Except for the **/usr/lib/x11/app-defaults** file, Resource Helper tries to write the new resources to the same directory and file where it first found SAS resources. This location must be a file that has Write access and a directory that has Write access. If Resource Helper cannot write to the file, the SAS resources in that file remain in effect and any new or modified resources generated by Resource Helper will not take effect. If this situation happens, Resource Helper displays an error dialog box that contains the file or directory and suggests a way to fix the problem.

Customizing Toolboxes and Toolsets in UNIX Environments

Techniques for Customizing Toolboxes

You can customize toolboxes in the following ways:

- Through the Preferences dialog box. The Preferences dialog box enables you to customize the appearance and behavior of toolboxes. See “Modifying X Resources Through the Preferences Dialog Box” on page 167 and “Modifying the SAS ToolBox Settings” on page 170 for information about using the Preferences dialog box.
- bBy specifying SAS resources in your resource file. See “X Resources That Control Toolbox Behavior” on page 178 for a description of the SAS resources that affect toolboxes.
- Through the Tool Editor. The Tool Editor enables you to customize the individual tools in a toolbox. See “Using the Tool Editor” on page 179 for more information.

X Resources That Control Toolbox Behavior

You can control the behavior of toolboxes with the following SAS resources:

SAS.autoComplete: **True | False**

specifies whether SAS automatically fills in the remaining letters of a command as you type a command in the command window that begins with the same letter as a command that you have entered previously. The default value is True.

SAS.commandsSaved: *number-of-commands-saved*

specifies whether SAS saves the commands that you enter in the command window and how many commands are saved. You can specify a number from 0 to 50. If you specify 0, no commands will be saved. If you specify 1 or more, that number of commands is saved in the file **commands.hist** in your Sasuser directory. If you specify 1 or more for this resource and **SAS.autoComplete** is True, then SAS can automatically fill in commands that were entered in previous sessions. The default value is 25.

SAS.defaultToolBox: **True | False**

controls opening the default toolbox when SAS is invoked. The default is True.

SAS.isToolBoxPersistent: **True | False**

controls whether the toolbox that is associated with the Program Editor stays open when you close the Program Editor. The default value is True.

SAS.toolboxAlwaysOnTop: **True | False**

controls whether the toolbox is always on top of the window stack. The default value of True might cause problems with window managers that are not Motif interface window managers or other applications that want to be on top of the window stack. If you have such a situation, set this resource to False.

SAS.toolboxTipDelay: *delay-in-milliseconds*

sets the delay in milliseconds before displaying the toolbox tip. The default is 750.

SAS.useCommandToolBoxCombo: **True | False**

controls whether the command window and toolbox are joined or separated. The **SAS.defaultToolBox** and **SAS.defaultCommandWindow** resources control whether the toolbox and command window are displayed. If both are displayed, this resource controls whether they are joined or separated. The default value is True.

SAS.useLargeToolBox: **True | False**

controls whether tool icons in the toolbox are displayed as 24x24 pixels or 48x48 pixels. The default is False (24x24 pixels).

SAS.useShowHideDecorations: **True | False**

controls whether the combined command window and toolbox window has arrows at the left and right. You can use these arrows to hide or show portions of the window as they are needed. The default value is False.

SAS.useToolBoxTips: **True | False**

determines whether toolbox tip text is displayed. Some window managers might place the toolbox tip behind the toolbox. If the toolbox tip is placed behind the toolbox in your environment, set this resource to False. The default is True.

Using the Tool Editor

What Is a Toolset?

The Tool Editor enables you to create custom toolsets for your SAS applications. A *toolset* is a set of predefined tools that is associated with an application. Toolsets make it easier for individual users to customize their application toolboxes. If you create a toolset for an application, users can select **Actions** in the Tool Editor and choose the tools that they want to appear in their toolboxes. Users do not have to define the icons, commands, tip text, and IDs for those tools.

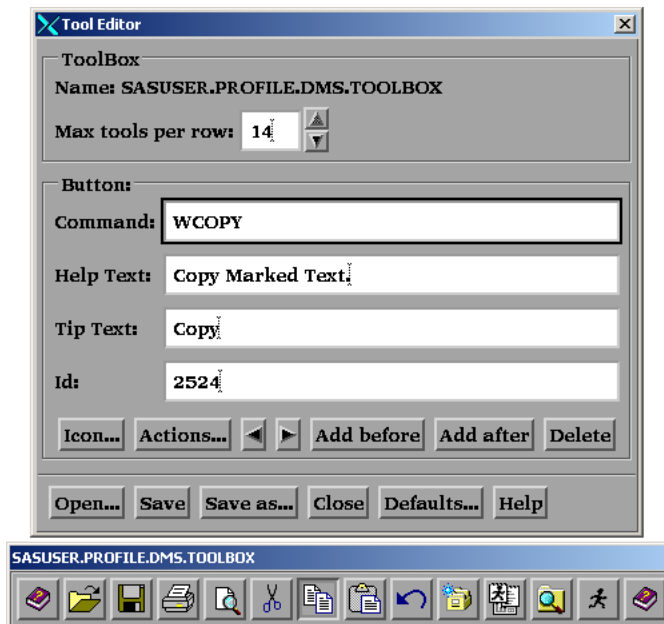
For example, you can define a default toolbox for your application that includes tools for opening files, cutting, copying, and pasting text, and saving files. You can define a toolset that includes those tools and tools for opening the Preferences dialog box, opening the Replace dialog box, and entering the RECALL command. These additional tools will not appear in users toolboxes unless users add them to their toolboxes with the Tool Editor. See “Changing the Attributes of an Existing Tool” on page 180 and “Create or Customize an Application- or Window-Specific Toolset” on page 183 for more information.

Invoking the Tool Editor

You can change the appearance and contents of a toolbox using the Tool Editor. To invoke the Tool Editor, select **Tools ► Options ► Edit Toolbox**. Alternatively, you can issue the TOOLEEDIT command as described in “TOOLEEDIT Command” on page 234.

The following display shows an example of a Tool Editor dialog box that was opened from the **Tools** menu in the Program Editor window:

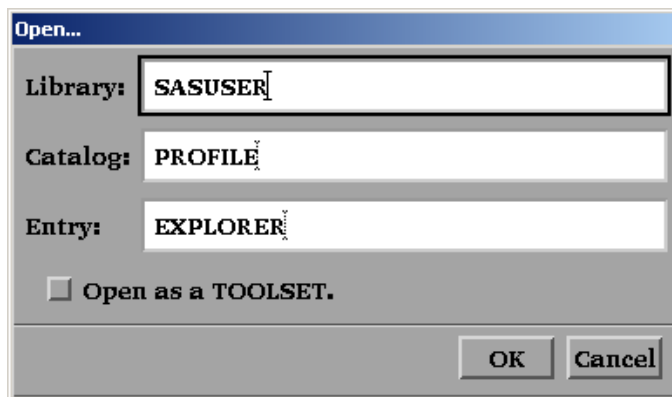
Display 8.6 Tool Editor Dialog Box



By default, the Tool Editor enables you to edit the current toolbox. To edit a different toolbox, click the **Open** button in the Tool Editor dialog box. Specify a library, catalog,

and entry name for the toolbox you want to edit. The following display shows the Open dialog box:

Display 8.7 The Open Dialog Box



After You Invoke the Tool Editor

After you invoke the Tool Editor, the toolbox displays in preview mode. In preview mode, clicking a tool icon to select a tool makes that icon the current icon. Its associated commands are displayed in the **Command** field. Attributes in the **Help Text**, **Tip Text**, and **Id** fields might also display, depending on whether this information was added when the tool was created or updated. For more information about the fields and buttons in the Tool Editor dialog box, click the **Help** button.

Changing the Appearance of the Entire Toolbox

The items in the **ToolBox** area of the Tool Editor affect the entire toolbox:

Name

displays the catalog entry that you are editing. The default toolbox is named SASUSER.PROFILE.DMS.TOOLBOX.

Max tools per row

specifies how the icons in the toolbox are arranged. The default value creates a horizontal toolbox. One tool per row creates a vertical toolbox.

Changing the Attributes of an Existing Tool

When you open the Tool Editor, the first tool in the toolset is selected, and the attributes for this tool appear in the **Button** area of the Tool Editor dialog box. If you click another icon in the toolset, the Tool Editor displays the attributes of that tool.

Alternatively, you can select a tool from the toolset that is displayed when you click the **Actions** button. When you click **OK** after you select a tool, the attributes in the **Button** area of the Tool Editor are updated to correspond to the new tool.

Note: Clicking the **Actions** button displays a toolset only if a toolset is associated with (meaning, has the same entry name as) the toolbox that you are editing. See “Saving Changes to the Toolbox or Toolset” on page 182 for more information. Δ

When you have selected the tool you want to change, you can then select an attribute field in the Tool Editor and type the changes you want to make.

To modify the attributes of a tool, follow these steps:

- 1 From the toolset, select the tool that you want to change.

- 2 In the **Button** area, select an attribute field associated with a button and change the text as appropriate:

Command

specifies the command or commands that you want executed when you click the icon. You can use any windowing environment command that is available under UNIX. For information about commands that are valid in all operating environments, see SAS Help and Documentation. Separate commands with a semicolon (;). For example, you could create an icon to open the Change Working Directory dialog box by using the DLGCDIR command.

Help Text

is used for applications that are designed to be run under Windows. The help text is displayed in the AWS status bar on Windows when a toolbox is ported to and loaded on those windows.

Tip Text

specifies the text that is displayed when you position the cursor over the icon.

Id

is useful if you are creating toolboxes for SAS/AF applications. The ID is the identifier of the corresponding menu item in the application. This number is the value assigned to the item in the ID option of the ITEM statement in PROC PMENU. If you specify an ID, then the application can set the state of the PMENU item to match the state of the tool in the toolbox. It can make the PMENU item active or inactive to match whether the tool in the toolbox is active or inactive. If you do not specify an ID, the ID defaults to 0.

- 3 Change the icon if necessary.

- a Click the **Icon** button or double-click an icon in the preview toolbox. The Tool Editor opens the Select a pixmap dialog box, which displays the icons that are provided with SAS. These icons are divided into several categories such as SAS windows, data, analysis, numbers and symbols, files, folders, and reports, and so on. To change categories, select the arrow to the right of the **Icon Category** field and select a new category.
- b Select the icon you want to use, and then select **OK**.

The following display shows the Select a pixmap dialog box:



- 4 Save your changes as described in “Saving Changes to the Toolbox or Toolset” on page 182.

Adding Tools to the Toolbox

To add a tool to the toolbox, follow these steps:

- 1 Select the icon next to where you want to add the new tool.
- 2 Select **Add before** or **Add after**. The Tool Editor adds a new icon to the toolbox and clears the Button fields.
- 3 Enter the appropriate information in the Button fields as described in “Changing the Attributes of an Existing Tool” on page 180.
- 4 Change the attributes of the icon, if necessary, as described in “Changing the Attributes of an Existing Tool” on page 180.
- 5 Save your changes as described in “Saving Changes to the Toolbox or Toolset” on page 182.

Change the Order of the Tools in the Toolbox

To change the position of a tool in the toolbox, select the tool icon, and then click on the left or right arrows to move the tool.

Deleting Tools from the Toolbox

To delete a tool from the toolbox, follow these steps:

- 1 Select the tool you want to delete.
- 2 Click **Delete**.
- 3 Save your changes as described in “Saving Changes to the Toolbox or Toolset” on page 182.

Returning to the Default Settings

To return all tools in the current Toolbox to their default settings, click **Defaults**. The Tool Editor asks you to verify your request. Click **Yes**, **No**, or **Cancel**.

Saving Changes to the Toolbox or Toolset

You can save the changes to the catalog entry shown in the **Name** field or create a new toolbox with a different name.

If you are customizing a window- or application-specific toolbox or toolset for your own use, you should save the customized toolbox or toolset in your Sasuser.Profile catalog using the same entry name as the PMENU entry for the window or application. SAS searches for toolboxes and toolsets first in the Sasuser.Profile catalog, and then in the application catalog.

If you are a SAS/AF application developer or site administrator and you are editing a window- or application-specific toolbox that you want to be accessible to all users, you must save the TOOLBOX entry with the same library, catalog, and entry name as the PMENU entry for the window or application. To associate a toolset with a specific toolbox, save the TOOLSET entry with the same library, catalog, and entry name as the TOOLBOX entry. You will need write permissions to the appropriate location. For example, to store a customized toolbox for the graphics editor, the site administrator needs to store the toolbox in SASHELP.GI.GEDIT.TOOLBOX.

Clicking the **Save** button saves the toolbox information to the catalog entry shown in the **Name** field. Clicking the **Save As** button prompts you to enter a different library,

catalog, and entry name. You can also choose to save the toolbox as a toolset. If you save the toolbox as a toolset, the entry type will be TOOLSET; otherwise, the entry type is always TOOLBOX. (Saving a set of tools as a TOOLSET does not change your TOOLBOX entry. See “Customizing Toolboxes and Toolsets in UNIX Environments” on page 177 and “Create or Customize an Application- or Window-Specific Toolset” on page 183 for information about toolsets.)

If you click the **Close** or the **Open** button without first saving your changes, the Tool Editor prompts you to save the changes to the current toolbox or toolset before continuing.

After you save the toolbox or toolset, the Tool Editor remains open for additional editing, and the **Name** field changes to the name of the new entry (if you entered a new name).

Creating a New Toolbox

To create an entirely new toolbox, follow these steps:

- Edit an existing toolbox using the Tool Editor and then save the toolbox by clicking the **Save as** button as described in “Saving Changes to the Toolbox or Toolset” on page 182.
- Open the Sasuser.Profile catalog in the Explorer window and add a new toolbox by selecting **File** \blacktriangleright **New** \blacktriangleright **Toolbox**.

Create or Customize an Application- or Window-Specific Toolbox

If you are an application developer and want to create or edit an existing application toolbox, follow these steps:

- 1 Delete any existing TOOLBOX entry in your Sasuser.Profile for the window or application that you want to customize.

Deleting the copy of the toolbox in your Sasuser.Profile enables you to pick up a copy of the toolbox that is supplied with SAS when you invoke the Tool Editor.

- 2 Create or edit the application toolbox as described in “Creating a New Toolbox” on page 183 or “Using the Tool Editor” on page 179.
- 3 Save the edited toolbox as described in “Saving Changes to the Toolbox or Toolset” on page 182.
- 4 Inform your users that you have changed the window or application toolbox.

If users want to use the new toolbox, they must delete the corresponding TOOLBOX entry from their Sasuser.Profile. The new toolbox will then be automatically loaded when the window or application is invoked. If a user does not delete the corresponding TOOLBOX entry from their Sasuser.Profile, that copy of the toolbox will be loaded instead of the new toolbox.

The TOOLLOAD and TOOLCLOSE commands are most useful when you are developing SAS/AF applications. You can use the EXECCMDI routine with these commands to enable your application to open and close the toolbox and to give users of your applications access to several toolboxes during the course of their work. See *SAS Component Language: Reference* for a description of the EXECCMDI routine.

Create or Customize an Application- or Window-Specific Toolset

You define application- or window-specific toolsets in the same way that you create an application- or window-specific toolbox. There are only two differences:

- To create a new toolset, start by defining a toolbox as described in “Creating a New Toolbox” on page 183.
- After you have defined the toolbox, save it as a TOOLSET entry, not as a TOOLBOX entry.

Note: If you are an application developer, make sure that you delete any existing TOOLSET entry for your application as described in “Create or Customize an Application- or Window-Specific Toolbox” on page 183 before you modify your application’s toolset. \triangle

Customizing Key Definitions in UNIX Environments

Techniques for Customizing Your Key Definitions

There are four ways to customize your key definitions:

- Through the Keys window.

To open the Keys window, issue the KEYS command or select **Tools** \blacktriangleright **Options** \blacktriangleright **Keys**. If you change any key definitions through the Keys window for the primary SAS windowing environment windows, the definitions are stored in the Sasuser.Profile catalog in the entry DMKEYS.KEYS. Key definitions for other SAS windows are stored in catalog entries named BUILD.KEYS, FSEEDIT.KEYS, and so on.

See the online SAS Help and Documentation for more information about the KEYS command and the Keys window.

- With the KEYDEF command. The KEYDEF command enables you to redefine individual function keys:

```
keydef keyname <command|-text-string>
```

For example, if you specify **keydef F8 dlgpref**, then the **F8** key will open the Preferences dialog box.

For more information about the KEYDEF command, see the Base SAS section in the online SAS Help and Documentation.

- Through the Resource Helper (reshelper).

Resource Helper generates SAS resource specifications based on keys and functions that you select. You can use Resource Helper to change the function of any key that is listed in the Keys window. See “Setting X Resources with the Resource Helper” on page 172 and “Defining Keys with the Resource Helper” on page 173 for more information about the Resource Helper.

In most cases, Resource Helper is much easier and faster than defining the resources yourself. However, because the X Window System searches for resources in several places, it is possible for Resource Helper to pick up the wrong key symbol for the key you are trying to define. Also, unless the action routine that you assign to your keys is the **sas-function-key** routine, then Resource Helper does not provide a way to change the key labels in the Keys window. In both of these cases, you will need to define your key resources yourself.

- By defining the **SAS.keyboardTranslations** and **SAS.keyboardWindowLabels** resources in your resources file as described in “Defining Key Translations” on page 185.

You can define most of the keys on your keyboard. However, a few keys have dedicated functions that are associated with them. For example, the mouse buttons are dedicated to the cursor and cut-and-paste operations and are not available for user customization.

Defining Key Translations

What Is a Key Translation?

Key customization for the X Window System consists of defining a key sequence and an action to be executed when that key sequence is typed on the keyboard. This customization is known as *binding keys to actions*; together they are referred to as a *translation*.

What Is the SAS.keyboardTranslations Resource?

The **SAS.keyboardTranslations** resource specifies the set of key bindings that SAS uses in all SAS windows. The default value for the **SAS.keyboardTranslations** resource is determined at run time based on the vendor identification string reported by the X server that you are using as the display. These defaults are listed in the files contained in **!SASROOT/X11/resource_files**. To modify the default bindings that are supplied by SAS, you must modify the **SAS.keyboardTranslations** resource.

Note: The X Toolkit Intrinsic translations that are specified in this resource apply to both the user area and the command line of all SAS windows that are affected by this resource. This resource does not affect windows that are controlled by Motif interface resources, such as the Command window, the Open or Import dialog boxes, and some other drop-down menu dialog boxes. Δ

Steps for Creating a Key Definition

To create a key definition, follow these steps:

- 1 Determine the keysyms for the keys that you want to define.

Keysyms are the symbols that are recognized by the X Window System for each key on a keyboard. See “Determining Keysyms” on page 186 for more information.

- 2 Modify or add the **SAS.keyboardTranslations** resource in your resource file to include the definitions of the keys that you want to define.

Use a keyboard action routine to define which action you want the key to perform. The definition in the right column in the Keys window will no longer control the function of any keys that are defined with a keyboard action routine other than **sas-function-key**. The definitions of those keys in the Keys window become labels that have no effect. See “Syntax of the SAS.keyboardTranslations Resource” on page 187 for more information.

- 3 Modify or add the **SAS.keysWindowLabels** resource in your resource file.

The **SAS.keysWindowLabels** resource specifies the set of valid labels that will appear in the SAS Keys window. Modify this resource only if you want to add new labels or modify existing labels in the left column in the Keys window.

The **SAS.keysWindowLabels** resource defines only the mnemonics used in the Keys window. For a specific key to perform an action, you must specify a **SAS.keyboardTranslations** definition for the key. See “Syntax of the SAS.keysWindowLabels Resource” on page 188 for more information.

- 4 Start a SAS session and open the Keys window.
- 5 In the right column in the Keys window, type a command name or other description of each key that you have defined.

See “Examples: Defining Keys Using SAS Resources” on page 191 for examples of key definitions.

Determining Keysyms

You can use the **xev** utility to determine the keysyms that are associated with the keys on your keyboard. The **xev** utility is distributed with most UNIX operating systems, but if **xev** is not installed in your operating environment, contact your UNIX system administrator for information about other methods that are available in your UNIX environment. The **xev** utility writes a message for each X event that occurs. The **KeyPress** event specifies the keysym for each key that is pressed.

To define keys, follow these steps:

- 1 Start the **xev** utility on the X server for which you want to define keys.

The **xev** client displays a small Event Tester window that lists the X events that occur. (The **xev** client generates a large amount of output, so you might want to save the output to a file for later review. You can issue the UNIX **script** command to save the output to a file.)
- 2 Give keyboard focus to the Event Tester window by clicking the mouse pointer on the window, if necessary.
- 3 Press the key that you want to define, and watch for the **KeyPress** event to be listed.

The listing contains a number of items that are separated by commas. One of the fields in the **KeyPress** event lists the keysym name that is associated with the key that was pressed.

For example, when the 0 key on the keypad of a Dell PC 105 keyboard is pressed, with the NumLock modifier toggled on, it generates the following output:

```
KeyPress event, serial 32, synthetic NO,
  window 0x1a00001, root 0x5d, subw 0x1a00002,
  time 600120687, (37,41), root:(240,458),
  state 0x10, keycode 90 (keysym 0xffb0, KP_0),
  same_screen YES,
  XLookupString gives 1 bytes: (30) "0"
  XmbLookupString gives 1 bytes: (30) "0"
  XFilterEvent returns: False
```

In this example, the keysym name is **KP_0**.

Note: SAS defines a set of *virtual keysyms* with the **SAS.defaultVirtualBindings** resource. Virtual keysyms all begin with **osf**, such as **osfPageDown**, **osfClear**, and **osfPrimaryPaste**. If you remap these virtual bindings instead of using the defaults supplied by SAS, you might get unexpected results. If you specify a key translation that does not work, you might be trying to redefine a key that is bound to a virtual keysym. In this case, you must specify the virtual keysym in the **SAS.keyboardTranslations** resource instead of the keysym that is displayed by the **xev** utility. To determine the virtual keysym that is bound to a key, you can start the Resource Helper, click **Keys**, and press the key or key combination that you want to define. Resource Helper will display the virtual keysym name. You can also refer to the key definition files in **/X11/resource_files** in the directory where SAS is installed (!**SASROOT**) and to the UNIX man pages for **VirtualBinding** or **xmbind**. \triangle

Syntax of the SAS.keyboardTranslations Resource

Note: Most SAS documentation uses angle brackets (<>) to indicate optional syntax. However, in this topic, optional syntax is shown with square brackets ([]). The angle brackets that are shown in this topic are part of the syntax and should be entered exactly as shown. △

Here is the syntax of the **SAS.keyboardTranslations** resource:

```
SAS.keyboardTranslations: #override \  
[modifier] <Key>keysym : action-routine \n\  
[modifier] <Key>keysym : action-routine
```

#override

indicates that this definition should override any existing bindings for the specific keys that you define without affecting any other keys. If you omit the *#override* directive, the new bindings replace all of the default bindings, and none of the other keys on the keyboard will be available.*

modifier

can be one of the following:

Alt

Ctrl

Meta

Shift

Lock

Mod1

Mod2

Mod3

Mod4

Mod5

None

blank space

The list of valid modifiers varies depending on your keyboard. To display a list of valid modifiers for your keyboard, enter the **xmodmap** UNIX command. Refer to the UNIX man page for **xmodmap** for more information.

<Key>

is required. It signals the beginning of the keysym.

keysym

is the key symbol recognized by X for the key that you are defining. See “Determining Keysyms” on page 186 for more information.

action-routine

is what you want the key to do. You can specify any action routine described in “SAS Keyboard Action Names” on page 188.

\n

enables the X translation manager to determine where one translation sequence ends and the next one begins. Do not enter \n after the end of the last translation.

* For information about the *#augment* and *#replace* directives, see the documentation for the X Window System.

\backslash
prevents the newline character at the end of the line from being interpreted as part of the definition. Using this character is a stylistic convention that allows each translation to be listed on a separate line. Do not enter a backslash after the end of the last translation.

Note: SAS does not prevent you from specifying invalid keys in the **SAS.keyboardTranslations** resource. In some cases, invalid keys will produce warnings in the shell window. Δ

Syntax of the SAS.keysWindowLabels Resource

Note: The square brackets ([]) in the following syntax indicate that (*InternalKeyName*) is optional. Δ

Here is the syntax of the **SAS.keysWindowLabels** resource:

```
SAS.keyWindowLabels:  $\backslash$   
KeyWindowLabel [(InternalKeyName)]  $\backslash$  $n$  $\backslash$   
KeyWindowLabel [(InternalKeyName)]
```

KeyWindowLabel
is the label (1 to 8 characters) that you want to appear in the Keys window.

InternalKeyName
is the character string that is passed to the **sas-function-key** action routine in the corresponding **SAS.keyboardTranslations** key binding. (*InternalKeyName* is used by SAS to correlate Keys window entries to key definitions in the KEYS modules loaded from SAS catalogs or defined in the SAS Keys window.) If the *InternalKeyName* is not specified, SAS uses the *KeyWindowLabel* as the *InternalKeyName*.

\backslash n and \backslash
serves the same purpose as in the **SAS.keyboardTranslations** resource. See “Syntax of the SAS.keyboardTranslations Resource” on page 187 for more information.

SAS Keyboard Action Names

Note: Most SAS documentation uses angle brackets (<>) to indicate optional syntax. However, in this topic optional syntax is shown with square brackets ([]). The angle brackets that are shown in this topic are part of the syntax and should be entered exactly as shown. Δ

SAS declares a set of keyboard actions during X initialization. You can think of these keyboard actions as simple functions. When the actions are executed, they act on the window that currently has keyboard input focus.

The following list of keyboard actions represents action routines registered by the Motif interface for use with X Toolkit keyboard event translations.

sas-cursor-down()

moves the cursor down one line in the SAS window. The cursor does not wrap when it reaches the bottom of the SAS window interior.

sas-cursor-left()

moves the cursor left one character in the SAS window. The cursor does not wrap when it reaches the left side of the SAS window interior.

sas-cursor-right()

moves the cursor right one character in the SAS window. The cursor does not wrap when it reaches the right side of the SAS window interior.

sas-cursor-up()

moves the cursor up one line in the SAS window. The cursor does not wrap when it reaches the top of the SAS window interior.

sas-delete()

deletes all text in the current field.

sas-delete-begin()

deletes text from the current cursor position to the beginning of the current text field.

sas-delete-char()

deletes the character under the text cursor and leaves the cursor in place.

sas-delete-end()

deletes text from the current cursor position to the end of the current text field.

sas-delete-prev-chr()

deletes the character to the left of the text cursor and moves the cursor back one space.

sas-delete-prev-word()

deletes text to the start of the previous word from the current cursor position. If the cursor is in the interior of a word when the action is invoked, the text from the cursor position to the start of the word is deleted.

sas-delete-word()

deletes text from the current cursor position to the end of the current or next word.

sas-do-command()

accepts one or more text string parameters that are interpreted as SAS commands to be executed when the action is invoked. The action might be invoked with multiple parameters. The parameters are concatenated with semicolon delimiters supplied by the **sas-do-command** action between the parameters. The assembled SAS command string is then submitted for execution. For example, the following translation syntax can be used to define a HOME, SUBMIT key sequence for all SAS windowing environment windows:

```
<Key>KP_F3: sas-do-command(HOME;SUBMIT)
```

sas-function-key("InternalKeyName")

invokes the SAS commands associated with the function key identified by the *InternalKeyName* label. *InternalKeyName* is the character string (1 to 8 characters long) that is passed to the keysWindowLabels resource. Enclose *InternalKeyName* in quotation marks. See “Defining Key Translations” on page 185 for a description of internal key names.

sas-home-cursor()

is the equivalent of the HOME command. It is provided for convenience so that the HOME action could be defined for all SAS windowing environment windows.

sas-insert-char(["InsertionString"])

inserts or overwrites the character typed into the input field under the text cursor. Insert or overstrike behavior is determined by the **sas-toggle-insert** action, which has a mode that is reflected by the text cursor style displayed; the block cursor indicates overstrike mode, and the underline cursor indicates insert mode. Normally, **sas-insert-char** translates the XKeyEvent into the appropriate character and inserts it at the SAS text cursor location. If you specify the parameter, the text string represented by this parameter is inserted at the SAS text cursor location. White space in the string is interpreted by the X Toolkit as a parameter delimiter unless you enclose the string in double quotation marks. See your X Window System documentation for information on embedding quotation marks in the string parameter. To include an escaped quotation mark, use the following syntax:

```
Shift<Key>KP_1: sas-insert-char("One\\"1\\")
```

This syntax produces the text string **One"1"** at the SAS text cursor location.

sas-kp-application()

sets the workstation's numeric keypad to allow function key translations to be reinstated. This action only works for those keypad keys that are bound to **sas-function-key()** actions. Keypad bindings to other actions are not affected by this translation.

sas-kp-numeric()

sets the workstation's keypad to generate numeric characters instead of its previous function key assignment. This action only works for keypad keys that are bound to **sas-function-key()** actions. Keypad bindings to other actions are not affected by this translation.

sas-move-begin()

moves the cursor to the beginning of the current text field.

sas-move-end()

moves the cursor to the end of the current text field.

sas-new-line()

generates an end-of-line event when invoked. This action is context sensitive. If the action is typed on the SAS command line, the text entered will be submitted for execution. If invoked in the SAS application client area, the action depends on the attributes of the text area under the text cursor. In simplest terms, this action is the general line terminator for an input field.

sas-next-field()

advances the SAS application to the next field in the SAS window client area.

sas-next-word()

skips the text cursor forward to the beginning of the next word in the current text field. If **sas-next-word** does not find the beginning of a word in the current text field, it advances to the next SAS application field. If you are typing in the SAS command line area of the window, the cursor will not wrap into the SAS window client area.

sas-page-down()

scrolls the current window contents forward by one page.

sas-page-end()

moves the text cursor to the end of the current page.

sas-page-top()

moves the text cursor to the top of the current page.

sas-page-up()

scrolls the window contents backward by one page.

sas-prev-field()

returns the SAS application to the previous field in the SAS window client area.

sas-prev-word()

skips the text cursor backward to the beginning of the previous word in the current text field. If **sas-prev-word** does not find the beginning of a previous word in the current text field, it returns to the end of the previous SAS application field. If you are typing in the SAS command line area of the window, the cursor will not wrap into the SAS window client area.

sas-to-bottom()

Moves the text cursor to the absolute bottom of the window's text range.

sas-to-top()

Moves the text cursor to the absolute top of the window's text range.

sas-toggle-insert()

switches the associated window line-editing behavior between insert and overstrike modes. This switching applies only to the SAS command line and the SAS window client area. The current mode is indicated by the cursor style in use. The block cursor indicates overstrike mode, and the underline cursor indicates insert mode.

sas-xattr-key(<KeyType>[, <KeyParam>])

processes SAS extended attribute keys. The *KeyType* parameter must be one of the following values: XACOLOR, XAATTR, XACLEAR. For *KeyType* XACOLOR, the 12 DMS color names are valid parameters; for *KeyType* XAATTR, the valid values are HIGHLIGHT, REVERSE, BLINK, and UNDERLINE; for XACLEAR, no parameter is required. The BLINK attribute is not supported in the Motif interface. However, if you specify the BLINK attribute, it will be displayed when the catalog is ported to other operating environments.

Examples: Defining Keys Using SAS Resources

Note: Most SAS documentation uses angle brackets (<>) to indicate optional syntax. However, in these examples, optional syntax is shown with square brackets ([]). The angle brackets that are shown in these examples are part of the syntax and should be entered exactly as shown. Δ

In the following example, the **sas-do-command** action routine specifies that the COMMAND command is to override any existing definition for **KP_0**.

```
SAS.keyboardTranslations: #override \n\
  None<Key>KP_0: sas-do-command(COMMAND)
```

All other keys retain their current definitions.

The following example binds the key sequence CTRL-K to the KEYS command and specifies that CTRL-D deletes the character under the cursor. Commands entered in the Keys window for CTRL-K and CTRL-D will have no effect.

```
SAS.keyboardTranslations: #override\
  Ctrl<Key>k: sas-do-command(keys)\n\
  Ctrl<Key>d: sas-delete-char()
```

The following example specifies that the key associated with the keySYM **hpClearLine** performs the command entered beside the **MyClrLn** label in the Keys window.

```
SAS.keyboardTranslations: #override \
    <Key>hpClearLine : sas-function-key("ClearLn")
SAS.keysWindowLabels: MyClrLn(ClearLn)
```

The character string that appears inside the parentheses in the **SAS.keysWindowLabels** resource must match the string entered as the parameter to the **sas-function-key** routine. The label (**MyClrLn**) can be any character string, and the keysym **hpClearLine** must be a valid keysym for your keyboard.

Customizing Fonts in UNIX Environments

Difference between the System Font and Fonts That Are Used in the Windowing Environment

SAS uses two main types of fonts:

- The system font is used in most dialog boxes and menus. SAS inherits the system font defined by the CDE ***.systemFont** resource. If this resource is not defined, SAS uses a Helvetica font.
- DMS fonts are used in SAS windows. You can change the SAS font either through the Fonts dialog box or by specifying the resources in your resources file. The font must be a fixed or monospace font.

Note: It is best to change fonts before invoking any applications. Changing fonts while applications are running might result in unexpected behavior. \triangle

How SAS Determines Which Font to Use

SAS determines the normal (not bold) default font as follows:

- 1 If you saved a font in SASUSER.PROFILE.DMSFONT.UNXPREFS through the Font dialog box, this font is used as the default normal font.
- 2 If you did not save a font through the Font dialog box, but you set the **SAS.DMSFont** resource, SAS uses the font specified by this resource as the default font.
- 3 If you did not set the **SAS.DMSFont** resource, SAS uses any font that matches the pattern ***Font** which might be defined or inherited.
- 4 If you did not specify or inherit any resources matching ***Font**, but you did set the **SAS.DMSFontPattern** resource, SAS uses this resource to determine which font to use. The **SAS.DMSfontPattern** resource has no effect if any resources matching ***Font** are inherited or defined.
- 5 If no resources were set, SAS chooses a font from the fonts that are available on your X server.

If you did not specify a value for the **SAS.DMSboldFont** resource, SAS uses the default normal font to determine the default bold font. If the normal **SAS.DMSFont** has an XLFD name associated with it, then SAS selects the matching bold font and loads it. If SAS cannot automatically select or load a bold font, the normal font is also used for the bold font.

In many cases, font names are given aliases so that a shorter name can be used to refer to a font that has an XLFD name associated with it. The name used in determining a bold font is based on the XA_FONT font property for the normal font.

Customizing Fonts by Using the Fonts Dialog Box

Introduction to the Fonts Dialog Box

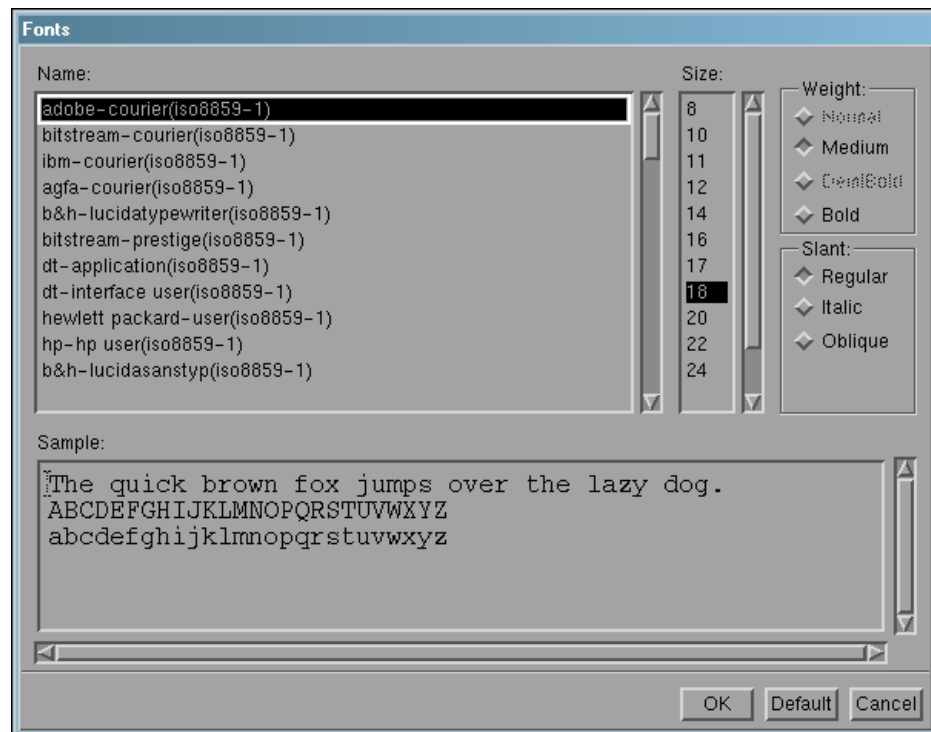
The Fonts dialog box enables you to change the windowing environment font for the entire SAS session. If you change the font, the font that you select is stored in SASUSER.PROFILE.DMSFONT.UNXPREFS and will be used in future SAS sessions.

How to Change the Default Font

To change the default font, complete the following steps:

- 1 To open the Fonts dialog box, use one of the following methods:
 - Issue the DLGFONT command in the command window.
 - Select **Tools ► Options ► Fonts**.

Display 8.8 Fonts Dialog Box



- Select a font name and, if desired, a size, weight, and slant. (Not all fonts are available in all sizes, weights, or slants.) The **Sample** field shows what the selected font looks like.
 - Click **OK** to change the existing font to the selected font.
- 2 To return to the default font, click **Default**.
 - 3 To cancel any changes and exit the Fonts dialog box, click **Cancel**.

Specifying Font Resources

You can customize the fonts that are used in the SAS windowing environment with the following resources:

SAS.DMSFont: *font-name*

specifies the font that you want to be used as the default normal font. The default normal font is Courier.

SAS.DMSboldFont: *font-name*

specifies the font that you want to be used as the default bold font.

SAS.DMSDBfont: *font-name*

specifies the multi-byte normal character set font used by the SAS windowing system for operating environments that support multi-byte character sets.

SAS.DMSDBboldFont: *font-name*

specifies the multi-byte bold character set font used by the SAS windowing system for operating environments that support multi-byte character sets.

SAS.DMSfontPattern: *XLFD-pattern*

specifies an X Logical Font Description, or XLFD pattern that you want SAS to use to determine the windowing environment font. Most fonts in the X Window System are associated with an XLFD, which contains a number of different fields delimited by a dash (-) character. The fields in the XLFD indicate properties such as the font family name, weight, size, resolution, and whether the font is proportional or monospaced. Refer to your X Window documentation for more information on the XLFD and font names used with X.

The *XLFD-pattern* that you specify for **SAS.DMSfontPattern** must contain the same number of fields as an XLFD. An asterisk (*) character means that any value is acceptable for that particular field. For example, the following pattern matches any font that has a regular slant, is not bold, is monospaced, and is an iso8859 font:

```
SAS.DMSFontPattern: -*-*-r-***-***-m*-iso8859-1
```

SAS uses the *XLFD-pattern* to choose a font as follows:

- 1 SAS queries the X server for the list of fonts that match the **SAS.DMSfontPattern** resource.
- 2 SAS excludes all fonts that have X and Y resolution values different from the current X display, all fonts that have variable character cell sizing (such as proportional fonts), and all fonts that have point sizes smaller than 8 points or larger than 15 points. If this step results in an empty list, SAS chooses a generic (and usually fixed) font.
- 3 The font with the largest point size is chosen from the remaining list.

SAS.fontPattern: *XLFD-pattern*

specifies an XLFD font pattern that describes the candidate fonts used to resolve SAS graphics font requests. Using this pattern allows the user to optimize or control the use of X fonts within the context of various SAS graphics applications. The default value of * usually does not affect performance to a significant degree. You might want to restrict the font search if you are running SAS on a server with an excessive number of fonts or that is operating in performance-limited environment.

SAS.systemFont: *font-name*

specifies the system font. The SAS font is used in SAS windows. The system font is used in most dialog boxes and menus. SAS typically inherits the system font from the font resources set by the X window environment, such as the Common Desktop Environment (CDE), or K Desktop Environment (KDE). If the ***.systemFont** resource, SAS uses a 12-point Helvetica font.

Specifying Font Aliases

If your server does not provide fonts to match all of the fonts that are supplied by SAS, you can use font alias resources to substitute the fonts available on your system. (Ask your system administrator about the fonts that are available.) Use the following syntax to specify font aliases in your resource file:

```
SAS.supplied-fontAlias: substitute-family
```

supplied-font is the name of the font supplied by SAS. *substitute-family* is the family name of the font that you want to substitute.

CAUTION:

Do not specify a SAS font as a font alias. There might be a conflict if you specify a font supplied by SAS as a font alias, as in the following example:

```
SAS.timesRomanAlias: symbol
```

Assigning this value to a font alias prevents the selection of any symbol fonts through the font selection dialog box, because they are specified as the Times Roman alias. △

The following table lists SAS font alias resource names.

Table 8.1 SAS Font Alias Resources

Resource Name	Class Name
SAS.timesRomanAlias	TimesRomanAlias
SAS.helveticaAlias	HelveticaAlias
SAS.courierAlias	CourierAlias
SAS.symbolAlias	SymbolAlias
SAS.avantGardeAlias	AvantGardeAlias
SAS.bookmanAlias	BookmanAlias
SAS.newCenturySchoolbookAlias	NewCenturySchoolbookAlias
SAS.palatinoAlias	PalatinoAlias
SAS.zapfChanceryAlias	ZapfChanceryAlias
SAS.zapfDingbatsAlias	ZapfDingbatsAlias

Example: Substitute the Lucida Font for Palatino

Suppose that your system does not have a Palatino font, but has the following Lucida font:

```
b&h-lucida-bold-r-normal-sans-
10-100-75-75-p-66-iso8859-1
```

To substitute Lucida for Palatino, include the following line in your resource file:

```
SAS.palatinoAlias: lucida
```

Customizing Colors in UNIX Environments

Methods for Customizing the Color Settings in Your SAS Session

SAS provides a default set of colors and attribute settings for the elements of all SAS windows. You can customize the colors in your SAS session in the following ways:

- Through Resource Helper (reshelper).
Resource Helper enables you to customize any color. See “Setting X Resources with the Resource Helper” on page 172 and “Modifying the Color of a SAS Window Using the Resource Helper” on page 174 for more information.
- Through the SASCOLOR window, as described in “Customizing Colors by Using the SASCOLOR Window” on page 196.
You can customize any window element for most SAS windows with the SASCOLOR window.
- With the COLOR command as described in “Syntax of the COLOR Command” on page 197.
The COLOR command affects only the specified element of the active window. Changes made with the COLOR command override changes entered through any of the other methods described here.
- By entering the color resource specifications yourself.
You can enter specific RGB values or color names for any of the X resources that control color. See “Defining Color Resources” on page 197 for more information.

Customizing Colors by Using the SASCOLOR Window

You can use the SASCOLOR window to change the color and highlighting of specific elements of SAS windows. To open the SASCOLOR window, issue the SASCOLOR command or select **Tools** \blacktriangleright **Options** \blacktriangleright **Colors**.

Display 8.9 SASCOLOR Window



To change a color for a window element, select the element name, and then select color and attribute that you want assigned to the element.

The BLINK attribute is not supported. The HIGHLIGHT attribute causes text to be displayed in bold font.

When you click **save**, your changes are saved to the catalog entry SASUSER.PROFILE.SAS.CPARMS.

Note: Close and reopen any active windows for new color settings to take effect. Δ

For more information about the SASCOLOR window, see the online SAS Help and Documentation.

Syntax of the COLOR Command

You can use the COLOR command to set the color for specific elements of the active window:

```
color field-type <color|NEXT <highlight>>
```

field-type

specifies an area of the window such as background, command, border, message, and so on.

color

specifies a color such as blue (which can be abbreviated B), red (R), green (G), cyan (C), pink (P), yellow (Y), white (W), black (K), magenta (M), gray (A), brown (B), or orange (O).

NEXT

changes the color to the next available color.

highlight

can be H (which causes text to be displayed in a bold font), U (underlined), or R (reverse video). The BLINK attribute is not supported.

To save your changes, issue the WSAVE command. The changes are saved to SASUSER.PROFILE.*window*.WSAVE.

Note: The WSAVE command is not available for all SAS windows. For example, with SAS/FSP changes are saved either through the EDPARMS or the PARM window. (To determine whether WSAVE is available for a SAS window, refer to the product documentation.) Δ

For more information about the COLOR and WSAVE commands, see the online SAS Help and Documentation.

Defining Color Resources

Types of Color Resources

Color resources fall into two categories:

- foreground and background definitions

These resources allow you to customize the RGB values that are used to define the 12 DMS colors. Because each color could be used as either a background or a foreground color, you can specify different RGB values or color names for each color for each usage. For example, you can specify that when blue is used as a

foreground color, color #0046ED is used, and when blue is used as a background color, CornflowerBlue is used.

- window element definitions

These resources, which are referred to as CPARMS resources, enable you to specify which of the 12 DMS colors you want to use for each window element. For example, you can specify that message text is displayed in magenta.

These two types of resources work together. The CPARMS color values use the current foreground and background definitions. For example, the following resources specify that the background of your primary windows will be CornflowerBlue:

```
SAS.blueBackgroundColor: CornflowerBlue
SAS.cparmBackground: DmBlue
```

Specifying RGB Values or Color Names for Foreground and Background Resources

SAS uses **SAS.systemBackground**, **SAS.systemForeground**, and the resources listed in the following table to determine the colors to be used in its windows.

SAS.systemForeground: color

specifies the color for the foreground system color in the SASCOLOR window.

SAS.systemBackground: color

specifies the color for the background system color in the SASCOLOR window.

SAS.systemSecondaryBackground: color

sets the system secondary background color and specifies the color for the secondary background system color in the SASCOLOR window.

You can specify color names such as MediumVioletRed or RGB values such as #0000FF for all of the foreground and background resources. See your X Window System documentation for information about RGB color values.

The following table lists all of the foreground and background color resources and their class names. All of these resources are of the type String.

Table 8.2 Foreground and Background Color Resources

Resource Name	Class Name
SAS.systemForeground	SystemForeground
SAS.systemBackground	SystemBackground
SAS.systemSecondaryBackground	Background
SAS.blackForegroundColor	BlackForegroundColor
SAS.blueForegroundColor	BlueForegroundColor
SAS.brownForegroundColor	BrownForegroundColor
SAS.cyanForegroundColor	CyanForegroundColor
SAS.grayForegroundColor	GrayForegroundColor
SAS.greenForegroundColor	GreenForegroundColor
SAS.magentaForegroundColor	MagentaForegroundColor
SAS.orangeForegroundColor	OrangeForegroundColor
SAS.pinkForegroundColor	PinkForegroundColor

Resource Name	Class Name
SAS.redForegroundColor	RedForegroundColor
SAS.whiteForegroundColor	WhiteForegroundColor
SAS.yellowForegroundColor	YellowForegroundColor
SAS.blackBackgroundColor	BlackBackgroundColor
SAS.blueBackgroundColor	BlueBackgroundColor
SAS.brownBackgroundColor	BrownBackgroundColor
SAS.cyanBackgroundColor	CyanBackgroundColor
SAS.grayBackgroundColor	GrayBackgroundColor
SAS.greenBackgroundColor	GreenBackgroundColor
SAS.magentaBackgroundColor	MagentaBackgroundColor
SAS.orangeBackgroundColor	OrangeBackgroundColor
SAS.pinkBackgroundColor	PinkBackgroundColor
SAS.redBackgroundColor	RedBackgroundColor
SAS.whiteBackgroundColor	WhiteBackgroundColor
SAS.yellowBackgroundColor	YellowBackgroundColor

Defining Colors and Attributes for Window Elements (CPARMS)

You can define the colors and attributes for specific window elements by assigning values to SAS resources known as CPARMS. Each CPARMS resource defines the color and attribute of a specific window element, such as the background in a secondary window or the border of a primary window.

You can specify multiple color and attribute names in the same resource definition, but only the final color and attribute will be used:

```
SAS.cparmResource: DmColorName|DmAttrName\  
<+DmColorName|DmAttrName>
```

Resource can be any of the CPARMS resources listed in the following table. All of these resources are of type DmColor, and their default values are dynamic—that is, the default values are determined at run time.

Table 8.3 SAS CPARMS Resources

Resource Name	Color and attribute settings	Class Name	Default Color
SAS.cparmBackground	for backgrounds within all primary windows displayed in a SAS session.	CparmBackground	DmWhite
SAS.cparmBanner	for a banner within a window.	CparmForeground	DmBlack
SAS.cparmBorder	for the border of a primary window.	CparmBackground	DmBlack
SAS.cparmByline	for BY lines written to the Output window.	CparmForeground	DmBlue

Resource Name	Color and attribute settings	Class Name	Default Color
SAS.cparmColumn	for text labels for column information. You can use this resource within the SAS editor to identify editing lines and in spreadsheet windows to label spreadsheets.	CparmForeground	DmBlue/ Underline
SAS.cparmCommand	for the command data entry field when menus are disabled.	CparmForeground	DmBlack
SAS.cparmData	for general lines written to the Log window or the Output window.	CparmForeground	DmBlack
SAS.cparmError	for ERROR lines that are written to the Log window or Output window.	CparmForeground	DmRed
SAS.cparmFootnote	for FOOTNOTE lines written to the Output window.	CparmForeground	DmBlue
SAS.cparmForeground	for all text fields within a SAS windowing environment window that can be edited.	CparmBackground	DmBlack
SAS.cparmHeader	for HEADER lines written to the Output window.	CparmForeground	DmBlue
SAS.cparmHelpLink	for links to additional levels of information in the Help system.	CparmForeground	DmGreen/ Underline
SAS.cparmHelpMainTopic	for topic words or phrases in the Help system.	CparmForeground	DmBlack
SAS.cparmHelpSubTopic	for topic words or phrases in the Help system.	CparmForeground	DmBlack
SAS.cparmInfo	for text that is displayed in a window as an aid to the user. For example: Press Enter to continue	CparmForeground	DmBlack
SAS.cparmLabel	for text that precedes a widget. For example, the text Name: in the following example is a label: Name: _____	CparmForeground	DmBlack
SAS.cparmMark	for areas that have been selected for operations such as FIND, CUT, and COPY.	CparmForeground	DmBlack/ DmReverse
SAS.cparmMessage	for the message field.	CparmForeground	DmRed
SAS.cparmNote	for NOTE lines that are written to the Log window or the Output window.	CparmForeground	DmBlue

Resource Name	Color and attribute settings	Class Name	Default Color
SAS.cparmSecondaryBackground	for backgrounds in secondary windows.	CparmForeground	DmGray
SAS.cparmSecondaryBorder	for the border of a secondary window.	CparmForeground	DmBlack
SAS.cparmSource	for SAS source lines that are written to the Log window.	CparmForeground	DmBlack
SAS.cparmText	for text labels for row information. You can use this resource within the SAS editor to identify editing lines and in spreadsheet windows to label spreadsheet rows.	CparmForeground	DmBlue
SAS.cparmTitle	for TITLE lines written to the Output window.	CparmForeground	DmBlue
SAS.cparmWarning	for WARNING lines written to the Log window or the Output window.	CparmForeground	DmGreen

DmColorName can be any one of the following colors:

- DmBLUE
- DmRED
- DmPINK
- DmGREEN
- DmCYAN
- DmYELLOW
- DmWHITE
- DmORANGE
- DmBLACK
- DmMAGENTA
- DmGRAY
- DmBROWN

DmAttrName can be any one of the following attributes:

- DmHIGHLIGHT
- DmUNDERLINE
- DmREVERSE

For example, the following resources specify that all background colors are gray and all foreground colors are black:

```
SAS.cparmBackground: DmGRAY
SAS.cparmForeground: DmBLACK
```

These resources specify that errors should be displayed in red with reverse video, and warnings should be displayed in yellow with reverse video and a bold font:

```
SAS.cparmError: DmRED + DmREVERSE
SAS.cparmWarning: DmHIGHLIGHT + DmYELLOW + DmREVERSE
```

SAS looks for default CPARMS resources in two places:

- If your on-site SAS support personnel entered color and attribute settings in the SASHELP.BASE.SAS.CPARMS catalog entry, then these settings become the default for your site.
- If you saved settings in SASUSER.PROFILE.SAS.CPARMS, then these settings override the settings specified for your site.

Controlling Contrast

During interactive move or stretch operations, such as rubber banding and dragging rectangles in SAS/INSIGHT software, you might find it hard to see the outline of the graphics primitive because of the lack of contrast between the primitive and the background. The XCONTRAST command makes the primitive visible against the background. The rendering performance and the aesthetic appearance of the primitive is compromised for the sake of visibility. You can enter XCONTRAST to act as a toggle, or you can specify XCONTRAST ON or XCONTRAST OFF.

In some color combinations, text fields, buttons, check boxes, and other foreground categories might not be visible. The **SAS.dmsContrastCheck** resource makes these categories legible.

SAS.dmsContrastCheck: True | False

controls whether contrast mapping is applied to non-graphic foreground colors in a SAS window. The default value is False. A value of True specifies that DMS foreground colors will be remapped if necessary to produce a contrast. Some color usage based on graphic operations are not affected by this resource.

Controlling Drop-Down Menus in UNIX Environments

Drop-down menus are controlled by the following resources:

SAS.pmenuOn: True | False

forces the global PMENU state on regardless of the information stored by the WSAVE command. The WSAVE state of an individual window takes precedence over the global state. The default is True. (You can also use the PMENU ON and PMENU OFF commands to turn drop-down menus on and off.)

SAS.usePmenuMnemonics: True | False

specifies whether mnemonics are attached to the drop-down menus for the current SAS session. The default is True.

Customizing Cut and Paste in UNIX Environments

Instructions for Cutting and Pasting Text

For instructions about cutting and pasting text, see “Selecting (Marking) Text in UNIX Environments” on page 153 and “Copying or Cutting and Pasting Selected Text in UNIX Environments” on page 155.

Types of Paste Buffers

There are four SAS paste buffers. Each SAS paste buffer is associated with an X paste buffer:

XPRIMARY

is associated with X primary selection (PRIMARY).

XSCNDARY

is associated with the X secondary selection (SECONDARY).

XCLIPBRD

is associated with the X clipboard selection (CLIPBOARD). This paste buffer allows you to use the MIT X Consortium xclipboard client with SAS.

XTERM

is associated with the paste buffer used by the xterm client. XTERM is the default buffer. DEFAULT is an alias for XTERM. If you copy or cut text into the XTERM buffer, the text is actually copied or cut into all four of the paste buffers. When you paste text from the XTERM buffer, the text is pasted from the XPRIMARY buffer.

XCUT_n

is associated with X cut buffer_n where $0 \leq n \leq 7$.

Selecting a Paste Buffer

If you are not sure which X data exchange protocols your other X clients are using, you should use the XTERM paste buffer. You can specify your default paste buffer with the **SAS.defaultPasteBuffer** resource:

```
SAS.defaultPasteBuffer: XTERM
```

If you know that the X clients in your workstation environment all use the X PRIMARY selections to exchange data, you should use the XPRIMARY paste buffer:

```
SAS.defaultPasteBuffer: XPRIMARY
```

This specification uses both SAS and X resources more efficiently and provides for the on-demand transfer of data between clients.

Sun OpenWindows desktop clients use the CLIPBOARD selection as the basis for their copy-and-paste operations. If you use the SAS XCLIPBRD paste buffer, you can exchange text directly with these clients.

You can also use the SAS XCLIPBRD paste buffer to interact with Motif clients that use the Motif clipboard mechanism for text exchanges. This clipboard mechanism makes it unnecessary to have a dedicated client such as xclipboard. For example, you can use XCLIPBRD to exchange text directly with the Motif xmeditor application when you select the **Cut**, **Copy**, or **Paste** items from the xmeditor **Edit** drop-down menu.

The Motif quick-copy data exchange and Motif clipboard data exchange mechanisms are specific to the Motif interface toolkit and are not currently supported as SAS paste buffers. However some dialog boxes, such as the File Selection dialog box, use Motif interface text widgets. In these dialog boxes, the Motif quick copy and clipboard data exchange mechanisms are available.

Manipulating Text Using a Paste Buffer

If you want SAS to automatically copy selected text into your paste buffer every time you mark a region of text with the mouse, you should also specify your paste buffer name in the **SAS.markPasteBuffer** resource:

```
SAS.markPasteBuffer: XTERM
```

Alternatively, because `DEFAULT` is an alias for `XTERM`, you could specify the following:

```
SAS.markPasteBuffer: DEFAULT
```

The **SAS.markPasteBuffer** definition causes SAS to automatically issue a `STORE` command whenever you select text.

The `STORE` command, as well as the `CUT` and `PASTE` commands, support a `BUFFER=` option that specifies which buffer to use. When these commands are issued from function keys or drop-down menus whose definitions do not include the `BUFFER=` option, if the **SAS.markPasteBuffer** resource is not defined, these commands use `BUFFER=DEFAULT`. If this resource is defined, these commands use `BUFFER=buffer-name`.

You can customize your normal cut, copy, or paste keys to issue any of these commands with the `BUFFER=` option. For example, you can override the **SAS.keyboardTranslations** definition for the `osfCopy` and `osfPaste` keys with the following specifications:

```
SAS.keyboardTranslations: #override \
<Key>osfCopy: sas-do-command(\"STORE BUFFER=XCLIPBRD\") \n\
<Key>osfPaste: sas-do-command(\"PASTE BUFFER=XCLIPBRD\")
```

For more information about customizing keys, see “Customizing Key Definitions in UNIX Environments” on page 184.

Notes about Preserving Text and Attribute Information

When you cut or copy and paste text between SAS sessions using the `XTERM`, `XPRIMARY`, or `XSCNDARY` paste buffers, the color and attribute information is preserved. However, if you copy and paste the same text into an xterm window while using the vi editor, the color and attribute information is lost. If you change the definition for **SAS.defaultPasteBuffer** and **SAS.markPasteBuffer** to `XCUT0`, then you will not retain the text and color attributes when you copy and paste text between two SAS sessions.

When you use the `xclipboard` client, SAS text attributes are not preserved in exchanges made between SAS sessions. However, when you use the `XCLIPBRD` paste buffer without a clipboard manager such as the `xclipboard` client, SAS text attributes are preserved in exchanges between SAS sessions.

Customizing Session Workspace, Session Gravity, and Window Sizes in UNIX Environments

SAS uses the following resources to determine the size of the session workspace, the gravity of the workspace, and the size of the windows. The default values for these resources are listed in Table 8.4 on page 209.

SAS.awsResizePolicy: `grow` | `fixed`

controls the policy for resizing AWS windows as interior windows are added and removed. The following values are valid:

<code>grow</code>	the AWS window will attempt to grow any time an interior window is grown or moved, in order to show all interior windows, but it will not shrink to remove dead areas.
-------------------	--

fixed the AWS window will attempt to size itself to the size of the first interior window and will not attempt any further size changes.

SAS.maxWindowHeight: units

specifies the number of units for the maximum height of a window. The unit is specified by the **SAS.windowUnitType** resource.

SAS.maxWindowWidth: units

specifies the number of units for the maximum width of a window. The unit is specified by the **SAS.windowUnitType** resource.

SAS.noAWS: True | False

controls whether each of your application's windows appears in its own native window rather than in an application work space (AWS). The default is True; each application runs in its own native window.

SAS.scrollBarSize: pixels

specifies the default size of the scroll bar in pixels.

SAS.sessionGravity: value

controls the region of the screen where SAS will attempt to place its windows. This resource might be ignored by some window manager configurations. Possible values include the following:

CenterGravity

EastGravity

WestGravity

SouthGravity

NorthGravity

SouthEastGravity

NorthEastGravity

SouthWestGravity

NorthWestGravity

SAS.sessionGravityXOffset: offset

specifies an x offset to be added when SAS attempts to place a window in the gravity region.

SAS.sessionGravityYOffset: offset

specifies a y offset to be added when SAS attempts to place a window in the gravity region.

SAS.windowHeight: units

specifies the number of units for the default height of a window. The unit is specified by the **SAS.windowUnitType** resource.

SAS.windowUnitType: character | pixel | percentage

specifies the unit type for **SAS.windowWidth**, **SAS.windowHeight**, **SAS.maxWindowWidth**, and **SAS.maxWindowHeight**. Possible values include the following:

character

units specify the number of rows and columns.

pixel

units specify the number of pixels.

percentage

units specify the percentage of the screen.

SAS.windowWidth: *units*

specifies the number of units for the default width of a window. The unit is specified by the **SAS.windowUnitType** resource.

Specifying User-Defined Icons in UNIX Environments

Why Specify User-Defined Icons?

You can add your own icons to those icons that are supplied with SAS. For example, if you want to use your own color icons in the toolbox, define the **SAS.colorUiconPath**, **SAS.colorUiconCount**, and **SAS.sasUiconx** resources. Then, when you are defining tools in the Tool Editor, the Tool Editor will include your icons in the display of icons that you can choose for each tool.

How SAS Locates a User-Defined Icon

The resource name that is used to locate the icon bitmap filename for user icon number *x* is **SAS.sasUiconx**. For example, the following resource defines the filename **myicon** for the user icon 1:

```
SAS.sasUicon1: myicon
```

If the resource name is not defined, SAS generates a filename of the form **sasuinnn.xbm** or **sasuinnn.xpm**. The path elements from the **SAS.uiconPath** or **SAS.colorUiconpath** resource are searched in sequence until the icon file is found or until the search path is exhausted.

For example, the following set of X resources defines a collection of color icons.

```
SAS.colorUiconPath: /users/jackaroe/pixmaps/
SAS.colorUiconCount: 7
SAS.sasUicon1: adsetup
SAS.sasUicon2: adverse
SAS.sasUicon3: altmenu
SAS.sasUicon4: batch
SAS.sasUicon5: is
SAS.sasUicon6: patgrps
SAS.sasUicon7: pctchg
```

The Motif interface will search for icon **sasUicon1** in a file named **/users/jackaroe/pixmaps/adsetup.xpm**.

X Resources for Specifying User-Defined Icons

SAS uses the following resources to determine the number of user-defined icons that are available and their location.

SAS.colorUiconPath: *search-path*

specifies the file search path for locating user-defined color icon files. This string resource specifies the directory paths to be searched for an icon file. These files should be in X Pixmap (xpm) format. Use a comma to separate individual directory pathnames. For example, the following string first searches for icon files

in the `/usr/lib/X11/pixmaps` directory and then in the `/usr/lib/X11/pixmaps/SAS` directory:

```
SAS.colorUiconPath : /usr/lib/X11/pixmaps, \
/usr/lib/X11/pixmaps/SAS
```

SAS.colorUiconCount: *num-icons*

specifies the number of user-defined color icons that are available for SAS to use.

SAS.uiconCount: *num-icons*

specifies the number of user-defined icons that are available for use in the SAS session.

SAS.uiconPath: *search-path*

specifies the file search path for locating user-defined icon bitmap files. This string resource specifies the directory paths to be searched for an icon file. These files should be in X Bitmap (xbm) format. Use a comma to separate individual directory pathnames. For example, the following string will first search for bitmap files in the `/usr/lib/X11/bitmaps` directory and then in the `/usr/lib/X11/bitmaps/SAS` directory:

```
SAS.uiconPath : /usr/lib/X11/bitmaps, \
/usr/lib/X11/bitmaps/SAS
```

SAS.sasUiconx: *name*

associates a value with the filename of an X Bitmap or Pixmap file. *x* is a number assigned to the file. A file extension of `.xbm` or `.xpm` is automatically supplied.

Miscellaneous Resources in UNIX Environments

You can also customize the following resources:

SAS.altVisualId: *ID*

specifies a visual type ID.

SAS.autoSaveInterval: *minutes*

specifies how often (in number of minutes) that the data from the Program Editor window should be saved.

SAS.autoSaveOn: *True | False*

specifies that data from the Program Editor window should be saved to a file at intervals specified by the **SAS.autoSaveInterval** resource.

SAS.confirmSASExit: *True | False*

controls whether SAS displays the Exit dialog box when you enter the DLGENDR command or select **File** ► **Exit**. The default is True.

SAS.defaultCommandWindow: *True | False*

specifies whether the command window is invoked when you start your SAS session. The default is True.

SAS.directory: *directory-pathname*

specifies the directory that you want when you first invoke the Open dialog box. By default, the Open dialog box uses the current directory.

SAS.helpBrowser: *pathname*

specifies the pathname of the World Wide Web browser to use for viewing the online Help or when the WBROWSE command is issued. The default browser is Netscape.

SAS.htmlUsePassword: True | False

specifies whether SAS prompts you to enter your password before sending HTML files to your browser. The default value is True.

SAS.insertModeOn: True | False

controls the editing mode in SAS editor windows. The default is False (overtyping).

SAS.noDoCommandRecall: True | False

controls whether SAS commands that are submitted through the `sas-do-command()` action routine are recorded in the command recall buffer. The default value of True causes commands to be omitted from the command recall buffer; a value of False causes them to be recorded.

SAS.pattern: *default-pattern*

specifies the default pattern that you want to be used as the file filter when you first invoke the Open and Import Image dialog boxes. This pattern is displayed in the text field at the top of the dialog box. By default, the dialog box uses the first filter in the File type list. The pattern resource has no effect on the File type field.

SAS.selectTimeout: *seconds*

specifies the X Toolkit selection conversion timeout value in units of seconds. This timeout value determines the amount of time that SAS will wait for a request to convert an X Toolkit selection to complete. The default value should be adequate in most cases.

SAS.startupLogo: *xpm-filename* | None | ""

specifies the XPM file that you want SAS to display when it is initialized. If the string is empty, SAS uses the default logo.

SAS.startSessionManager: True | False

specifies whether SAS automatically starts the SAS Session Manager when a new SAS session is started. Using your own host editor with SAS requires that the SAS Session Manager be running. The default is True.

SAS.suppressMenuIcons: True | False

specifies whether SAS displays any menu icons other than the check box and toggle button icons in cascade or pop-up menus. Suppressing the icons reduces memory usage and improves how quickly the menus display on slower X servers. The default is False.

SAS.suppressTutorialDialog: True | False

specifies whether SAS displays the Getting Started Tutorial dialog box at the start of your SAS session. True suppresses the dialog box. You might want to suppress this dialog box if you have previously used SAS. The default is False.

SAS.useNativeXmTextTranslations: True | False

specifies whether any XmText widget translations are inherited by all instances of the Text, Combo Box, and Spin Box widgets used by the SAS X Motif user interface. When the value is False, the SAS keys windows translations supersede any user or system-supplied XmText translations. The default value is True.

The following example shows SAS XmText translations:

```
SAS*XmText*translations: #override \n\  
Ctrl<Key>e:end-of-line()\n\  
Ctrl<Key>u:delete-to-start-of-line()\n\  

```

```

Ctrl<Key>k:delete-to-end-of-line()\n\
Ctrl<Key>f:forward-character()\n\
Ctrl<Key>b:backward-character()\n\
Ctrl<Key>a:beginning-of-line()\n\
Ctrl<Key>c:copy-clipboard()\n\
Ctrl<Key>v:paste-clipboard()\n\

```

SAS.wsaveAllExit: True | False

specifies whether SAS should issue the WSAVE ALL command when you end your session. This command saves the global settings, such as window color and window position, that are in effect for all windows that are currently open. The default is False.

Note: For the WSAVE command to work, your window manager must support explicit window placement. See the documentation for your window manager to determine how to configure your window manager. For example, if you are running Exceed, open the Screen Definition Settings dialog box and deselect **Cascade Windows**. Δ

Summary of X Resources for SAS in UNIX Environments

The following table lists the instance and class names, type, and default values for many of the SAS resources. See the following tables for additional resources of specific types:


- “SAS Font Alias Resources,” Table 8.1 on page 195
- “Foreground and Background Color Resources,” Table 8.2 on page 198
- “SAS CPARM Resources,” Table 8.3 on page 199

Table 8.4 SAS Resources

Resource Name	Class Name	Type	Default
SAS.altVisualId	AltVisualId	Integer	NULL
SAS.autoComplete	AutoComplete	Boolean	True
SAS.autoSaveInterval	AutoSaveInterval	Integer	10
SAS.autoSaveOn	AutoSaveOn	Boolean	True
SAS.awsResizePolicy	AWSResizePolicy	String	grow
SAS.colorUiconCount	UiconCount	Integer	0
SAS.colorUiconPath	UiconPath	String	NULL
SAS.commandsSaved	CommandsSaved	Integer	25
SAS.confirmSASExit	ConfirmSASExit	Boolean	True
SAS.defaultCommandWindow	DefaultCommandWindow	Boolean	True
SAS.defaultPasteBuffer	DefaultPasteBuffer	String	XTERM
SAS.defaultToolBox	DefaultToolBox	Boolean	True
SAS.directory	Directory	String	NULL
SAS.dmsContrastCheck	DmsContrastCheck	Boolean	False
SAS.DMSDBFont	Font	String	<i>dynamic</i>


Resource Name	Class Name	Type	Default
SAS.DMSDBboldFont	Font	String	<i>dynamic</i>
SAS.DMSboldFont	Font	String	<i>dynamic</i>
SAS.DMSFont	Font	String	<i>dynamic</i>
SAS.DMSfontPattern	DMSFontPattern	String	_*_*_*-r-*_*-_*_*-m-*_iso8859-1
SAS.fontPattern	FontPattern	String	*
SAS.helpBrowser	HelpBrowser	String	netscape
SAS.htmlUsePassword	HtmlUsePassword	Boolean	True
SAS.insertModeOn	InsertModeOn	Boolean	False
SAS.isToolBoxPersistent	IsToolBoxPersistent	Boolean	True
SAS.keyboardTranslations	KeyboardTranslations	Translation	<i>dynamic</i>
SAS.keysWindowLabels	KeysWindowLabels	String	<i>dynamic</i>
SAS.markPasteBuffer	MarkPasteBuffer	String	XTERM
SAS.maxWindowHeight	WindowHeight	Dimension	95
SAS.maxWindowWidth	WindowWidth	Dimension	95
SAS.noAWS	NoAWS	Boolean	True
SAS.noDoCommandRecall	NoDoCommandRecall	Boolean	True
SAS.pattern	Pattern	String	NULL
SAS.pmenuOn	PmenuOn	Boolean	True
SAS.sasUicon	SasUicon	String	NULL
SAS.scrollBarSize	ScrollBarSize	Dimension	17
SAS.selectTimeout	SelectTimeout	Integer	60
SAS.sessionGravity	SASGravity	String	NorthWestGravity
SAS.sessionGravityXOffset	SASGravityOffset	Integer	0
SAS.sessionGravityYOffset	SASGravityOffset	Integer	0
SAS.startSessionManager	StartSessionManager	Boolean	True
SAS.startupLogo	StartUpLogo	String	NULL
SAS.suppressMenuIcons	SuppressMenuIcons	Boolean	False
SAS.suppressTutorialDialog	SuppressTutorialDialog	Boolean	False
SAS.systemFont	SystemFont	String	“-adobe-helvetica-medium-r-normal-12_*_*_*_*_*”
SAS.toolBoxAlwaysOnTop	ToolBoxAlwaysOnTop	Boolean	True
SAS.toolBoxTipDelay	ToolBoxTipDelay	Integer	750

Resource Name	Class Name	Type	Default
SAS.uiconCount	UiconCount	Integer	0
SAS.uiconPath	UiconPath	String	NULL
SAS.useCommandToolBoxCombo	UseCommandToolBoxCombo	Boolean	True
SAS.useLargeToolBox	UseLargeToolBox	Boolean	False
SAS.useNativeXmTextTranslations	UseNativeXmTextTranslations	Boolean	False
SAS.usePmenuMnemonics	UsePmenuMnemonics	Boolean	True
SAS.useShowHideDecorations	UseShowHideDecorations	Boolean	False
SAS.useToolBoxTips	UseToolBoxTips	Boolean	True
SAS.wsaveAllExit	WsaveAllExit	Boolean	False
SAS.windowHeight	WindowHeight	Dimension	50
SAS.windowWidth	WindowWidth	Dimension	67
SAS.windowUnitType	WindowUnitType	String	percentage



Application Considerations

Chapter 9 **Data Representation** 215



CHAPTER

9

Data Representation

Numeric Variable Length and Precision in UNIX Environments 215

Missing Values in UNIX Environments 216

Reading and Writing Binary Data in UNIX Environments 216

Numeric Variable Length and Precision in UNIX Environments

The default length of numeric variables in SAS data sets is 8 bytes. (You can control the length of SAS numeric variables with the `LENGTH` or `ATTRIB` statements in the `DATA` step.)

The issue of numeric precision affects the return values of almost all SAS math functions and many numeric values returned from SAS procedures. Numeric values in SAS for UNIX are represented as IEEE double-precision floating-point numbers. The decimal precision of a full 8-byte number is effectively 15 decimal digits.

The following table specifies the significant digits and largest integer that can be stored exactly in SAS numeric variables.

Table 9.1 Significant Digits and Largest Integer by Length for SAS Variables under UNIX

Length in Bytes	Significant Digits Retained	Largest Integer Represented Exactly
3	3	8,192
4	6	2,097,152
5	8	536,870,912
6	11	137,438,953,472
7	13	35,184,372,088,832
8	15	9,007,199,254,740,992

When you are specifying variable lengths, keep in mind that the length of a variable affects both the amount of disk space used and the number of I/O operations required to read and write the data set. See *SAS Language Reference: Dictionary* for more information about specifying variable lengths.

If you know that the value of a numeric variable will be an integer between -8192 and 8192 inclusive, you can use a length of 3 to store the number and thus save space in your data set. For example:

```
data mydata;
  length num 3;
```

```

    ...more SAS statements...
run;

```

Numeric *dummy variables* (variables whose only purpose is to hold 0 or 1) can be stored in a variable whose length is 3 bytes.

CAUTION:

Use the LENGTH statement to reduce length only for variables whose values are always integers. Fractional numbers lose precision if they are truncated. In addition, you must ensure that the values of your variable will always be represented exactly in the number of bytes that you specify. You can do this programmatically in a DATA step with the TRUNC function. No warnings or errors are issued when the length that you specify in the LENGTH statement results in the truncation of data. Δ

For more information about specifying variable lengths and optimizing system performance, see *SAS Language Reference: Concepts*.

Missing Values in UNIX Environments


In SAS on UNIX, missing values are represented by IEEE Not-a-Number values. An IEEE Not-a-Number value is an IEEE floating-point bit pattern that represents something other than a valid numeric value. These numbers are not computationally derivable.

Reading and Writing Binary Data in UNIX Environments

Different computers store numeric binary data in different forms. For more information about compatible computer types, see “Compatible Computer Types in UNIX Environments” on page 42. If you try to move binary data in flat files across systems that are incompatible, problems will occur. A safer way to move data is by using SAS data sets.


SAS provides several sets of informats and formats for handling binary data. Some of these informats and formats are host dependent. For example, the *IBw.d*, *PDw.d*, *PIBw.d*, and *RBw.d* informats and formats read and write data in native mode. That is, they use the byte-ordering system that is standard for the computer. If you create a file using the *IBw.d* format on a 64-bit HP-UX host and then use the *IBw.d* informat to read the same file on a 32-bit Linux host, you will get unpredictable results.

For more information about all of the informats and formats, see *SAS Language Reference: Dictionary*.



Host-Specific Features of the SAS Language

<i>Chapter 10</i>	Commands under UNIX	<i>219</i>
<i>Chapter 11</i>	Data Set Options under UNIX	<i>241</i>
<i>Chapter 12</i>	Formats under UNIX	<i>251</i>
<i>Chapter 13</i>	Functions and CALL Routines under UNIX	<i>257</i>
<i>Chapter 14</i>	Informats under UNIX	<i>279</i>
<i>Chapter 15</i>	Macro Facility under UNIX	<i>285</i>
<i>Chapter 16</i>	Procedures under UNIX	<i>291</i>
<i>Chapter 17</i>	Statements under UNIX	<i>315</i>
<i>Chapter 18</i>	System Options under UNIX	<i>341</i>



CHAPTER

10

Commands under UNIX

<i>SAS Commands under UNIX</i>	220
<i>AUTOSCROLL Command</i>	220
<i>CAPS Command</i>	221
<i>COLOR Command</i>	221
<i>DLGABOUT Command</i>	221
<i>DLGCDIR Command</i>	222
<i>DLGENDR Command</i>	222
<i>DLGFIND Command</i>	223
<i>DLGFONT Command</i>	223
<i>DLGOPEN Command</i>	224
<i>DLGPREF Command</i>	224
<i>DLGREPLACE Command</i>	225
<i>DLGSAVE Command</i>	225
<i>DLGSCRDUMP Command</i>	226
<i>DLGSMAIL Command</i>	227
<i>FILE Command</i>	227
<i>FILL Command</i>	229
<i>FONTLIST Command</i>	229
<i>GSUBMIT Command</i>	230
<i>HOME Command</i>	230
<i>HOSTEDIT Command</i>	231
<i>INCLUDE Command</i>	231
<i>SETAUTOSAVE Command</i>	233
<i>SETDMSFONT Command</i>	233
<i>TOOLCLOSE Command</i>	234
<i>TOOLEDIT Command</i>	234
<i>TOOLLARGE Command</i>	235
<i>TOOLLOAD Command</i>	235
<i>TOOLTIPS Command</i>	236
<i>WBROWSE Command</i>	236
<i>WCOPY Command</i>	237
<i>WCUT Command</i>	237
<i>WDEF Command</i>	238
<i>WPASTE Command</i>	238
<i>WUNDO Command</i>	239
<i>X Command</i>	239
<i>XSYNC Command</i>	240

SAS Commands under UNIX

This section describes commands that you can enter on the command line in the windowing environment of SAS. The commands that are described here have behavior or syntax that is specific to UNIX environments. Each command description includes a brief “UNIX specifics” section that explains which aspect of the command is specific to UNIX. If the information under “UNIX specifics” says “all,” then the command applies to the UNIX operating environment and is described only in this document.

The following commands are not supported in UNIX environments:

CASCADE
DCALC
ICON
PCLEAR
RESIZE
SCROLLBAR
SMARK
TILE
WGROW
WMOVE
WSHRINK
ZOOM

AUTOSCROLL Command

Specifies how often the Log and Output windows scroll to display output.

UNIX specifics: valid arguments and default values

Syntax

AUTOSCROLL *<n>*

n specifies the number of lines that the window should scroll when it receives a line of data that cannot fit.

Details

The AUTOSCROLL command controls the scrolling of lines as they are written to the Log and Output windows. The default value for AUTOSCROLL in the Log and Output windows is 1. Processing is slower when AUTOSCROLL displays one line at a time. To expedite processing, you can specify a greater AUTOSCROLL value in your autoexec.sas file. Specifying a value of 0 optimizes processing and results in the fastest scrolling (similar to jump scrolling in xterm windows). To add the AUTOSCROLL command to your autoexec.sas file, you must use the DM command. The following example maximizes scrolling in both the Log and Output windows:

```
dm 'output; autoscroll 0; log; autoscroll 0; pgm;';
```

CAPS Command

Changes the default case of text.

UNIX specifics: all

Syntax

CAPS <ON | OFF>

COLOR Command

Specifies the color and highlighting of selected portions of a window.

UNIX specifics: valid field types and attributes

Syntax

COLOR *field-type color* | NEXT *<highlight>*

Details

Under UNIX, you cannot use the COLOR command to change the colors in these field types: BORDER, MENU, MENUBORDER, SCROLLBAR, or TITLE. Also, the H (HIGHLIGHT) and B (BLINK) attributes are not supported. For more information about the COLOR command, see the online Help for the Program Editor window.

See Also

Online Help for the Program Editor window
“Syntax of the COLOR Command” on page 197

DLGABOUT Command

Opens the About SAS dialog box.

UNIX specifics: all

Syntax

DLGABOUT

Details

The About SAS dialog box displays information such as the release of SAS that you are running, your site number, the operating system, the version of Motif that you are using, and the color information from your PC.

To access this dialog box from the menu, select **Help ► About SAS 9**.

DLGCDIR Command

Opens the Change Working Directory dialog box.

UNIX specifics: all

Syntax

DLGCDIR

Details

The Change Working Directory dialog box enables you to select a new working directory. To access this dialog box from the menu, select **Tools ► Options ► Change Directory**.

DLGENDR Command

Opens the Exit dialog box.

UNIX specifics: all

Syntax

DLGENDR

Details

The Exit dialog box prompts you to confirm that you want to exit SAS. If you choose **OK**, the SAS session ends. If you have set the **SAS.confirmSASExit** resource to **False**, this command becomes equivalent to the **BYE** command. To access this dialog box from the menu, select **File ► Exit**.

See Also

“Miscellaneous Resources in UNIX Environments” on page 207

DLGFIND Command

Opens the Find dialog box.

UNIX specifics: all

Syntax

DLGFIND

Details

The Find dialog box enables you to search for text strings. To access this dialog box from the menu, select **Edit ► Find**.

See Also

“DLGREPLACE Command” on page 225

DLGFONT Command

Opens the Fonts dialog box.

UNIX specifics: all

Syntax

DLGFONT

Details

The Font dialog box enables you to dynamically change the SAS font. To access this dialog box from the menu, select **Tools ► Options ► Fonts**.

See Also

“Customizing Fonts in UNIX Environments” on page 192

“SETDMSFONT Command” on page 233

DLGOPEN Command

Opens the Open dialog box.

UNIX specifics: all

Syntax

DLGOPEN <FILTERS=*filters*' <IMPORT> <SUBMIT|NOSUBMIT> <VERIFY>>

FILTERS=*filters*'

specifies one or more file filters to use as search criteria when displaying files. For example, the following command displays all files in the current directory that have a **.sas** extension and adds ***.txt** to the **File type** box in the dialog box:

```
DLGOPEN FILTERS="*.sas *.txt"
```

You can specify multiple filters; they all appear in the box. If you do not specify any filters, the dialog box displays a default list. See the description of the **SAS.pattern** resource in “Miscellaneous Resources in UNIX Environments” on page 207 for information about specifying a default file pattern.

IMPORT

invokes the Import Image dialog box, which enables you to import graphic files to SAS/GRAPH applications.

SUBMIT|NOSUBMIT

specifies whether the SUBMIT command is pushed after the file is opened.

VERIFY

checks whether the DLGOPEN command is appropriate for the active window.

Details

The Open and Import Image dialog boxes enable you to select a file to read into the active window. If the active window is a SAS/GRAPH window, then the Import Image dialog box is displayed; otherwise, the Open dialog box is displayed. To access these dialog boxes from the menu, select **File ► Open** or **File ► Import Image**.

See Also

Information about image extensions in online documentation for SAS/GRAPH

DLGPREF Command

Opens the Preferences dialog box.

UNIX specifics: all

Syntax

DLGPREF

Details

The Preferences dialog box enables you to dynamically change certain resource settings. To access this dialog box from the menu, select **Tools ▶ Options ▶ Preferences**.

See Also

“Modifying X Resources Through the Preferences Dialog Box” on page 167

DLGREPLACE Command

Opens the Change dialog box.

UNIX specifics: all

Syntax

DLGREPLACE

Details

The Change dialog box enables you to search for and replace text strings. To access this dialog box from the menu, select **Edit ▶ Replace**.

See Also

“DLGFIND Command” on page 223

DLGSAVE Command

Opens the Save As or Export dialog box.

UNIX specifics: all

Syntax

DLGSAVE <FILTERS=*filters*' <EXPORT> <VERIFY>>

FILTERS='filters'

specifies one or more file filters to use as search criteria when displaying files. For example, the following command displays all files in the current directory that have a **.sas** extension and adds ***.txt** to the **File type** box in the dialog box:

```
DLGSAVE FILTERS="*.sas *.txt"
```

You can specify multiple filters; they all appear in the dialog box. If you do not specify any filters, the dialog box displays a default list.

EXPORT

invokes the Export dialog box, enabling you to export graphic files in your SAS session.

VERIFY

checks whether the DLGSAVE command is appropriate for the active window.

Details

To access this dialog box from the menu, select **File ► Save as** or **File ► Export as Image**.

See Also

Information about image extensions in online documentation for SAS/GRAPH

DLGSCRDUMP Command

Saves the active SAS/GRAPH window as an image file using the filename and file type that you specify.

UNIX specifics: all

Syntax

DLGSCRDUMP <'filename.ext' 'FORMAT=file-type'>

Details

DLGSCRDUMP saves the active GRAPH window as an image file by using the filename and file type that you specify. If you do not specify arguments, DLGSCRDUMP opens the Export dialog box and enables you to choose a filename and file type. You can save displays in any image format supported by SAS/GRAPH with image extensions. If your

site has not licensed SAS/GRAPH with image extensions, then displays can be saved only as XPM files.

See Also

Information about image extensions in online documentation for SAS/GRAPH

DLGSMail Command

Opens the Send Mail dialog box.

UNIX specifics: all

Syntax

DLGSMail

Details

The Send Mail dialog box lets you send electronic mail while working in SAS. To access this dialog box from the menu, select **File ► Send mail**.

See Also

“Sending Electronic Mail Using the FILENAME Statement (EMAIL)” on page 82 and
 “Sending Mail from within Your SAS Session in UNIX Environments” on page 158
 “EMAILSYS System Option” on page 366

FILE Command

Writes the contents of the current window to an external file.

UNIX specifics: valid values for *encoding-value* and *host-options*

Syntax

FILE <*file-specification*> <ENCODING=*encoding-value*'><*portable-options*>
 <*host-options*>

file-specification

can be any of the following:

single filename	SAS writes the file in the current directory. If you enclose the filename in quotation marks, SAS uses the filename exactly as you specify it. If you do not enclose the filename in quotation marks and if you do not specify a filename extension, SAS uses .sas, .log, or .lst, depending on whether you issue the command from the Program Editor, Log, or Output window.
entire pathname	SAS does not assume any filename extensions, even if you do not enclose the pathname in quotation marks.
fileref	SAS specifies a fileref to assign to an external file.

ENCODING='encoding-value'

specifies the encoding to use when writing to the output file. The value for ENCODING= indicates that the output file has a different encoding from the current session encoding.

When you write data to the output file, SAS transcodes the data from the session encoding to the specified encoding.

For valid encoding values, see “Encoding Values in SAS Language Elements” in the *SAS National Language Support (NLS): Reference Guide*.

portable-options

are options for the FILE command that are valid in all operating environments. See *SAS Language Reference: Dictionary* for information about these options.

host-options

are specific to UNIX environments. These options can be any of the following:

BLKSIZE=

BLK=

specifies the number of bytes that are physically written in one I/O operation. The default is 8K. The maximum is 1G–1.

LRECL=

specifies the logical record length. Its meaning depends on the record format in effect (RECFM). The default is 256. The maximum length is 1G.

- If RECFM=F, then the value for the LRECL= option determines the length of each output record. The output record is truncated or padded with spaces to fit the specified size.
- If RECFM=N, then the value for the LRECL= option must be at least 256.
- If RECFM=V, then the value for the LRECL= option determines the maximum record length. Records that are longer than the specified length are truncated.

NEW|OLD

indicates that a new file is to be opened for output. If the file already exists, then it is deleted and re-created. This is the default action.

RECFM=

specifies the record format. Values for the RECFM= option are listed below:

D	default format (same as variable).
F	fixed format. That is, each record has the same length. Do not use RECFM=F for external files that contain carriage-control characters.
N	binary format. The file consists of a stream of bytes with no record boundaries.
P	print format. SAS writes carriage-control characters.

V	variable format. Each record ends with a newline character.
S370V	variable S370 record format (V).
S370VB	variable block S370 record format (VB).
S370VBS	variable block with spanned records S370 record format (VBS).
UNBUF	tells SAS not to perform buffered writes to the file on any subsequent FILE or INCLUDE command. This option applies especially when you are writing to a data collection device.

Details

If you do not enter a file specification, SAS uses the filename from the previous FILE or INCLUDE command. In this case, SAS first asks if you want to overwrite the file. If you have not issued any FILE or INCLUDE commands, you receive an error message that indicates that no default file exists.

FILL Command

Specifies the fill character.

UNIX specifics: default character

Syntax

FILL <fill-character>

fill-character

specifies the character to be used to fill out a line.

Details

Under UNIX, the default fill character is an underscore (_).

FONTLIST Command

Opens the Select Font window, which lists available software fonts.

UNIX specifics: all

Syntax

FONTLIST

Details

The FONTLIST command opens windows that list all of the software fonts that are available in your operating environment. This feature is useful if you want to choose a font to use in a SAS program, typically with a FONT= or FTEXT= option.

Issuing the FONTLIST command from the SAS command line opens the Select Font window, which contains two buttons, **Copy** and **System**. Clicking **System** opens the Fonts window, from which you can select and preview all available system fonts. After you select the desired font and font attributes, click **OK**. The Select Font window reopens with your selected font name displayed. Clicking **Copy** places the font name in the copy buffer so that you can paste the selected font name into your SAS program.

GSUBMIT Command

Submits SAS code stored in a paste buffer.

UNIX specifics: valid buffer names

Syntax

GSUBMIT BUF=*buffername* | “*statement1;statementN...*”

buffername

can be XPRIMARY, XSCNDARY, XCLIPBRD, XTERM, or XCUT n where $0 \leq n \leq 7$. See “Customizing Cut and Paste in UNIX Environments” on page 202 for more information.

statementN

can be any SAS statement.

HOME Command

Toggles the cursor position between the current position and the command line.

UNIX specifics: keyboard equivalent

Syntax

HOME

Details

Keyboards vary among the different UNIX operating environments. To determine which key is assigned to the HOME command, look in the Keys window. To open the Keys window, issue the KEYS command.

See Also

Online Help for the Program Editor window
 “Customizing Key Definitions in UNIX Environments” on page 184

HOSTEDIT Command

Starts the UNIX editor, specified by the EDITCMD system option, in the current window.

UNIX specifics: all

Syntax

HOSTEDIT

Details

When you issue the HOSTEDIT command from a SAS text editor window, the contents of the buffer for that window are written to a temporary file in the `/tmp` directory. A command invoking the host editor that was specified in the EDITCMD system option is passed to the SAS Session Manager. The SAS Session Manager issues the command to the operating environment to invoke the editor for the temporary file.

The X display used with the HOSTEDIT command is the same one used with your SAS session.

HED is an alias for the HOSTEDIT command.

See Also

“Configuring SAS for Host Editor Support in UNIX Environments” on page 160
 “EDITCMD System Option” on page 365

INCLUDE Command

Copies the entire contents of an external file into the current window.

UNIX specifics: valid values for *encoding-value* and *portable-options*

Syntax

INCLUDE *<file-specification>* *<ENCODING='encoding-value'>* *<portable-options>*
<host-options>

file-specification

can be any of the following:

- a single filename. SAS searches for the file in the current directory. If you enclose the filename in quotation marks, then SAS uses the filename exactly as you specify it. If you do not enclose the filename in quotation marks and if you do not specify a filename extension, then SAS searches for *file-specification.sas*.
- an entire pathname. SAS does not assume any filename extensions, even if you do not enclose the pathname in quotation marks.
- a fileref.

ENCODING=*encoding-value*

specifies the encoding to use when reading from the external file. The value for ENCODING= indicates that the external file has a different encoding from the current session encoding.

When you read data from an external file, SAS transcodes the data from the specified encoding to the session encoding.

For valid encoding values, see “Encoding Values in SAS Language Elements” in *SAS National Language Support (NLS): Reference Guide*.

portable-options

are options for the INCLUDE command that are valid in all operating environments. See the *SAS Language Reference: Dictionary* for information about these options.

host-options

are specific to UNIX environments. These options can be any of the following:

BLKSIZE=

BLK=

specifies the number of bytes that are physically read in one I/O operation. The default is 8K. The maximum is 1G–1.

LRECL=

specifies the logical record length. Its meaning depends on the record format in effect (RECFM). The default is 256. The maximum length is 1G.

- If RECFM=F, then the value for the LRECL= option determines the number of bytes to be read as one record.
- If RECFM=N, then the value for the LRECL= option must be at least 256.
- If RECFM=V, then the value for the LRECL= option determines the maximum record length. Records that are longer than the specified length are truncated.

RECFM=

specifies the record format. Values for the RECFM= option are

D	default format (same as variable).
F	fixed format. That is, each record has the same length.
N	binary format. The file consists of a stream of bytes with no record boundaries.
P	print format.
V	variable format. Each record ends with a newline character.

Details

If you do not enter a file specification, then SAS uses the filename from the previous FILE or INCLUDE command. In this case, SAS first asks if you want to overwrite the

file. If you have not issued any FILE or INCLUDE commands, then you receive an error message to indicate that no default file exists.

SETAUTOSAVE Command

Turns autosave on and off.

UNIX specifics: all

Syntax

SETAUTOSAVE <ON|OFF>

Details

The SETAUTOSAVE command turns autosave on or off for the Program Editor. However, the value set for autosave in the Preferences dialog box has precedence. To open the Preferences dialog box, select **Tools ► Options ► Preferences**. Autosave is controlled by the **Backup Documents** check box on the **DMS** tab. On this tab, there is also a field in which you can specify the interval for these backups.

If you turn autosave on using the SETAUTOSAVE command and the **Backup Documents** check box is selected, then SAS automatically saves the contents of the Program Editor into a file named **pgm.asv** in your current directory at the interval specified on the **DMS** tab.

If you issue this command but do not specify ON or OFF, SAS displays the current autosave setting.

See Also

“Modifying the DMS Settings” on page 168

“Miscellaneous Resources in UNIX Environments” on page 207

SETDMSFONT Command

Specifies a windowing environment font for the current session.

UNIX specifics: all

Syntax

SETDMSFONT *“font-specification”*

font-specification

specifies an XLFD (X Logical Font Description) pattern that you want SAS to use in order to determine the windowing environment font.

Details

Most fonts in the X Window System are associated with an XLFD, which contains a number of different fields that are delimited by a dash (-) character. The fields in the XLFD indicate properties such as the font family name, weight, size, resolution, and whether the font is proportional or monospaced. See your X Window system documentation for more information about the XLFD and font names that are used with the X Window System.

See Also

“DLGFONT Command” on page 223

TOOLCLOSE Command

Closes the toolbox.

UNIX specifics: all

Syntax

TOOLCLOSE

Details

The TOOLCLOSE command closes the toolbox.

See Also

“TOOLLOAD Command” on page 235

TOOLEDIT Command

Opens the Tool Editor on the specified toolbox.

UNIX specifics: all

Syntax

TOOLEEDIT <library.catalog.entry>

Details

If you do not specify an entry name, the Tool Editor edits the toolbox for the active window.

TOOLLARGE Command

Toggles the size of the SAS ToolBox window.

UNIX specifics: all

Syntax

TOOLLARGE <ON|OFF>

ON

sets the size of the icons in the SAS ToolBox to 48x48.

OFF

sets the size of the icons in the SAS ToolBox to 24x24.

Details

If you do not specify ON or OFF, the TOOLLARGE command toggles the size of the SAS ToolBox. The size of the SAS ToolBox changes for your current session only; the new size is not saved.

You can also use the menu to change the size of the SAS ToolBox through the Preferences dialog box. Select **Tools ► Options ► Preferences**. Select the **ToolBox** tab, and then select **Use large tools**. If you change the size of the SAS ToolBox through the Preferences dialog box, the new size is saved, and SAS displays the large toolbox in subsequent sessions.

TOOLLOAD Command

Loads a specific toolbox.

UNIX specifics: all

Syntax

TOOLLOAD <library.catalog.entry>

Details

If you do not specify an entry name, TOOLLOAD loads the toolbox for the active window.

See Also

“TOOLCLOSE Command” on page 234

TOOLTIPS Command

Toggles the ToolTip text for an icon in the toolbox.

UNIX specifics: all

Syntax

TOOLTIPS <ON|OFF>

ON

specifies that the ToolTip text is displayed when you move the cursor over an icon in the toolbox.

OFF

specifies that the ToolTip text is not displayed.

Details

If you do not specify ON or OFF, the TOOLTIPS command turns the ToolTip text on or off, depending on the current setting.

You can also use the Preferences dialog box to specify whether ToolTip text is displayed by selecting **Tools ► Options ► Preferences**. Select the **ToolBox** tab, and then select **Use tip text**.

See Also

“Changing the Attributes of an Existing Tool” on page 180

WBROWSE Command

Opens a World Wide Web (WWW) browser.

UNIX specifics: all

Syntax

WBROWSE <“url”>

Details

WBROWSE invokes the Web browser that is specified by the resource **SAS.helpBrowser**. If you specify a URL, the document that the URL identifies is automatically displayed. If you do not specify a URL, the SAS home page is displayed.

See Also

“Miscellaneous Resources in UNIX Environments” on page 207

WCOPY Command

Copies the marked contents of the active window to the default buffer.

UNIX specifics: all

Syntax

WCOPY

Details

In Base SAS windows, this command executes the STORE command.

See Also

online SAS Help and Documentation for information about the STORE command

WCUT Command

Moves the marked contents of the active window to the default buffer.

UNIX specifics: all

Syntax

WCUT

Details

In Base SAS windows, this command executes the CUT command.

This command is valid only when the active window is a text editor window, such as Program Editor or Notepad.

See Also

online SAS Help and Documentation for information about the CUT and WCUT commands

WDEF Command

Redefines the active window.

UNIX specifics: behavior is controlled by the **SAS.awsResizePolicy** resource

Syntax

WDEF *starting-row starting-col nrows ncols*

Details

The WDEF command operates in the application workspace assigned to the SAS session. The WDEF command does not operate in the AWS container window, except when the container window needs to be enlarged so that you can view a SAS window contained in it. AWS resize behavior is controlled by the **SAS.awsResizePolicy** resource.

See Also

The description of the **SAS.awsResizePolicy** resource in “Miscellaneous Resources in UNIX Environments” on page 207

“X Window Managers” on page 141 or your window manager documentation for information about moving and resizing windows in the X environment

WPASTE Command

Pastes the contents of the default buffer into the active window.

UNIX specifics: all

Syntax

WPASTE

Details

In Base SAS windows, this command executes the PASTE command.

See Also

online SAS Help and Documentation for information about the PASTE and WPASTE commands

WUNDO Command

Undoes one line of text entry, or undoes the last cut, copy, or paste action.

UNIX specifics: all

Syntax

WUNDO

Details

In Base SAS windows, this command executes the UNDO command. In SAS/GRAPH windows, WUNDO is invalid.

One execution of the WUNDO command undoes text entry for only one line at a time. If you issue the WUNDO command again, the previous line of text is undone.

If you use the CC command to copy and paste a block of text, and then issue the WUNDO command, the block of text that you copied is deleted. If you use the DD command to delete a block of text, and then issue the WUNDO command, the block of text that you deleted is restored.

Note: The WUNDO command cannot replace lines that the SUBMIT command removes. It cannot reverse the effects of submitted SAS statements. △

X Command

Enables you to enter UNIX commands without ending the SAS session.

UNIX specifics: all

Syntax

X *UNIX-command*

X '*cmd1;cmd2....<;cmd-n>*'

Details

When you enter the X command, SAS starts a shell to execute the commands that you specified. The commands that you enter are processed differently, depending on whether you enter one command or more than one command.

See Also

“Executing Operating System Commands from Your SAS Session” on page 14

XSYNC Command

Changes X synchronization during a SAS session.

UNIX specifics: all


Syntax

XSYNC <ON|OFF>

Details

This command turns off the buffering that is normally done by the X Window System. X synchronization is off by default. Turning it on is useful when you are debugging applications, although it drastically reduces performance.

If you do not specify ON or OFF, XSYNC toggles the synchronization. The XSYNC command is valid from any SAS window.



CHAPTER

11

Data Set Options under UNIX

<i>SAS Data Set Options under UNIX</i>	241
<i>Summary of SAS Data Set Options in UNIX Environments</i>	241
<i>ALTER= Data Set Option</i>	245
<i>BUFNO= Data Set Option</i>	245
<i>BUFSIZE= Data Set Option</i>	246
<i>FILECLOSE= Data Set Option</i>	247
<i>PW= Data Set Option</i>	248
<i>USEDIRECTIO= Data Set Option</i>	248

SAS Data Set Options under UNIX

This section describes SAS data set options that exist only in the UNIX environment, as well as options whose behavior or syntax is specific to UNIX. Each data set option description includes a brief “UNIX specifics” section that explains which aspect of the data set option is specific to UNIX. For data set options that have behavior or syntax specific to UNIX, see *SAS Language Reference: Dictionary* for a complete description of the option.

Specify data set options following the data set name in SAS statements as follows:

```
..data-set-name(option-1=value-1 option-2=value-2,..)
```

A few data set options are also SAS system options (for example, `BUFSIZE=`). If the same option is specified both as a system option and as a data set option, SAS uses the value given with the data set option. See “Customizing Your SAS Session by Using System Options” on page 17 and Chapter 18, “System Options under UNIX,” on page 341 for more information about SAS system options.

See “Summary of SAS Data Set Options in UNIX Environments” on page 241 for a table of all of the data set options available under UNIX.

Summary of SAS Data Set Options in UNIX Environments

SAS data set options are listed in the following table. The table lists the name of each option, a brief description, whether the option can be used for a data set opened for input, output, or update, and a list of engines for which the option is valid. The **See** column tells you where to look for more information about an option. Use the following legend to locate the additional information.

COMP	See the description of the data set option in this section.
------	---

LR	See <i>SAS Language Reference: Dictionary</i> .
NLS	See the <i>SAS National Language Support (NLS): Reference Guide</i> .

Table 11.1 Summary of SAS Data Set Options

Option Name	Description	When Used	Engines	See
ALTER=	specifies a password for a SAS file that prevents users from replacing or deleting the file, but permits Read and Write access.	output, update	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR, COMP
BUFNO=	specifies the number of buffers to be allocated for processing a SAS data set.	input, output, update	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR, COMP
BUFSIZE=	specifies the size of a permanent buffer page for an output SAS data set.	output	V9, V8, V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR, COMP
CNTLLEV=	specifies the level of shared access to SAS data sets.	input, update	V9, V8	LR
COMPRESS=	controls the compression of observations in a new output SAS data set.	output	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE	LR
DLDMGACTION=	specifies the action to take when a SAS data set in a SAS library is detected as damaged.	input, output, update	V9, V8	LR
DROP=	for an input data set, excludes the specified variables from processing; for an output data set, excludes the specified variables from being written to the data set.	input, output, update	all	LR
ENCODING=	overrides the encoding for the input or output SAS data set.	input, output	V9, V8, V9TAPE, V8TAPE, V7TAPE	NLS
ENCRYPT=	specifies whether to encrypt an output SAS data set.	output	all	LR
FILECLOSE=	specifies how a tape is positioned when a SAS data set is closed.	input, output	V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR, COMP
FIRSTOBS=	specifies the first observation that SAS processes in a SAS data set.	input, update	all	LR

Option Name	Description	When Used	Engines	See
GENMAX=	requests generations for a SAS data set, and specifies the maximum number of versions.	output, update	V9, V8	LR
GENNUM=	specifies a particular generation of a SAS data set.	input, output, update	V9, V8	LR
IDXNAME=	directs SAS to use a specific index to meet the conditions of a WHERE expression.	input, update	V9, V8, V6	LR
IDXWHERE=	specifies whether SAS uses an index or uses a sequential search, to match the conditions of a WHERE expression.	input, update	V9, V8, V6	LR
IN=	creates a Boolean variable that indicates whether the data set contributed data to the current observation.	input, update	all	LR
INDEX=	defines an index for a new output SAS data set.	output	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE	LR
KEEP=	for an input data set, specifies the variables to process; for an output data set, specifies the variables to write to the data set.	input, output, update	all	LR
LABEL=	specifies a label for a SAS data set.	input, output, update	all	LR
OBS=	specifies the last observation that SAS processes in a data set.	input, update	all	LR
OBSBUF=	determines the size of the view buffer for processing a DATA step view.	input	V9, V8	LR
OUTREP=	specifies the data representation for the output SAS data set.	output	V9, V8, V9TAPE, V8TAPE, V7TAPE	LR , NLS
POINTOBS=	controls whether to process a compressed SAS data set by observation number or by sequential access.	output	V9, V8	LR
PW=	assigns a READ, WRITE, or ALTER password to a SAS file and enables access to a password-protected file.	input, output, update	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR, COMP
PWREQ=	specifies whether to display a dialog box for a SAS data set password.	input, output, update	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR

Option Name	Description	When Used	Engines	See
READ=	assigns a password to a SAS file and enables access to a read-protected SAS file.	input, output, update	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR
RENAME=	changes the name of a variable.	input, output, update	all	LR
REPEMPTY=	specifies whether a new, empty data set can overwrite an existing SAS data set that has the same name.	output	V9, V8	LR
REPLACE=	specifies whether a new SAS data set that contains data can overwrite an existing data set that has the same name.	output	all	LR
REUSE=	specifies whether new observations can be written to freed space in compressed SAS data sets.	output	V9, V8, V6	LR
SORTEDBY=	indicates how the SAS data set is currently sorted.	input, output, update	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR
SPILL=	specifies whether to create a spill file for non-sequential processing of a DATA step view.	output	V9, V8	LR
TOBSNO= ¹	specifies the number of observations to send in a client/server transfer.	input, output, update	V9, V8	LR
TYPE=	specifies the data set type for a specially structured SAS data set.	input, output, update	all	LR
USEDIRECTIO=	turns on direct file I/O for the file that you specify. To use this data set option, you must specify the ENABLEDIRECTIO statement option in the LIBNAME statement where the libref was assigned.	input, output, update	V9, V8	COMP
WHERE=	selects observations in a SAS data set that match the specified conditions.	input, output, update	all	LR
WHEREUP=	specifies whether to evaluate new observations and updated observations against a WHERE expression.	output, update	V9, V8, V6	LR
WRITE=	assigns a WRITE password to a SAS data set and enables access to a write-protected SAS file.	output, update	V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE	LR

¹ The TOBSNO= option is valid only for data sets that are accessed through a SAS server via the REMOTE engine.

ALTER= Data Set Option

Specifies a password for a SAS file that prevents users from replacing or deleting the file, but permits Read and Write access.

Default: none

Valid in: DATA step and PROC steps

Category: Data Set Control

Engines: V9, V8, V6

UNIX specifics: TAPE engines ignore the *alter-password*

See: ALTER= Data Set Option in *SAS Language Reference: Dictionary*

Syntax

ALTER=*alter-password*

alter-password

must be a valid SAS name. See “Rules for Words and Names in the SAS Language” in *SAS Language Reference: Concepts*.

Details

The ALTER= option applies to all types of SAS files except catalogs. You can use this option to assign an *alter-password* to a SAS file or to access a read-protected, write-protected, or alter-protected SAS file.

Note: Under UNIX, TAPE engines ignore the *alter-password*. △

BUFNO= Data Set Option

Specifies the number of buffers to be allocated for processing a SAS data set.

Default: 1

Valid in: DATA step and PROC steps

Category: Data Set Control

Engines: V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE

UNIX specifics: default value

See: BUFNO= Data Set Option in *SAS Language Reference: Dictionary*

Syntax

BUFNO=*n* | *nK* | *hexX* | MIN | MAX

n* | *nK

specifies the number of buffers in multiples of 1 (bytes); 1,024 (kilobytes). For example, a value of **8** specifies 8 buffers, and a value of **1k** specifies 1024 buffers.

hexX

specifies the number of buffers as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example the value **2dx** specifies 45 buffers.

MIN

sets the minimum number of buffers to 0, which causes SAS to use the minimum optimal value for the operating environment.

MAX

sets the number of buffers to the maximum possible number in your operating environment, up to the largest four-byte signed integer, which is $2^{31}-1$, or approximately 2 billion.

Details

The buffer number is not a permanent attribute of the data set; it is valid only for the current SAS step. BUFSIZE= applies to SAS data sets that are opened for input, output, or update.

See Also

- “BUFSIZE= Data Set Option” on page 246
- “BUFNO System Option” on page 359

BUFSIZE= Data Set Option

Specifies the size of a permanent buffer page for an output SAS data set.

Default: 0

Valid in: DATA step and PROC steps

Category: Data Set Control

Engines: V9, V8, V9TAPE, V8TAPE, V7TAPE, V6TAPE

UNIX specifics: valid range

See: BUFSIZE= Data Set Option in *SAS Language Reference: Dictionary*

Syntax

BUFSIZE=*n* | *nK* | *nM* | *nG* | *hexX* | MAX

n* | *nK* | *nM* | *nG

specifies the buffer size in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); or 1,073,741,824 (gigabytes). For example, a value of **8** specifies 8 bytes, and a value of **3m** specifies 3,145,728 bytes.

The buffer size can range from 1K to 2G–1. For values greater than 1G, use the *nM* option.

hexX

specifies the page size as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, **2dx** sets the page size to 45 bytes.

MAX

sets the buffer page size to the maximum possible number in your operating environment, up to the largest four-byte, signed integer, which is $2^{31}-1$, or approximately 2 billion bytes.

Details

The BUFSIZE= data set option specifies the buffer size for data sets you are creating. This option is valid only for output data sets.

If you use the default value (0) when you create a SAS data set, the engine calculates a buffer size to optimize CPU and I/O use. This size is the smallest multiple of 8K that can hold 80 observations, but is not larger than 64K.

If you specify a nonzero value when you create a SAS data set, the engine uses that value. If that value cannot hold at least one observation or is not a valid buffer size, the engine rounds the value up to a multiple of 1K.

See Also

- “BUFSIZE System Option” on page 360

FILECLOSE= Data Set Option

Specifies how a tape is positioned when a SAS data set is closed.

Default: REREAD

Valid in: DATA step and PROC steps

Category: Miscellaneous

Engines: V9TAPE, V8TAPE, V7TAPE, V6TAPE

UNIX specifics: list of valid values

See: FILECLOSE= Data Set Option in *SAS Language Reference: Dictionary*

Syntax

FILECLOSE= FREE | LEAVE | REREAD | REWIND

FREE

rewinds and dismounts the tape. If the device cannot dismount the tape, then the tape will be rewound only.

LEAVE

positions the tape at the end of the file that was just processed. Use FILECLOSE=LEAVE if you are not repeatedly accessing the same files in a SAS program, but you are accessing one or more subsequent SAS files on the same tape.

REREAD

positions the tape volume at the beginning of the file that was just processed. Use FILECLOSE=REREAD if you are accessing the same SAS data set on tape several times in a SAS program.

REWIND

rewinds the tape volume to the beginning. Use FILECLOSE=REWIND if you are accessing one or more previous SAS files on the same tape, but you are not repeatedly accessing the same files in a SAS program.

PW= Data Set Option

Assigns a READ, WRITE, or ALTER password to a SAS file, and enables access to a password-protected SAS file.

Default: none

Valid in: DATA step and PROC steps

Category: Data Set Control

Engines: V9, V8, V6, V9TAPE, V8TAPE, V7TAPE, V6TAPE

UNIX specifics: TAPE engines ignore the *alter-password*

See: PW= Data Set Option in *SAS Language Reference: Dictionary*

Syntax

PW=*password*

password

must be a valid SAS name. See “Rules for Words and Names in the SAS Language” in *SAS Language Reference: Concepts*.

Details

The PW= option applies to all types of SAS files except catalogs. You can use this option to assign a password to a SAS file or to access a password-protected SAS file.

Note: Under UNIX, TAPE engines ignore the *alter-password*. Δ

USEDIRECTIO= Data Set Option

Turns on direct file I/O for a library that contains the file to which the ENABLEDIRECTIO option has been applied.

Default: Off

Valid in: DATA step

Category: Data Set Control

Engines: V9, V8

UNIX specifics: To use this option, you must also use the ENABLEDIRECTIO option in the LIBNAME statement where the libref was assigned.

Syntax

USEDIRECTIO=

Details

The Basics

The USEDIRECTIO= data set option turns on direct file I/O for a data set that is listed on a DATA statement. The associated libref must have been defined with the ENABLEDIRECTIO option in the LIBNAME statement.

Using ENABLEDIRECTIO on a LIBNAME statement makes direct file I/O possible for data sets in that library. Direct I/O itself is not turned on. You must use the USEDIRECTIO= option to produce direct file I/O.

Ways to Turn on Direct File I/O You can turn on direct file I/O in two ways:

□

Use both the ENABLEDIRECTIO and USEDIRECTIO= options in the LIBNAME statement:

```
libname libref-name '.' ENABLEDIRECTIO USEDIRECTIO=yes;
```

In this case, SAS uses direct file I/O on all SAS I/O data sets that are opened using the libref *libref-name*.

□

Use ENABLEDIRECTIO in the LIBNAME statement and use USEDIRECTIO= in a DATA statement:

```
libname libref-name '.' ENABLEDIRECTIO;
data libref-name.data-set-name (USEDIRECTIO=yes);
```

In this case, *libref-name.data-set-name*.DATA will be opened for direct file I/O. Other SAS I/O data sets referenced by *libref-name* will not use direct file I/O.

USEDIRECTIO= by itself has no effect. Neither of the following statements open a data set for direct file I/O:

```
libname libref-name '.' USEDIRECTIO=yes;
data libref-name.data-set-name (USEDIRECTIO=yes);
```

Examples

The following example uses the ENABLEDIRECTIO LIBNAME option to enable files that are associated with the libref **test** to be opened for direct I/O. The USEDIRECTIO= data set option opens **test.file1** for direct I/O. **test.file2** is not opened for direct I/O.


```
LIBNAME test '.' ENABLEDIRECTIO;
data test.file1(USEDIRECTIO=yes);
```

```
    ... more SAS statements ...  
run;  
data test.file2;  
    ... more SAS statements ...  
run;
```

See Also

Statements:

“LIBNAME Statement” on page 328



CHAPTER

12

Formats under UNIX

SAS Formats under UNIX 251

HEXw. Format 251

\$HEXw. Format 252

IBw.d Format 252

PDw.d Format 252

PIBw.d Format 253

RBw.d Format 253

ZDw.d Format 254

SAS Formats under UNIX

This section describes SAS formats that have behavior or syntax that is specific to UNIX environments. Each format description includes a brief “UNIX specifics” section that explains which aspect of the data set option is specific to UNIX. Each format is described in this documentation and in *SAS Language Reference: Dictionary*.

HEXw. Format

Converts real binary (floating-point) numbers to hexadecimal representation.

Category: Numeric

Width range: 1 to 16

Default width: 8

Alignment: left

UNIX specifics: floating-point representation

See: HEXw. format in *SAS Language Reference: Dictionary*

Details

The HEXw. format converts a real (floating-point) binary number to its hexadecimal representation. When you specify a width value of 1 through 15, the real binary number is truncated to a fixed-point integer before being converted to a hexadecimal number. When you specify 16 for the width, SAS writes the floating-point value of the number, but does not truncate it.

Note: UNIX systems vary widely in their floating-point representation. See “Reading and Writing Binary Data in UNIX Environments” on page 216 for more information. Δ

\$HEXw. Format

Converts character values to hexadecimal representation.

Category: Character

Width range: 1 to 32767

Default width: 4

Alignment: left

UNIX specifics: produces ASCII codes

See: \$HEXw. format in *SAS Language Reference: Dictionary*

Details

Under UNIX, the \$HEXw. format produces hexadecimal representations of ASCII codes for characters, with each byte requiring two columns. Therefore, you need twice as many columns to output a value with the \$HEXw. format.

IBw.d Format

Writes integer binary (fixed-point) values.

Category: Numeric

Width range: 1 to 8

Default width: 4

Decimal Range: 0–10

Alignment: left

UNIX specifics: byte order

See: IBw.d format in *SAS Language Reference: Dictionary*

Details

The IBw.d format writes integer binary (fixed-point) values. Integers are stored in integer-binary, or fixed-point, form. For example, the number 2 is stored as 00000002. If the format includes a *d* value, the data value is multiplied by 10^d .

For more details, see “Reading and Writing Binary Data in UNIX Environments” on page 216.

PDw.d Format

Writes data in packed decimal format.

Category: Numeric
Width range: 1 to 16
Default width: 1
Decimal Range: 0–31
Alignment: left
UNIX specifics: data representation
See: PDw.d format in *SAS Language Reference: Dictionary*

Details

The PDw.d format writes values in packed decimal format. In packed decimal data, each byte contains two digits. The *w* value represents the number of bytes, not the number of digits. The value's sign is the first byte. Because the entire first byte is used for the sign, you should specify at least a width of 2.

The PDw.d format writes missing numerical data as -0 . When the PDw.d informat reads a value of -0 , the result is a value of 0.

For more information, see “Reading and Writing Binary Data in UNIX Environments” on page 216.

PIBw.d Format

Writes positive integer binary (fixed-point) values.

Category: Numeric
Width range: 1 to 8
Default width: 1
Decimal Range: 0–10
Alignment: left
UNIX specifics: byte order
See: PIBw.d format in *SAS Language Reference: Dictionary*

Details

The PIBw.d format writes fixed-point binary values, treating all values as positive. Thus, the high-order bit is part of the value, rather than the value's sign. If a *d* value is specified, the data value is multiplied by 10^d .

For more information, see “Reading and Writing Binary Data in UNIX Environments” on page 216.

RBw.d Format

Writes real binary (floating-point) data in real binary format.

Category: Numeric

Width range: 2 to 8**Default width:** 4**Decimal Range:** 0–10**Alignment:** left**UNIX specifics:** floating-point representation**See:** RBw.d format in *SAS Language Reference: Dictionary*

Details

The RBw.d format writes numeric data in real binary (floating-point) notation. SAS stores all numeric values in floating-point.

Real binary is the most efficient format for representing numeric values because SAS already represents numbers this way and no conversion is needed.

For more information, see “RBw.d Informat” on page 281 and “Reading and Writing Binary Data in UNIX Environments” on page 216.

ZDw.d Format

Writes numeric data in zoned decimal format.**Category:** Numeric**Width range:** 1 to 32**Default width:** 1**Alignment:** left**UNIX specifics:** data representation**See:** ZDw.d format in *SAS Language Reference: Dictionary*


Details

The ZDw.d format writes zoned decimal data. This format is also known as overprint trailing numeric format. Under UNIX, the last byte of the field includes the sign with the last digit. The conversion table for the last byte is as follows:

Digit	ASCII Character	Digit	ASCII Character
0	{	–0	}
1	A	–1	J
2	B	–2	K
3	C	–3	L
4	D	–4	M
5	E	–5	N
6	F	–6	O
7	G	–7	P

8	H	-8	Q
9	I	-9	R

For more information, see “ZDw.d Informat” on page 283 and “Reading and Writing Binary Data in UNIX Environments” on page 216.



CHAPTER

13

Functions and CALL Routines under UNIX

SAS Functions and CALL Routines under UNIX 257

<i>BYTE Function</i>	257
<i>CALL MODULE Routine</i>	258
<i>CALL SLEEP Routine</i>	260
<i>CALL SYSTEM Routine</i>	261
<i>COLLATE Function</i>	262
<i>DINFO Function</i>	263
<i>DOPEN Function</i>	264
<i>DOPTNAME Function</i>	265
<i>DOPTNUM Function</i>	265
<i>FDELETE Function</i>	266
<i>FEXIST Function</i>	266
<i>FILEEXIST Function</i>	267
<i>FILENAME Function</i>	267
<i>FILeref Function</i>	269
<i>FINFO Function</i>	269
<i>FOPTNAME Function</i>	270
<i>FOPTNUM Function</i>	272
<i>MODEXIST Function</i>	273
<i>MOPEX Function</i>	273
<i>PATHNAME Function</i>	274
<i>PEEKLONG Function</i>	275
<i>RANK Function</i>	276
<i>SYSGET Function</i>	276
<i>TRANSLATE Function</i>	277

SAS Functions and CALL Routines under UNIX

This section describes SAS functions and CALL routines whose behavior is specific to UNIX environments. Each function and CALL routine description includes a brief “UNIX specifics” section that explains which aspect of the function and CALL routine is specific to UNIX. For more information about all of these CALL routines and functions, except SYSGET, see *SAS Language Reference: Dictionary*.

BYTE Function

Returns one character in the ASCII collating sequence.

Category: Character

UNIX specifics: Uses the ASCII collating sequence

See: BYTE Function in *SAS Language Reference: Dictionary*

Syntax

BYTE(*n*)

n

specifies an integer that represents a specific ASCII character. The value of *n* can range from 0 to 255.

Details

If the BYTE function returns a value to a variable that has not yet been assigned a length, by default the variable is assigned a length of 1.

CALL MODULE Routine

Calls a specific routine or module that resides in a shared executable library.

Category: External Files

UNIX specifics: All

Syntax

CALL MODULE(<*cntl*>, *module*, *arg-1*, *arg-2*..., *arg-n*);

num=**MODULEN**(<*cntl*>, *module*, *arg-1*, *arg-2*..., *arg-n*);

char=**MODULEC**(<*cntl*>, *module*, *arg-1*..., *arg-2*, *arg-n*);

Note: The following functions permit vector and matrix arguments. You can use them only within the IML procedure (See the documentation for SAS/IML for more information.): Δ

CALL MODULEI (<*cntl*>, *module*, *arg-1*, *arg-2*..., *arg-n*);

num=**MODULEIN**(<*cntl*>, *module*, *arg-1*, *arg-2*..., *arg-n*)

char=**MODULEIC**(<*cntl*>, *module*, *arg-1*, *arg-2*..., *arg-n*);

cntl

is an optional control string whose first character must be an asterisk (*), followed by any combination of the following characters:

I prints the hexadecimal representations of all arguments to the MODULE function and to the requested shared library routine before and after the shared library routine is called. You can use this option to help diagnose problems that are caused by incorrect

arguments or attribute tables. If you specify the I option, the E option is implied.

- E prints detailed error messages. Without the E option (or the I option, which supersedes it), the only error message that the MODULE function generates is "Invalid argument to function," which is usually not enough information to determine the cause of the error.
- S*x* uses *x* as a separator character to separate field definitions. You can then specify *x* in the argument list as its own character argument to serve as a delimiter for a list of arguments that you want to group together as a single structure. Use this option only if you do not supply an entry in the SASCBTBL attribute table. If you do supply an entry for this module in the SASCBTBL attribute table, you should use the FDSTART option in the ARG statement in the table to separate structures.
- H provides brief help information about the syntax of the MODULE routines, the attribute file format, and the suggested SAS formats and informats.

For example, the control string '*IS/' specifies that parameter lists be printed and that the string '/' is to be treated as a separator character in the argument list.

module

is the name of the external module to use. The *module* can be specified as a shared library and the routine name or ordinal value, separated by a comma. You do not need to specify the shared library name if you specified the MODULE attribute for the routine in the SASCBTBL attribute table, as long as the routine name is unique (that is, no other routines have the same name in the attribute file).

The module must reside in a shared library, and it must be externally callable. Note that while the shared library name is not case sensitive, the routine name is based on the restraints of the routine's implementation language, so the routine name is case sensitive.

If the shared library supports ordinal-value naming, you can provide the shared library name followed by a decimal number, such as 'XYZ,30'.

You can specify *module* as a SAS character expression instead of as a constant; most often, though, you will pass it as a constant.

arg-1, arg-2, ...arg-n

are the arguments to pass to the requested routine. Use the proper attributes for the arguments (that is, numeric arguments for numeric attributes and character arguments for character attributes).

CAUTION:

Be sure to use the correct arguments and attributes. If you use incorrect arguments or attributes for a shared library function, you can cause SAS to crash, or you will see unexpected results. △

Details

The MODULE functions execute a routine *module* that resides in an external (outside SAS) shared library with the specified arguments *arg-1* through *arg-n*.

The MODULE call routine does not return a value, while the MODULEN and MODULEC functions return a number *num* or a character *char*, respectively. Which routine you use depends on the expected return value of the shared library function that you want to execute.

MODULEI, MODULEIC, and MODULEIN are special versions of the MODULE functions that permit vector and matrix arguments. Their return values are still scalar. You can invoke these functions only from PROC IML.

Other than this name difference, the syntax for all six routines is the same.

The MODULE function builds a parameter list by using the information in *arg-1* to *arg-n* and by using a routine description and argument attribute table that you define in a separate file. Before you invoke the MODULE routine, you must define the fileref of SASCBTBL to point to this external file. You can name the file whatever you want when you create it.

This way, you can use SAS variables and formats as arguments to the MODULE function and ensure that these arguments are properly converted before being passed to the shared library routine.

CAUTION:

Using the MODULE function without defining an attribute table can cause SAS to crash, produce unexpected results, or result in severe errors. You need to use an attribute table for all external functions that you want to invoke. \triangle

See Also

- “The SASCBTBL Attribute Table” on page 109
- “PEEKLONG Function” on page 275

CALL SLEEP Routine

For a specified period of time, suspends the execution of a program that invokes this CALL routine.

Category: Special

UNIX specifics: All

See: CALL SLEEP Routine in *SAS Language Reference: Dictionary*

Syntax

CALL SLEEP(*n*<,unit>);

n

is a numeric constant that specifies the number of units of time for which you want to suspend execution of a program.

unit

specifies the unit of time, as a power of 10, which is applied to *n*. For example, 1 corresponds to a second, and .001 corresponds to a millisecond. The default is .001.

Details

CALL SLEEP puts the DATA step in which it is invoked into a nonactive wait state, using no CPU time and performing no input or output. If you are running multiple SAS processes, each process can execute CALL SLEEP independently without affecting the other processes.

Note: Extended sleep periods can trigger automatic host session termination based on timeout values set at your site. Contact your host system administrator to determine the timeout values used at your site. △

CALL SYSTEM Routine

Submits an operating environment command for execution.

Category: Special

UNIX specifics: *Command* must evaluate to a valid UNIX command

See: CALL SYSTEM Routine in *SAS Language Reference: Dictionary*

Syntax

CALL SYSTEM(*command*);

command

specifies any of the following:

- a UNIX command enclosed in quotation marks
- an expression whose value is a UNIX command
- the name of a character variable whose value is a UNIX command

Details

The CALL SYSTEM routine issues operating system commands. The output of the command appears in the window from which you invoked SAS.

The value of the XSYNC system option affects how the CALL SYSTEM routine works.

Note: The CALL SYSTEM routine can be executed within a DATA step. However, neither the X statement nor the %SYSEXEC macro program statement is intended for use during the execution of a DATA step. △

In the following example, for each record in **answer.week**, if the **resp** variable is **y**, the CALL SYSTEM routine will mail a message:

```
data _null_;
  set answer.week;
  if resp='y' then
    do;
      call system('mail mgr < $HOME/msg');
    end;
run;
```

See Also

- “Executing Operating System Commands from Your SAS Session” on page 14

COLLATE Function

Returns a character string in an ASCII collating sequence.

Category: Character

UNIX specifics: Uses ASCII collating sequence

See: COLLATE Function in *SAS Language Reference: Dictionary*

Syntax

COLLATE(*start-position* <*end-position*>) | (*start-position*<,*length*>)

start-position

specifies the numeric position in the collating sequence of the first character to be returned.

end-position

specifies the numeric position in the collating sequence of the last character to be returned.

length

specifies the number of characters in the collating sequence.

Details

The COLLATE function returns a string of ASCII characters. The ASCII collating sequence contains 256 positions, referenced with the numbers 0 through 255. Characters above 127 correspond to characters used in European languages as defined in the ISO 8859 character set.

Unless you assign the return value of the COLLATE function to a variable with a defined length less than 200, the ASCII collating sequence string is padded with spaces to a length of 200. If the ASCII collating sequence is greater than 200 characters, you must specify the length for the return string in a LENGTH statement; otherwise, the returned string will be truncated to a length of 200 characters. For more information, see the following examples.

Examples: How SAS Determines the Length of the Return String

Example 1: Truncating the Variable Length to 200 Characters Because the following code does not include a LENGTH statement, the length attribute for the ADDRESS variable is truncated to 200 characters:

```
data sales;
  Address=collate(1,241);
run;

proc contents;
run;
```

Output 13.1 Portion of PROC CONTENTS Output

Alphabetic List of Variables and Attributes			
#	Variable	Type	Len
1	Address	Char	200

Because length of ADDRESS is limited to 200 characters, the returned string from the COLLATE function will be limited to 200 characters.

Example 2: Specifying a Length Greater than 200 Characters To specify a length greater than 200 characters for a specific variable, you can use the LENGTH statement. In the following code, the length of ADDRESS is specified as 240 characters:

```
data sales;
  length Address $240;
  Address=collate(1,241);
run;

proc contents;
run;
```

Output 13.2 Portion of PROC CONTENTS Output

Alphabetic List of Variables and Attributes			
#	Variable	Type	Len
1	Address	Char	240

Because the length of ADDRESS is set to 240 characters, the returned string from the COLLATE function will contain 240 characters.

See Also

- “LENGTH Statement” on page 327

DINFO Function

Returns information about a directory.

Category: External Files

UNIX specifics: Directory pathname is the only information available

See: DINFO Function in *SAS Language Reference: Dictionary*

Syntax

DINFO(*directory-id*, *info-item*)

directory-id

specifies the identifier that was assigned when the directory was opened, generally by the DOPEN function.

info-item

specifies the information item to be retrieved. DINFO returns a blank if the value of *info-item* is invalid.

Details

Directories that are opened with the DOPEN function are identified by a *directory-id*. Use DOPTNAME to determine the names of the available system-dependent information items. Use DOPTNUM to determine the number of directory information items available.

Under UNIX, the only *info-item* available is Directory, which is the pathname of *directory-id*. If *directory-id* points to a list of concatenated directories, then Directory is the list of concatenated directory names.

See Also

- “DOPEN Function” on page 264
- “DOPTNAME Function” on page 265
- “DOPTNUM Function” on page 265

DOPEN Function

Opens a directory, and returns a directory identifier value.

Category: External Files

UNIX specifics: *fileref* can be assigned with an environment variable

See: DOPEN Function in *SAS Language Reference: Dictionary*

Syntax

DOPEN(*fileref*)

fileref

specifies the *fileref* assigned to the directory. In a DATA step, *fileref* can be a character expression, a string enclosed in quotation marks, or a DATA step variable whose value contains the *fileref*. In a macro, *fileref* can be any expression.

Details

DOPEN opens a directory and returns a directory identifier value (a number greater than 0) that is used to identify the open directory in other SAS external file access functions. If the directory could not be opened, DOPEN returns 0. The directory to be opened must be identified by a *fileref*.

DOPTNAME Function

Returns directory attribute information.

Category: External Files

UNIX specifics: Directory is the only item available

See: DOPTNAME Function in *SAS Language Reference: Dictionary*

Syntax

DOPTNAME(*directory-id*, *nval*)

directory-id

specifies the identifier that was assigned when the directory was opened, generally by the DOPEN function.

nval

specifies the sequence number of the information item.

Details

Under UNIX, the only directory information item available is Directory, which is the pathname of the *directory-id*. The *nval*, or sequence number, of Directory is 1. If *directory-id* points to a list of concatenated directories, then Directory is the list of concatenated directory names.

DOPTNUM Function

Returns the number of information items that are available for a directory.

Category: External Files

UNIX specifics: Directory is the only item available

See: DOPTNUM Function in *SAS Language Reference: Dictionary*

Syntax

DOPTNUM(*directory-id*)

directory-id

specifies the identifier that was assigned when the directory was opened, generally by the DOPEN function.

Details

Under UNIX, only one information item is available for a directory. The name of the item is Directory; its value is the pathname or list of pathnames for *directory-id*, and its

sequence number is 1. Because only one information item is available for a directory, this function will return a value of 1.

FDELETE Function

Deletes an external file or an empty directory.

Category: External Files

UNIX specifics: *fileref* can be assigned with an environment variable

See: FDELETE Function in *SAS Language Reference: Dictionary*

Syntax

FDELETE(*fileref*)

fileref

specifies the fileref that is assigned to the external file or directory. The fileref cannot be associated with a list of concatenated filenames or directories. If the fileref is associated with a directory, the directory must be empty. You must have permission to delete the file. See the UNIX man page for **chmod** for more information about permissions.

Under UNIX, *fileref* can also be an environment variable. The *fileref* must be enclosed in double quotation marks.

Details

FDELETE returns 0 if the operation was successful, or a nonzero number if it was not successful.

FEXIST Function

Verifies the existence of an external file that is associated with a fileref.

Category: External Files

UNIX specifics: *fileref* can be assigned with an environment variable

See: FEXIST Function in *SAS Language Reference: Dictionary*

Syntax

FEXIST(*fileref*)

fileref

specifies the fileref assigned to the external file or directory. In a DATA step, *fileref* can be a character expression, a string enclosed in quotation marks, or a DATA step variable whose value contains the fileref. In a macro, *fileref* can be any expression.

Under UNIX, *fileref* can also be an environment variable. The *fileref* or the environment variable that you specify must be enclosed in double quotation marks.

Details

The FEXIST function returns a value of 1 if the external file that is associated with *fileref* exists, and a value of 0 if the file does not exist.

FILEEXIST Function

Verifies the existence of an external file by its physical name.

Category: External Files

UNIX specifics: *filename* can be assigned with an environment variable

See: FILEEXIST Function in *SAS Language Reference: Dictionary*

Syntax

FILEEXIST(*filename*)

filename

specifies a fully qualified physical filename of the external file. In a DATA step, *filename* can be a character expression, a string in quotation marks, or a DATA step variable. In a macro, *filename* can be any expression.

Under UNIX, *filename* can also be an environment variable. The *filename* or the environment variable that you specify must be enclosed in double quotation marks.

Details

FILEEXIST returns 1 if the external file exists, and 0 if the external file does not exist. You can check for the existence of a directory by using FILEEXIST.

FILENAME Function

Assigns or deassigns a fileref for an external file, directory, or output device.

Category: External Files

UNIX specifics: *fileref* can be assigned with an environment variable; valid values of *device-type* and *host-options*

See: FILENAME Function in *SAS Language Reference: Dictionary*

Syntax

FILENAME(*fileref*, *filename* <,*device-type*<,"*host-options*"<,<*dir-ref*>>>)

fileref

specifies the *fileref* to assign to an external file. In a DATA step, *fileref* can be a character expression, a string enclosed in quotation marks, or a DATA step variable whose value contains the fileref. In a macro (for example, in the %SYSFUNC function), *fileref* is the name of a macro variable (without an ampersand) whose value contains the fileref to assign to the external file. (For information, see the FILENAME function in *SAS Language Reference: Dictionary*.)

Under UNIX, the *fileref* can be a UNIX environment variable. The *fileref* or the environment variable that you specify must be enclosed in double quotation marks.

filename

specifies the external file. Specifying a blank filename (" ") deassigns a *fileref* that was previously assigned.

Under UNIX, the filename differs according to the device type. See "Device Information in the FILENAME Statement" on page 322 for information that is appropriate for each device. Remember that UNIX filenames are case sensitive.

In a DATA step, *filename* can be a character expression, a string enclosed in quotation marks, or a DATA step variable whose value contains the filename. In a macro, *filename* can be any expression.

device-type

specifies the type of device or the access method that is used if the fileref points to an input or output device or location that is not a physical file. It can be any one of the devices listed in "Device Information in the FILENAME Statement" on page 322. DISK is the default device type.

host-options

are options that are specific to UNIX. You can use any of the options that are available in the FILENAME statement. See "FILENAME Statement" on page 318 for a description of the host options.

Requirement: Enclose host options in quotation marks. If you have multiple host options, then all of the host options must be enclosed in one set of quotation marks. The following example shows the syntax:

```
rc=filename("try","MISCHL.FLAT.FILE1","ftp",
           'user="mischl1",host="sdcunx",prompt');
```

dir-ref

specifies the fileref that is assigned to the directory in which the external file resides.

Details

FILENAME returns a 0 if the operation is successful, and a nonzero number if it was not successful.

If you use the FTP access method to communicate with a remote system, SAS might return the following error message:

```
ERROR: Physical file does not exist.
```

This error is likely to occur when the fully qualified data set name is specified within single quotation marks. For example:

```
FILENAME fileref FTP 'system.dataset.name' USER='username'
                    PASS='password' HOST='ip_address';
```

By default, SAS appends the profile prefix to the beginning of the data set name. To prevent the profile prefix from being appended, enclose the data set name with both double and single quotation marks:

```
FILENAME fileref FTP "'external_file'" USER='username'
                    PASS='password' HOST='ip_address';
```

FILEREF Function

Verifies whether a fileref has been assigned for the current SAS session.

Category: External Files

UNIX specifics: *fileref* can be assigned with an environment variable

See: FILEREF Function in *SAS Language Reference: Dictionary*

Syntax

FILEREF(*fileref*)

fileref

specifies the fileref assigned to be validated. In a DATA step, *fileref* can be a character expression, a string enclosed in quotation marks, or a DATA step variable whose value contains the fileref. In a macro, *fileref* can be any expression.

Under UNIX, *fileref* can also be a UNIX environment variable. The *fileref* or the environment variable you specify must be enclosed in double quotation marks.

Details

A negative return code indicates that the fileref exists, but the physical file associated with the fileref does not exist. A positive value indicates that the fileref is not assigned. A value of zero indicates that the fileref and external file both exist.

See “FILENAME Function” on page 267 for more information.

FINFO Function

Returns the value of a file information item for an external file.

Category: External Files

UNIX specifics: *info-item* is available

See: FINFO Function in *SAS Language Reference: Dictionary*

Syntax

FINFO(*file-id*, *info-item*)

file-id

specifies the identifier that was assigned when the file was opened, generally by the FOPEN function.

info-item

specifies the name of the file information item to be retrieved. This value is a character value. *Info-item* is either a variable containing a valid value or the valid value in quotation marks.

Under UNIX, *info-item* for disk files can have one of the following values:

- Filename
- Owner Name
- Group Name
- Access Permission
- File Size (bytes)

If you concatenate filenames, then an additional *info-item* is available: File List.

If you are using pipe files, then the only valid value for *info-item* is PIPE Command.

Details

The FINFO function returns the value of a system-dependent information item for an external file that was previously opened and assigned a *file-id* by the FOPEN function. FINFO returns a blank if the value given for *info-item* is invalid.

For an example of how to use the FINFO function, see “Example: File Attributes When Using the Pipe Device Type” on page 271.

See Also

- “FOPEN Function” in *SAS Language Reference: Dictionary*

FOPTNAME Function

Returns the name of an item of information about an external file.

Category: External Files

UNIX specifics: Information items available

See: FOPTNAME Function in *SAS Language Reference: Dictionary*

Syntax

FOPTNAME(*file-id*, *nval*)

file-id

specifies the identifier that was assigned when the file was opened, generally by the FOPEN function.

nval

specifies the number of the file information item to be retrieved. The following table shows the values that *nval* can have in UNIX operating environments for single, pipe, and concatenated files:

<i>nval</i>	Information Items Available For...		
	Single File	Pipe Files	Concatenated Files
1	File Name	PIPE Command	File Name
2	Owner Name		File List
3	Group Name		Owner Name
4	Access Permission		Group Name
5	File Size (bytes)		Access Permission
6			File Size (bytes)

Details

FOPTNAME returns a blank if an error occurs.

Example: File Attributes When Using the Pipe Device Type

The following example creates a data set that contains the NAME and VALUE attributes returned by the FOPTNAME function when you are using pipes:

```
data fileatt;
  length name $ 20 value $ 40;
  drop fid j infonum;
  filename mypipe pipe 'UNIX-command';
  fid=fopen("mypipe","s");
  infonum=foptnum(fid);
  do j=1 to infonum;
    name=foptname(fid,j);
    value=finfo(fid,name);
    put 'File attribute' name 'has a value of ' value;
  output;
  end;
run;
```

The following statement should appear in the SAS log:

```
File attribute Pipe Command has a value of UNIX-command
```

UNIX-command is the UNIX command or program where you are piping your output or where you are reading your input. This command or program must be either fully qualified or defined in your PATH environment variable.

See Also

- “FINFO Function” on page 269
- “FOPTNUM Function” on page 272
- “FOPEN Function” in *SAS Language Reference: Dictionary*

FOPTNUM Function

Returns the number of information items that are available for an external file.

Category: External Files

UNIX specifics: Information items available

See: FOPTNUM Function in *SAS Language Reference: Dictionary*

Syntax

FOPTNUM(*file-id*)

file-id

specifies the identifier that was assigned when the file was opened, generally by the FOPEN function.

Details

Under UNIX, five information items are available for all types of files:

- File Name
- Owner Name
- Group Name
- Access Permission
- File Size (bytes)

If you concatenate filenames, then an additional information item is available: File List. If you are using pipe files, then the only information item available is PIPE Command.

The *open-mode* specified in the FOPEN function determines the value that FOPTNUM returns.

Open Mode	FOPTNUM Value	Information Items Available
Append	6 for concatenated files	All information items available.
Input	5 for single files	
Update		
Output	5 for concatenated files 4 for single files	Because the file is open for output, the File Size information type is unavailable.
Sequential (using Pipe Device Type)	1	The only information item available is PIPE Command.

For an example of how to use the FOPTNUM function, see “Example: File Attributes When Using the Pipe Device Type” on page 271.

See Also

- “FINFO Function” on page 269
- “FOPTNAME Function” on page 270
- “FOPEN Function” in *SAS Language Reference: Dictionary*

MODEXIST Function

Determines whether a product image exists in the release of SAS that you have installed.

Category: Numeric

UNIX specifics: *pathname* is available

See: MODEXIST Function in *SAS Language Reference: Dictionary*

Syntax

MODEXIST '*product-name*' | '*pathname*'

Arguments

'*product-name*'

specifies a character constant, variable, or expression that is the name of the product image you are checking.

'*pathname*'

specifies the pathname for the product image you are checking.

Details

The MODEXIST function searches the directories that are listed in the *pathname* argument for an executable module. The name of the executable module is passed to MODEXIST. MODEXIST returns 1 if the module is found, and 0 if the module is not found.

MOPEN Function

Opens a file by directory ID and member name, and returns either the file identifier or a 0.

Category: External Files

UNIX specifics: OPEN modes

See: MOPEN Function in *SAS Language Reference: Dictionary*

Syntax

MOPEN(*directory-id*,*member-name*<,&i>open-mode<,&i>record-length<,&i>record-format>>>)

Note: This version is a simplified version of the MOPEN function syntax. For the complete syntax and its explanation, see the MOPEN function in *SAS Language Reference: Dictionary*. Δ

open-mode

specifies the type of access to the file:

A	APPEND mode allows writing new records after the current end of the file.
I	INPUT mode allows reading only (default).
O	OUTPUT mode defaults to the OPEN mode specified in the host option in the FILENAME statement or function. If no host option is specified, it allows writing new records at the beginning of the file.
S	Sequential input mode is used for pipes and other sequential devices such as hardware ports.
U	UPDATE mode allows both reading and writing.
W	Sequential update mode is used for pipes and other sequential devices such as ports.

Details

MOPEN returns the identifier for the file, or 0 if the file could not be opened.

PATHNAME Function

Returns the physical name of a SAS library or an external file, or returns a blank.

Category: SAS File I/O

UNIX specifics: *fileref* or *libref* argument can also specify a UNIX environment variable

See: PATHNAME Function in *SAS Language Reference: Dictionary*

Syntax

PATHNAME((*fileref* | *libref*) <,*search-ref*>)

fileref

specifies the *fileref* that is assigned to the external file. In a DATA step, *fileref* can be a character expression, a string enclosed in quotation marks, or a DATA step variable whose value contains the *fileref*. In a macro, *fileref* can be any expression.

The value of *fileref* can be a UNIX environment variable.

libref

specifies the *libref* that is assigned to a SAS library. In a DATA step, *libref* can be a character expression, a string enclosed in quotation marks, or a DATA step variable whose value contains the *libref*. In a macro, *libref* can be any expression.

The value of *libref* can be a UNIX environment variable.

search-ref

specifies whether to search for a fileref or a libref.

F specifies a search for a fileref.

L specifies a search for a libref.

Details

PATHNAME returns the physical name of an external file or SAS library, or a blank if *fileref* or *libref* is invalid.

For more information about using a UNIX environment variable for *fileref* or *libref*, see “FILENAME Function” on page 267.

PEEKLONG Function

Stores the contents of a memory address in a numeric variable on 32-bit and 64-bit platforms.

Category: Special

UNIX specifics: All

See: PEEKLONG Function in *SAS Language Reference: Dictionary*

Syntax

PEEKCLONG(*address,length*);

PEEKLONG(*address,length*);

address

specifies the character string that is the memory address.

length

specifies the data length.

Details

CAUTION:

Use the PEEKLONG functions only to access information returned by one of the MODULE functions. Δ

The PEEKLONG function returns a value of length *length* that contains the data that starts at memory address *address*.

Here are the variations of the PEEKLONG functions:

PEEKCLONG accesses character strings.

PEEKLONG accesses numeric values.

Usually, when you need to use one of the PEEKLONG functions, you will use PEEKCLONG to access a character string. The PEEKLONG function is mentioned for completeness.

RANK Function

Returns the position of a character in the ASCII collating sequence.

Category: Character

UNIX specifics: Uses ASCII collating sequence

See: RANK Function in *SAS Language Reference: Dictionary*

Syntax

RANK(*x*)

x

specifies a character constant, variable, or expression that contains a character in the ASCII collating sequence. If the length of *x* is greater than 1, you receive the rank of the first character in the string.

Details

Because UNIX uses the ASCII character set, the RANK function returns an integer that represents the position of a character in the ASCII collating sequence.

SYSGET Function

Returns the value of the specified operating environment variable.

Category: Special

UNIX specifics: *environment-variable* is a UNIX environment variable

See: SYSGET Function in *SAS Language Reference: Dictionary*

Syntax

SYSGET('environment-variable')

environment-variable

is the name of a UNIX environment variable.

Details

The SYSGET function returns the value of an environment variable as a character string. For example, this statement returns the value of the HOME environment variable:

```
here=sysget('HOME');
```

TRANSLATE Function

Replaces specific characters in a character expression.

Category: Character

UNIX specifics: *to* and *from* arguments are required

See: TRANSLATE Function in *SAS Language Reference: Dictionary*

Syntax

TRANSLATE(*source,to-1,from-1<,...to-n,from-n>*)

Note: This version is a simplified version of the TRANSLATE function syntax. For the complete syntax and its explanation, see the TRANSLATE function in *SAS Language Reference: Dictionary*. Δ

source

specifies a constant, variable, or expression that contains the original character value.

to


specifies the characters that you want TRANSLATE to use as substitutes.

from

specifies the characters that you want TRANSLATE to replace.

Details

Under UNIX, you must specify pairs of *to* and *from* arguments, and you can use a comma as a placeholder.



CHAPTER

14

Informats under UNIX

SAS Informats under UNIX 279

HEXw. Informat 279

\$HEXw. Informat 280

IBw.d Informat 280

PDw.d Informat 280

PIBw.d Informat 281

RBw.d Informat 281

ZDw.d Informat 283

SAS Informats under UNIX

This section describes SAS informats that have behavior or syntax that is specific to UNIX environments. Each informat description includes a brief “UNIX specifics” section that explains which aspect of the informat is specific to UNIX. All of these informats are described in this documentation and in *SAS Language Reference: Dictionary*.

HEXw. Informat

Converts hexadecimal positive binary values to either fixed-point or floating-point binary values.

Category: Numeric

Width range: 1 to 16

Default width: 8

UNIX specifics: floating-point representation

See: HEXw. informat in *SAS Language Reference: Dictionary*

Details

The HEXw. informat converts the hexadecimal representation of positive binary numbers to real floating-point binary values. The width value of the HEXw. informat determines whether the input represents an integer (fixed-point) or real (floating-point) binary number. When you specify a width of 1 through 15, the informat interprets the input hexadecimal as an integer binary number. When you specify 16 for the width value, the informat interprets the input hexadecimal as a floating-point value.

For more information, see “Reading and Writing Binary Data in UNIX Environments” on page 216.

\$HEXw. Informat

Converts hexadecimal data to character data.

Category: Character

Width range: 1 to 32,767

Default width: 2

UNIX specifics: values are interpreted as ASCII values

See: \$HEXw. informat in *SAS Language Reference: Dictionary*

Details

The \$HEXw. informat converts every two digits of hexadecimal data into one byte of character data. Use the \$HEXw. informat to encode hexadecimal values into a character variable when your input data is limited to printable characters. SAS under UNIX interprets values that are read with this informat as ASCII values.

IBw.d Informat

Reads integer binary (fixed-point) values.

Category: Numeric

Width range: 1 to 8

Default width: 4

Decimal range: 0 to 10

UNIX specifics: byte values

See: IBw.d informat in *SAS Language Reference: Dictionary*

Details

The IBw.d informat reads fixed-point binary values. For integer binary data, the high-order bit is the value’s sign: 0 for positive values, 1 for negative. Negative values are represented in two’s-complement notation. If the informat includes a *d* value, the data value is divided by 10^d .

For more details, see “Reading and Writing Binary Data in UNIX Environments” on page 216.

PDw.d Informat

Reads data that is stored in packed decimal format.

Category: Numeric
Width range: 1 to 16
Default width: 1
Decimal range: 0 to 31
UNIX specifics: data representation
See: PDw.d informat in *SAS Language Reference: Dictionary*

Details

The PDw.d informat reads packed decimal data. Although it is usually impossible to type packed decimal data directly from a console, many programs write packed decimal data.

Each byte contains two digits in packed decimal data. The value's sign is the first byte. Because the entire first byte is used for the sign, you should specify at least a width of 2.

The PDw.d format writes missing numerical data as -0. When the PDw.d informat reads a value of -0, the result is a value of 0.

For more information, see "Reading and Writing Binary Data in UNIX Environments" on page 216.

PIBw.d Informat

Reads positive integer binary (fixed-point) values.

Category: Numeric
Width range: 1 to 8
Default width: 1
Decimal range: 0 to 10
UNIX specifics: byte order
See: PIBw.d informat in *SAS Language Reference: Dictionary*

Details

The PIBw.d informat reads integer binary (fixed-point) values. Positive integer binary values are the same as integer binary (see "IBw.d Informat" on page 280), except that all values are treated as positive. Thus, the high-order bit is part of the value rather than the value's sign. If the informat includes a *d* value, the data value is divided by 10^d .

For more information, see "Reading and Writing Binary Data in UNIX Environments" on page 216.

RBw.d Informat

Reads numeric data that is stored in real binary (floating-point) notation.

Category: Numeric

Width range: 2 to 8

Default width: 4

Decimal range: 0 to 10

UNIX specifics: floating-point representation; supports single-precision numbers only for those applications that truncate numeric data

See: RBw.d informat in *SAS Language Reference: Dictionary*

Details

The RBw.d informat reads numeric data that is stored in real binary (floating-point) notation. SAS stores all numeric values in floating-point.

It is usually impossible to type floating-point binary data directly from a console, but many programs write floating-point binary data. Use caution if you are using the RBw.d informat to read floating-point data created by programs other than SAS because the RBw.d informat is designed to read only double-precision data.

All UNIX systems that are currently supported by SAS use the IEEE standard for floating-point representation. This representation supports both single-precision and double-precision floating-point numbers. Double-precision representation has more bytes of precision, and the data within the representation is interpreted differently. For example, for single-precision, the value of 1 in hexadecimal representation is **3F800000**. For double-precision, the hexadecimal representation of 1 is **3FF0000000000000**.

The RBw.d informat is designed to read only double-precision data. It supports widths less than 8 only for applications that truncate numeric data for space-saving purposes. RB4. does *not* expect a single-precision floating-point number; it expects a double-precision number truncated to four bytes. Using the example of 1 above, RB4. expects **3FF00000** to be the hexadecimal representation of the four bytes of data to be interpreted as 1. If given **3F800000**, the single-precision value of 1, a different number results.

External programs such as those programs that are written in the C and FORTRAN languages can produce only single- or double-precision floating-point numbers. No length other than four or eight bytes is allowed. RBw.d allows a length of 3 through 8, depending on the storage you need to save.

The FLOAT4. informat has been created to read a single-precision floating-point number. If you read 3F800000 with FLOAT4., the result is a value of 1.

To read data created by a C or FORTRAN program, you need to decide on the proper informat to use. If the floating-point numbers require an eight-byte width, you should use the RB8. informat. If the floating point numbers require a four-byte width, you should use FLOAT4.

Consider the following C example:

```
#include <stdio.h>

main() {

FILE *fp;
float x[3];

fp = fopen("test.dat","wb");
x[0] = 1; x[1] = 2; x[2] = 3;

fwrite((char *)x,sizeof(float),3,fp);
fclose(fp);
}
```


The file Test.dat contains **3f8000004000000040400000** in hexadecimal representation. The following statements read Test.dat correctly:

```
data _null_;
  infile 'test.dat';
  input (x y z) (float4.);
run;
```

Also available is the *IEEEw.d* informat, which reads IEEE floating-point data. On UNIX systems, *IEEE8.* is equivalent to *RB8.*, and *IEEE4.* is equivalent to *FLOAT4.* *IEEEw.d* can be used on any platform, as long as the original IEEE binary data originated on a platform that uses the IEEE representation.

For more information, see “Reading and Writing Binary Data in UNIX Environments” on page 216.

ZDw.d Informat

Reads zoned decimal data.

Category: Numeric

Width range: 1 to 32

Default width: 1

UNIX specifics: last byte includes the sign; data representation


See: ZDw.d informat in *SAS Language Reference: Dictionary*

Details

The *ZDw.d* informat reads zoned decimal data; it is also known as overprint trailing numeric format. Under UNIX, the last byte of the field includes the sign along with the last digit. The conversion table for the last byte is as follows:

Digit	ASCII Character	Digit	ASCII Character
0	{	-0	}
1	A	-1	J
2	B	-2	K
3	C	-3	L
4	D	-4	M
5	E	-5	N
6	F	-6	O
7	G	-7	P
8	H	-8	Q
9	I	-9	R

For more information, see “ZDw.d Format” on page 254 and “Reading and Writing Binary Data in UNIX Environments” on page 216.



CHAPTER

15

Macro Facility under UNIX

<i>About the Macro Facility under UNIX</i>	285
<i>Automatic Macro Variables in UNIX Environments</i>	285
<i>Macro Statements in UNIX Environments</i>	287
<i>Macro Functions in UNIX Environments</i>	287
<i>SAS System Options Used by the Macro Facility in UNIX Environments</i>	288
<i>Using Autocall Libraries in UNIX Environments</i>	288
<i>What Is an Autocall Library?</i>	288
<i>Available Autocall Macros</i>	288
<i>Guidelines for Naming Macro Files</i>	288
<i>The SASAUTOS System Option</i>	289
<i>Example: Setting Up and Testing a Macro in an Autocall Library</i>	289

About the Macro Facility under UNIX

Most features of the SAS macro facility are valid in all operating environments. This documentation discusses only those components of the macro facility that depend on the UNIX environment. For more information, see the following documentation:

- *SAS Macro Language: Reference*
- *SAS Macro Facility Tips and Techniques*
- the online Help for the macro facility

Automatic Macro Variables in UNIX Environments

The following automatic macro variables are valid in all operating environments, but their values are determined by the operating environment:

SYSCC

contains the current SAS condition code. Upon exit, SAS translates this condition code to a return code that has a meaningful value for the operating environment.

Note: The value of SYSCC might not match the return code returned by the operating system. Δ

Under UNIX, the following codes can be returned:

0	Normal completion
1	SAS issued warnings
2	SAS issued errors
3	ABORT;

- 4 ABORT RETURN *n*;
- 5 ABORT ABEND *n*;
- 6 Internal error

Note: When ERRORCHECK=NORMAL, the return code will be 0, even if an error exists in a LIBNAME or FILENAME statement, or in a LOCK statement in SAS/SHARE software. Also, the SAS job or session will not abort when the %INCLUDE statement fails due to a nonexistent file. For more information, see the “ERRORCHECK= System Option” in *SAS Language Reference: Dictionary*. Δ

SYSDEVIC

contains the name of the current graphics device. The current graphics device is determined by the DEVICE system option. Contact your on-site SAS support personnel to determine which graphics devices are available at your site. (See “DEVICE System Option” on page 363 and *SAS Language Reference: Dictionary* for information about the DEVICE system option.)

SYSENV

reports whether SAS is running interactively. Values for SYSENV are **FORE** when the TERMINAL system option is in effect, and **BACK** when the NOTERMINAL system option is in effect.

SYSJOBID

lists the process identification number (PID) of the process that is executing SAS (for example, 00024).

SYSMAXLONG

returns the maximum long integer value allowed under UNIX, which is 9,007,199,254,740,992. On 32-bit systems, the maximum is 2,147,483,647.

SYSRC

holds the decimal value of the exit status code that is returned by the last UNIX command executed from your SAS session. The following output shows an interactive line mode SAS session that shows two sample SYSRC values:

Output 15.1 Sample SYSRC Values

```
1? x 'data';
/bin/ksh: data: not found
2? %put UNIX exit status code is &sysrc;
UNIX exit status code is 256
3? x 'date';
Tue Mar 19 09:41:27 CST 2008
4? %put UNIX exit status code is now &sysrc;
UNIX exit status code is now 0
```

SYSSCP

returns the abbreviation for your operating environment, such as **HP 800**, **SUN 4**, or **RS6000**.

SYSSCPL

returns the name of the specific UNIX environment that you are using, such as **HP-UX**, **SunOS**, or **AIX**. This variable returns the same value that is returned by the UNIX command **uname**.

Macro Statements in UNIX Environments

The arguments that can be entered with the following statements depend on the operating environment:

%SYSEXEC

executes UNIX commands. It is similar to the X statement described in “Executing Operating System Commands from Your SAS Session” on page 14. The %SYSEXEC statement enables you to execute operating environment commands immediately and, if necessary, determine whether they executed successfully by examining the value of the automatic macro variable SYSRC. You can use the %SYSEXEC statement inside a macro or in open code. The form of the %SYSEXEC statement is as follows, where *command* can be any UNIX command:

```
%SYSEXEC <command>;
```

For example, the following code writes the status of the default printer to your UNIX shell:

```
%sysexec lpstat;
```

Entering %SYSEXEC without a UNIX command starts a new shell, except under the X interface to SAS. See “Executing Operating System Commands from Your SAS Session” on page 14 for information.

Macro Functions in UNIX Environments

The following functions have operating environment dependencies:

%SCAN

searches for a word that is specified by its position in a string. Here is the form of the %SCAN function:

```
%SCAN(argument,n,<delimiters>;
```

On ASCII systems, the default delimiters are the following:

```
blank . < ( + & ! $ * ) ; ^ - / , % |
```

%SYSGET

returns the character string that is the value of the environment variable passed as the argument. Both UNIX and SAS environment variables can be translated using the %SYSGET function. A warning message is written if the global variable does not exist. Here is the form of the %SYSGET function:

```
%SYSGET(environment-variable);
```

For example, the following code writes the value of the HOME environment variable to the SAS log:

```
%let var1=%sysget(HOME);
%put &var1;
```

SAS System Options Used by the Macro Facility in UNIX Environments

The following system options have operating environment dependencies:

MSYMTABMAX

specifies the maximum amount of memory available to all symbol tables (global and local, combined). Under UNIX, the default value for this option is 4M. For more information, see *SAS Language Reference: Dictionary*.

MVARSIZE

specifies the maximum number of bytes for any macro variable stored in memory. Under UNIX, the default value for this option is 32K. For more information, see *SAS Language Reference: Dictionary*.

SASAUTOS

specifies the AUTOCALL library. For more information, see “The SASAUTOS System Option” on page 289.

Using Autocall Libraries in UNIX Environments

What Is an Autocall Library?

An autocall library contains files that define SAS macros. The following sections discuss aspects of autocall libraries that are dependent on the operating environment. For more information, see *SAS Macro Language: Reference*.

Available Autocall Macros

There are two types of autocall macros, those macros that are provided by SAS, and those macros that you define yourself. To use the autocall facility, you must have the MAUTOSOURCE system option set.

When SAS is installed, the SASAUTOS system option is defined in the configuration file to refer to the location of the default macros supplied by SAS. The products licensed at your site determine the autocall macros you have available. You can also define your own autocall macros and store them in one or more directories. SAS does not recognize autocall macros if their filenames are written in uppercase or in mixed case. Use only filenames that are lowercase.

Guidelines for Naming Macro Files

Macro names in SAS are case insensitive, but they all map to a lowercase filename. If you store autocall macros in a UNIX directory, the file extension must be **.sas**, and the filename must be entirely in lowercase. In the UNIX environment, each macro file in the directory must contain a macro definition with a macro name that matches the filename. For example, a file named **prtdata.sas** should define a macro named **prtdata**.

The SASAUTOS System Option

To use your own autocall macros in your SAS program, specify their directories with the SASAUTOS system option. For more information, see “SASAUTOS System Option” on page 398.

Note: The SASAUTOS system option under UNIX does not recognize filenames that are in uppercase or mixed case. Δ

You can set the SASAUTOS system option when you start SAS, or you can use it in an OPTIONS statement during your SAS session. However, autocall libraries specified with the OPTIONS statement override any previous specification.

If you use the CONFIG system option to specify a configuration file, add your autocall library to the library concatenation supplied by SAS. If you use the default configuration files (sasv9.cfg), specify your autocall library there.

Autocall libraries are searched in the order in which you specify them.


Example: Setting Up and Testing a Macro in an Autocall Library

This example shows how to set up and test a macro in an autocall library.

The following output shows the results of executing two UNIX (**cat**) commands to display the contents of two files, and a SAS command to run the autocall.sas program:

Output 15.2 AUTOCALL Library Example

```
$ cat maclib/testauto.sas
%macro testauto;
x echo 'Autocall library is working.';
%mend testauto;
$ cat source/autocall.sas
filename sysautos ('!SASROOT/sasautos' '$HOME/test/sasautos');
options mautosource sasautos=(sysautos '$HOME/macros/maclib');
%testauto
%TestAuto
%TESTAUTO
$ sas source/autocall.sas
Autocall library is working.
Autocall library is working.
Autocall library is working.
```

CHAPTER 16

Procedures under UNIX

<i>SAS Procedures under UNIX</i>	291
<i>CATALOG Procedure</i>	291
<i>CIMPORT Procedure</i>	292
<i>CONTENTS Procedure</i>	293
<i>CONVERT Procedure</i>	295
<i>CPORT Procedure</i>	298
<i>DATASETS Procedure</i>	299
<i>OPTIONS Procedure</i>	302
<i>PMENU Procedure</i>	303
<i>PRINTTO Procedure</i>	304
<i>SORT Procedure</i>	305

SAS Procedures under UNIX

This section describes SAS procedures that have behavior or syntax that is specific to UNIX environments. Each procedure description includes a brief “UNIX specifics” section that explains which aspect of the procedure is specific to UNIX. Each procedure is described in both this documentation and in the *Base SAS Procedures Guide*.

CATALOG Procedure

Manages entries in SAS catalogs.

UNIX specifics: FILE= option in the CONTENTS statement

See: CATALOG Procedure in *Base SAS Procedures Guide*

Syntax

```
PROC CATALOG CATALOG=<libref.>catalog <ENTRYTYPE=etype> <KILL>;
  CONTENTS <OUT=SAS-data-set> <FILE=fileref>;
```

Note: This version is a simplified version of the CATALOG procedure syntax. For the complete syntax and its explanation, see the CATALOG procedure in the *Base SAS Procedures Guide*. Δ

fileref

names a file specification that is specific to the UNIX operating environment.

Details

The FILE= option in the CONTENTS statement of the CATALOG procedure accepts a fileref. If the name specified does not correspond to a fileref, a file with that name and an extension of **.lst** is created in the current directory. For example, if **myfile** is not a fileref, the following code creates the file **myfile.lst** in your current directory:

```
proc catalog catalog=sasuser.profile;
  contents file=myfile;
run;
```

SAS writes the following output to the Log:

```
NOTE: 6 entries have been written to the output file /users/userid/MYFILE.lst.
```

Note: The filename that is created is always stored in lowercase, even if you specified it in uppercase. In the SAS log, however, the filename is listed in uppercase. Δ

CIMPORT Procedure

Restores a transport file that was created by the CPORT procedure.

UNIX specifics: name and location of transport file

See: CIMPORT Procedure in *Base SAS Procedures Guide*

Syntax

PROC CIMPORT *destination=libref* | *<libref.>member-name* *<option(s)>*;

Note: This version is a simplified version of the CIMPORT procedure syntax. For the complete syntax and its explanation, see the CIMPORT procedure in the *Base SAS Procedures Guide*. Δ

destination

identifies the files in the transport file as a single SAS data set, single SAS catalog, or multiple members of a SAS library.

libref* | *<libref.>member-name

specifies the name of the SAS data set, catalog, or library to be created from the transport file.

Details

Note: Starting in SAS 9.1, you can use the MIGRATE procedure to migrate a SAS library from a previous release. For more information, see “Migrating 32-Bit SAS Files to 64-Bit in UNIX Environments” on page 44, and Migration at support.sas.com/migration. Δ

The CIMPORT procedure *imports* a transport file that was created (*exported*) by the CPORT procedure. The transport file can contain a SAS data set, a SAS catalog, or an entire SAS library.

Typically, the INFILE= option is used to designate the source of the transport file. If this option is omitted, CIMPORT uses the default file Sascat.dat in the current directory as the transport file.

Note: CIMPORT works only with transport files created by the CPORT procedure. If the transport file was created using the XPORT engine with the COPY procedure, then another PROC COPY must be used to restore the transport file. For more information about PROC COPY, see the *Base SAS Procedures Guide*. △

Example

For this example, a SAS library that contains multiple SAS data sets was exported to a file (called **transport-file**) using the CPORT procedure on a foreign host. The transport file is then moved by a binary transfer to the receiving host.

The following code extracts all of the SAS data sets and catalogs stored within the transport file and restores them to their original state in the new library, called **SAS-library**.

```
libname newlib 'SAS-library';
filename tranfile 'transport-file';

proc cimport lib=newlib infile=tranfile;
run;
```

See Also

- “CPORT Procedure” on page 298
- “Migrating 32-Bit SAS Files to 64-Bit in UNIX Environments” on page 44
- The MIGRATE Procedure and Cross-Release Compatibility at support.sas.com/rnd/migration
- *Moving and Accessing SAS Files*

CONTENTS Procedure

Prints the description of the contents of one or more files from a SAS library.

UNIX specifics: information displayed in the SAS output

See: CONTENTS Procedure in *Base SAS Procedures Guide*

Syntax

PROC CONTENTS<option(s)>;

Details

The CONTENTS procedure produces the same information as the CONTENTS statement in the DATASETS procedure. (See “DATASETS Procedure” on page 299 for a comparison.)

Example

The following SAS code creates two data sets, **classes.grades** and **classes.majors**, and executes PROC CONTENTS to describe the **classes.majors** data set:

```
options nodate pageno=1;
libname classes '.';

data classes.grades (label='First Data Set');
  input student year state $ grade1 grade2;
  label year='Year of Birth';
  format grade1 4.1;
  datalines;
1000 1980 NC 85 87
1042 1981 MD 92 92
1095 1979 PA 78 72
1187 1980 MA 87 94
;

data classes.majors(label='Second Data Set');
  input student $ year state $ grade1 grade2 major $;
  label state='Home State';
  format grade1 5.2;
  datalines;
1000 1980 NC 84 87 Math
1042 1981 MD 92 92 History
1095 1979 PA 79 73 Physics
1187 1980 MA 87 74 Dance
1204 1981 NC 82 96 French
;

proc contents data=classes.majors;
run;
```

Output 16.1 Output from the CONTENTS Procedure

The SAS System				1	
The CONTENTS Procedure					
Data Set Name	CLASSES.MAJORS	Observations	5		
Member Type	DATA	Variables	6		
Engine	V9	Indexes	0		
Created	Tuesday, August 22, 2006 03:10:46 PM	Observation Length	48		
Last Modified	Tuesday, August 22, 2006 03:10:46 PM	Deleted Observations	0		
Protection		Compressed	NO		
Data Set Type		Sorted	NO		
Label	Second Data Set				
Data Representation	HP_UX_64, RS_6000_AIX---64, SOLARIS_64, HP_IA64				
Encoding	latin1 Western (ISO)				
Engine/Host Dependent Information					
Data Set Page Size	8192				
Number of Data Set Pages	1				
First Data Page	1				
Max Obs per Page	169				
Obs in First Data Page	5				
Number of Data Set Repairs	0				
File Name	/user/xxxxxx/majors.sas7bdat				
Release Created	9.0201B0				
Host Created	HP_UX				
Inode Number	592				
Access Permission	rw-r- -r- -				
Owner Name	xxxxxx				
File Size (bytes)	16384				
Alphabetic List of Variables and Attributes					
#	Variable	Type	Len	Format	Label
4	grade1	Num	8	5.2	
5	grade2	Num	8		
6	major	Char	8		
3	state	Char	8		Home State
1	student	Char	8		
2	year	Num	8		

CONVERT Procedure

Converts BMDP and OSIRIS system files and SPSS export files to SAS data sets.

UNIX specifics: all

Syntax

PROC CONVERT *product-specification* <option-list>;

Details

The CONVERT procedure converts BMDP and OSIRIS system files, and SPSS export files to SAS data sets. The procedure is supplied for compatibility. The procedure invokes the appropriate engine to convert files.

PROC CONVERT produces one output data set, but no printed output. The new data set contains the same information as the input system file. Exceptions are noted in “How Missing Values Are Handled” on page 297.

The procedure converts system files from these products:

- BMDP saves files up to and including the most recent release of BMDP (available for AIX, HP-UX, and Solaris only).
- OSIRIS saves files through and including OSIRIS IV. (Hierarchical file structures are not supported.)

Because the BMDP, OSIRIS, and SPSS products are maintained by other organizations, changes might be made that make new files incompatible with the current version of PROC CONVERT. SAS upgrades PROC CONVERT to support changes to these products only when a new version of SAS is released.

In the PROC CONVERT statement, *product-specification* is required and can be one of the following:

BMDP=*fileref* <(CODE=*code* CONTENT=*content-type*)>

converts the first member of a BMDP save file created under UNIX (AIX) into a SAS data set. Here is an example:

```
filename save '/usr/mydir/bmdp.dat';
proc convert bmdp=save;
run;
```

If you have more than one save file in the BMDP file referenced by the *fileref* argument, you can use two options in parentheses after *fileref*. The CODE= option specifies the code of the save file that you want, and the CONTENT= option specifies the content of the save file. For example, if a file with **code=judges** has a content of DATA, you can use the following statements:

```
filename save '/usr/mydir/bmdp.dat';
proc convert bmdp=save(code=judges
                      content=data);
run;
```

OSIRIS=*fileref*|*libref*

specifies a *fileref* or *libref* for the OSIRIS file to be converted into a SAS data set. You must also include the DICT= option.

SPSS=*fileref*|*libref*

specifies a *fileref* or *libref* for the SPSS export file that is to be converted into a SAS data set. The SPSS file must be created by using the SPSS EXPORT command, but it can be from any operating environment.

The *option-list* can be one or more of the following:

DICT=*fileref*|*libref*

specifies a *fileref* or *libref* of the dictionary file for the OSIRIS file. DICT= is valid only when used with the OSIRIS product specification.

FIRSTOBS=*n*

gives the number of the observation where the conversion is to begin, so that you can skip observations at the beginning of the BMDP, OSIRIS, or SPSS file.

OBS=*n*

specifies the number of the last observation to be converted. This option enables you to exclude observations at the end of the file.

OUT=SAS-data-set

names the SAS data set that will hold the converted data. If *OUT=* is omitted, SAS still creates a Work data set and automatically names it *DATA n* , just as if you had omitted a data set name in a *DATA* statement. See Chapter 2, “Using SAS Files,” on page 31 for more information.

How Missing Values Are Handled

If a numeric variable in the input data set has no value or a system missing value, PROC CONVERT assigns it a missing value.

How Variable Names Are Assigned

The following sections explain how names are assigned to the SAS variables created by the CONVERT procedure.

CAUTION:

Make sure that the translated names will be unique. Variable names are translated as indicated in the following sections. Δ

Variable Names in BMDP Output Variable names from the BMDP save file are used in the SAS data set, but nontrailing blanks and all special characters are converted to underscores in the SAS variable names. The subscript in BMDP variable names, such as *x(1)*, becomes part of the SAS variable name with the parentheses omitted: *X1*. Alphabetic BMDP variables become SAS character variables of corresponding length. Category records from BMDP are not accepted.

Variable Names in OSIRIS Output For single-response variables, the *V1* through *V9999* name becomes the SAS variable name. For multiple-response variables, the suffix *R n* is added to the variable name where *n* is the response. For example, *V25R1* would be the first response of the multiple-response *V25*. If the variable after *V1000* has 100 or more responses, responses above 99 are eliminated. Numeric variables that OSIRIS stores in character, fixed-point binary, or floating-point binary mode become SAS numeric variables. Alphabetic variables become SAS character variables; any alphabetic variable of length greater than 200 is truncated to 200. The OSIRIS variable description becomes a SAS variable label, and OSIRIS print format information becomes a SAS format.

Variable Names in SPSS Output SPSS variable names and variable labels become variable names and labels without change. SPSS alphabetic variables become SAS character variables of the same length. SPSS blank values are converted to SAS missing values. SPSS print formats become SAS formats, and the SPSS default precision of no decimal places becomes part of the variables' formats. The SPSS DOCUMENT data is copied so that the CONTENTS procedure can display it. SPSS value labels are not copied.

File Conversion Examples

These three examples show how to convert BMDP, OSIRIS, and SPSS files to SAS data sets.

Converting a BMDP save file

The following statements convert a BMDP save file and produce the temporary SAS data set *ttemp*, which contains the converted data:

```
filename bmdpfile 'bmdp.savefile';
proc convert bmdp=bmdpfile out=temp;
run;
```

Converting an OSIRIS file

The following statements convert an OSIRIS file and produce the temporary SAS data set **temp**, which contains the converted data:

```
filename osirfile 'osirdata';
filename dictfile 'osirdict';
proc convert osiris=osirfile dict=dictfile
            out=temp;
run;
```

Converting an SPSS file

The following statements convert an SPSS file and produce the temporary SAS data set **temp**, which contains the converted data:

```
filename spssfile 'spssfile.num1';
proc convert spss=spssfile out=temp;
run;
```

Comparison with Interface Library Engines

The CONVERT procedure is closely related to the interface library engines BMDP, OSIRIS, and SPSS. (In fact, the CONVERT procedure uses these engines.) For example, the following two sections of code provide identical results:

```
filename myfile 'mybmdp.dat';
proc convert bmdp=myfile out=temp;
run;

libname myfile bmdp 'mybmdp.dat';
data temp;
  set myfile._first_;
run;
```

However, the BMDP, OSIRIS, and SPSS engines provide more extensive capability than PROC CONVERT. For example, PROC CONVERT converts only the first BMDP member in a save file. The BMDP engine, in conjunction with the COPY procedure, copies all members.

See Also

- “Accessing BMDP, OSIRIS, or SPSS Files in UNIX Environments” on page 62

CPORT Procedure

Writes SAS data sets and catalogs into a transport file.

UNIX specifics: name and location of transport file

See: CPORT Procedure in *Base SAS Procedures Guide*

Syntax

PROC CPORT *source-type=libref* | *<libref.>member-name* *<option(s)>*;

Note: This version is a simplified version of the CPORT procedure syntax. For the complete syntax and its explanation, see the “CPORT Procedure” in the *Base SAS Procedures Guide*. Δ

source-type

identifies the files to export as either a single SAS data set, single SAS catalog, or multiple members of a SAS library.

libref | *<libref.> member-name*

specifies the name of the SAS data set, catalog, or library to be exported.

Details

Note: Starting in SAS 9.1, you can use the MIGRATE procedure to migrate a SAS library from a previous release. For more information, see “Migrating 32-Bit SAS Files to 64-Bit in UNIX Environments” on page 44, and Migration at support.sas.com/migration. Δ

The CPORT procedure creates a transport file to later be restored (*imported*) by the CIMPORT procedure. The transport file can contain a SAS data set, SAS catalog, or an entire SAS library.

Typically, the FILE= option is used to specify the path of the transport file. The value of the FILE= option can be a fileref defined in a FILENAME statement or an environment variable. If this option is omitted, CPORT creates the default file Sascat.dat in the current directory as the transport file.

Examples

In this example, a SAS library (called **oldlib**) that contains multiple SAS data sets is being exported to the file, called **transport-file**:

```
libname oldlib 'SAS-data-library';
filename tranfile 'transport-file';

proc cport lib=oldlib file=tranfile;
run;
```

This transport file is then typically moved by binary transfer to a different host, where the CIMPORT procedure will be used to restore the SAS library.

See Also

- “CIMPORT Procedure” on page 292
- “Migrating 32-Bit SAS Files to 64-Bit in UNIX Environments” on page 44
- The MIGRATE Procedure and Cross-Release Compatibility at support.sas.com/rnd/migration
- Moving and Accessing SAS Files*

DATASETS Procedure

Manages SAS files, and creates and deletes indexes and integrity constraints for SAS data sets.

UNIX specifics: Directory information, CONTENTS statement output

See: DATASETS Procedure in the *Base SAS Procedures Guide*

Syntax

```
PROC DATASETS< option(s)>;
  CONTENTS <option(s)>
```

Note: This version is a simplified version of the DATASETS procedure syntax. For the complete syntax and its explanation, see the DATASETS procedure in *Base SAS Procedures Guide*. Δ

CONTENTS *option(s)*

the value for *option(s)* can be the following:

DIRECTORY

prints a list of information specific to the UNIX operating environment.

Details

The output from the DATASETS procedure shows you the libref, engine, and physical name that are associated with the library, as well as the names and other properties of the SAS files that are contained in the library. Some of the SAS library information, such as the filenames and access permissions, that is displayed in the SAS log by the DATASETS procedure depends on the operating environment and the engine. The information generated by the CONTENTS statement also varies according to the device type or access method associated with the data set.

If you specify the DIRECTORY option in the CONTENTS statement, the directory information is displayed in both the Log and Output windows.

The CONTENTS statement in the DATASETS procedure generates the same engine and host dependent information as the CONTENTS procedure.

Example

The following SAS code creates two data sets, **classes.grades** and **classes.majors**, and executes PROC DATASETS using **classes.majors** as the input data set.

The first page of output from this example is produced by the DIRECTORY option in the CONTENTS statement. This information also appears in the SAS log. Page 2 in this output describes the data set **classes.majors** and appears only in the SAS output:

```
options nodate pageno=1;
libname classes '.';

data classes.grades (label='First Data Set');
  input student year state $ gradel grade2;
  label year='Year of Birth';
  format gradel 4.1;
  datalines;
1000 1980 NC 85 87
1042 1981 MD 92 92
1095 1979 PA 78 72
1187 1980 MA 87 94
;
```

```

data classes.majors(label='Second Data Set');
  input student $ year state $ gradel grade2 major $;
  label state='Home State';
  format gradel 5.2;
  datalines;
1000 1980 NC 84 87 Math
1042 1981 MD 92 92 History
1095 1979 PA 79 73 Physics
1187 1980 MA 87 74 Dance
1204 1981 NC 82 96 French
;

proc datasets library=classes;
  contents data=majors directory;
run;

```

Output 16.2 Output from the DATASETS Procedure

The SAS System					1
The DATASETS Procedure					
Directory					
Libref			CLASSES		
Engine			V9		
Physical Name			/users/xxxx		
File Name			/users/xxxx		
Inode Number			4259		
Access Permission			rw-rw-rw-		
Owner Name			xxxx		
File Size (bytes)			512		
		Member	File		
#	Name	Type	Size	Last Modified	
1	GRADES	DATA	16384	02May07:07:42:10	
2	MAJORS	DATA	16384	02May07:07:42:11	

The SAS System				2	
The DATASETS Procedure					
Data Set Name	CLASSES.MAJORS	Observations	5		
Member Type	DATA	Variables	6		
Engine	V9	Indexes	0		
Created	Tuesday, August 22, 2006 02:31:32 PM	Observation Length	48		
Last Modified	Tuesday, August 22, 2006 02:31:32 PM	Deleted Observations	0		
Protection		Compressed	NO		
Data Set Type		Sorted	NO		
Label	Second ds				
Data Representation	SOLARIS_X86_64, LINUX_X86_64, ALPHA_TRU64, LINUXIA_64				
Encoding	latin1 Western (ISO)				
Engine/Host Dependent Information					
Data Set Page Size	8198				
Number of Data Set Pages	1				
First Data Page	1				
Max Obs per Page	169				
Obs in First Data Page	5				
Number of Data Set Repairs	0				
File Name	/users/xxxxxx/majors.sas7bdat				
Release Created	9.0201B0				
Host Created	SunOS				
Alphabetic List of Variables and Attributes					
#	Variable	Type	Len	Format	Label
4	grade1	Num	8	5.2	
5	grade2	Num	8		
6	major	Char	8		
3	state	Char	8		Home State
1	student	Char	8		
2	year	Num	8		

See Also

- “CONTENTS Procedure” in the *Base SAS Procedures Guide*

OPTIONS Procedure

Lists the current settings of SAS system options.

UNIX specifics: options available only under UNIX

See: OPTIONS Procedure in *Base SAS Procedures Guide*

Syntax

PROC OPTIONS=<option(s)>

Note: This version is a simplified version of the OPTIONS procedure syntax. For the complete syntax and its explanation, see the OPTIONS procedure in the *Base SAS Procedures Guide*. △

option(s)

HOST | NOHOST

displays only host options (HOST) or only portable options (NOHOST).
PORTABLE is an alias for NOHOST.

Details

PROC OPTIONS lists the current settings of the system options that are available in all operating environments, as well as the system options that are available in the UNIX environment. If you specify the HOST option in the PROC OPTIONS statement, it lists those options that are available only under UNIX (host options). The option values that are displayed by PROC OPTIONS depend on the default values of SAS, the default values specified by your site administrator, the default values in your own configuration file, any changes made in your current session through the System Options window or OPTIONS statement, and possibly, the device on which you are running SAS.

For more information about a specific option, see Chapter 18, “System Options under UNIX,” on page 341.

See Also

- For more information about restricted options, see “Order of Precedence for Processing SAS Configuration Files” on page 21.

PMENU Procedure

Defines menu facilities for windows that are created with SAS software.

UNIX specifics: ATTR= and COLOR= options in the TEXT statement have no effect; ACCELERATE= and MNEMONIC= options in the ITEM statement are ignored

See: PMENU Procedure in *Base SAS Procedures Guide*

Syntax

```
PROC PMENU <CATALOG=<libref.>catalog>
    <DESC 'entry-description'>;
```

Note: This version is a simplified version of the PMENU procedure syntax. For the complete syntax and its explanation, see the PMENU procedure in the *Base SAS Procedures Guide*. △

CATALOG=<libref.>catalog

specifies the catalog in which you want to store PMENU entries. If you omit *libref*, the PMENU entries are stored in a catalog in the Sasuser library. If you omit CATALOG=, the entries are stored in the Sasuser.Profile catalog.

DESC 'entry-description'

provides a description of the PMENU catalog entries created in the step.

Details

The PMENU procedure defines PMENU facilities for windows created by using the WINDOW statement in Base SAS software, the %WINDOW macro statement, the BUILD procedure of SAS/AF software, or the SAS Component Language (SCL) PMENU function with SAS/AF and SAS/FSP software.

Under UNIX, the following options are ignored:

- ATTR= and COLOR= options in the TEXT statement. The colors and attributes for text and input fields are controlled by the CPARMS colors specified in the SASCOLOR window. See “Customizing Colors in UNIX Environments” on page 196 for more information.
- ACCELERATE= and the MNEMONIC= options in the ITEM statement.

PRINTTO Procedure

Defines destinations for SAS procedure output and the SAS log.

UNIX specifics: Valid values of *file specification*

See: PRINTTO Procedure in the *Base SAS Procedures Guide*

Syntax

PROC PRINTTO <option(s)>

Note: This version is a simplified version of the PRINTTO procedure syntax. For the complete syntax and its explanation, see the PRINTTO procedure in *Base SAS Procedures Guide*. Δ

LOG=file-specification

specifies a fully qualified pathname (in quotation marks), an environment variable, a fileref, or a file in the current directory (without extension).

PRINT=file-specification

specifies a fully qualified pathname (in quotation marks), an environment variable, a fileref, or a file in the current directory (without extension). If you specify a fileref that is defined with the PRINTER device-type keyword, output is sent directly to the printer.

Examples

The following statements send any SAS log entries that are generated after the RUN statement to the external file that is associated with the fileref **myfile**:

```
filename myfile '/users/myid/mydir/mylog';
proc printto log=myfile;
run;
```

If **myfile** has not been defined as a fileref, then PROC PRINTTO creates the file **myfile.log** in the current directory.

The following statements send any procedure output that is generated after the RUN statement to the file `/users/myid/mydir/myout`:

```
proc printto print='/users/myid/mydir/myout';
run;
```

The following statements send the procedure output from the CONTENTS procedure directly to the system printer:

```
filename myfile printer;
proc printto print=myfile;
run;

proc contents data=oranges;
run;
```

To redirect the SAS log and procedure output to their original default destinations, run PROC PRINTTO without any options:

```
proc printto;
run;
```

If filerefs `myprint` and `mylog` have not been defined, then the following statements send any SAS procedure output to `myprint.lst` and any log output to `mylog.log` in the current directory:

```
proc printto print=myprint log=mylog;
run;
```

If filerefs `myprint` and `mylog` had been defined, the output would have gone to the files associated with these filerefs.

See Also

- Chapter 4, “Printing and Routing Output,” on page 91

SORT Procedure

Sorts observations in a SAS data set by one or more variables, and then stores the resulting sorted observations in a new SAS data set or replaces the original data set.

UNIX specifics: sort utilities available

See: SORT Procedure in *Base SAS Procedures Guide*

Syntax

PROC SORT<option(s)><collating-sequence-option>

Note: This version is a simplified version of the SORT procedure syntax. For the complete syntax and its explanation, see the SORT procedure in the *Base SAS Procedures Guide*. △

option

SORTSIZE=*memory-specification*

specifies the maximum amount of memory available to the SORT procedure. For more information about the SORTSIZE= option, see “Details” on page 306 .

TAGSORT

stores only the BY variables and the observation numbers in temporary files. The TAGSORT option has no effect on a UNIX host that uses SyncSort.

For more information about the TAGSORT option, see “Details” on page 306.

DETAILS

specifies that PROC SORT write messages to the SAS log detailing whether the sort was performed in memory. (This option is a statement option.)

If the sort was not performed in memory, then the details that are written to the SAS log include the number of utility files that were used and their sizes.

Tip: Using the DETAILS option can help determine an ideal SORTSIZE value.

Details

The SORT procedure sorts observations in a SAS data set by one or more character or numeric variables, either replacing the original data set or creating a new, sorted data set. By default under UNIX, the SORT procedure uses the ASCII collating sequence.

The SORT procedure uses the sort utility that is specified by the SORTPGM system option. Sorting can be done by SAS or by the **syncsort** utility. You can use all of the options that are available to the SAS sort utility, such as the SORTSEQ and NODUPKEY options. In some situations, you can improve your performance by using the NOEQUALS option. If you specify an option that is not supported by the host sort, then the SAS sort will be used instead. For more information about all of the options that are available, see the SORT procedure in the *Base SAS Procedures Guide*.

SORTSIZE= Option

Limiting the Amount of Memory Available to PROC SORT You can use the SORTSIZE= option in the PROC SORT statement to limit the amount of memory that is available to the SORT procedure. This option can reduce the amount of swapping SAS must do to sort the data set.

Note: If you do not specify the SORTSIZE= option, PROC SORT uses the value of the SORTSIZE system option. The SORTSIZE system option can be defined in the command line or in the SAS configuration file. △

Syntax of the SORTSIZE= Option The syntax of the SORTSIZE= option is as follows:

SORTSIZE=*memory-specification*

memory-specification can be one of the following:

<i>n</i>	specifies the amount of memory in bytes.
<i>nK</i>	specifies the amount of memory in 1-kilobyte multiples.
<i>nM</i>	specifies the amount of memory in 1-megabyte multiples.
<i>nG</i>	specifies the amount of memory in 1-gigabyte multiples.

Default Value of the SORTSIZE= Option

The default SAS configuration file sets this option based on the value of the SORTSIZE system option. To view the default value for your operating environment, execute the following code:


```
proc options option=sortsize;
run;
```

You can override the default value of the SORTSIZE system option in one of the following ways:

- by specifying a different SORTSIZE= value in the PROC SORT statement
- by submitting an OPTIONS statement that sets the SORTSIZE system option to a new value
- by setting the SORTSIZE system option in the command line during the invocation of SAS

Improving Performance with the SORTSIZE= Option The SORTSIZE system option limits the amount of memory that is available to PROC SORT. In general, you should set the SORTSIZE= option no larger than the amount of physical memory that is available to the SAS process.

Setting SORTSIZE memory below the default range might adversely affect sorting and SAS performance in general. Setting SORTSIZE memory to a value that is greater than the default might not necessarily improve performance. The key indicator as to whether additional memory can improve performance is whether the sort will fit in memory.

When the SORTSIZE= value is large enough to sort the entire data set in memory, you can achieve optimal sort performance. If the entire data set to be sorted will not fit in memory, SAS creates a temporary utility file to store the data. In this case, SAS uses a sort algorithm that is tuned to sort using disk space instead of memory.

If the SORTSIZE= value is larger than the amount of available memory, then the operating system will be forced to page excessively. If the SORTSIZE= value is too small, then not all of the sorting can be done in memory, which also results in more disk I/O.

If the data set to be sorted does not fit in memory, then setting a SORTSIZE value in the 64–128M range is generally the optimal value. SORTSIZE should always be set to a value that is at least 8M smaller than MEMSIZE.

Note: You can also use the SORTSIZE system option, which has the same effect as the SORTSIZE= option in the PROC SORT statement. △

TAGSORT Option

The TAGSORT option in the PROC SORT statement is useful when there might not be enough disk space to sort a large SAS data set. When you specify the TAGSORT option, only the sort keys (that is, the variables specified in the BY statement) and the observation number for each observation are stored in the temporary utility files. The sort keys, together with the observation number, are referred to as *tags*. At the completion of the sorting process, the tags are used to retrieve the records from the input data set in sorted order. Thus, in cases where the total number of bytes of the sort keys is small compared with the length of the record, temporary disk use is reduced considerably.

You must have enough disk space to hold an additional copy of the data set (the output data set) and the utility file that contains the tags. By default, this utility file is stored in the Work library. If this directory is too small, you can change this directory by using the WORK system option. For more information, see “WORK System Option” on page 415.

Note that while using the TAGSORT option might reduce temporary disk use, the processing time could be higher. However, on systems with limited available disk space, the TAGSORT option might enable data sets to be sorted in situations where that would otherwise not be possible.

Disk Space Considerations for PROC SORT

You need to consider the following information when determining the amount of disk space needed to run PROC SORT:

input SAS data set

PROC SORT uses the SAS input data set specified by the DATA= option.

output SAS data set

PROC SORT stores the output SAS data set in the location specified by the OUT= option. If the OUT= option is not specified, PROC SORT stores the output SAS data set in the Work library.

utility file

If the UTILLOC system option is not set, the utility file is stored in the Work directory. This utility file is approximately the size of the input SAS data set.

Note: You can use the UTILLOC system option to specify a location in which applications can store utility files. Δ

temporary output SAS data set

During the sort, PROC SORT creates its output in the directory specified in the OUT= option (or directory of the input SAS data set if the OUT= option is not specified). The temporary data set has the same filename as the original data set, except it has an extension of .lck. After the sort completes successfully, the original data set is deleted, and the temporary data set is renamed to match the original data set. Therefore, you need to have enough available disk space in the target directory to hold two copies of the data set.

You can reduce the amount of disk space needed by specifying the OVERWRITE option in the PROC SORT statement. When you specify this option, SAS replaces the input data set with the sorted data set. This option should be used only with a data set that is backed up, or with a data set that you can reconstruct. For more information, see the SORT procedure in the *Base SAS Procedures Guide*.

Performance Tuning for PROC SORT

How SAS Determines the Amount of Memory to Use The MEMSIZE system option controls the upper limit for the maximum amount of memory that is available to the SAS process. The SORTSIZE system option limits the amount of memory that is available to PROC SORT. The REALMEMSIZE system option specifies the amount of real (not virtual) memory that will be made available to SAS.

While memory settings below the default values for MEMSIZE and SORTSIZE might adversely affect sorting and SAS performance, making large amounts of memory available might be of no benefit. The key for determining whether additional memory might improve performance is whether the sort will fit in memory. If the sorted file requires more memory than is allocated, then a SORTSIZE value in the range of 64–128M is generally the optimal value. SORTSIZE should always be set to a value that is at least 8M smaller than MEMSIZE.

For information about setting the REALMEMSIZE system option, see “Guidelines for Setting the REALMEMSIZE System Option” on page 308.

Note: If you receive an out of memory error, then increase the value of MEMSIZE. For more information, see “MEMSIZE System Option” on page 384. Δ

Guidelines for Setting the REALMEMSIZE System Option Because PROC SORT might use the REALMEMSIZE system option to determine how much memory to use, it is important that the REALMEMSIZE value reflects the amount of memory that is

available on your system. The default value is 80% of the MEMSIZE setting. If REALMEMSIZE is set too high, then PROC SORT might use more memory than is actually available. Using too much memory will cause excessive paging and adversely impact system performance.

In general, REALMEMSIZE should be set to the amount of physical memory (not including swap space) that you expect to be available to SAS at run time. A good starting value is the amount of physical memory installed on the computer less the amount that is being used by running applications and the operating system. You can experiment with the REALMEMSIZE value until you reach optimum performance for your environment. In some cases, optimum performance can be achieved with a very low REALMEMSIZE value. A low value could cause SAS to use less memory and leave more memory for the operating system to perform I/O caching.

For more information, see “REALMEMSIZE System Option” on page 394.

Using Other Options that Affect Performance The THREADS system option controls whether threaded procedures will use threads. It is available as both a system option and as a procedural override in PROC SORT.

The CPUCOUNT option is directly related to the THREADS option and defaults to the number of CPUs on your computer. Depending on your file system and the number of concurrent users, you might benefit from lowering the CPUCOUNT on machines that have many CPUs. When the value of CPUCOUNT equals ACTUAL, SAS returns the number of physical CPUs that are associated with the operating environment where SAS is executing.

The UTILLOC system option allows for the spreading of utility files, and is a good option for balancing I/O.

The DETAILS option, specified in the PROC SORT statement, causes PROC SORT to write messages to the SAS log detailing whether the sort was performed in memory. If the sort was not performed in memory, then the details that are written include the number of utility files and their sizes.

For more information about the THREADS, CPUCOUNT, and UTILLOC system options see the *SAS Language Reference: Dictionary*.

Creating Your Own Collating Sequences

If you want to provide your own collating sequences or change a collating sequence provided for you, use the TRANTAB procedure to create or modify translation tables. For more information, see the TRANTAB procedure in the *SAS National Language Support (NLS): Reference Guide*. When you create your own translation tables, they are stored in your Sasuser.Profile catalog, and they override any translation tables by the same name that are stored in the Host catalog.

Note: System managers can modify the Host catalog by copying newly created tables from the Profile catalog to the Host catalog. Then, all users can access the new or modified translation table. △

If you are using the SAS windowing environment and want to see the names of the collating sequences that are stored in the Host catalog, issue the following command from any window:

```
catalog sashelp.host
```

If you are not using the SAS windowing environment, then issue the following statements to generate a list of the contents in the Host catalog:

```
proc catalog catalog=sashelp.host;
  contents;
run;
```

Entries of type TRANTAB are the collating sequences.

To see the contents of a particular translation table, use the following statements:

```
proc trantab table=table-name;
list;
run;
```

The contents of collating sequences are displayed in the SAS log.

Specifying the Host Sort Utility

Introduction to Using the Host Sort SAS supports one host sort utility on UNIX called **syncsort**. You can use this sorting application as an alternative sorting algorithm to the SAS sort. SAS determines which sort to use by the values that are set for the SORTNAME, SORTPGM, SORTCUT, and SORTCUTP system options.

Setting the Host Sort Utility as the Sort Algorithm To specify a host sort utility as the sort algorithm, complete the following steps:

- 1 Specify the name of the host utility (**syncsort**) in the SORTNAME system option.
- 2 Set the SORTPGM system option to tell SAS when to use the host sort utility.
 - If you specify SORTPGM=HOST, then SAS always prefers to use the host sort utility.
 - If you specify SORTPGM=BEST, then SAS chooses the best sorting method (either the SAS sort or the host sort) for the situation.

Sorting Based on Size or Observations The sort routine that SAS uses can be based on either the number of observations in a data set or on the size of the data set. When the SORTPGM system option is set to BEST, SAS uses the first available and pertinent sorting algorithm based on this order of precedence:

- host sort utility
- SAS sort utility

SAS looks at the values for the SORTCUT and SORTCUTP system options to determine which sort to use.

The SORTCUT option specifies the number of observations above which the host sort utility is used instead of the SAS sort. The SORTCUTP option specifies the number of bytes in the data set above which the host sort utility is used.

If SORTCUT and SORTCUTP are set to zero, SAS uses the SAS sort routine. If you specify both options and either condition is met, SAS uses the host sort utility.

When the following OPTIONS statement is in effect, the host sort utility is used when the number of observations is 501 or greater:

```
options sortpgm=best sortcut=500;
```

In this example, the host sort utility is used when the size of the data set is greater than 40M:

```
options sortpgm=best sortcutp=40M;
```

For more information about these sort options, see “SORTCUT System Option” on page 403, “SORTCUTP System Option” on page 404, and “SORTPGM System Option” on page 407.

Changing the Location of Temporary Files Used by the Host Sort Utility By default, the host sort utilities use the location that is specified in the -WORK option for temporary files. To change the location of these temporary files, specify a location by using the SORTDEV system option. Here is an example:

```
options sortdev="/tmp/host";
```

For more information, see “SORTDEV System Option” on page 406.

Passing Options to the Host Sort Utility To specify options for the sort utility, use the SORTANOM system option. For a list of valid options, see “SORTANOM System Option” on page 403.

Passing Parameters to the Host Sort Utility To pass parameters to the sort utility, use the SORTPARM system option. The parameters that you can specify depend on the host sort utility. For more information, see “SORTPARM System Option” on page 407.

Specifying the SORTSEQ= Option with a Host Sort Utility

The SORTSEQ= option enables you to specify the collating sequence for your sort. For a list of valid values, see the SORT procedure in *Base SAS Procedures Guide*.

CAUTION:

If you are using a host sort utility to sort your data, then specifying the SORTSEQ= option might corrupt the character BY variables if the sort sequence translation table and its inverse are not one-to-one mappings. In other words, for the sort to work, the translation table must map each character to a unique weight, and the inverse table must map each weight to a unique character variable. △

If your translation tables do not map one to one, then you can use one of the following methods to perform your sort:

- Create a translation table that maps one to one. Once you create a translation table that maps one to one, you can easily create a corresponding inverse table using the TRANTAB procedure. If your translation table is not mapped one to one, then you will receive the following note in the SAS log when you try to create an inverse table:

```
NOTE: This table cannot be mapped one to one.
```

For more information, see TRANTAB Procedure in the *SAS National Language Support (NLS): Reference Guide*.

- Use the SAS sort. You can specify the SAS sort using the SORTPGM system option. For more information, see “SORTPGM System Option” on page 407.
- Specify the collation order options of your host sort utility. See the documentation for your host sort utility for more information.
- Create a view with a dummy BY variable. For an example, see “Example: Creating a View with a Dummy BY Variable” on page 311.

Note: After using one of these methods, you might need to perform subsequent BY processing using either the NOTSORTED option or the NOBYSORTED system option. For more information about the NOTSORTED option, see BY Statement in *SAS Language Reference: Dictionary*. For more information about the NOBYSORTED system option, see BYSORTED System Option in *SAS Language Reference: Dictionary*. △

Example: Creating a View with a Dummy BY Variable The following example shows how to create a view using a dummy BY variable. SAS uses the BEST argument in the SORTPGM system option to sort the data. By using BEST, SAS selects either the host sort or the SAS sort.

```
options nodate pageno=1 nostimer ls=78 ps=60;
options sortpgm=best msglevel=i;
```

```

data one;
  input name $ age;
  datalines;
anne 35
ALBERT 10
JUAN 90
janet 5
bridget 23
BRIAN 45
;

data oneview / view=oneview;
  set one;
  name1=upcase(name);
run;

proc sort data=oneview out=final(drop=name1);
  by name1;
run;

proc print data=final;
run;

```

SAS writes the following note about sort to the log:

NOTE: SAS threaded sort was used.

SAS creates the following output:

Output 16.3 Creating a View with a Dummy BY Variable


The SAS System		
Obs	name	age
1	ALBERT	10
2	anne	35
3	BRIAN	45
4	bridget	23
5	janet	5
6	JUAN	90

1

See Also

- TRANTAB Procedure in *SAS National Language Support (NLS): Reference Guide*
- “MEMSIZE System Option” on page 384
- “REALMEMSIZE System Option” on page 394
- “SORTANOM System Option” on page 403
- “SORTCUT System Option” on page 403
- “SORTCUTP System Option” on page 404
- “SORTDEV System Option” on page 406
- “SORTNAME System Option” on page 406

- “SORTPARM System Option” on page 407
- “SORTPGM System Option” on page 407
- “SORTSIZE System Option” on page 408
- UTILLOC System Option in *SAS Language Reference: Dictionary*



CHAPTER

17

Statements under UNIX

<i>SAS Statements under UNIX</i>	315
<i>ABORT Statement</i>	315
<i>ATTRIB Statement</i>	316
<i>FILE Statement</i>	316
<i>FILENAME Statement</i>	318
<i>FOOTNOTE Statement</i>	324
<i>%INCLUDE Statement</i>	324
<i>INFILE Statement</i>	326
<i>LENGTH Statement</i>	327
<i>LIBNAME Statement</i>	328
<i>SYSTASK Statement</i>	334
<i>TITLE Statement</i>	337
<i>WAITFOR Statement</i>	338
<i>X Statement</i>	339

SAS Statements under UNIX

This section describes SAS statements that exhibit behavior or syntax that is specific to UNIX environments. Each statement description includes a brief “UNIX specifics” section that explains which aspect of the statement is specific to UNIX. If the information under the “UNIX specifics” says “all,” then the statement is described only in this documentation. Otherwise, the statement is described in this documentation and in *SAS Language Reference: Dictionary*.

ABORT Statement

Stops executing the current DATA step, SAS job, or SAS session.

Valid: in a DATA step

UNIX specifics: values of *n*

See: ABORT Statement in *SAS Language Reference: Dictionary*

Syntax

ABORT <ABEND | RETURN><*n*>;

Details

The n option allows you to specify the value of the exit status code that SAS returns to the shell when it stops executing. The value of n can range from 0 to 255. Normally, a return code of 0 is used to indicate that the program ran with no errors, and return codes greater than 0 are used to indicate progressively more serious error conditions. Return codes of 0–6 and those codes that are greater than 977 are reserved for use by SAS.

See Also

- “Determining the Completion Status of a SAS Job in UNIX Environments” on page 24

ATTRIB Statement

Associates a format, informat, label, or length with one or more variables.

Valid: in a DATA step

UNIX specifics: length specification

See: ATTRIB Statement in *SAS Language Reference: Dictionary*

Syntax

ATTRIB *variable-list-1 attribute-list-1* <...*variable-list-n attribute-list-n*>;

Note: Here is a simplified explanation of the ATTRIB statement syntax. For complete syntax and its explanation, see the ATTRIB statement in *SAS Language Reference: Dictionary*. Δ

attribute-list

LENGTH=<\${}>*length*

specifies the length of the variables in *variable-list*. The minimum length that you can specify for a numeric variable depends on the floating-point format used by your system. Because most systems use the IEEE floating-point format, the minimum is 3 bytes.

See Also

- Chapter 9, “Data Representation,” on page 215

FILE Statement

Specifies the current output file for PUT statements.

Valid: in a DATA step

UNIX specifics: valid values for *file-specification*, *host-options*, and *encoding-value*

See: FILE Statement in *SAS Language Reference: Dictionary*

Syntax

FILE *file-specification* <ENCODING=*'encoding-value'*> <*options*> <*host-options*>;

file-specification

can be any of the file specification forms that are discussed in the “Accessing an External File or Device in UNIX Environments” on page 69.

ENCODING=*'encoding-value'*

specifies the encoding to use when writing to the output file. The value for ENCODING= indicates that the output file has a different encoding from the current session encoding.

When you write data to the output file, SAS transcodes the data from the session encoding to the specified encoding.

For valid encoding values, see “Encoding Values in SAS Language Elements” in the *SAS National Language Support (NLS): Reference Guide*.

options

can be any of the options for the FILE statement that are valid in all operating environments. See the *SAS Language Reference: Dictionary* for a description of these options.

host-options

are specific to UNIX environments. These options can be any of the following:

BLKSIZE=

BLK=

specifies the number of bytes that are physically written in one I/O operation. The default is 8K. The maximum is 1G–1.

TERMSTR=

controls the end-of-line or record delimiters in PC- and UNIX-formatted files. This option enables the sharing of UNIX- and PC-formatted files between the two hosts. The following are values for the TERMSTR= option:

CRLF Carriage Return Line Feed. This parameter is used to create PC format files.

NL Newline. This parameter is used to create UNIX format files. NL is the default format.

Use TERMSTR=CRLF when you are writing to a file that you want to read on a PC. If you use this option when creating the file, then you do not need to use TERMSTR=NL when reading the file on the PC.

LRECL=

specifies the logical record length. Its meaning depends on the record format in effect (RECFM). The default is 256. The maximum length is 1G.

- If RECFM=F, then the value for the LRECL= option determines the length of each output record. The output record is truncated or padded with blanks to fit the specified size.
- If RECFM=N, then the value for the LRECL= option must be at least 256.

- If RECFM=V, then the value for the LRECL= option determines the maximum record length. Records that are longer than the specified length are divided into multiple records.

MOD

indicates that data written to the file should be appended to the file.

NEW | OLD

indicates that a new file is to be opened for output. If the file already exists, then it is deleted and re-created. If you specify OLD, then the previous contents of the file are replaced. NEW is the default.

RECFM=

specifies the record format. Values for the RECFM= option are the following:

D	default format (same as variable).
F	fixed format. That is, each record has the same length. Do not use RECFM=F for external files that contain carriage-control characters.
N	binary format. The file consists of a stream of bytes with no record boundaries.
P	print format. SAS writes carriage-control characters.
V	variable format. Each record ends with a newline character.
S370V	variable S370 record format (V).
S370VB	variable block S370 record format (VB).
S370VBS	variable block with spanned records S370 record format (VBS).

UNBUF

tells SAS not to perform buffered writes to the file on any subsequent FILE statement. This option applies especially when you are writing to a data collection device.

Details

The ENCODING= option is valid only when the FILE statement includes a file specification that is not a reserved fileref. If the FILE statement includes the ENCODING= argument and the reserved filerefs Log or Print as the *file-specification*, then SAS issues an error message. The ENCODING= value in the FILE statement overrides the value of the ENCODING= system option.

You can set the permissions of the output file by issuing the **umask** command from within the SAS session. For more information, see “Executing Operating System Commands from Your SAS Session” on page 14.

See Also

- Chapter 3, “Using External Files and Devices,” on page 67

FILENAME Statement

Associates a SAS fileref with an external file or output device; disassociates a fileref and external file; lists attributes of external files.

Valid: anywhere

UNIX specifics: *device-type*, *external-file*, *host-options*, and *encoding-value*

See: FILENAME Statement in *SAS Language Reference: Dictionary*

Syntax

FILENAME *fileref* <*device-type*> '*external-file*' <ENCODING=*encoding-value*'>
<*host-options*'><LOCKINTERNAL= AUTO | SHARED>;

FILENAME *fileref* *device-type* <'external-file'> <ENCODING=*encoding-value*'>
<*host-options*'><LOCKINTERNAL= AUTO | SHARED>;

FILENAME *fileref* ('*pathname-1*' ... '*pathname-n*') <ENCODING=*encoding-value*'>
<*host-options*'><LOCKINTERNAL= AUTO | SHARED>;

FILENAME *fileref* *directory-name*
<ENCODING=*encoding-value*'><LOCKINTERNAL= AUTO | SHARED>;

FILENAME *fileref* <*access-method*> '*external-file*' *access-information*;

FILENAME *fileref* CLEAR | _ALL_ CLEAR;

FILENAME *fileref* LIST | _ALL_ LIST;

fileref

is the name by which you reference the file. Under UNIX, the value of *fileref* can be an environment variable.

Note: You cannot clear a fileref that is defined by an environment variable.

Filerefs that are defined by an environment variable are assigned for the entire SAS session. △

See “Using Environment Variables to Assign Filerefs in UNIX Environments” on page 76 for more information.

device-type

specifies a device for the output, such as a disk, terminal, printer, pipe, and so on. The *device-type* keyword must follow *fileref* and precede *pathname*. “Device Information in the FILENAME Statement” on page 322 describes the valid device types. DISK is the default device type. If you are associating the fileref with a DISK file, then you do not need to specify the device type.

'*external-file*'

differs according to device type. “Device Information in the FILENAME Statement” on page 322 shows the information appropriate to each device. Remember that UNIX filenames are case sensitive. See “Specifying Pathnames in UNIX Environments” on page 70 for more information.

Note: If a filename has leading blanks, then they will be trimmed. △

ENCODING=*encoding-value*'

specifies the encoding to use when reading from or writing to the external file. The value for ENCODING= indicates that the external file has a different encoding from the current session encoding.

When you read data from an external file, SAS transcodes the data from the specified encoding to the session encoding. When you write data to an external file, SAS transcodes the data from the session encoding to the specified encoding.

For valid encoding values, see “Encoding Values in SAS Language Elements” in the *SAS National Language Support (NLS): Reference Guide*.

'host-options'

are specific to UNIX environments. These options can be any of the following:

BLKSIZE=

BLK=

specifies the number of bytes that are physically written or read in one I/O operation. The default is 8K. The maximum is 1G–1. If you specify RECFM=S370VBS, then you should specify BLKSIZE=32760 in order to avoid errors with records longer than 255 characters.

TERMSTR=

controls the end of line/record delimiters in PC and UNIX formatted files. This option enables the sharing of UNIX and PC formatted files between the two hosts. The following values for the TERMSTR= option are valid:

CRLF Carriage Return Line Feed. This parameter is used to read and write PC format files.

NL Newline. This parameter is used to read and write UNIX format files. NL is the default format.

If you are working on UNIX and reading a file that was created on a PC, specify TERMSTR=CRLF unless the file was created with the TERMSTR=NL option. If you are writing a file that will be read on a PC, specify TERMSTR=CRLF.

If you are working on a PC and reading a file that was created on UNIX, specify TERMSTR=NL unless the file was created with the TERMSTR=CRLF option. If you are writing a file that will be read on UNIX, specify TERMSTR=NL.

LRECL=

specifies the logical record length. Its meaning depends on the record format in effect (RECFM). The default is 256. The maximum length is 1G.

- If RECFM=F, then the value for the LRECL= option determines either the number of bytes to be read as one record or the length of each output record. The output record is truncated or padded with blanks to fit the specified size.
- If RECFM=N, then the value for the LRECL= option must be at least 256.
- If RECFM=V, then the value for the LRECL= option determines the maximum record length. Records that are longer than the specified length are divided into multiple records on output and truncated on input.
- If RECFM=S370VBS, then you should specify LRECL=32760 in order to avoid errors with records longer than 255 characters.

MOD

indicates that data written to the file should be appended to the file.

NEW | OLD

indicates that a new file is to be opened for output. If the file already exists, then it is deleted and re-created. If you specify OLD, then the previous contents of the file are replaced. NEW is the default.

RECFM=

specifies the record format. Values for the RECFM= option are the following:

D default format (same as variable).

F fixed format. That is, each record has the same length. Do not use RECFM=F for external files that contain carriage-control characters.

N	binary format. The file consists of a stream of bytes with no record boundaries. N is not valid for the PIPE device type. If you do not specify the LRECL option, then by default SAS reads 256 bytes at a time from the file.
P	print format. On output, SAS writes carriage-control characters.
V	variable format. Each record ends with a newline character.
S370V	variable S370 record format (V).
S370VB	variable block S370 record format (VB).
S370VBS	variable block with spanned records S370 record format (VBS). If you specify RECFM=S3270VBS, then you should specify BLDSIZE=32760 and LRECL=32760 in order to avoid errors with records longer than 255 characters.

The RECFM= option is used for both input and output.

LOCKINTERNAL=AUTO | SHARED

specifies the SAS system locking that is to be used for the files that are listed in a FILENAME statement. LOCKINTERNAL can have one of the following values:

AUTO

locks a file so that in a SAS session, if a user has Write access to a file, then no other users can have Read or Write access to the file. If a user has Read access to a file, no other user can have Write access to the file, but multiple users can have Read access.

SHARED

locks a file so that in a SAS session, two users do not have simultaneous Write access to the file. The file can be shared simultaneously by one user who has Write access and multiple users who have Read access.

Default: AUTO

UNBUF

tells SAS not to perform buffered writes to the file on any subsequent FILE statement. This option applies especially when you are reading from or writing to a data collection device. As explained in *SAS Language Reference: Dictionary*, it also prevents buffered reads on INFILE statements.

'pathname-1'...'pathname-n'

are pathnames for the files that you want to access with the same fileref. Use this form of the FILENAME statement when you want to concatenate filenames. Concatenation of filenames is available only for DISK files, so you do not have to specify the *device-type*. Separate the pathnames with either commas or blank spaces. Enclose each pathname in quotation marks. Table 2.6 on page 51 shows character substitutions you can use when specifying a pathname. If the fileref that you are defining is to be used for input, then you can also use wildcards as described in "Using Wildcards in Pathnames (Input Only)" on page 71. Remember that UNIX filenames are case-sensitive.

directory-name

specifies the directory that contains the files that you want to access. For more information, see "Assigning a Fileref to a Directory (Using Aggregate Syntax)" on page 75.

access-method

can be CATALOG, SOCKET, FTP, SFTP, or URL. "Device Information in the FILENAME Statement" on page 322 describes the information expected by these access methods.

access-information

differs according to the access method. “Device Information in the FILENAME Statement” on page 322 shows the information appropriate to each access method.

CLEAR

clears the specified fileref or, if you specify `_ALL_`, clears all filerefs that are currently defined.

Note: You cannot clear a fileref that is defined by an environment variable. Filerefs that are defined by environment variables are assigned for the entire SAS session. Δ

ALL

refers to all filerefs currently defined. You can use this keyword when you are listing or clearing filerefs.

LIST

writes to the SAS log the pathname of the specified fileref or, if you specify `_ALL_`, lists the definition for all filerefs that are currently defined. Filerefs defined as environment variables appear only if you have already used those filerefs in a SAS statement. If you are using the Bourne shell or the Korn shell, SAS cannot determine the name of a pre-opened file, so it displays the following string instead of a filename:

<File Descriptor number>

See “Using Environment Variables to Assign Filerefs in UNIX Environments” on page 76 for more information.

Details

File Locking File locking of external files is controlled at the FILENAME statement level by the LOCKINTERNAL option. If you use the AUTO (default) value for LOCKINTERNAL, then SAS locks a file exclusively for one user who has Write access, or SAS locks a file non-exclusively for multiple users who have Read access. For example, if a file is opened in UPDATE or OUTPUT mode, then all other access from internal processes will be blocked. If a file is opened in INPUT mode, then multiple users can read the file, but UPDATE and OUTPUT functions are blocked.

If you use the SHARED value for LOCKINTERNAL, then SAS allows one user Write access to a file as well as allowing multiple users to read the file.

Device Information in the FILENAME Statement The following table lists the relationship between device type or access method and the related external file.

Table 17.1 Device Information in the FILENAME Statement

Device or Access Method	Function	External File
CATALOG	references a SAS catalog as an external file.	is a valid two-, three-, or four-part SAS catalog name followed by catalog options (if needed). See <i>SAS Language Reference: Dictionary</i> for information.
DISK	associates the fileref with a DISK file.	is either the pathname for a single file or, if you are concatenating filenames, a list of pathnames separated by spaces or commas and enclosed in parentheses. The level of specification depends on your location in the file system. Table 2.6 on page 51 shows character substitutions that you can use when specifying a UNIX pathname.

Device or Access Method	Function	External File
DUMMY	associates a fileref with a null device.	None. DUMMY enables you to debug your application without reading from or writing to a device. Output to this device is discarded.
EMAIL	sends electronic mail to an address.	is an address and e-mail options. See “Sending Electronic Mail Using the FILENAME Statement (EMAIL)” on page 82 for information.
FTP	reads from or writes to a file from any computer on a network that is running an FTP server.	is the pathname of the external file on the remote computer followed by FTP options. See <i>SAS Language Reference: Dictionary</i> and “Assigning Filerefs to Files on Other Systems (FTP, SFTP, and SOCKET Access Types)” on page 74 for information. If you are transferring a file to UNIX from the z/OS operating environment and you want to use either the S370V or S370VB format to access that file, then the file must be of type RECFM=U and BLKSIZE=32760 before you transfer it. If you FTP to an z/OS computer, only one member of an z/OS PDS can be written to at a time. If you need to write to multiple members at the same time, an z/OS PDSE or a UNIX System Services directory should be used.
PIPE	reads input from or writes output to a UNIX command.	is a UNIX command. See Chapter 4, “Printing and Routing Output,” on page 91 for information.
PLOTTER	sends output to a plotter.	is a device name and plotter options. See “Printing the Contents of a Window” on page 98 and “Using the PRINTTO Procedure in UNIX Environments” on page 100 for information.
PRINTER	sends output to a printer.	is a device name and printer option. See “Printing the Contents of a Window” on page 98 and “Using the PRINTTO Procedure in UNIX Environments” on page 100 for information.
SFTP	reads from or writes to a file from any host computer that you can connect to on a network with an SSHD server running.	is the pathname of the external file on the remote computer, followed by SFTP options. See FILENAME Statement, SFTP Access Method in <i>SAS Language Reference: Dictionary</i> and “Assigning Filerefs to Files on Other Systems (FTP, SFTP, and SOCKET Access Types)” on page 74 for more information.
SOCKET	reads and writes information over a TCP/IP socket.	depends on whether the SAS application is a server application or a client application. In a client application, <i>external-file</i> is the name or IP address of the host and the TCP/IP port number to connect to, followed by any TCP/IP options. In a server application, <i>external-file</i> is the port number to create for listening, followed by the SERVER keyword, and then any TCP/IP options. See <i>SAS Language Reference: Dictionary</i> for information.
TEMP	associates a fileref with an external file stored in the Work library.	None
TERMINAL	associates a fileref with a terminal.	is the pathname of a terminal.
UPRINTER	sends output to the default printer that was set up through the Printer Setup dialog box.	None

Device or Access Method	Function	External File
URL	enables you to use the URL of a file to access it remotely.	is the name of the file that you want to read from or write to on a URL server. The URL must be in one of these forms: http://hostname/file http://hostname:portno/file
WebDAV	enables you to use WebDAV (Web Distributed Authoring and Versioning) to read from or write to a file from any host machine that you can connect to on a network with a WebDAV server running.	is the name of the file that you want to read from or write to a WebDAV server. The external file must be in one of these forms: http://hostname/path-to-the-file https://hostname/path-to-the-file http://hostname:port/path-to-the-file https://hostname:port/path-to-the-file

See Also

- Chapter 3, “Using External Files and Devices,” on page 67
- Chapter 4, “Printing and Routing Output,” on page 91

FOOTNOTE Statement

Writes up to 10 lines of text at the bottom of the procedure or DATA step output.

Valid: anywhere

UNIX specifics: maximum length of footnote

See: FOOTNOTE Statement in *SAS Language Reference: Dictionary*

Syntax

FOOTNOTE <n> <'text' | "text">;

Details

The maximum footnote length is 255 characters. If the length of the specified footnote is greater than the value of the LINESIZE option, SAS truncates the footnote to the line size.

%INCLUDE Statement

Brings a SAS programming statement, data lines, or both into a current SAS program.

Valid: anywhere

UNIX specifics: *source*, if a file specification is used; valid values for *encoding-value*

See: %INCLUDE Statement in *SAS Language Reference: Dictionary*

Syntax

```
%INCLUDE source-1 <...source-n> </<SOURCE2>
      <S2=length><ENCODING='encoding-value'><host-options>>;
```

source

describes the location you want to access with the %INCLUDE statement. The three possible sources are a file specification, internal lines, or keyboard entry. The file specification can be any of the file specification forms that are discussed in “Accessing an External File or Device in UNIX Environments” on page 69.

Note: When using aggregate syntax, if the member name contains a leading digit, enclose the member name in quotation marks. If the member name contains a macro variable reference, use double quotation marks. △

ENCODING='encoding-value'

specifies the encoding to use when reading from the specified source. The value for ENCODING= indicates that the specified source has a different encoding from the current session encoding.

When you read data from the specified source, SAS transcodes the data from the specified encoding to the session encoding.

For valid encoding values, see “Encoding Values in SAS Language Elements” in *SAS National Language Support (NLS): Reference Guide*.

host-options

consists of statement options that are valid under UNIX. The following options are available:

BLKSIZE=*block-size*

BLK=*block-size*

specifies the number of bytes that are physically read or written in an I/O operation. The default is 8K. The maximum is 1M.

LRECL=*record-length*

specifies the record length (in bytes). Under UNIX, the default is 256. The value of *record-length* can range from 1 to 1,048,576 (1 megabyte).

RECFM=*record-format*

specifies the record format. The following values are valid under UNIX:

D default format (same as variable).

F fixed format. That is, each record has the same length.

N binary format. The file consists of a stream of bytes with no record boundaries.

P print format.

V variable format. Each record ends with a newline character.

S370V variable S370 record format (V).

S370VB variable block S370 record format (VB).

S370VBS variable block with spanned records S370 record format (VBS).

The S370 values are valid with files laid out as z/OS files only. That is, files are binary, have variable-length records, and are in EBCDIC format. If you want to use a fixed-format z/OS file, first copy it to a variable-length, binary z/OS file.

Details

If you specify any options in the %INCLUDE statement, remember to precede the options list with a forward slash (/).

See Also

- Chapter 3, “Using External Files and Devices,” on page 67

INFILE Statement

Identifies an external file to read with an INPUT statement.

Valid: in a DATA step

UNIX specifics: valid values for *encoding-value*, *file-specification*, and *host-options*

See: INFILE Statement in *SAS Language Reference: Dictionary*

Syntax

INFILE *file-specification* <ENCODING=*'encoding-value'*> <*options*> <*host-options*>;

file-specification

can be any of the file specification forms that are discussed in the “Accessing an External File or Device in UNIX Environments” on page 69.

ENCODING=*'encoding-value'*

specifies the encoding to use when reading from the external file. The value for ENCODING= indicates that the external file has a different encoding from the current session encoding.

When you read data from an external file, SAS transcodes the data from the specified encoding to the session encoding.

For valid encoding values, see “Encoding Values in SAS Language Elements” in *SAS National Language Support (NLS): Reference Guide*.

host-options

are specific to UNIX environments. These options can be any of the following:

BLKSIZE=

BLK=

specifies the number of bytes that are physically read in one I/O operation. The default is 8K. The maximum is 1G–1.

TERMSTR=

controls the end of line or record delimiters in PC- and UNIX-formatted files. This option enables the sharing of UNIX- and PC-formatted files between the two hosts. The following are values for the TERMSTR= option:

CR	Carriage return. Use TERMSTR=CR to read files formatted by Apple Macintosh.
----	---

CRLF Carriage Return Line Feed. Use `TERMSTR=CRLF` to read files formatted by a PC.

LF LineFeed. This parameter is used to read files formatted by UNIX. LF is the default.

Use `TERMSTR=CRLF` to read a file that was created on a PC. If the PC-formatted file was created using `TERMSTR=LF`, then the `TERMSTR=` option is unnecessary.

LRECL=

specifies the logical record length. Its meaning depends on the record format in effect (RECFM). The default is 256. The maximum length is 1G.

- If `RECFM=F`, then the value for the `LRECL=` option determines the number of bytes to be read as one record.
- If `RECFM=N`, then the value for the `LRECL=` option must be at least 256.
- If `RECFM=V`, then the value for the `LRECL=` option determines the maximum record length. Records that are longer than the specified length are truncated.

RECFM=

specifies the record format. The following values are valid under UNIX:

D	default format (same as variable).
F	fixed format. That is, each record has the same length.
N	binary format. The file consists of a stream of bytes with no record boundaries. If you do not specify the <code>LRECL</code> option, then, by default, SAS reads 256 bytes at a time from the file.
P	print format.
V	variable format. Each record ends with a newline character.
S370V	variable S370 record format (V).
S370VB	variable block S370 record format (VB).
S370VBS	variable block with spanned records S370 record format (VBS).

Details

The `ENCODING=` option is valid only when the `INFILE` statement includes a file specification that is not a reserved fileref. If the `INFILE` statement includes the `ENCODING=` argument and the reserved filerefs `DATALINES` or `DATALINES4` as a *file-specification*, then SAS issues an error message. The `ENCODING=` value in the `INFILE` statement overrides the value of the `ENCODING=` system option.

See Also

- Chapter 3, “Using External Files and Devices,” on page 67

LENGTH Statement

Specifies the number of bytes for storing variables.

Valid: in a DATA step

UNIX specifics: valid numeric variable lengths

See: LENGTH Statement in *SAS Language Reference: Dictionary*

Syntax

LENGTH <variable-1><...variable-n> <\$> length <DEFAULT=*n*>

length

can range from 3 to 8 for numeric variables under UNIX. The minimum length you can specify for a numeric variable depends on the floating-point format used by your system. Because most systems use the IEEE floating-point format, the minimum is 3 bytes.

DEFAULT=*n*

changes the default number of bytes that are used for storing the values of newly created numeric variables from 8 to the value of *n*. Under UNIX, *n* can range from 3 to 8.

See Also

- Chapter 9, “Data Representation,” on page 215

LIBNAME Statement

Associates or disassociates a SAS library with a libref (a shortcut name); clears one or all librefs; lists the characteristics of a SAS library; concatenates SAS libraries; implicitly concatenates SAS catalogs; turns off file locking.

Valid: anywhere

UNIX specifics: *engine*, *library*, and *engine/host-options*

See: LIBNAME Statement in *SAS Language Reference: Dictionary*

Syntax

LIBNAME libref <engine> 'SAS-library' <options> <engine/host-options>;

LIBNAME libref <engine> ('library-1'<,...'library-n') <options>;

LIBNAME libref ('library-1' | libref-1,...,'library-n' | libref-n);

LIBNAME libref CLEAR | _ALL _ CLEAR;

LIBNAME libref LIST | _ALL _ LIST;

libref

is any valid libref as documented in *SAS Language Reference: Dictionary*. SAS reserves some librefs for special system libraries. See “Librefs Assigned by SAS in UNIX Environments” on page 54 for more information.

engine

is one of the library engines supported under UNIX. See “Details” on page 329 for a description of the engines. If no engine name is specified, SAS determines which engine to use as described in “Omitting Engine Names from the LIBNAME Statement” on page 331.

'SAS-library'

differs based on the engine that you specify and based on your current working directory. Table 17.2 on page 330 describes what each engine expects for this argument. Specify directory pathnames as described in “Specifying Pathnames in UNIX Environments” on page 50. You cannot create directories with the LIBNAME statement. The directory that you specify must already exist, and you must have permissions to it. Enclose the library name in quotation marks. Remember that UNIX pathnames are casesensitive.

'library-n' | libref-n

are pathnames or librefs (that have been assigned) for the libraries that you want to access with one libref. Use these forms of the LIBNAME statement when you want to concatenate libraries. Separate the pathnames with either commas or blank spaces. Enclose library pathnames in quotation marks. Do not enclose librefs in quotation marks. See “Assigning a Libref to Several Directories (Concatenating Directories)” on page 51 for more information.

options

are LIBNAME statement options that are available in all operating environments. See *SAS Language Reference: Dictionary* for information about these options.

engine/host-options

can be any of the options described in “Engine/Host Options” on page 331.

ALL

refers to all librefs currently defined. You can use this keyword when you are listing or clearing librefs.

CLEAR

clears the specified libref, or, if you specify ALL, clears all librefs that are currently defined. Sasuser, Sashelp, and Work remain assigned.

Note: When you clear a libref defined by an environment variable, the variable remains defined, but it is no longer considered a libref. You can still reuse it, either as a libref or a fileref. See “Using Environment Variables as Librefs in UNIX Environments” on page 53 for more information. Δ

SAS automatically clears the association between librefs and their respective libraries at the end of your job or session. If you want to associate an existing libref with a different SAS library during the current session, you do not have to end the session or clear the libref. SAS automatically reassigns the libref when you issue a LIBNAME statement for the new SAS library.

LIST

prints to the SAS log the engine, pathname, file format, access permissions, and so on that are associated with the specified libref, or, if you specify ALL, prints this information for all librefs that are currently defined. Librefs defined as environment variables appear only if you have already used those librefs in a SAS statement.

Details

There are two main types of engines:

View engines

enable SAS to read SAS views that are described by SAS/ACCESS software, the SQL procedure, and DATA step views. The use of SAS view engines is automatic because the name of the view engine is stored as part of the descriptor portion of the SAS data set.

Library engines

control access at the SAS library level. Every SAS library has an associated library engine, and the files in that library can be accessed only through that engine. There are two types of library engines:

native engines

access SAS files created and maintained by SAS. See the following table for a description of these engines.

interface engines

treat other vendors' files as if they were SAS files. See the following table and "Accessing BMDP, OSIRIS, or SPSS Files in UNIX Environments" on page 62 for more information.

Table 17.2 Engine Names and Descriptions

Engine Type	Name (Alias)	Description	SAS Library
default	V9 (BASE) V8	enables you to create new SAS data files and to access existing SAS data files that were created with Version 8 or SAS 9. The V8 and V9 engines are identical. This engine enables Read access to data files that were created with some earlier releases of SAS, but this engine is the only one that supports SAS 9 catalogs. This engine allows for data set indexing and compression, and is also documented in <i>SAS Language Reference: Dictionary</i> .	is the pathname of the directory containing the library.
sequential	V9TAPE	enables you to access SAS files that were created in a sequential format, whether on tape or on disk.	is the name of the special file (see "Introduction to External Files and Devices in UNIX Environments" on page 68) associated with the sequential device, such as /dev/rmt/Omn .
	V8TAPE	requires less overhead than the default engine because sequential access is simpler than random access. This engine is also documented in <i>SAS Language Reference: Dictionary</i> .	
	V6TAPE	accesses V6 SAS data files that were created in a sequential format. This engine is read-only.	is the name of the special file that is associated with the sequential device, such as /dev/rmt/Omn .
compatibility	V6	accesses any data file that was created by Releases 6.09 through 6.12. This engine is read-only.	is the pathname of the directory containing the library.
servers	SPDS	enables communication between a client session and a data server. You must have the Scalable Performance Data Server licensed on your client computer to use this engine. See the <i>Scalable Performance Data Server User's Guide</i> for more information.	is the logical LIBNAME domain name for a Scalable Performance Data Server (SPDS) library on the server. The name server resolves the domain name into the physical path for the library.

Engine Type	Name (Alias)	Description	SAS Library
	MDDDB	enables communication between a client session and an MDDDB server. You must have SAS/MDDDB Server licensed on your client computer or on your server to use this engine. See the <i>SAS MDDDB Server Software: Administration Guide</i> for more information.	
transport	XPORT	accesses transport data sets. This engine creates computer-independent SAS transport files that can be used under all hosts running Release 6.06 or later of SAS. This engine is documented in <i>Moving and Accessing SAS Files</i> .	is the pathname of either a sequential device or a disk file.
XML	XML	generates (writes) and processes (reads) any XML document, which is an application- and computer-independent file.	is the pathname of the XML document.
interface	BMDP	provides read-only access to BMDP files. This engine is available only on AIX, HP-UX, and Solaris.	is the pathname of the data file.
	OSIRIS	provides read-only access to OSIRIS files.	is the pathname of the data file.
	SPSS	provides read-only access to SPSS files	is the pathname of the data file.

Omitting Engine Names from the LIBNAME Statement It is always more efficient to specify the engine name than to have SAS determine the correct engine. However, if you omit an engine name in the LIBNAME statement or if you define an environment variable to serve as a libref, SAS determines the appropriate engine.

If you have specified the ENGINE= system option, SAS uses the engine name that you specified. See “ENGINE System Option” on page 367 for a discussion of the ENGINE= system option.

Note: The ENGINE= system option specifies the default engine for libraries on disk only. △

If you did not specify the ENGINE= system option, SAS looks at the extensions of the files in the given directory and uses these rules to determine an engine:

- If all the SAS data sets in the library were created by the same engine, the libref is assigned using that engine.

Note: If the engine used to create the data sets is not the same as the default engine, then you will not be able to create a view or stored program. See “Using Multiple Engines for a Library in UNIX Environments” on page 53 for more information. △

- If there are no SAS data sets in the given directory, the libref is assigned using the default engine.
- If there are SAS data sets from more than one engine, the system issues a message about finding mixed engine types and assigns the libref using the default engine.

Engine/Host Options The LIBNAME statement accepts the following options:

ENABLEDIRECTIO

specifies that direct file I/O can be available for all files that are opened in the library that is identified in the LIBNAME statement.

Tip: A libref that is assigned to a directory with the ENABLEDIRECTIO option will not match another libref that is assigned to the same directory without ENABLEDIRECTIO. The two librefs will point to the same directory, but the files that are opened using one libref will be read from and written to using Direct I/O. Files that are opened using the other libref will be read from and written to using the regular disk I/O calls.

Tip: You must use the ENABLEDIRECTIO option with the USEDIRECTIO= option to turn on direct I/O for the file or files whose libref is listed in the LIBNAME statement.

The following example uses the ENABLEDIRECTIO and USEDIRECTIO= LIBNAME options. In this case, all files that are referenced with libref **test** will be opened for direct I/O:

```
LIBNAME test'.'ENABLEDIRECTIO USEDIRECTIO=yes;
```

The following example uses the ENABLEDIRECTIO LIBNAME option to enable files that are associated with the libref **test** to be opened for direct I/O. The USEDIRECTIO= data set option opens **test.file1** for direct I/O. **test.file2** is not opened for direct I/O, although it is enabled for direct I/O:

```
LIBNAME test'.'ENABLEDIRECTIO;
data test.file1(USEDIRECTIO=yes);
    ... more SAS statements ...
run;
data test.file2;
    ... more SAS statements ...
run;
```

FILELOCKS=NONE | FAIL | CONTINUE

specifies whether file locking is turned on or off for the files that are opened under the libref in the LIBNAME statement. The FILELOCKS statement option works like the FILELOCKS system option, except that it applies only to the files that are associated with the libref. The following values for the FILELOCKS statement option are available:

NONE

turns file locking off. NONE specifies that SAS attempts to open the file without checking for an existing lock on the file. NONE does not place an operating system lock on the file. These files are not protected from shared Update access.

FAIL

turns file locking on. FAIL specifies that SAS attempts to place an operating system lock on the file. Access to the file is denied if the file is already locked, or if it cannot be locked.

CONTINUE

turns file locking on. CONTINUE specifies that SAS attempts to place an operating system lock on the file. If the file is already locked by someone else, an attempt to open it fails. If the file cannot be locked for some other reason (for example, if the file system does not support locking), the file is opened and a warning message is sent to the log.

The FILELOCKS option in the LIBNAME statement applies to most (but not all) of the SAS I/O files (for example, data sets and catalogs) opened under the libref that is listed in the LIBNAME statement.

Tip: For the FILELOCKS statement option, RESET is not a valid value as it is when you use the FILELOCKS system option.

Tip: Use the FILELOCKS system option instead of the FILELOCKS statement option to set the locking behavior for your files. (The FILELOCKS statement option will be deprecated in a future release of SAS.)

Note: The FILELOCKS option in the LIBNAME statement overrides the FILELOCKS system option. Δ

See “FILELOCKS System Option” on page 368 for more information.

You can also specify any of the options supported by the Scalable Performance Data Server. Refer to *Scalable Performance Data Server User’s Guide* at support.sas.com for a description of these options.

FILELOCKWAIT=*n*

specifies the number of seconds SAS will wait for a locked file to become available to another process.

If the locked file is released before the number of seconds specified by *n*, then SAS locks the file for the current process and continues. If the file is still locked when the number of seconds has been reached, then SAS writes a “Locked File” error to the log and the DATA step fails.

Interaction: Specifying the FILELOCKWAIT= option can have an adverse effect on one or more SAS/SHARE server and client sessions that are waiting for the release of a SAS file that is locked by another process. One or more wait conditions could lead to failed processes for a SAS/SHARE server and clients.

To prevent the possibility of a failed SAS/SHARE process, you can set FILELOCKWAIT=0, which cancels the amount of time that a SAS/SHARE server and clients would wait for the release of a locked file. Canceling the wait time would prevent a failed process. For more information, see the “FILELOCKWAITMAX= System Option” on page 370. In addition, see the section about predefining a server library by using the LIBNAME statement in the *SAS/SHARE User’s Guide*.

Range: 0–600

Default: 0

TRANSFERSIZE=*n*K | *n*M

specifies the size of blocks of data in units of kilobytes or megabytes.

Requirement: To use the TRANSFERSIZE option, you must have files open for direct I/O. That is, both the ENABLEDIRECTIO and USEDIRECTIO= options must be in effect. If you use TRANSFERSIZE without the ENABLEDIRECTIO and USEDIRECTIO= options, the option is accepted, but it has no effect.

In the following example, 128k blocks of data are read from the **test.file1** file because this file is opened for direct I/O. **test.file2** is not open for direct I/O, and the TRANSFERSIZE option has no effect on this file:

```
LIBNAME test'.'ENABLEDIRECTIO TRANSFERSIZE=128k;
data test.file1(USEDIRECTIO=yes);
    ... more SAS statements ...
run;
data test.file2;
    ... more SAS statements ...
run;
```

In the following example, all the files that are listed in the DATA statements read 128k blocks of data because all the files are affected by the ENABLEDIRECTIO, USEDIRECTIO=, and TRANSFERSIZE options:

```
LIBNAME test'.'ENABLEDIRECTIO USEDIRECTIO=yes TRANSFERSIZE=128k;
data test.file1;
```

```

        ... more SAS statements ...
run;
data test.file2;
    ... more SAS statements ...
run;
data test.file3;
    ... more SAS statements ...
run;

```

USEDIRECTIO= YES | NO

if used with the ENABLEDIRECTIO statement option, turns on or turns off direct file I/O for all the files associated with the libref listed in the LIBNAME statement. (See ENABLEDIRECTIO in “Engine/Host Options” on page 331.)

Requirement: Use USEDIRECTIO= with the ENABLEDIRECTIO statement option to turn on direct file I/O.

See Also

- System Options:
 - “FILELOCKS System Option” on page 368
- Chapter 2, “Using SAS Files,” on page 31

SYSTASK Statement

Executes asynchronous tasks.

Valid: anywhere

UNIX specifics: all

Syntax

SYSTASK COMMAND *“operating-environment-command”*

```

<WAIT | NOWAIT>
<TASKNAME=taskname>
<MNAME=name-variable>
<STATUS=status-variable>
<SHELL<=“shell-command”>>
<CLEANUP>;

```

SYSTASK LIST *<_ALL_ | taskname> <STATE> <STATVAR>;*

SYSTASK KILL *taskname <taskname...>;*

COMMAND

executes the *operating-environment-command*.

LIST

lists either a specific active task or all of the active tasks in the system. A task is *active* if it is running or if it has completed and has not been waited for using the WAITFOR statement.

KILL

forces the termination of the specified tasks.

operating-environment-command

specifies the name of a UNIX command (including any command-specific options) or the name of an X Windows or Motif application. Enclose the command in either single or double quotation marks. If the command-specific options require quotation marks, repeat them for each option. For example:

```
SYSTASK COMMAND "xdialog -m ""There was an error."" -t ""Error"" -o";
```

Note: If the command name is a shell alias, or if you use the shell special characters tilde (~) and asterisk (*) in a pathname in a command, you need to specify the SHELL option so that the shell will process the alias or special characters:

```
SYSTASK COMMAND "mv ~usr/file.txt /tmp/file.txt" shell;
```

In this example, by using the SHELL option, the ~usr path is expanded on execution and is not executed directly. Δ

Note: The *operating-environment-command* that you specify cannot require input from the keyboard. Δ

Tip: If using a shell alias results in an error even though the SHELL option is used, then the shell is not processing your shell initialization files. Use the actual SHELL command instead of the SHELL alias.

WAIT | NOWAIT

determines whether SYSTASK COMMAND suspends execution of the current SAS session until the task has completed. NOWAIT is the default. For tasks that start with the NOWAIT option, you can use the WAITFOR statement when necessary to suspend execution of the SAS session until the task has finished. See “WAITFOR Statement” on page 338.

TASKNAME=*taskname*

specifies a name that identifies the task. Task names must be unique among all active tasks. A task is *active* if it is running or if it has completed and has not been waited for using the WAITFOR statement. Duplicate task names generate an error in the SAS log. If you do not specify a task name, SYSTASK will automatically generate a name. If the task name contains a blank character, enclose it in quotation marks.

Task names cannot be reused, even if the task has completed, unless you either issue the WAITFOR statement for the task or you specify the CLEANUP option.

MNAME=*name-variable*

specifies a macro variable in which you want SYSTASK to store the task name that it automatically generated for the task. If you specify both the TASKNAME option and the MNAME option, SYSTASK copies the name that you specified with TASKNAME into the variable that you specified with MNAME.

STATUS=*status-variable*

specifies a macro variable in which you want SYSTASK to store the status of the task. Status variable names must be unique among all active tasks.

SHELL<=*shell-command*>

specifies that the *operating-environment-command* should be executed with the operating system shell command. The shell will expand shell special characters that are contained in the *operating-environment-command*. If you specify a *shell-command*, SYSTASK uses the shell command that you specify to invoke the

shell; otherwise, SYSTASK uses the default shell. Enclose the shell command in quotation marks.

Note: The SHELL option assumes that the shell command that you specify uses the **-i** option to pass statements. Usually, your shell command will be **sh**, **cs**, **ks**, or **ba**. Δ

CLEANUP

specifies that the task should be removed from the LISTTASK output when the task completes. You can then reuse the task name without issuing the WAITFOR statement.

Details

The SYSTASK statement enables you to execute host-specific commands from within your SAS session or application. Unlike the X statement, the SYSTASK statement runs these commands as *asynchronous* tasks, which means that these tasks execute independently of all other tasks that are currently running. Asynchronous tasks run in the background, so you can perform additional tasks while the asynchronous task is still running.

For example, to start a new shell and execute the UNIX **cp** command in that shell, you might use this statement:

```
systask command "cp /tmp/sas* ~/archive/" taskname="copyjob1"
                status=copysts1 shell;
```

The return code from the **cp** command is saved in the macro variable COPYSTS1.

The output from the command is displayed in the SAS log.

If you convert PC SAS jobs to run on UNIX, you might encounter an error in the conversion process. Entering the following command results in an error:

```
systask command "md directory-name" taskname="mytask";
```

SAS writes the following error message to the log:

```
ERROR: Could not create a new process.
```

This error message indicates that the launch of the command failed. The most common problem is that the command cannot be found or was not executable. It is also possible that there was a failure in obtaining resources to launch the command.

Note: Program steps that follow the SYSTASK statements in SAS applications usually depend on the successful execution of the SYSTASK statements. Therefore, syntax errors in some SYSTASK statements will cause your SAS application to abort. Δ

There are two types of asynchronous processes that can be started from SAS:

Task

All tasks started with SYSTASK COMMAND are of type Task. For these tasks, if you do not specify STATVAR or STATE, then SYSTASK LIST displays the task name, type, and state, and the name of the status macro variable. You can use SYSTASK KILL to kill only tasks of type Task.

SAS/CONNECT Process

Tasks started from SAS/CONNECT with the SIGNON statement or command and RSUBMIT statement are of type SAS/CONNECT Process. To display SAS/CONNECT processes, use the LISTTASK statement to displays the task name, type, and state. To terminate a SAS/CONNECT process, use the KILLTASK

statement. For information about SAS/CONNECT processes, see the *SAS/CONNECT User's Guide*.

Note: The preferred method for displaying any task (not just SAS/CONNECT processes) is to use the LISTTASK statement instead of SYSTASK LIST.

The preferred method for ending a task is using the KILLTASK statement instead of SYSTASK KILL. △

The SYSRC macro variable contains the return code for the SYSTASK statement. The status variable that you specify with the STATUS option contains the return code of the process started with SYSTASK COMMAND. To ensure that a task executes successfully, you should monitor both the status of the SYSTASK statement and the status of the process that is started by the SYSTASK statement.

If a SYSTASK statement cannot execute successfully, the SYSRC macro variable will contain a non-zero value. For example, there might be insufficient resources to complete a task or the SYSTASK statement might contain syntax errors. With the SYSTASK KILL statement, if one or more of the processes cannot be killed, SYSRC is set to a non-zero value.

When a task is started, its status variable is set to NULL. You can use the status variables for each task to determine which tasks failed to complete. Any task whose status variable is NULL did not complete execution. If a task terminates abnormally, then its status variable will be set to --1. See “WAITFOR Statement” on page 338 for more information about the status variables.

Unlike the X statement, you cannot use the SYSTASK statement to start a new interactive session.

See Also

- “WAITFOR Statement” on page 338
- “X Statement” on page 339
- “Executing Operating System Commands from Your SAS Session” on page 14

TITLE Statement

Specifies title lines for SAS output.

Valid: anywhere

UNIX specifics: maximum length of title

See: TITLE Statement in *SAS Language Reference: Dictionary*

Syntax

TITLE <n> <'text' | "text">;

Details

In interactive modes, the maximum title length is 254 characters; otherwise, the maximum length is 200 characters. If the length of the specified title is greater than the value of the LINESIZE option, the title is truncated to the line size.

WAITFOR Statement

Suspends execution of the current SAS session until the specified tasks finish executing.

Valid: anywhere

UNIX specifics: all

Syntax

```
WAITFOR <_ANY_ | _ALL_> taskname <taskname...> <TIMEOUT=seconds>;
```

taskname

specifies the name of the tasks that you want to wait for. See “SYSTASK Statement” on page 334 for information about task names. The task names that you specify must match exactly the task names assigned through the SYSTASK COMMAND statement. You cannot use wildcards to specify task names.

ANY | _ALL_

suspends execution of the current SAS session until either one or all of the specified tasks finishes executing. The default setting is _ANY_, which means that as soon as one of the specified tasks completes executing, the WAITFOR statement will finish executing.

TIMEOUT=*seconds*

specifies the maximum number of seconds that WAITFOR should suspend the current SAS session. If you do not specify the TIMEOUT option, WAITFOR will suspend execution of the SAS session indefinitely.

Details

The WAITFOR statements suspends execution of the current SAS session until the specified tasks finish executing or until the TIMEOUT= interval has elapsed. If the specified task was started with the WAIT option, then the WAITFOR statement ignores that task. See “SYSTASK Statement” on page 334 for a description of the WAIT option.

For example, the following statement starts three different X Windows programs and waits for them to complete:

```
systask command "xv" taskname=pgm1;
systask command "xterm" taskname=pgm2;
systask command "xcalc" taskname=pgm3;
waitfor _all_ pgm1 pgm2 pgm3;
```

The WAITFOR statement can be used to execute multiple concurrent SAS sessions. The following statements start three different SAS jobs and suspend the execution of the current SAS session until those three jobs have finished executing:

```
systask command "sas myprog1.sas" taskname=sas1;
systask command "sas myprog2.sas" taskname=sas2;
systask command "sas myprog3.sas" taskname=sas3;
waitfor _all_ sas1 sas2 sas3;
```

Note: In this method, SAS terminates after each command, which can result in reduced performance. SAS/CONNECT can also be used for executing parallel SAS sessions. See the *SAS/CONNECT User's Guide* for more information. \triangle

The SYSRC macro variable contains the return code for the WAITFOR statement. If a WAITFOR statement cannot execute successfully, the SYSRC macro variable will contain a nonzero value. For example, the WAITFOR statement might contain syntax errors. If the number of seconds specified with the TIMEOUT option elapses, then the WAITFOR statement finishes executing, and SYSRC is set to a nonzero value if one of the following occurs:

- you specify a single task that does not finish executing
- you specify more than one task and the `_ANY_` option (which is the default setting), but none of the tasks finishes executing
- you specify more than one task and the `_ALL_` option, and any one of the tasks does not finish executing

Any task whose status variable is still NULL after the WAITFOR statement has executed did not complete execution. See “SYSTASK Statement” on page 334 for a description of status variables for individual tasks.

See Also

- “SYSTASK Statement” on page 334
- “X Statement” on page 339
- *SAS/CONNECT User’s Guide*
- “Executing Operating System Commands from Your SAS Session” on page 14

X Statement

Issues an operating environment command from within a SAS session.

Valid: anywhere

UNIX specifics: valid operating system command

See: X Statement in *SAS Language Reference: Dictionary*

Syntax

X <’operating system command’>;

operating system command

specifies the UNIX command. If you specify only one UNIX command, you do not need to enclose it in quotation marks. Also, if you are running SAS from the Korn shell, you cannot use aliases.

Details

The X statement issues a UNIX command from within a SAS session. SAS executes the X statement immediately.


Neither the X statement nor the %SYSEXEC macro program statement is intended for use during the execution of a DATA step. The CALL SYSTEM routine, however, can

be executed within a DATA step. See “CALL SYSTEM Routine” on page 261 for an example.

Note: The X statement is not supported without arguments under the X Window System. Δ

See Also

- “Executing Operating System Commands from Your SAS Session” on page 14



CHAPTER

18

System Options under UNIX

<i>SAS System Options under UNIX</i>	342
<i>Behavior or Syntax That Is Specific to UNIX Environments</i>	342
<i>When You Use Parentheses in a Command Line</i>	343
<i>List of SAS System Options for UNIX</i>	343
<i>Determining How a System Option Was Set</i>	343
<i>Summary of All SAS System Options in UNIX Environments</i>	343
<i>ALTLOG System Option</i>	354
<i>ALTPRINT System Option</i>	355
<i>APPEND System Option</i>	356
<i>AUTHPROVIDERDOMAIN</i>	357
<i>AUTOEXEC System Option</i>	357
<i>AUTOSAVELOC System Option</i>	358
<i>BUFNO System Option</i>	359
<i>BUFSIZE System Option</i>	360
<i>CATCACHE System Option</i>	361
<i>CLEANUP System Option</i>	362
<i>CONFIG System Option</i>	363
<i>DEVICE System Option</i>	363
<i>ECHO System Option</i>	364
<i>EDITCMD System Option</i>	365
<i>EMAILSYS System Option</i>	366
<i>ENGINE System Option</i>	367
<i>FILELOCKS System Option</i>	368
<i>FILELOCKWAITMAX= System Option</i>	370
<i>FMTSEARCH System Option</i>	371
<i>FONTSLLOC System Option</i>	371
<i>FULLSTIMER System Option</i>	372
<i>HELPINDEX System Option</i>	375
<i>HELPLLOC System Option</i>	376
<i>HELPTOC System Option</i>	377
<i>INSERT System Option</i>	378
<i>JREOPTIONS System Option</i>	379
<i>LINESIZE System Option</i>	380
<i>LOG System Option</i>	381
<i>LPTYPE System Option</i>	382
<i>MAPS System Option</i>	383
<i>MAXMEMQUERY System Option</i>	383
<i>MEMSIZE System Option</i>	384
<i>MSG System Option</i>	386
<i>MSGCASE System Option</i>	386
<i>MSYMTABMAX System Option</i>	387

<i>MVARSIZE</i> System Option	388
<i>NEWS</i> System Option	389
<i>OBS</i> System Option	389
<i>OPLIST</i> System Option	390
<i>PAGESIZE</i> System Option	391
<i>PATH</i> System Option	392
<i>PRIMARYPROVIDERDOMAIN=</i> System Option	392
<i>PRINT</i> System Option	393
<i>PRINTCMD</i> System Option	393
<i>REALMEMSIZE</i> System Option	394
<i>RSASUSER</i> System Option	396
<i>RTRACE</i> System Option	397
<i>RTRACELOC</i> System Option	397
<i>SASAUTOS</i> System Option	398
<i>SASHELP</i> System Option	400
<i>SASSCRIPT</i> System Option	401
<i>SASUSER</i> System Option	401
<i>SET</i> System Option	402
<i>SORTANOM</i> System Option	403
<i>SORTCUT</i> System Option	403
<i>SORTCUTP</i> System Option	404
<i>SORTDEV</i> System Option	406
<i>SORTNAME</i> System Option	406
<i>SORTPARM</i> System Option	407
<i>SORTPGM</i> System Option	407
<i>SORTSIZE</i> System Option	408
<i>STDIO</i> System Option	409
<i>STIMEFMT</i> System Option	410
<i>STIMER</i> System Option	411
<i>SYSIN</i> System Option	412
<i>SYSPRINT</i> System Option	413
<i>USER</i> System Option	413
<i>VERBOSE</i> System Option	414
<i>WORK</i> System Option	415
<i>WORKINIT</i> System Option	416
<i>WORKPERMS</i> System Option	417
<i>XCMD</i> System Option	418

SAS System Options under UNIX

Behavior or Syntax That Is Specific to UNIX Environments

This section describes SAS system options that have behavior or syntax that is specific to UNIX environments. Each system option description includes a brief “UNIX specifics” section that explains which aspect of the system option is specific to UNIX. If the information under “UNIX specifics” is “all,” then the system option is described only in this documentation. Otherwise, the system option is described in this documentation and in *SAS Language Reference: Dictionary*.

When You Use Parentheses in a Command Line

On a command line, if arguments are enclosed in quotation marks, then you must use a backslash before the open parenthesis and close parenthesis so that UNIX can interpret the arguments correctly.

List of SAS System Options for UNIX

See “Summary of All SAS System Options in UNIX Environments” on page 343 for a table of all of the system options available under UNIX.

Determining How a System Option Was Set

Because of the relationship between some SAS system options, SAS might modify an option’s value. This modification might change your results.

To determine how an option was set, enter the following code in the SAS Program Editor:

```
proc options option=option value;
run;
```

After you submit this code, the SAS log will explain how the option was set. For example, the following output is displayed in the SAS log when you enter the MEMSIZE system option:

```
proc options option=MEMSIZE value;
run;
```

Output 18.1 Log Output for the MEMSIZE System Option

```
Option Value Information For SAS Option MEMSIZE
Option Value: 100663296
Option Scope: SAS Session
How option value set: Config File(s)
```

Options that are set by SAS will often say **Internal** in the **How option value was set** field. Some SAS options are internal only. You cannot specify an internal option as the **option=** value in the preceding code. If you do, SAS will return an error stating that this value is an unrecognized option value.

Summary of All SAS System Options in UNIX Environments

The following table lists every SAS system option available under UNIX. Many of these options have no host-specific behavior and are described in *SAS Language Reference: Dictionary*. If an option is available only under UNIX, it is described in this documentation. If an option is available under all environments, but has some environment-specific behavior, it is described in *SAS Language Reference: Dictionary* and this documentation. Use the following legend to know how to locate more information about an option:

Access	<i>SAS/ACCESS for Relational Databases: Reference</i>
Comp	indicates that the option is completely described in this section
Conn	<i>SAS/CONNECT User's Guide</i>
DQ	<i>SAS Data Quality Server: Reference</i>
DST	<i>Encryption in SAS</i>
ARM	<i>SAS Interface to Application Response Measurement (ARM): Reference</i>
LR	<i>SAS Language Reference: Dictionary</i>
Macro	<i>SAS Macro Language: Reference</i>
NLS	<i>SAS National Language Support (NLS): Reference Guide</i>
SAS/SHARE	<i>SAS/SHARE User's Guide</i>
SPDE	<i>SAS Scalable Performance Data Engine: Reference</i>
Web	indicates that the option is described in documentation on the SAS Web site (support.sas.com)

The table also shows the default value for each option and where you can specify the option:

- at initialization: in the SAS command, in the SASV9_OPTIONS environment variable, or in the configuration file
- in the OPTIONS statement
- in the System Options window

Table 18.1 Summary of All SAS System Options

Name	Default	Can Be Specified In			See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window	
ALTLOG	NOALTLOG	X			Comp
ALTPRINT	NOALTPRINT	X			Comp
APPEND	none	X	X		Comp /LR
APPLETLOC=	none	X	X	X	LR
ARMAGENT=	none	X	X	X	ARM
ARMLOC=	ARMLOC.LOG	X	X	X	ARM
ARMSUBSYS=	ARM_NONE	X	X	X	ARM
AUTHPROVIDERDOMAIN	NULL	X			LR/ Comp
AUTOEXEC	see description	X			Comp
AUTOSAVELOC=	none	X	X	X	Comp/ LR

Name	Default	Can Be Specified In				See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window		
AUTOSIGNON	NOAUTOSIGNON	X	X	X	Conn	
BINDING=	DEFAULT	X	X	X	LR	
BLKSIZE	256	X	X	X	LR	
BOTTOMMARGIN=	0.000	X	X	X	LR	
BUFNO=	1	X	X	X	Comp/ LR	
BUFSIZE=	0	X	X	X	Comp/ LR	
BYERR	BYERR	X	X	X	LR	
BYLINE	BYLINE	X	X	X	LR	
BYSORTED	BYSORTED	X			LR	
CAPS	NOCAPS	X	X	X	LR	
CARDIMAGE	NOCARDIMAGE	X	X	X	LR	
CATCACHE=	0	X			Comp/ LR	
CBUFNO=	0	X	X	X	LR	
CENTER	CENTER	X	X	X	LR	
CGOPTIMIZE		X	X	X	LR	
CHARCODE	NOCHARCODE	X	X	X	LR	
CLEANUP	see description	X	X	X	Comp/ LR	
CMDMAC	NOCMDMAC	X	X	X	Macro	
CMPLIB=	none	X	X	X	LR	
CMPMODEL=	BOTH	X	X	X	LR	
CMPOPT=	ALL	X	X		LR	
COLLATE	NOCOLLATE	X	X	X	LR	
COLORPRINTING	COLORPRINTING	X	X	X	LR	
COMAMID	TCP/IP	X	X	X	Conn/ Share	
COMAUX1	none	X			Share	
COMPRESS=	NO	X	X	X	LR/ SPDE	
CONFIG	see description	X (also SASV9_CONFIG environment variable)			COMP	
CONNECTPERSIST	CONNECTPERSIST	X	X	X	Conn	

Name	Default	Can Be Specified In				See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window		
CONNECTREMOTE	none	X	X	X	Conn	
CONNECTSTATUS	CONNECTSTATUS	X	X	X	Conn	
CONNECTWAIT	CONNECTWAIT	X	X	X	Conn	
COPIES=	1	X	X	X	LR	
CPUCOUNT=	1-1024 or ACTUAL	X	X	X	LR	
CPUID	CPUID	X			LR	
DATASTMTCHK=	COREKEYWORDS	X	X	X	LR	
DATE	DATE	X	X	X	LR	
DATESTYLE=	MDY	X	X	X	LR	
DBCS	NODBCS	X			NLS	
DBCSLANG	none	X			NLS	
DBCSTYPE	see description	X			NLS	
DBSLICEPARM	(THREADED_APPS, 2)	X	X	X	Access	
DBSRVTP	NONE	X			Access	
DEFLATION=	6	X	X	X	LR	
DETAILS	NODETAILS	X	X	X	LR	
DEVICE=	none	X (also, SASV9_OPTIONS environment variable, GOPTIONS statement	X	X	Comp/ LR	
DFLANG	ENGLISH	X	X	X	NLS	
DKRICOND=	ERROR	X	X	X	LR	
DKROCOND=	WARN	X	X	X	LR	
DLDMGACTION=	REPAIR	X	X	X	LR	
DMR	NODMR	X			LR/ Conn	
DMS	DMS	X			LR	
DMSEXP	DMSEXP	X			LR	
DMSLOGSIZE=	99999	X			LR	
DMSOUTSIZE=	99999	X			LR	
DMSPGMLINESIZE=	136	X			LR	
DMSSYNCHK	NODMSSYNCHK	X	X	X	LR	
DQLOCALE	none	X	X	X	DQ	
DQSETUPLOC	none	X	X	X	DQ	

Name	Default	Can Be Specified In		SAS System Options Window	See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement		
DSNFERR	DSNFERR	X	X	X	LR
DTRESET	NODTRESET	X	X	X	LR
DUPLEX	NODUPLEX	X	X	X	LR
ECHO	none	X			Comp
ECHOAUTO	NOECHOAUTO	X			LR
EDITCMD	none	X	X		Comp
EMAILAUTHPROTOCOL=	none	X	X	X	LR
EMAILFROM	none	X	X	X	LR
EMAILHOST	LOCALHOST	X			LR
EMAILID=	none	X	X	X	LR
EMAILPORT	25	X	X	X	LR
EMAILPW=	none	X	X	X	LR
EMAILSYS	SMTP	X (except SASV9_OPTIONS environment variable)	X	X	COMP
ENCODING	LATIN1	X			NLS
ENGINE=	V9	X			Comp/ LR
ERRORABEND	NOERRORABEND	X	X	X	LR
ERRORBYABEND	NOERRORBYABEND	X	X	X	LR
ERRORCHECK=	NORMAL	X	X	X	LR
ERRORS=	20	X	X	X	LR
EXPLORER	NOEXPLORER	X			LR
FILELOCKS	FAIL	X	X		Comp
FILELOCKWAITMAX=	600	X		X	Comp
FILESYNC=	HOST	X			LR
FIRSTOBS=	1	X	X	X	LR
FMTERR	FMTERR	X	X	X	LR
FMTSEARCH=	WORK LIBRARY	X	X	X	Comp/ LR
FONTEMBEDDING	FONTEMBEDDING	X	X	X	LR
FONTRENDERING=	FREETYPE_POINTS	X	X	X	LR
FONTSLC=	!SASROOT/misc/fonts	X			Comp/ LR

Name	Default	Can Be Specified In				See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window		
FORMCHAR=	---- + ----+= -/\<>*	X	X	X	LR	
FORMDLIM=	none	X	X	X	LR	
FORMS=	DEFAULT	X	X	X	LR	
FSDBTYPE	DEFAULT	X			NLS	
FSIMM	none	X			NLS	
FSIMMOPT	none	X			NLS	
FULLSTIMER	NOFULLSTIMER	X	X		Comp	
GSTYLE	GSTYLE	X	X	X	LR	
GWINDOW	GWINDOW	X	X	X	LR	
HELPBROWSER	REMOTE	X	X	X	LR	
HELPCMD	X (except SASV9_OPTIONS)				LR	
HELPHOST	null	X	X	X	Comp, LR	
HELPINDEX	/help/common.hlp/index.txt /help/common.hlp/keywords.htm common.hhk	X			Comp	
HELPLC	!SASROOT/X11/native_help/en	X			Comp	
HELPPORT	0	X	X	X	LR	
HELPTOC	/help/helpnav.help/config.txt /help/common.hlp /toc.htm, common.hhc	X			Comp	
HTTPSERVERPORTMAX		X	X		LR	
HTTPSERVERPORTMIN		X	X		LR	
IBUFNO=	0	X	X	X	LR	
IBUFSIZE=	0	X	X	X	LR	
IMPLMAC	NOIMPLMAC	X	X	X	Macro	
INITCMD	none	X			LR	
INITSTMT=	none	X			LR	
INSERT	none		X	X	Comp/ LR	
INTERVALDS	none	X	X	X	LR	
INVALIDDATA=	.	X	X	X	LR	
JPEGQUALITY=	75	X	X	X	LR	
JREOPTIONS	none	X			Comp	
LABEL	LABEL	X	X	X	LR	

Name	Default	Can Be Specified In				See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window		
LAST=	_NULL_	X	X	X	LR	
LEFTMARGIN=	0	X	X	X	LR	
LINESIZE	see description	X	X	X	Comp/ LR	
LOCALE	ENGLISH_UNITEDSTATES	X	X	X	NLS	
LOG	see description	X			Comp	
LOGPARM=	none	X			LR	
LPTYPE	none	X	X		Comp	
LRECL=	256	X	X	X	LR	
MACRO	MACRO	X			Macro	
MAPS=	!SASROOT/maps	X	X	X	Comp/ LR	
MAUTOLOCDISPLAY	NOMAUTOLOCDISPLAY	X	X	X	Macro	
MAUTOSOURCE	MAUTOSOURCE	X	X	X	Macro	
MAXMEMQUERY	256M	X	X		Comp	
MAXSEGRATIO	75	X	X	X	SPDE	
MCOMPILENOTE	NONE	X	X	X	Macro	
MEMSIZE	!SASROOT/sasv9.cfg	X			Comp	
MERGENOBY	NOWARN	X	X	X	LR	
MERROR	MERROR	X	X	X	Macro	
MEXECNOTE	NOMEXECNOTE	X	X	X	Macro	
MEXECSIZE	65536	X	X	X	Macro	
MFILE	NOMFILE	X	X	X	Macro	
MINDELIMITER	none	X	X	X	Macro	
MINPARTSIZE	16M	X			SPDE	
MISSING=	.	X	X	X	LR	
MLOGIC	NOMLOGIC	X	X	X	Macro	
MLOGICNEST	NOMLOGICNEST	X	X	X	Macro	
MPRINT	NOMPRINT	X	X	X	Macro	
MPRINTNEST	NOMPRINTNEST	X	X	X	Macro	
MRECALL	NOMRECALL	X	X	X	Macro	
MSG	!SASROOT/sasmsg	X			Comp	
MSGCASE	NOMSGCASE	X			Comp	
MSGLEVEL=	N	X	X	X	LR	
MSTORED	NOMSTORED	X	X	X	Macro	

Name	Default	Can Be Specified In			See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window	
MSYMTABMAX	4M	X	X	X	Comp/ Macro
MULTENVAPPLE	NOMULTENVAPPLE	X	X		LR
MVARSIZE	32K	X	X	X	Comp/ Macro
NETENCRYPT	NONETENCRYPT	X	X	X	DST
NETENCRYPTALGORITHM	none	X	X	X	DST
NETENCRYPTKEYLEN	0	X	X	X	DST
NEWS=	!SASROOT/misc/base/news	X			Comp/ LR
NLSCOMPATMODE	NONLSCOMPATMODE	X			NLS
NOTES	NOTES	X	X	X	LR
NUMBER	NUMBER	X	X	X	LR
OBS	MAX	X	X	X	Comp/ LR
OPLIST	NOOPLIST	X			Comp
ORIENTATION=	PORTRAIT	X	X	X	LR
OVP	NOOVP	X	X	X	LR
PAGEBREAKINITIAL	NOPAGEBREAKINITIAL	X			LR
PAGENO=	1	X	X	X	LR
PAGESIZE=	see description	X	X	X	Comp/ LR
PAPERDEST=	none	X	X	X	LR
PAPERSIZE=	LETTER	X	X	X	LR
PAPERSOURCE=	none	X	X	X	LR
PAPERTYPE=	PLAIN	X	X	X	LR
PARAM=	none	X	X	X	LR
PARMCARDS=	FT15F001	X	X	X	LR
PATH	!SASROOT/sasexe	X			Comp
PDFACCESS	none	X	X	X	LR
PDFASSEMBLY	NOPDFASSEMBLY	X	X	X	LR
PDFCOMMENT	NOPDFCOMMENT	X	X	X	LR
PDFCONTENT	none	X	X	X	LR
PDFCOPY	PDFCOPY	X	X	X	LR
PDFFILLIN	none	X	X	X	LR

Name	Default	Can Be Specified In			See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window	
PDFPAGELAYOUT=	DEFAULT	X	X	X	LR
PDFPAGEVIEW=	DEFAULT	X	X	X	LR
PDFPASSWORD=		X	X		LR
PDFPRINT=	HRES	X	X	X	LR
PDFSECURITY=	NONE	X	X	X	LR
PRIMARYPROVIDERDOMAIN=	none	X (except SASV9_OPTIONS)			COMPLR
PRINT	see description	X			Comp
PRINTCMD	none	X	X		Comp
PRINTERPATH=	PostScript Level 1	X			LR
		X			LR
PRINTINIT	NOPRINTINIT	X			LR
PRINTMSGLIST	PRINTMSGLIST	X	X	X	LR
QUOTELENMAX	QUOTELENMAX	X	X	X	LR
REALMEMSIZE	0	X			Comp
REPLACE	REPLACE	X	X	X	LR
REUSE=	NO	X	X	X	LR
RIGHTMARGIN=	0.000	X	X	X	LR
RSASUSER	NORSASUSER	X			Comp/ LR
RTRACE	NONE	X			Comp
RTRACELOC	none	X			Comp
S=	none	X	X	X	LR
S2=	none	X	X	X	LR
S2V=	S2	X	X	X	LR
SASAUTOS	SASAUTOS fileref	X	X	X	Comp/ Macro
SASCMD	none	X	X	X	Conn
SASFRSCR	none	Valid in SAS Component Language			Conn
SASHELP=	!SASROOT/sashelp	X			Comp/ LR
SASMSTORE	none	X	X	X	Macro
SASSCRIPT	!SASROOT/misc/connect	X	X	X	Comp/ Conn

Name	Default	Can Be Specified In		SAS System Options Window	See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement		
SASUSER	~SASUSER.v92	X			Comp/ LR
SEQ=	8	X	X	X	LR
SERROR	SERROR	X	X	X	Macro
SET	none	X	X		Comp
SETINIT	NOSETINIT	X			LR
SHARESESSIONCNTL	SERVER	X	X	X	Share
SIGNONWAIT	SIGNONWAIT	X	X	X	Conn
SKIP=	0	X	X	X	LR
SOLUTIONS	SOLUTIONS	X			LR
SORTANOM	none	X	X		Comp
SORTCUT	0	X	X		Comp
SORTCUTP	0	X	X		Comp
SORTDEV	see description	X	X		Comp
SORTDUP=	PHYSICAL	X	X	X	LR
SORTEQUALS	SORTEQUALS	X	X	X	LR
SORTNAME	none	X	X		Comp
SORTPARM	none	X	X		Comp
SORTPGM	BEST	X	X		Comp
SORTSEQ	none	X	X	X	NLS
SORTSIZE	value of MAX	X	X	X	Comp/ LR
SORTVALIDATE	NOSORTVALIDATE	X	X	X	LR
SOURCE	SOURCE	X	X	X	LR
SOURCE2	NOSOURCE2	X	X	X	LR
SPDEINDEXSORTSIZE	32M	X	X	X	SPDE
SPDEMAXTHREADS	0	X			SPDE
SPDESORTSIZE	32M	X	X	X	SPDE
SPDEUTILLOC	none	X			SPDE
SPDEWHEVAL	COST	X			SPDE
SPOOL	NOSPOOL	X	X	X	LR
SQLCONSTDATETIME	none	X	X	X	LR
SQLMAPPUTTO	SAS_PUT	X	X		Access
SQLREDUCEPUT=	DBMS	X	X	X	LR
SQLREDUCEPUTOBS=	0	X	X	X	LR

Name	Default	Can Be Specified In				See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window		
SQLREDUCEPUTVALUES=	0	X	X	X		LR
SQLREMERGE	none	X	X	X		LR
SQLLUNDOPOLICY=	REQUIRED	X	X			LR
SSLCALISTLOC	none	X	X	X		DST
SSLCERTLOC	none	X	X	X		DST
SSLCLIENTAUTH	NOSSLCLIENTAUTH	X	X	X		DST
SSLCRLCHECK	NOSSLCRLCHECK	X	X	X		DST
SSLCRLLOC	none	X	X	X		DST
SSLPVTKEYLOC	none	X	X	X		DST
SSLPVTKEYPASS	none	X	X	X		DST
STARTLIB	STARTLIB	X				LR
STDIO	NOSTDIO	X				Comp
STEPCHKPT	NOSTEPCHKPT	X				LR
STEPCHKPTLIB=	WORK	X				LR
STEPRESTART	none	X				LR
STIMEFMT	M	X	X			Comp
STIMER	STIMER	X	X			Comp
SUMSIZE=	0	X	X	X		LR
SVGCONTROLBUTTONS	NOSVGCONTROLBUTTONS	X	X	X		LR
SVGHEIGHT=	none	X	X	X		LR
SVGPRESERVEASPECTRATIO=	none	X	X	X		LR
SVGTITLE=	none	X	X	X		LR
SVGVIEWBOX=	none	X	X	X		LR
SVGWIDTH=	none	X	X	X		LR
SVGX=	none	X	X	X		LR
SVGY=	none	X	X	X		LR
SYMBOLGEN	NOSYMBOLGEN	X	X	X		Macro
SYNTAXCHECK	SYNTAXCHECK	X	X	X		LR
SYSIN	none	X				Comp
SYSPARM	none	X	X	X		Macro
SYSPRINT	default system printer	X	X			Comp
SYSPRINTFONT=	none	X	X	X		LR
TBUFSIZE	0	X	X	X		Conn/ Share

Name	Default	Can Be Specified In			See
		SAS Invocation, SASV9_OPTIONS, Configuration File	OPTIONS Statement	SAS System Options Window	
TCPPORTFIRST	0	SAS invocation on the remote host			Conn
TCPPORTLAST	0	SAS invocation on the remote host			Conn
TERMINAL	TERMINAL	X			LR
TERMSTMT=	none	X			LR
TEXTURELOC=	!SASROOT/misc/textures	X	X	X	LR
THREADS	THREADS	X	X	X	LR
TOOLSMENU	TOOLSMENU	X			LR
TOPMARGIN=	0.000	X	X	X	LR
TRAINLOC=	none	X			LR
TRANTAB	none	X	X	X	NLS
UPRINTCOMPRESSION	UPRINTCOMPRESSION	X	X	X	LR
UNIVERSALPRINT	UNIVERSALPRINT	X			LR
USER=	none	X	X	X	Comp/ LR
UTILLOC=	WORK	X			LR
UIDCOUNT=	100	X	X	X	LR
UIDGENDHOST=	none	X			LR
V6CREATEUPDATE=	NOTE	X			LR
VALIDFMTNAME=	LONG	X	X	X	LR
VALIDVARNAME=	V7	X	X	X	LR
VARLENCHR		X	X	X	LR
VERBOSE	NOVERBOSE	X			Comp
VIEWMENU	VIEWMENU	X			LR
VNFERR	VNFERR	X	X	X	LR
WORK=	see description	X			Comp/ LR
WORKINIT	WORKINIT	X			Comp/ LR
WORKPERMS	700	X			Comp
WORKTERM	WORKTERM	X	X	X	LR
YEARCUTOFF=	1920	X	X	X	LR

ALTLOG System Option

Specifies the destination for the SAS log.

Default: NOALTLOG

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES, LOGCONTROL

UNIX specifics: all

Syntax

`-ALTLOG file-specification | -NOALTLOG`

-ALTLOG *file-specification*

specifies the location where an alternate SAS log is to be written. The *file-specification* argument can be any valid UNIX path to a directory, a filename, or an environment variable that is associated with a path. If you specify only the path to a directory, the SAS log is placed in a file in the specified directory. The name of the file will be *filename.log*, where *filename* is the name of your SAS job. If you are running SAS interactively and specify only the path to a directory, the log is written to a file named *sas.log* within that path.

-NOALTLOG

specifies that the SAS log is not copied.

Details

The ALTLOG system option specifies a destination to which a copy of the SAS log is written. All messages that are written to the SAS log are also written to the location specified in *file-specification*. You can use this option to capture log output for printing.

Note: You can use the LOG option in the PRINTTO procedure to redirect any portion of the log to an external file. The code for PROC PRINTTO will not appear in the SAS log for the current session, but it will appear in the SAS log that you created with the ALTLOG system option. △

Note: When SAS is started with the OBJECTSERVER and NOTERMINAL system options and no log is specified, SAS discards all log and alternate log messages. △

See Also

- “ALTPRINT System Option” on page 355
- “PRINTTO Procedure” on page 304
- “Using SAS System Options to Route Output” on page 102
- “The SAS Log” in *SAS Language Reference: Concepts*

ALTPRINT System Option

Specifies the destination for the output files from SAS procedures.

Default: NOALTPRINT

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

-ALTPRINT *file-specification* | **-NOALTPRINT**

-ALTPRINT *file-specification*

specifies the location for copies of procedure output to be written. The *file-specification* argument can be any valid UNIX path to a directory, a filename, or an environment variable that is associated with a path. If you specify only the path to a directory, the copy is placed in a file in the specified directory. The name of the file will be *filename.lst*, where *filename* is the name of your SAS job. If you are running SAS interactively and specify only the path to a directory, the output is written to a file named *sas.lst*.

-NOALTPRINT

causes any previous ALTPRINT specifications to be ignored.

Details

The ALTPRINT system option specifies a destination to which copies of the SAS procedure output file are written. All messages that are written to the SAS procedure output file are also written to the location specified in *file-specification*. You can use this option to capture the procedure output for printing.

See Also

- “ALTLOG System Option” on page 354
- “Using SAS System Options to Route Output” on page 102

APPEND System Option

Used when SAS starts; appends the specified value to the existing value at the end of the specified system option.

Default: none

Valid in: configuration file, SAS invocation

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

-APPEND *system-option new-option-value*

system-option

can be FMTSEARCH, HELPLOC, MAPS, MSG, SASAUTOS, or SASHELP.

new-option-value

is the new value that you want to append to the current value of *system-option*.

Details

By default, if you specify the FMTSEARCH, HELPLOC, MAPS, MSG, SASAUTOS, or SASHELP system option more than one time, the last value that is specified is the value that SAS uses. If you want to add additional pathnames to the pathnames already specified by one of these options, you must use the APPEND system option to add the new pathname. For example, if you entered the following SAS command, the only location that SAS will look for help files is **/apps/help**, and the output of PROC OPTIONS will show only **/apps/help**:

```
sas -helploc /sas/help -helploc /apps/help
```

If you want SAS to look in both locations for help files, you must use the APPEND option:

```
sas -helploc /sas/help -append helploc /apps/help
```

For the value of the HELPLOC option, PROC OPTIONS will now show the following:

```
(' /sas/help' '/apps/help')
```

See Also

- See the APPEND system option in *SAS Language Reference: Dictionary* to append a value to a system option after SAS starts
- “INSERT System Option” on page 378

AUTHPROVIDERDOMAIN

Associates a domain suffix with an authentication provider.

Default: NULL

Valid in: configuration file, SAS invocation

Alias: AUTHPD

Category: Environment control: Initialization and operation

PROC OPTIONS GROUP: EXECMODES

See: AUTHPROVIDERDOMAIN in *SAS Language Reference: Dictionary*

AUTOEXEC System Option

Specifies the SAS autoexec file.

Default: `autoexec.sas` (see Details)

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

`-AUTOEXEC file-specification | -NOAUTOEXEC`

-AUTOEXEC *file-specification*

specifies the SAS autoexec file. The *file-specification* must resolve to a valid UNIX pathname.

-NOAUTOEXEC

specifies that SAS is not to process any autoexec files.

Details

The AUTOEXEC system option specifies the autoexec file. The autoexec file contains SAS statements that are executed automatically when you invoke SAS or when you start another SAS process. The autoexec file can contain any SAS statements. For example, your autoexec file can contain LIBNAME statements for SAS libraries that you access routinely in SAS sessions.

SAS looks for this option in the following order:

- 1 in the command line
- 2 in the SASV9_OPTIONS environment variable
- 3 in the configuration file

It uses the first AUTOEXEC option it encounters and ignores all others.

If neither AUTOEXEC nor NOAUTOEXEC is specified, SAS searches three directories for an autoexec.sas file in the following order:

- 1 your current directory
- 2 your home directory
- 3 the **!SASROOT** directory (see Appendix 1, “The !SASROOT Directory,” on page 421)

SAS uses the first file it finds to initialize the session.

If you want to see the contents of the autoexec file for your session, use the ECHOAUTO system option when you invoke SAS.

See Also

- “Customizing Your SAS Session by Using System Options” on page 17

AUTOSAVELOC System Option

Specifies the location of the Program Editor autosave file.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window

Category: Environment control: Display

PROC OPTIONS GROUP= ENVDISPLAY

UNIX specifics: valid values of *pathname*

See: AUTOSAVELOC= System Option in *SAS Language Reference: Dictionary*

Syntax

-AUTOSAVELOC *fileref* | *pathname*

AUTOSAVELOC= *fileref* | *pathname*

fileref

specifies a fileref to the location where the autosave file is saved.

pathname

specifies the pathname of the autosave file. The *pathname* must be a valid UNIX pathname.

Details

By default, SAS saves the Program Editor autosave file, *pgm.asv*, in the open folder. You can use the AUTOSAVELOC system option to specify a different location for the autosave file.

See Also

- “SETAUTOSAVE Command” on page 233

BUFNO System Option

Specifies the number of buffers to be allocated for processing a SAS data set.

Default: 1

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Files: SAS Files

PROC OPTIONS GROUP= SASFILES, PERFORMANCE

UNIX specifics: default value

See: BUFNO System Option in *SAS Language Reference: Dictionary*

Syntax

-BUFNO *n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

BUFNO=*n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

n* | *nK* | *nM* | *nG

specifies the number of buffers in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); or 1,073,741,824 (gigabytes). You can specify decimal values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **8** specifies 8 buffers, a value of **.782k** specifies 801 buffers, and a value of **3m** specifies 3,145,728 buffers.

hexX

specifies the number of buffers as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, **2dx** specifies 45 buffers.

MIN

sets the number of buffers to 0, and requires SAS to use the default value of 1.

MAX

sets the number of buffers to 2,147,483,647.

Details

The number of buffers is not a permanent attribute of the data set; it is valid only for the current SAS session or job.

BUFNO= applies to SAS data sets that are opened for input, output, or update.

Using BUFNO= can improve execution time by limiting the number of input/output operations that are required for a particular SAS data set. The improvement in execution time, however, comes at the expense of increased memory consumption.

Under UNIX, the maximum number of buffers that you can allocate is determined by the amount of memory available.

BUFSIZE System Option

Specifies the size of a permanent buffer page for an output SAS data set.

Default: 0

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Files: SAS Files

PROC OPTIONS GROUP= SASFILES, PERFORMANCE

UNIX specifics: valid range

See: BUFSIZE System Option in *SAS Language Reference: Dictionary*

Syntax

-BUFSIZE *n* | *nK* | *nM* | *nG* | *hexX* | MAX

BUFSIZE=*n* | *nK* | *nM* | *nG* | *hexX* | MAX

n* | *nK* | *nM* | *nG

specifies the buffer page size in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); or 1,073,741,824 (gigabytes). You can specify decimal values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **8** specifies 8 bytes, a value of **.782k** specifies 801 bytes, and a value of **3m** specifies 3,145,728 bytes.

hexX

specifies the buffer page size as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, **2dx** sets the buffer page size to 45 bytes.

MAX

sets the buffer page size to 2,147,483,647.

Details

The buffer page size can range from 1K to 2G–1.

If you specify a nonzero value when you create a SAS data set, the BASE engine uses that value. If that value cannot hold at least one observation or is not a multiple of 1K, the engine rounds the value up to a multiple of 1K.

CATCACHE System Option

Specifies the number of SAS catalogs to keep open.

Default: 0

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Files: SAS Files

PROC OPTIONS GROUP= SASFILES

UNIX specifics: Valid values for *n*

See: CATCACHE System Option in *SAS Language Reference: Dictionary*

Syntax

-CATCACHE *n* | *nK* | MIN | MAX

n* | *nK

specifies the number of open-file descriptors to keep in cache memory in multiples of 1 (*n*) or 1,024 (*nK*). You can specify decimal values for the number of kilobytes. For example, a value of **8** specifies 8 open-file descriptors, a value of **.782k** specifies 801 open-file descriptors, and a value of **3k** specifies 3,072 open-file descriptors.

If *n* > 0, SAS places up to that number of open-file descriptors in cache memory, instead of closing the catalogs.

MIN

sets the number of open-file descriptors that are kept in cache memory to 0.

MAX

sets the number of open-file descriptors that are kept in cache memory to 32,767.

Details

By using the CATCACHE system option to specify the number of SAS catalogs to keep open, you can avoid repeatedly opening and closing the same catalogs.

See Also

- The section on optimizing system performance in *SAS Language Reference: Concepts*

CLEANUP System Option

Specifies how to handle out-of-resource conditions.

Default: CLEANUP for interactive modes; NOCLEANUP otherwise

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Environment control: Error handling

PROC OPTIONS GROUP= ERRORHANDLING

UNIX specifics: behavior when running in interactive line mode and batch mode

See: CLEANUP System Option in *SAS Language Reference: Dictionary*

Syntax

-CLEANUP | -NOCLEANUP

CLEANUP | NOCLEANUP

CLEANUP

specifies that during the entire session, SAS attempts to perform automatic, continuous cleanup of resources that are not essential for execution. Nonessential resources include those resources that are not visible to the user (for example, cache memory) and those resources that are visible to the user (for example, the Keys window).

CLEANUP does not prompt you before SAS attempts to clean up your disk. However, when an out-of-disk-space condition occurs and your display is attached to the process, you are prompted with a menu selection even if the CLEANUP option is on. If you do not want to be prompted for out-of-disk-space conditions, use the CLEANUP option with the NOTERMINAL option.

When the CLEANUP option is on, SAS performs automatic continuous cleanup. If not enough resources are recovered, the request for the resource fails, and an appropriate error message is written to the SAS log.

CLEANUP is the default in batch mode because there is no display attached to the process to accommodate prompting.

NOCLEANUP

specifies that SAS allows the user to choose how to handle an out-of-resource condition. When NOCLEANUP is in effect and SAS cannot execute because of a lack of resources, SAS automatically attempts to clean up resources that are not visible to the user (for example, cache memory). However, resources that are visible to the user

(for example, the Keys window) are not automatically cleaned up. Instead, SAS prompts you before attempting to clean up your disk.

Details

The CLEANUP system option indicates whether you should be prompted with a menu of items to be cleaned up when SAS encounters an out-of-resource condition. In batch mode, SAS ignores this option, and if an out-of-resource condition occurs, the SAS session terminates.

CONFIG System Option

Specifies the configuration file that is used when initializing or overriding the values of SAS system options.

Default: `sasv9.cfg` (see “Order of Precedence for Processing SAS Configuration Files” on page 21)

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable, SASV9_CONFIG environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

`-CONFIG file-specification | -NOCONFIG`

-CONFIG *file-specification*

specifies a configuration file to be read. The *file-specification* must resolve to a valid UNIX filename.

-NOCONFIG

specifies that any previous CONFIG specification should be ignored and that the default system options should be used.

Details

Configuration files contain system option specifications that execute automatically whenever SAS is invoked.

If you specify the CONFIG system option in a configuration file, the option is ignored.

See Also

- “Customizing Your SAS Session by Using System Options” on page 17

DEVICE System Option

Specifies a device driver for graphics output for SAS/GRAPH software.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable, GOPTIONS statement

Category: Graphics: Driver settings

PROC OPTIONS GROUP= GRAPHICS

UNIX specifics: valid device drivers

See: DEVICE System Option in *SAS Language Reference: Dictionary*

Syntax

`-DEVICE device-driver-name`

`DEVICE=device-driver-name`

device-driver-name

specifies the name of a device driver for graphics output.

Details

To see the list of device drivers that are available under UNIX, you can use the GDEVICE procedure. If you are using the SAS windowing environment, submit the following statements:

```
proc gdevice catalog=sashelp.devices;
run;
```

If you are running SAS in interactive line mode or batch mode, submit the following statements:

```
proc gdevice catalog=sashelp.devices nofs;
  list _all_;
run;
```

See Also

- *SAS/GRAPH: Reference*

ECHO System Option

Specifies a message to be echoed to the computer.

Default: none

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log

PROC OPTIONS GROUP= LOGCONTROL

UNIX specifics: all

Syntax

`-ECHO "message" | -NOECHO`

-ECHO “message”

specifies the text of the message to be echoed to the computer. The text must be enclosed in single or double quotation marks if the message is more than one word. Otherwise, the quotation marks are not needed.

-NOECHO

specifies that no messages are to be echoed to the computer.

Details

Messages that result from errors in the autoexec file are printed in the SAS log regardless of how the ECHO system option is set.

You can specify multiple ECHO options. The strings are displayed in the order in which SAS encounters them. See “How SAS Processes System Options Set in Multiple Places” on page 20 for information on how that order is determined.

Example

For example, you can specify the following:

```
-echo "SAS 9.2 under UNIX is initializing."
```

The message appears in the Log window as SAS initializes.

See Also

- “ECHOAUTO System Option” □ in *SAS Language Reference: Dictionary*

EDITCMD System Option

Specifies the host editor to be used with the HOSTEDIT command.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Environment control: Display

PROC OPTIONS GROUP= ENVDISPLAY

UNIX specifics: all

Syntax

```
-EDITCMD "host-editor-pathname editor-options"
```

```
EDITCMD="host-editor-pathname editor-options"
```

Details

The EDITCMD system option specifies the command that is issued to the operating environment. If you are using a terminal-based editor, such as vi, you must specify a command that runs the editor inside a terminal emulator window.

You can define the EDITCMD option using the SASV9_OPTIONS environment variable as part of a configuration file or on the command line to make the definition available automatically to SAS. The option must be specified as a quoted string. You can use either single or double quotation marks. You can change the value for the EDITCMD option during a SAS session by issuing an OPTIONS statement.

The host editor that you specify is used when you issue the HOSTEDIT command. The HOSTEDIT command is valid only when you are running SAS in a windowing environment.

If you do not specify the full pathname, SAS searches the pathnames specified in the \$PATH environment variable. For example, to use vi, you would specify the following:

```
sas -editcmd "/usr/bin/x11/xterm -e /usr/bin/vi"
```

See Also

- “Configuring SAS for Host Editor Support in UNIX Environments” on page 160

EMAILSYS System Option

Specifies the e-mail protocol to use for sending electronic mail.

Default: SMTP

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window

Category: Communications: Email

PROC OPTIONS GROUP= EMAIL

UNIX specifics: all

Syntax

-EMAILSYS SMTP | *name-of-script*

EMAILSYS=SMTP | *name-of-script*

SMTP

specifies the Simple Mail Transfer Protocol (SMTP) electronic mail interface.

name-of-script

specifies which script to use for sending electronic mail from within SAS. Some external scripts do not support sending e-mail attachments. These scripts are not supported by SAS.

Details

The EMAILSYS system option specifies which e-mail protocol to use for sending electronic mail from within SAS. Specifying SMTP supports sending e-mail attachments on UNIX, but might require changing the values of the EMAILHOST= and EMAILPORT=, depending on your site configuration.

You can set the EMAILSYS option at any time in your SAS session.

See Also

- “Sending Mail from within Your SAS Session in UNIX Environments” on page 158
- “Sending Electronic Mail Using the FILENAME Statement (EMAIL)” on page 82
- “EMAILHOST System Option” in *SAS Language Reference: Dictionary*
- “EMAILPORT System Option” in *SAS Language Reference: Dictionary*
- “The SMTP E-mail Interface” in *SAS Language Reference: Concepts*

ENGINE System Option

Specifies the default access method to use for SAS libraries.

Default: V9

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Files: SAS Files

PROC OPTIONS GROUP= SASFILES

UNIX specifics: valid values of *engine-name*

See: ENGINE= System Option in *SAS Language Reference: Dictionary*

Syntax

-ENGINE *engine-name*

engine-name

can be one of the following under UNIX:

BASE | V9

specifies the default SAS engine for SAS 9 through SAS 9.2 files.

V8

specifies the SAS engine for all SAS Version 9 files.

V7

specifies the SAS engine for all Version 7 files.

V6

specifies the SAS engine for Release 6.09 through Release 6.12. This engine is read-only.

V9TAPE

specifies the default sequential engine for SAS 9 and SAS 9.1 files.

V8TAPE | V7TAPE

specifies the SAS sequential engine for all Version 8 and Version 7 files. These engines are identical to the V9TAPE engine.

V6TAPE

specifies the SAS sequential engine for Version 6 files. This engine is read-only.

See Also

- “Compatible Computer Types in UNIX Environments” on page 42
- “ENGINE System Option” in *SAS Language Reference: Dictionary*
- *SAS Language Reference: Concepts*

FILELOCKS System Option

Specifies whether file locking is turned on or off and what action should be taken if a file cannot be locked.

Default: FAIL

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window

Category: Files: External Files | SAS Files

PROC OPTIONS GROUP= EXTFILES, SASFILES, ENVFILES

UNIX specifics: all

Syntax

-FILELOCKS *setting path* | *path setting*

-FILELOCKS NONE | FAIL | CONTINUE | RESET

FILELOCKS=(*setting path* | *path setting*)

FILELOCKS=NONE | FAIL | CONTINUE | RESET

setting

specifies the operating environment locking value for the specified path. The following values are valid:

- NONE
- FAIL
- CONTINUE
- RESET

path

specifies a path to a UNIX directory. Enclose the path in single or double quotation marks.

Tip: The *path* argument can contain an environment variable.

NONE

turns file locking off. NONE specifies that SAS attempts to open the file without checking for an existing lock on the file. NONE does not place an operating system lock on the file. These files are not protected from shared Update access.

Tip: NONE does not suppress internal locking.

FAIL

turns file locking on. FAIL specifies that SAS attempts to place an operating system lock on the file. Access to the file is denied if the file is already locked, or if it cannot be locked. FAIL is the default value for FILELOCKS.

CONTINUE

turns file locking on. CONTINUE specifies that SAS attempts to place an operating system lock on the file. If a file is already locked by someone else, an attempt to open it fails. If the file cannot be locked for some other reason (for example, if the file system does not support locking), the file is opened and a warning message is sent to the log.

Tip: CONTINUE does not suppress internal locking.

RESET

specifies that all previous FILELOCKS settings will be deleted, and resets the global setting to the default value of FAIL. If you use the **FILELOCKS=(*setting path* | *path setting*)** syntax, then RESET resets only those files that are in *path*.

Details

The Basics of File Locking In previous releases of SAS, the FILELOCKS system option was able to lock only SAS files. In SAS 9.2, the FILELOCKS system option is able to lock external files as well.

The FILELOCKS system option enables you to lock both external files and SAS files based on global settings that you set in the FILELOCKS system option. External file locking applies to all files that are opened.

You can use multiple instances of the FILELOCKS option to establish different settings for different paths. One path can be a subdirectory of another path. In this case, the most specific matching path currently in effect governs operating system file locking. The following example shows how you can specify multiple instances of the FILELOCKS option in a configuration file:

```
filelocks = ('/u/myuserid/temp' NONE)
filelocks = ('/tmp' CONTINUE)
```

When the value of the FILELOCKS option is a set of *path* and *setting*, the path must be enclosed in quotation marks. If you use FILELOCKS on the command line, then quotation marks are not needed.

Note: To prevent data corruption, setting FILELOCKS to NONE or CONTINUE is not recommended. △

Resetting Paths by Using the *Path* and *Setting* Arguments The *path* and *setting* arguments enable you to apply a setting to a particular directory and its subtrees. If you set the value of *setting* to RESET, then the *path* and *setting* values are deleted.

For example, in the following case, **filelocks=('/' reset)**, the current values for *path* and *setting* are deleted, and FILELOCKS resets the values to the following default: **('/' fail)**.

When FILELOCKS Is Set to FAIL When FILELOCKS is set to FAIL (the default value), the following actions occur:

- SAS prevents two sessions from simultaneously opening the same SAS file for update or output.
- SAS prevents one session from reading a SAS file that another SAS session has open for update or output.
- SAS prevents one session from writing to a file that another SAS session has open in read mode.

See Also

- “WORKINIT System Option” on page 416

FILELOCKWAITMAX= System Option

Sets an upper limit on the time SAS will wait for a locked file.

Default: 600

Valid in: configuration file, SAS invocation

Category: Files: SAS Files

PROC OPTIONS GROUP= SASFILES

UNIX specifics: all

Syntax

FILELOCKWAITMAX = *wait-time*

wait-time

specifies the amount of time, in seconds, that SAS will wait for a locked file to become available.

Interaction: Specifying the FILELOCKWAITMAX= system option can have an adverse effect on one or more SAS/SHARE server and client sessions that are waiting for the release of a SAS file that is locked by another process. One or more wait conditions could lead to failed processes for a SAS/SHARE server and clients.

To prevent the possibility of a failed SAS/SHARE process, you can set FILELOCKWAITMAX=0, which cancels the amount of time that a SAS/SHARE server and clients would wait for the release of a locked file. Canceling the wait time would prevent a failed process. For more information, see the LIBNAME statement option FILELOCKWAIT=.

Range: 0–600

Default: 600

Details

The FILELOCKWAITMAX= system option enables you to limit or turn off the amount of time SAS will wait for a locked file. SAS uses the FILELOCKWAIT= LIBNAME option to wait for the file to become available. Using the FILELOCKWAITMAX= system option, an administrator can limit or turn off this behavior. Normally, SAS returns an error if the file it attempts to access is locked. If you set FILELOCKWAITMAX= to 0, SAS fails immediately upon encountering a locked file. This option is used primarily by a system administrator.

For more information about the FILELOCKWAIT= option in the LIBNAME statement, see the “LIBNAME Statement” on page 328.

See Also

System Option:

“FILELOCKS System Option” on page 368

FMTSEARCH System Option

Specifies the order in which format catalogs are searched.

Default: (WORK LIBRARY)

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVDISPLAY

UNIX specifics: valid values for *catalog-specification*

See: FMTSEARCH= System Option in *SAS Language Reference: Dictionary*

Syntax

-FMTSEARCH (*catalog-specification-1... catalog-specification-n*)

FMTSEARCH=(*catalog-specification-1... catalog-specification-n*)

catalog-specification

specifies the order in which format catalogs are searched until the desired member is found. The value of *libref* can be either *libref* or *libref.catalog*. If only the *libref* is given, SAS assumes that FORMATS is the catalog name.

Note: The value of *libref* must be in uppercase characters. Δ

Details

To add additional *catalog-specification* entries, use the INSERT or APPEND system options. For more information, see the INSERT System Option and APPEND System Option in *SAS Language Reference: Dictionary*.

FONTSLOC System Option

Specifies the location of the SAS fonts that are loaded during the SAS session.

Default: !SASROOT/misc/fonts

Valid in: configuration file, SAS invocation

Category: Environment control: Display

PROC OPTIONS GROUP= ENVDISPLAY

UNIX specifics: valid pathname

See: FONTSLOC= System Option in *SAS Language Reference: Dictionary*

Syntax

-FONTSLOC "*directory-specification*"

"directory-specification"

specifies the directory that contains the SAS fonts that are loaded during the SAS session. The *directory-specification* must be enclosed in double quotation marks.

Details

The directory must be a valid operating environment pathname.

FULLSTIMER System Option

Specifies whether to write all available system performance statistics and the datetime stamp to the SAS log.

Default: NOFULLSTIMER

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log

PROC OPTIONS GROUP= LOGCONTROL

UNIX specifics: all

Syntax

-FULLSTIMER | -NOFULLSTIMER

FULLSTIMER | NOFULLSTIMER

FULLSTIMER

writes to the SAS log a list of the host-dependent resources that were used for each step and for the entire SAS session. A datetime stamp is included in the output.

NOFULLSTIMER

does not write to the SAS log a complete list of resources or a datetime stamp.

Details

SAS uses the `getrusage()` and `times()` UNIX system calls for your operating environment to obtain the statistics presented with FULLSTIMER. The datetime stamp is also listed in the output. The following is an example of FULLSTIMER output:

Output 18.2 FULLSTIMER Output

```

NOTE: DATA statement used (Total process time):
      real time           0.01 seconds
      user cpu time       0.00 seconds
      system cpu time     0.00 seconds
      Memory              236k
      OS Memory           5672k
      Timestamp           3/16/2006  9:13:39 AM
      Page Faults        3
      Page Reclaims      24
      Page Swaps         0
      Voluntary Context Switches 11
      Involuntary Context Switches 1
      Block Input Operations 2
      Block Output Operations 0

```

Note: If both FULLSTIMER and STIMER system options are set, the FULLSTIMER statistics are printed. Δ

FULLSTIMER displays the following statistics:

Table 18.2 Description of FULLSTIMER Statistics

Statistic	Description
Real Time	the amount of real time (clock time) that is spent to process the SAS job. Real time is also referred to as elapsed time.
User CPU Time	the CPU time that is spent in the user program.
System CPU Time	the CPU time that is spent to perform operating system tasks (system overhead tasks) that support the execution of your SAS code.
Memory	the amount of memory required to run a step.
OS Memory	the largest amount of operating system memory that is available to SAS during the step.
Timestamp	the date and time that a step was executed.
Page Faults	the number of pages that SAS tried to access but were not in main memory and required I/O activity.
Page Reclaims	the number of pages that were accessed without I/O activity.
Page Swaps	the number of times a process was swapped out of main memory.
Voluntary Context Switches	the number of times that the SAS process had to pause because of a resource constraint such as a disk drive.
Involuntary Context Switches	the number of times that the operating system forced the SAS session to pause processing to allow other process to run.
Block Input Operations	the number of I/O operations that are performed to read the data into memory.
Block Output Operations	the number of I/O operations that are performed to write the data to a file.

For more information about these statistics, see the man pages for the `getrusage()` and `times()` UNIX system calls.

Note: Starting in SAS 9, some procedures use multiple threads. On computers with multiple CPUs, the operating system can run more than one thread simultaneously. Consequently, CPU time might exceed real time in your FULLSTIMER output.

For example, a SAS procedure could use two threads that run on two separate CPUs simultaneously. The value of CPU time would be calculated as the following:

```
CPU1 time + CPU2 time = total CPU time
1 second + 1 second = 2 seconds
```

Because CPU1 can run a thread at the same time that CPU2 runs a separate thread for the same SAS process, you can theoretically consume 2 CPU seconds in 1 second of real time. \triangle

See Also

- \square “STIMER System Option” on page 411

HELPHOST System Option

Specifies the name of the computer where the remote browsing system is to be displayed.

Default: NULL

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, VMS_SAS_OPTIONS DCL symbol, or SASV9_OPTIONS environment variable

Category: Environment control: Help

PROC OPTIONS GROUP= HELP

UNIX specifics: all

Syntax

```
HELPHOST="host"
-HELPHOST ("host")
```

“host”

specifies the name of the computer where the remote browsing system is to be displayed. Quotation marks or parentheses are required. The maximum number of characters is 2048.

Details

If you do not specify the HELPHOST option, the remote browsing system is displayed on the host that is specified in the X Windows display setting.

Example 1: SAS Invocation

The syntax for specifying the HELPHOST system option differs for the OpenVMS environment and the UNIX environments. The syntax for specifying the HELPHOST system option for the UNIX environment is shown in the following example:

```
sas92 -helphost "my.computer.com"
```

The syntax for specifying the HELPHOST system option for the OpenVMS environment is shown in the following example:

```
sas92/helpost="my.computer.com"
```

Example 2: OPTIONS Statement

The syntax for the HELPHOST system option is identical for both the OpenVMS and UNIX operating environments when it is specified in the OPTIONS statement as shown in the following example:

```
options helpost="my.computer.com";
```

See Also

- “Installing the Remote Browser Server” on page 134

HELPINDEX System Option

Specifies one or more index files for the online SAS Help and Documentation.

Default: /help/common.hlp/index.txt, /help/common.hlp/keywords.htm, common.hhk

Valid in: configuration file, SAS invocation

Category: Environment control: Help

PROC OPTIONS GROUP= HELP

UNIX specifics: applet and HTML files must reside in the path specified by the HELPLOC option

Syntax

```
-HELPINDEX index-pathname-1 < index-pathname-2 < index-pathname-3>>
```

index-pathname

specifies the partial pathname for the index that is to be used by the online SAS Help and Documentation. The *index-pathname* can be any or all of the following:

/help/applet-index-filename

specifies the partial pathname of the index file that is to be used by the SAS Documentation Java applet in a UNIX environment. *applet-index-filename* must have a file extension of .txt, and it must reside in a path that is specified by the HELPLOC system option. The default is /help/common.hlp/index.txt.

See the default index file for the format that is required for an index file.

/help/accessible-index-filename

specifies the partial pathname of an accessible index file that is to be used by the online SAS Help and Documentation in UNIX, OpenVMS, or z/OS environments. An accessible index file is an HTML file that can be used by Web browsers.

accessible-index-filename must have a file extension of .htm, and it must reside in a path that is specified by the HELPLOC system option. The default pathname is /help/common.hlp/keywords.htm.

See the default index file for the format that is required for an index file.

HTML-Help-index-pathname

specifies the pathname of the Microsoft HTML Help index that is to be used by the online SAS Help and Documentation in Windows environments. The default pathname is **common.hhk**. For information about creating an index for Microsoft HTML Help, see your Microsoft HTML Help documentation.

Details

Use the HELPINDEX option if you have a customized index that you want to use instead of the index that SAS supplies. If you use one configuration file to start SAS in more than one operating environment, you can specify all of the partial pathnames in the HELPINDEX option. The order of the pathnames is not important, although only one pathname of each type can be specified.

When the HELPINDEX option specifies a pathname for UNIX, OpenVMS, or z/OS operating environments, SAS determines the complete path by replacing **/help/** in the partial pathname with the pathname specified in the HELPLOC option. If the HELPLOC option contains more than one pathname, SAS searches each path for the specified index.

For example, when the value of HELPINDEX is **/help/common.hlp/myindex.htm** and the value of HELPLOC is **/u/myhome/myhelp**, the complete path to the index is **/u/myhome/myhelp/common.hlp/myindex.htm**.

See Also

- “HELPLOC System Option” on page 376

HELPLOC System Option

Specifies the location of the text and index files for the facility that is used to view the online SAS Help and Documentation.

Default: `!SASROOT/X11/native_help/en`

Valid in: configuration file, SAS invocation

Category: Environment control: Help

PROC OPTIONS GROUP= HELP

UNIX specifics: default *pathname*

Syntax

`-HELPLOC (pathname<,pathname-2...pathname-n>)`

pathname

specifies one or more directory pathnames in which the online SAS Help and Documentation files are located.

Details

Specifying a value for the HELPLOC system option causes SAS to insert that value at the start of a list of concatenated values, the last of which is the default value. This

behavior enables you to access help for your site without losing access to SAS Help and Documentation.

To add pathnames, use the INSERT or APPEND system options. For more information, see the INSERT System Option and APPEND System Option in *SAS Language Reference: Dictionary*.

Example

The following command contains two specifications of HELPLOC:

```
sas -insert helploc /app2/help -insert helploc /app1/help -append
```

The value of the system option is the following:

```
/app1/help, /app2/help, !SASROOT/X11/native_help
```

See Also

- INSERT System Option in *SAS Language Reference: Dictionary*
- APPEND System Option in *SAS Language Reference: Dictionary*

HELPTOC System Option

Specifies the table of contents files for the online SAS Help and Documentation.

Default: `/help/helpnav.hlp/config.txt, /help/common.hlp/toc.htm, common.hhc`

Valid in: configuration file, SAS invocation

Category: Environment control: Help

PROC OPTIONS GROUP= HELP

UNIX specifics: applet and HTML files must reside in the path specified by the HELPLOC option

Syntax

```
-HELPTOC TOC-pathname-1 < TOC-pathname-2 < TOC-pathname-3>>
```

TOC-pathname

specifies a partial pathname for the table of contents that is to be used by the online SAS Help and Documentation. *TOC-pathname* can be any or all of the following:

/help/applet-TOC-filename

specifies the partial pathname of the table of contents file that is to be used by the SAS Documentation Java applet in a UNIX environment. The *applet-TOC-filename* must have a file extension of .txt, and it must reside in a path that is specified by the HELPLOC system option. The default is `/help/helpnav.hlp/config.txt`.

See the default table of contents file for the format that is required for an index file.

/help/accessible-TOC-filename

specifies the partial pathname of an accessible table of contents file that is to be used by the online SAS Help and Documentation in UNIX, OpenVMS, or z/OS

environments. An accessible table of contents file is an HTML file that can be used by Web browsers. The *accessible-TOC-filename* must have a file extension of .htm, and it must reside in a path that is specified by the HELPLOC system option. The default pathname is `/help/common.hlp/toc.htm`.

See the default table of contents file for the format that is required for a table of contents.

HTML-Help-TOC-pathname

specifies the complete pathname to the Microsoft HTML Help table of contents that is to be used by the online SAS Help and Documentation in Windows environments. The default pathname is `common.hhc`. For information about creating an index for Microsoft HTML Help, see your Microsoft HTML Help documentation.

Details

Use the HELPTOC if you have a customized table of contents that you want to use, instead of the table of contents that SAS provides. If you use one configuration file to start SAS in more than one operating environment, you can specify all of the partial pathnames in the HELPTOC option. The order of the pathnames is not important, although only one pathname of each type can be specified.

When the HELPTOC option specifies the pathname for UNIX, OpenVMS, or z/OS operating environments, SAS determines the complete path by replacing `/help/` in the partial pathname with the pathname specified in the HELPLOC option. If the HELPLOC option contains more than one pathname, SAS searches each path for the table of contents.

For example, when HELPTOC is `/help/common.hlp/mytoc.htm`, and the value of HELPLOC is `/u/myhome/myhelp`, the complete path to the table of contents is `/u/myhome/myhelp/common.hlp/mytoc.htm`.

See Also

- “HELPLOC System Option” on page 376

INSERT System Option

Used when SAS starts; inserts the specified value at the beginning of the specified system option.

Default: none

Valid in: configuration file, SAS invocation

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

`-INSERT system-option new-option-value`

system-option

can be FMTSEARCH, HELPLOC, MAPS, MSG, SASAUTOS, or SASHELP.

new-option-value

is the new value that you want to insert at the beginning of the current value of *system-option*.

Details

By default, if you specify the FMTSEARCH, HELPLOC, MAPS, MSG, SASAUTOS, or SASHELP system option more than one time, the last specification is the one that SAS uses. If you want to insert additional pathnames at the beginning of the search paths already specified by one of these options, you must use the INSERT system option to add the new pathname. For example, if you entered the following SAS command, the only location that SAS will look for help files is **/apps/help**, and the output of PROC OPTIONS will show only **/apps/help**:

```
sas -helploc /apps/help
```

If you want SAS to look in both the current path for help files and in **/apps/help**, and if you want SAS to search **/apps/help** first, then you must use the INSERT option:

```
sas -insert helploc /apps/help
```

If your current path for help files is **!SASROOT/X11/native_help**, then for the value of the HELPLOC option, PROC OPTIONS will now show:

```
(' /apps/help' '!SASROOT/X11/native_help')
```

See Also

- “APPEND System Option” on page 356

JREOPTIONS System Option

Identifies the Java Runtime Environment (JRE) options for SAS.

Default: none

Valid in: configuration file, SAS invocation

Category: Environment control: Initialization and operation

PROC OPTIONS GROUP= EXECMODES

UNIX specifics: all

Syntax

-JREOPTIONS (*-JRE-option-1*<*-JRE-option-n*>)

-JRE-option

specifies one or more JRE options. JRE options must begin with a hyphen (-). Use a space to separate multiple JRE options. Valid values for *JRE-option* depend on your installation's JRE. For information about JRE options, see your installation's Java documentation.

Details

JRE options must be enclosed in parentheses. If you issue JREOPTIONS options on the command line, then you must put a backslash (\) before the open parenthesis and close parenthesis, as shown in the Examples. If you specify multiple JREOPTIONS options, then SAS appends JRE options to JRE options that are currently defined. Incorrect JRE options are ignored.

Examples

```
-jreoptions \(-verbose\)
```

```
-jreoptions \(-Djava.class.path=myjava/classes/myclasses.jar:myjava2/
classes/myclasses.jar -oss600k\)
```

LINESIZE System Option

Specifies the line size of the SAS Log and Output windows.

Default: the display width setting for the interactive modes; 132 for batch mode

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log and procedure output

PROC OPTIONS GROUP= LOG_LISTCONTROL, LOGCONTROL

UNIX specifics: default values

See: LINESIZE= System Option in *SAS Language Reference: Dictionary*

Syntax

```
-LINESIZE n | hexX | MIN | MAX
```

```
LINESIZE=n | hexX | MIN | MAX
```

n

specifies the line size in characters. Valid values range between 64 and 256.

hexX

specifies the line size as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, **2dx** specifies 45 characters.

MIN

sets the line size to 64 characters.

MAX

sets the line size to 256 characters.

See Also

- “Controlling the Content and Appearance of Output in UNIX Environments” on page 104

LOG System Option

Specifies a destination for the SAS log when running in batch mode.

Default: a file in the current directory with the same filename as the SAS source file and an extension of .log

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES, LOGCONTROL

UNIX specifics: all

Syntax

`-LOG file-specification | -NOLOG`

-LOG *file-specification*

specifies the destination for the SAS log. The *file-specification* can be any valid UNIX path to a directory, a filename, or an environment variable that is associated with a path. If you specify only the path to a directory, the log file is created in the specified directory. The default name for this file is *filename.log*, where *filename* is the name of your SAS job.

-NOLOG

suppresses the creation of the SAS log. Do not use this value unless your SAS program is thoroughly debugged.

Details

The LOG system option specifies a destination for the SAS log when running in batch mode. The LOG system option is valid in batch mode; it is ignored in interactive modes.

Using directives in the value of the LOG system option enables you to control when logs are open and closed and how they are named, based on real-time events, such as time, month, day of week, and so on. For a valid list of directives, see the LOGPARM= System Option in *SAS Language Reference: Dictionary*.

If you start SAS in batch mode or server mode and the LOGCONFIGLOC= option is specified, logging is performed by the SAS logging facility. The traditional SAS log option LOGPARM= is ignored. The traditional SAS log option LOG= is honored only when the %S{App.Log} conversion character is specified in the logging configuration file. For more information, see SAS Logging Facility in *SAS Logging: Configuration and Programming Reference*.

Note: When SAS is started with the OBJECTSERVER and NOTERMINAL system options and no log is specified, SAS discards all log and alternate log messages. △

See Also

- LOGPARM= System Option
- “The SAS Log” in *SAS Language Reference: Concepts*
- “Using SAS System Options to Route Output” on page 102

LPTYPE System Option

Specifies which UNIX command and options settings will be used to route files to the printer.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Log and procedure output: Procedure output

PROC OPTIONS GROUP= LISTCONTROL

UNIX specifics: all

Syntax

-LPTYPE BSD | SYSV

LPTYPE=BSD | SYSV

Details

The LPTYPE option determines whether SAS is to use the **lpr** or the **lp** UNIX command to print files.

The LPTYPE option has two forms:

-LPTYPE BSD

causes SAS to use the **lpr** command to send files to the printer. The **lpr** command is usually supported on UNIX operating systems developed at the University of California, Berkeley, such as HP-UX.

-LPTYPE SYSV

causes SAS to use the **lp** command to send files to the printer. The **lp** command is usually supported on operating systems derived from UNIX System V, such as Solaris.

If you do not know whether to specify BSD or SYSV, check with your system administrator.

By default, SAS uses the **lpr** command if your operating system is derived from Berkeley’s version; otherwise, it uses the **lp** command.

See Also

- “PRINTCMD System Option” on page 393

MAPS System Option

Specifies the name of the SAS library containing the SAS/GRAPH map data sets.

Default: !SASROOT/maps (set in the installed !SASROOT/sasv9.cfg file)

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Graphics: Driver settings

PROC OPTIONS GROUP= GRAPHICS

UNIX specifics: default value and *location-of-maps*

See: MAPS= System Option in *SAS Language Reference: Dictionary*

Syntax

-MAPS *location-of-maps*

MAPS=*location-of-maps*

location-of-maps

specifies a libref, a valid UNIX pathname, or an environment variable associated with a pathname. Do not use a specific filename.

Details

You can reassign the MAPS libref, but you cannot clear it.

Map files might have to be uncompressed before they are used. Use the CONTENTS statement in the DATASETS procedure to determine whether they are compressed.

See Also

- INSERT System Option in *SAS Language Reference: Dictionary*
- APPEND System Option in *SAS Language Reference: Dictionary*

MAXMEMQUERY System Option

Specifies the maximum amount of memory that is allocated per request for certain procedures.

Default: 256M

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: System administration: Memory

PROC OPTIONS GROUP= MEMORY

UNIX specifics: all

Syntax

-MAXMEMQUERY *n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

MAXMEMQUERY=*n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

n* | *nK* | *nM* | *nG

specifies the limit in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); or 1,073,741,824 (gigabytes). You can specify decimal values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **8** specifies 8 bytes, a value of **.782k** specifies 801 bytes, and a value of **3m** specifies 3,145,728 bytes.

hexX

specifies the amount of memory as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, **2dx** sets the amount of memory to 45 bytes.

MIN

specifies 0 bytes, which indicates that there is no limit on the total amount of memory that can be allocated per request by each SAS procedure. These memory allocations are limited by the value of MEMSIZE.

MAX

specifies a limit of 9,007,199,254,740,992 bytes on 64-bit machines.

Details

Some SAS procedures use the MAXMEMQUERY option to specify the largest block of virtual memory that a procedure can request at one time. By contrast, the MEMSIZE option places a limit on the total amount of virtual memory that SAS dynamically allocates at any time. This virtual memory is supported by a combination of real memory and paging space. The operating environment begins paging when the amount of virtual memory that is required exceeds the real memory that is available. To prevent paging and the associated performance problems, the MAXMEMQUERY and MEMSIZE system options should be set to a subset of real memory.

MEMSIZE System Option

Specifies the limit on the total amount of virtual memory that can be used by a SAS session.

Default: value set in SAS configuration file **!SASROOT/sasv9.cfg**

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: System administration: Memory

PROC OPTIONS GROUP= MEMORY, PERFORMANCE

UNIX specifics: all

Syntax

-MEMSIZE *n* | *nK* | *nM* | *nG* | *nT* | *hexX* | MAX

n* | *nK* | *nM* | *nG* | *nT

specifies the limit in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); 1,073,741,824 (gigabytes); or 1,099,511,627,776 (terabytes). You can specify decimal

values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **16.5M** specifies 17,301,504 bytes, and a value of **.25G** specifies 268,435,456 bytes.

hexX

specifies the amount of memory as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, **0F00000x** sets the value of the MEMSIZE option to 15,728,640 bytes. A value of **0x** is equivalent to using the MAX value.

MAX

specifies to set the memory size to the largest reasonable value depending on the amount of physical memory and paging space that is available at the time that SAS is started.

Details

The Basics The MEMSIZE system option specifies the total amount of memory that is available to each SAS session, and places an enforced limit on the amount of virtual memory that SAS can dynamically allocate at any one time. Too low a value will result in out-of-memory conditions. By contrast, the REALMEMSIZE and MAXMEMQUERY system options, the SORTSIZE= option in the SORT procedure, and the SUMSIZE= option in the SUMMARY procedure all provide for procedure tuning.

If you specify an unreasonably small numeric value (for example, 6K) for MEMSIZE, the setting of this option automatically increases to a minimum reasonable value that will enable SAS to start. If you specify a numeric value in excess of 4,294,967,295 on a 32-bit version of SAS, then the setting of MEMSIZE will be reduced to 4,294,967,295.

Numeric values in excess of 9,223,372,036,854,775,807 bytes will be rejected as invalid and will prevent SAS from starting.

SAS does not automatically reserve or allocate the amount of memory that you specify in the MEMSIZE system option. SAS will use only as much memory as it needs to complete a process. For example, a DATA step might require only 20M of memory, so even though MEMSIZE is set to 500M, SAS will use only 20M of memory. While your SAS jobs are running, you can monitor the effect of larger memory settings by using system monitoring tools, such as VMSTAT and **top**.

Setting MEMSIZE to MAX

Setting MEMSIZE to MAX is reasonable only if no processes that consume large amounts of memory are likely to become active after SAS has started. For example, if multiple instances of SAS are running concurrently, and all of the sessions were started with a MEMSIZE value of MAX, then one or more of these sessions can encounter out-of-memory conditions, or the operating system can run out of available paging space. MEMSIZE=MAX calculates a value that would help prevent the system from paging if all of the memory were allocated.

If you set MEMSIZE to the maximum amount of memory that is reasonably attainable, some procedures scale themselves to the available memory.

Note: Setting MEMSIZE to MAX is the same as setting MEMSIZE to 0. Δ

Setting MEMSIZE to 0 The optimal setting for this option depends on the other applications that are running and the system resources available at your site. The amount of memory available to SAS processes can also be limited by your system administrator.

To determine the optimal setting of MEMSIZE, run the SAS procedure or DATA step with MEMSIZE=0 with the FULLTIMER option. Note the amount of memory that is used by the process, and then set MEMSIZE to a larger amount.

Note: Setting MEMSIZE to 0 is the same as setting MEMSIZE to MAX. Δ

See Also

- “REALMEMSIZE System Option” on page 394
- “SORT Procedure” on page 305

MSG System Option

Specifies the library that contains the SAS error messages.

Alias: SASMSG

Default: !SASROOT/sasmsg (set in the installed !SASROOT/sasv9.cfg file)

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

-MSG *pathname*

-MSG (*'pathname' 'pathname' ...*)

pathname

must resolve to a valid UNIX pathname. You can use an environment variable that resolves to a valid pathname.

Details

The MSG system option specifies the library that contains the SAS error messages. This option is set during the installation process and is not normally changed after installation.

To add additional pathnames, use the INSERT or APPEND system options. For more information, see the INSERT System Option and APPEND System Option in *SAS Language Reference: Dictionary*.

See Also

- INSERT System Option in *SAS Language Reference: Dictionary*
- APPEND System Option in *SAS Language Reference: Dictionary*

MSGCASE System Option

Specifies whether notes, warnings, and error messages that are generated by SAS are displayed in uppercase characters.

Default: NOMSGCASE

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Log and procedure output: SAS log

PROC OPTIONS GROUP= LOGCONTROL

UNIX specifics: all

Syntax

-MSGCASE | -NOMSGCASE

-MSGCASE

displays notes, warnings, and error messages in uppercase characters.

-NOMSGCASE

displays notes, warnings, and error messages in uppercase and lowercase characters.

Details

The MSGCASE system option specifies whether notes, warnings, and error messages that are generated by SAS are displayed in uppercase characters. User-generated messages and source lines are not affected by the MSGCASE system option.

MSYMTABMAX System Option

Specifies the maximum amount of memory available to the macro variable symbol tables.

Default: 4M (set in the installed `!SASROOT/sasv9.cfg` file)

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Macro: SAS macro

PROC OPTIONS GROUP= MACRO

UNIX specifics: default value

See: MSYMTABMAX= System Option in *SAS Macro Language: Reference*

Syntax

-MSYMTABMAX *n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

MSYMTABMAX=*n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

n* | *nK* | *nM* | *nG

specifies the maximum amount of memory that is available in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); or 1,073,741,824 (gigabytes). You can specify decimal values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **8** specifies 8 bytes, a value of **.782k** specifies 801 bytes, and a value of **3m** specifies 3,145,728 bytes.

hexX

specifies the maximum amount of memory that is available as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, **2dx** sets the maximum amount of memory to 45 bytes.

MIN

sets the amount of memory that is available to the minimum setting, which is 0 bytes. Setting the amount of memory to the minimum setting causes all macro symbol tables to be written to disk.

MAX

sets the amount of memory that is available to the maximum setting. On 64-bit computers, this value is 9,007,199,254,740,992 bytes.

MVARSIZE System Option

Specifies the maximum size for in-memory macro variables.

Default: 32K (set in the installed **!SASROOT/sasv9.cfg** file)

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Macro: SAS macro

PROC OPTIONS GROUP= MACRO

UNIX specifics: default value

See: MVARSIZE= System Option in *SAS Macro Language: Reference*

Syntax

-MVARSIZE *n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

MVARSIZE=*n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

n* | *nK* | *nM* | *nG

specifies the maximum macro variable size in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); or 1,073,741,824 (gigabytes). You can specify decimal values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **8** specifies 8 bytes, a value of **.782k** specifies 801 bytes, and a value of **3m** specifies 3,145,728 bytes.

hexX

specifies the maximum macro variable size as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, **2dx** sets the maximum macro variable size to 45 bytes.

MIN

sets the macro variable size to the minimum setting, which is 0 bytes. Setting the macro variable size to the minimum setting causes all macro variable values to be written to disk.

MAX

sets the macro variable size to the maximum setting, which is 65,534 bytes.

NEWS System Option

Specifies a file that contains messages to be written to the SAS log.

Default: `!SASROOT/misc/base/news` (set in the installed `!SASROOT/sasv9.cfg` file)

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES, LOGCONTROL

UNIX specifics: -NONEWS option

See: NEWS= System Option in *SAS Language Reference: Dictionary*

Syntax

`-NEWS file-specification | -NONEWS`

-NEWS file-specification

specifies an external file. This file contains the messages for the SAS log.

-NONEWS

specifies that the contents of the NEWS file are not displayed in the SAS log, even if the file exists. This option causes any previous NEWS specifications to be ignored.

Details

The contents of the NEWS file are displayed in the SAS log immediately after the SAS header.

See Also

- “The SAS Log” in *SAS Language Reference: Concepts*

OBS System Option

Specifies the last observation that SAS processes in a data set.

Default: MAX

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Files: SAS Files

PROC OPTIONS GROUP= SASFILES

UNIX specifics: default value

See: OBS= System Option in *SAS Language Reference: Dictionary*

Syntax

-OBS *n* | *nK* | *nM* | *nG* | *nT* | *hexX* | MIN | MAX

OBS=*n* | *nK* | *nM* | *nG* | *nT* | *hexX* | MIN | MAX

n | *nK* | *nM* | *nG* | *nT*

specifies a number to indicate when to stop processing. Using one of the letter notations results in multiplying the integer by a specific value. That is, specifying K (kilo) multiplies the integer by 1,024, M (mega) multiplies by 1,048,576, G (giga) multiplies by 1,073,741,824, or T (tera) multiplies by 1,099,511,627,776. You can specify a decimal value for *n* when it is used to specify a K, M, G, or T value. For example, a value of **20** specifies 20 observations or records, a value of **.782k** specifies 801 observations or records, and a value of **3m** specifies 3,145,728 observations or records.

hexX

specifies a number to indicate when to stop processing as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, the hexadecimal value F8 must be specified as **0F8x** in order to specify the decimal equivalent of 248. For example, the value **2dx** specifies the decimal equivalent of 45.

MIN

sets the number to 0 to indicate when to stop processing.

If OBS=0 and the NOREPLACE option is in effect, SAS might still be able to take certain actions. See *SAS Language Reference: Dictionary* for more information.

MAX

sets the number 9,223,372,036,854,775,807 to indicate when to stop processing.

OPLIST System Option

Specifies whether the settings of the SAS system options are written to the SAS log.

Default: NOOPLIST

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log

PROC OPTIONS GROUP= LOGCONTROL

UNIX specifics: all

Syntax

-OPLIST | -NOOPLIST

Details

The OPLIST system option echoes only the system options specified on the command line; it does not echo any system options specified in the configuration file or in the SASV9_OPTIONS environment variable. (If you want to echo the contents of the configuration file, use the VERBOSE option.) For example, invoke SAS with the following command:

```
sas -nodms -fullstimer -nonews -oplist
```

SAS writes this line to the SAS log:

```
NOTE: SAS command line: -nodms -fullstimer -nonews -oplist
```

See Also

- “VERBOSE System Option” on page 414

PAGESIZE System Option

Specifies the number of lines that compose a page of SAS output.

Default: number of lines on your display for interactive modes; 60 for batch mode

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log and procedure output

PROC OPTIONS GROUP= LOG_LISTCONTROL, LOGCONTROL

UNIX specifics: default values and range

See: PAGESIZE= System Option in *SAS Language Reference: Dictionary*

Syntax

-PAGESIZE *n* | *nK* | *hexX* | MIN | MAX

PAGESIZE=*n* | *nK* | *hexX* | MIN | MAX

n | *nK*

specifies the number of lines that compose a page in multiples of 1 (*n*) or 1,024 (*nK*). You can specify decimal values for the number of kilobytes. For example, a value of **800** specifies 800 lines, a value of **.782k** specifies 801 lines, and a value of **3k** specifies 3,072 lines.

hexX

specifies the number of lines that compose a page as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, the value **2dx** specifies 45 lines.

MIN

sets the number of lines that compose a page to the minimum setting, which is 15.

MAX

sets the number of lines that compose a page to the maximum setting, which is 32,767.

Details

The default for interactive modes is the number of lines on your display. For batch mode, the default is 60.

See Also

- “Controlling the Content and Appearance of Output in UNIX Environments” on page 104
- “The SAS Log” in *SAS Language Reference: Concepts*.

PATH System Option

Specifies one or more search paths for SAS executable files.

Default: `!SASROOT/sasexe` (set in the installed `!SASROOT/sasv9.cfg` file)

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

`-PATH directory-specification`

Details

The PATH system option identifies the search paths for SAS executable files. You can specify multiple PATH options to define the search order. The paths are searched in the order in which SAS encounters them; therefore, specify at the beginning of the list the paths for the products that you run most frequently. See “How SAS Processes System Options Set in Multiple Places” on page 20 for information about how that order is determined when you specify the PATH system option more than once.

PRIMARYPROVIDERDOMAIN= System Option

Specifies the domain name of the primary authentication provider.

Valid in: configuration file, SAS invocation

Alias: PRIMPD=

Category: Environment control: Initialization and operation

PROC OPTIONS GROUP: EXECMODES

See: PRIMARYPROVIDERDOMAIN= System Option in *SAS Language Reference: Dictionary*

PRINT System Option

Specifies a destination for SAS output when running in batch mode.

Default: the SAS output from a batch SAS program is written to a file in the current directory with the same filename as the SAS source file, with an extension of .lst.

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

-PRINT *file-specification* | -NOPRINT

-PRINT *file-specification*

specifies the location for the SAS procedure output file. The *file-specification* can be any valid UNIX path to a directory, a filename, or an environment variable that is associated with a path. If you specify only the path to a directory, the procedure output file is created in the specified directory. The default name for this file is *filename.lst*, where *filename* is the name of your SAS job.

-NOPRINT

suppresses the creation of the SAS procedure output file.

Details

The PRINT system option specifies a destination for SAS output when running in batch mode. The PRINT system option is valid in batch mode; it is ignored in interactive modes.

See Also

- “Using SAS System Options to Route Output” on page 102

PRINTCMD System Option

Specifies the print command SAS is to use.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Log and procedure output control: Procedure output

PROC OPTIONS GROUP= LISTCONTROL

UNIX specifics: all

Syntax

-PRINTCMD "*print-command*"

PRINTCMD="*print-command*"

Details

The syntax of the options passed to the print command is controlled by the LPTYPE system option. If LPTYPE is set to BSD, the command uses **lpr** command options; if LPTYPE is set to SYSV, the command uses **lp** command options.

If your site uses a print command (spooler) other than **lp** or **lpr**, *print-command* specifies its name. The PRINTCMD option overrides the LPTYPE setting.

When specified in an OPTIONS statement, the PRINTCMD option will not change the print commands assigned to previously defined filenames. For example, consider the following code:

```
filename pc1 printer;
proc printto print=pc1;
run;
proc print data=sales.week;
run;

options printcmd="netlp";

filename pc2 printer;
proc printto print=pc2;
run;
proc print data=sales.month;
run;
```

Output associated with PC2 will use the **netlp** command; output associated with PC1 will use the default print command.

See Also

- Chapter 4, "Printing and Routing Output," on page 91
- "LPTYPE System Option" on page 382

REALMEMSIZE System Option

Specifies the amount of real (physical) memory SAS can expect to allocate.

Default: 0

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: System administration: Memory

PROC OPTIONS GROUP= MEMORY**UNIX specifics:** valid values**Syntax****-REALMEMSIZE** *n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX***n* | *nK* | *nM* | *nG***

specifies the amount of memory to reserve in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes), or 1,073,741,824 (gigabytes). The value of *n* can be a decimal value. For example, a value of **8** specifies 8 bytes, a value of **.782k** specifies 801 bytes, and a value of **3m** specifies 3,145,728 bytes.

hexX

specifies the amount of memory as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, the value **2dx** sets the amount of memory to 45 bytes.

MIN

specifies a value of 0, which indicates that the memory usage is determined by SAS when SAS starts.

MAX

specifies to set the memory size to the largest permissible value. This value depends on the system limit.

Details

The REALMEMSIZE system option sets a recommended upper limit on working memory for procedures that can use both memory and utility disk space, such as PROC SUMMARY and PROC SORT, so that they can avoid virtual memory thrashing.

The REALMEMSIZE option should never be set beyond the amount of real memory. If the amount of real memory is insufficient for a job to run, then setting the MEMSIZE option beyond the real memory might enable the job to run using a combination of real and virtual memory.

Comparisons

Some SAS procedures use the REALMEMSIZE system option to specify how much real memory the procedure can allocate and use without inducing excessive page swapping. By contrast, the MEMSIZE system option places a limit on the total amount of virtual memory that SAS dynamically allocates at any time. This virtual memory is supported by a combination of real memory and paging space.

The operating environment begins paging when the amount of virtual memory that is required exceeds the real memory that is available. To prevent paging and the associated performance problems, the REALMEMSIZE and MEMSIZE system options should be set to a subset of real memory.

See Also

- “MEMSIZE System Option” on page 384

- “SORT Procedure” on page 305

RSASUSER System Option

Controls whether members of the Sasuser library can be opened for update or for read-only access.

Default: NORSASUSER

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: network considerations

See: RSASUSER System Option in *SAS Language Reference: Dictionary*

Syntax

-RSASUSER | -NORSASUSER

-RSASUSER

limits access to the Sasuser library to read-only access. (If the Sasuser library is being shared by multiple users or the same user is running SAS multiple times simultaneously, the Sasuser library is often shared.) By default, if one user has a member of the Sasuser library open for update, all other users are denied access to that SAS library member. For example, if one user is writing to the Sasuser.Profile catalog, no other user can even read data from the Profile catalog.

Specifying RSASUSER enables a group of users to share Sasuser library members by allowing all users read-only access to members. In the Profile catalog example, if RSASUSER is in effect, all users can open the Profile catalog for read-only access, allowing other users to concurrently read from the Profile catalog. However, no user can write information out to the Profile catalog; you receive an error message if you try to do so.

Specifying RSASUSER from the command line affects only that session’s access to files. To enable a group of users to share members in the Sasuser library, the system manager should set RSASUSER in a common SAS configuration file, which is shared by all users who will be sharing the Sasuser library.

If you specify RSASUSER but no Profile catalog exists in the Sasuser library, the Profile catalog is created in the Work library.

Note: While the RSASUSER option is extremely useful for sharing information (such as the Profile catalog) stored in the Sasuser library, it is less practical when used in conjunction with SAS/ASSIST software or other SAS modules that require Update access to the Sasuser library. Δ

-NORSASUSER

prevents users from sharing members of the Sasuser library because it allows a user to open a file in the Sasuser library for Update access. Update access requires exclusive rights to the library member.

See Also

- “Sharing SAS Files in a UNIX Environment” on page 38

RTRACE System Option

Produces a list of resources that are read or loaded during a SAS session.

Default: NONE

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log

PROC OPTIONS GROUP= LOGCONTROL

UNIX specifics: all

Syntax

-RTRACE ALL | NONE

ALL

produces a list of resources that are read or loaded during a SAS session.

NONE

turns off RTRACE on all files.

Details

The RTRACE system option produces a list of resources that are read or loaded during the execution of SAS. If you specify -RTRACE ALL but do not specify the RTRACELOC system option, the output is written to the SAS log.

See Also

- “RTRACELOC System Option” on page 397

RTRACELOC System Option

Specifies the pathname of the file to which the list of resources that are read or loaded during a SAS session is written.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

-RTRACELOC *pathname*

RTRACELOC=*pathname*

pathname

specifies the file to which RTRACE information is written. The *pathname* must include the path and the filename for the RTRACE output.

Details

The RTRACELOC system option specifies the pathname of the file to which RTRACE information is written. If the *pathname* does not include a filename, the output will be directed to standard output. If you specify -RTRACE ALL, but do not specify the RTRACELOC system option, the output is written to the SAS log.

See Also

- “RTRACE System Option” on page 397

SASAUTOS System Option

Specifies the autocall library.

Default: SASAUTOS fileref

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Environment control: Files
Macro: SAS macro

PROC OPTIONS GROUP= ENVFILES
MACRO

UNIX specifics: syntax for specifying multiple *directory-specifications*

See: SASAUTOS= System Option in *SAS Macro Language: Reference*

Syntax

-SASAUTOS '*directory-specification*' | *fileref*

-SASAUTOS ('*directory-specification1*' | *fileref1*,...,'*directory-specification-n*' | *filerefn*)

-NOSASAUTOS

SASAUTOS=*directory-specification*' | *fileref*

SASAUTOS =('*directory-specification1*' | *fileref1*,..., '*directory-specification-n*' | *filerefn*)

NOSASAUTOS

directory-specification

specifies a pathname to an autocall macro library.

fileref

specifies a name (shorthand reference) that has been assigned to an autocall macro library.

Note: The SASAUTOS option uses filerefs, not librefs. △

Details

Each autocall macro library consists of files in a UNIX directory. The *directory-specification* can be the pathname of a UNIX directory, a fileref, or an environment variable.

If you specify the pathname of a directory, you must enclose the name in quotation marks. You can omit the quotation marks only if you are specifying the option in the configuration file, in the SAS command, or in the SASV9_OPTIONS environment variable, and if the name cannot be taken to be a fileref.

If you specify a fileref, you must define it before attempting to use any of the autocall macros. You can define the fileref in a FILENAME statement, in an environment variable, or with the FILENAME function (see “Assigning Filerefs to External Files or Devices with the FILENAME Statement” on page 73).

How you specify multiple directory names, filerefs, or environment variables depends on where you specify the SASAUTOS option:

- If you specify the SASAUTOS option in the configuration file or in the SASV9_OPTIONS environment variable, use either multiple SASAUTOS options, or enclose the directory names in parentheses. Separate the names with a comma or a blank space.
- If you specify the SASAUTOS option in the SAS command, use the APPEND or INSERT system options to append to the end or insert at the beginning of the current SASAUTOS value. For example, the following code adds `/users/userid/also` to the end of the current SASAUTOS value, `/users/userid/here`:

```
sas -sasautos /users/userid/here -append sasautos /users/userid/also
```

For more information, see the and in *SAS Language Reference: Dictionary*.

- If you specify the SASAUTOS option in the OPTIONS statement or in the SAS System Options window, you must enclose the directory names in parentheses. Separate the names with a comma or a blank space.

At configuration time, SAS concatenates all directories specified for SASAUTOS. However, after the session starts, any new directories you specify override any current autocall libraries.

The NOSASAUTOS option causes SAS to ignore all previous SASAUTOS specifications (whether specified in the SAS command, in the configuration file, or in the SASV9_OPTIONS environment variable).

The default value of the SASAUTOS option is the SASAUTOS fileref. There is no UNIX directory assigned to the fileref, so you must define the SASAUTOS fileref if you want to use it as your autocall library.

Example: Specifying Multiple Environment Variables in the OPTIONS Statement

The following example shows the syntax to use if you are specifying multiple environment variables in the OPTIONS statement:

```
options sasautos=(AUTODIR, SASAUTOS);
```

The environment variables that you specify must be defined. For example, you could define the AUTODIR environment variable at SAS invocation by using the following code:

```
-set AUTODIR /tmp/sasautos
```

For more information about how to define an environment variable, see “SET System Option” on page 402.

Example: Specifying a Fileref in the OPTIONS Statement

The fileref you specify must be defined. For example, you could define the AUTODIR fileref using a FILENAME statement:

```
filename AUTODIR '/tmp/sasautos';
```

Once the fileref is defined, you can use it in an OPTIONS statement to set the autocall library.

```
options sasautos=AUTODIR;
```

See Also

- in *SAS Language Reference: Dictionary*
- in *SAS Language Reference: Dictionary*
- MAUTOSOURCE System Option in *SAS Macro Language: Reference*
- MRECALL System Option in *SAS Macro Language: Reference*
- *SAS Macro Language: Reference*

SASHELP System Option

Specifies the locations of Sashelp libraries.

Default: `!SASROOT/sashelp` (set in the installed `!SASROOT/sasv9.cfg` file)

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: *directory-specification* can also be an environment variable

See: SASHELP= System Option in *SAS Language Reference: Dictionary*

Syntax

```
-SASHELP directory-specification
```

```
-SASHELP ('directory-specification', 'directory-specification'...)
```

Details

This option is set in the installation process and is not normally changed after installation. An environment variable can be specified as the value of SASHELP.

To add additional directory specifications, use the INSERT or APPEND system option. For more information, see the and in *SAS Language Reference: Dictionary*.

See Also

- in *SAS Language Reference: Dictionary*
- in *SAS Language Reference: Dictionary*

SASSCRIPT System Option

Specifies one or more storage locations of SAS/CONNECT script files.

Default: `!SASROOT/misc/connect`

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Communications: Networking and encryption

PROC OPTIONS GROUP= COMMUNICATIONS

UNIX specifics: syntax for specifying multiple directory names

Syntax

`-SASSCRIPT 'dir-name' | ('dir-name-1',..., 'dir-name-n')`

`SASSCRIPT='dir-name' | ('dir-name-1',..., 'dir-name-n')`

Details

How you specify multiple directory names in the same SASSCRIPT option depends on where you specify the SASSCRIPT option:

- If you specify the option in the configuration file or in the SASV9_OPTIONS environment variable, use either multiple SASSCRIPT options, or enclose the directory names in parentheses. Separate the names with a comma or a blank space.
- If you specify the option in the SAS command, use multiple SASSCRIPT options because parentheses cause syntax errors.
- If you specify the option in the OPTIONS statement or in the SAS System Options window, you must enclose the directory names in parentheses. Separate the names with a comma or a blank space.

See Also

- “SASSCRIPT System Option” in *SAS/CONNECT User’s Guide*

SASUSER System Option

Specifies the name of the Sasuser library.

Default: `~/sasuser.v92` (set in the installed `!SASROOT/sasv9.cfg` file)

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: *pathname* can be an environment variable

See: SASUSER= System Option in *SAS Language Reference: Dictionary*

Syntax

`-SASUSER pathname`

Details

The *pathname* identifies the directory for the Sasuser library that contains a user's Profile catalog. You can use an environment variable to specify the pathname, for example:

```
sas -sasuser $HOME
```

SET System Option

Defines an environment variable.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS Statement, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

`-SET variable-name value`

`SET=variable-name value`

Details

The SET option lets you define an environment variable that is valid within the SAS session and any shell started from within the SAS session. Using the SET option is similar to using the SAS `setenv` command. See “Executing Operating System Commands from Your SAS Session” on page 14 for information .

A special use for the SET option is to specify the name of the `!SASROOT` directory:

```
-set SASROOT pathname
```


The pathname specified can then be used to expand **!SASROOT** (as shown in Table 2.6 on page 51).

After exiting your SAS session, environment variables that are set with the SET option no longer exist.

See Also

- \square Appendix 1, “The !SASROOT Directory,” on page 421
- \square “Defining Environment Variables in UNIX Environments” on page 23

SORTANOM System Option

Specifies certain options for the host sort utility.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Sort: Procedure options

PROC OPTIONS GROUP= SORT

UNIX specifics: all

Syntax

SORTANOM=*option(s)*

–SORTANOM *option(s)*

options

can be any one or more of the following:

B

tells SyncSort to run in multi-call mode, instead of single-call mode. (Refer to the documentation for SyncSort for more information.)

Note: This option is available for SyncSort only. Δ

T

writes to the SAS log statistics about the external sorting process.

V

writes to the SAS log all of the commands that are passed to the host sort utility.

SORTCUT System Option

Specifies the number of observations that SAS sorts; if the number of observations in the data set is greater than the specified number, the host sort program sorts the remaining observations.

Default: 0

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Sort: Procedure options

PROC OPTIONS GROUP= SORT

UNIX specifics: all

Syntax

-SORTCUT *n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

SORTCUT=*n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

n* | *nK* | *nM* | *nG

specifies the number of observations in multiples of 1 (*n*); 1,024 (*nK*); 1,048,576 (*nM*); or 1,073,741,824 (*nG*). You can specify decimal values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **800** specifies 800 observations, a value of **.782k** specifies 801 observations, and a value of **3m** specifies 3,145,728 observations.

hexX

specifies the number of observations as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, the value **2ffx** specifies 767 observations.

MIN

specifies 0 observations.

MAX

specifies 9,007,199,254,740,992 observations.

Details

When you specify SORTPGM=BEST, SAS uses the value of the SORTCUT and SORTCUTP options to determine whether to use the host sort or the SAS sort. If the number of observations in the data set is greater than the number that you specify with SORTCUT, the host sort will be used. If both SORTCUT and SORTCUTP are either not defined or are set to 0, the SAS sort is used. If you specify both options and either condition is true, SAS chooses the host sort.

See Also

- “SORTCUTP System Option” on page 404
- “SORTPGM System Option” on page 407

SORTCUTP System Option

Specifies the number of bytes that SAS sorts; if the number of bytes in the data set is greater than the specified number, the host sort program sorts the remaining data set.

Default: 0

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Sort: Procedure options

PROC OPTIONS GROUP= SORT

UNIX specifics: all

Syntax

-SORTCUTP *n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

SORTCUTP=*n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

n* | *nK* | *nM* | *nG

specifies the number of bytes in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); or 1,073,741,824 (gigabytes). You can specify decimal values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **8** specifies 8 bytes, a value of **.782k** specifies 801 bytes, and a value of **3m** specifies 3,145,728 bytes.

hexX

specifies the number of bytes as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, the value **2dx** specifies 45 bytes.

MIN

specifies 0 bytes.

MAX

specifies 9,007,199,254,740,992 bytes.

Details

When you specify SORTPGM=BEST, SAS uses the value of the SORTCUT and SORTCUTP options to determine whether to use the host sort or the SAS sort. If the data set to be sorted is larger than the number of bytes (or kilobytes or megabytes) that you specify with SORTCUTP, the host sort will be used instead of the SAS sort. The value you specify must be less than or equal to 2,147,483,647 bytes. If both SORTCUT and SORTCUTP are either not defined or are set to 0, the SAS sort is used. If you specify both options and either condition is true, SAS chooses the host sort.

The following equation computes the number of bytes to be sorted:

$$\begin{aligned} \text{number-of-bytes} &= ((\text{length-of-obs})+(\text{length-of-all-keys})) \\ & * \text{number-of-obs} \end{aligned}$$

See Also

- “SORTANOM System Option” on page 403
- “SORTCUT System Option” on page 403
- “SORTPGM System Option” on page 407

SORTDEV System Option

Specifies the pathname used for temporary files created by the host sort utility.

Default: same location as -WORK, which is set in the installed **!SASROOT/sasv9.cfg** file

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Sort: Procedure options

PROC OPTIONS GROUP= SORT

UNIX specifics: all

Syntax

SORTDEV=*'directory-specification'*

-SORTDEV *directory-specification*

Details

The SORTDEV option specifies an alternative directory for temporary files created by the host sort program.

SORTNAME System Option

Specifies the name of the host sort utility.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Sort: Procedure options

PROC OPTIONS GROUP= SORT

UNIX specifics: all

Syntax

SORTNAME=*'host-sort-utility-name'*

-SORTNAME *host-sort-utility-name*

Details

The SORTNAME option specifies the name of the default host sort utility, **syncsort**.

See Also

- “SORTPGM System Option” on page 407

SORTPARM System Option

Specifies parameters for the host sort utility.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Sort: Procedure options

PROC OPTIONS GROUP= SORT

UNIX specifics: all

Syntax

SORTPARM=*parameters*'

-SORTPARM *parameters*'

The *parameters* are any parameters that you want to pass to the sort utility. For a description of these parameters, see the documentation for the sort that you are using.

SORTPGM System Option

Specifies whether SAS sorts using the SAS sort utility or the host sort utility.

Default: BEST

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Sort: Procedure options

PROC OPTIONS GROUP= SORT

UNIX specifics: all

Syntax

-SORTPGM SAS | HOST | BEST

SORTPGM=SAS | HOST | BEST

SAS

tells SAS to use the SAS sort.

HOST

tells SAS to use the sort that is specified by the SORTNAME system option.

BEST

tells SAS to determine the best routine to sort the data set: the SAS sort or the host sort that is specified by the SORTNAME system option. The settings of the

SORTCUT and SORTCUTP system options determine whether SAS chooses the SAS sort or the host sort.

Details

The SORTPGM system option tells SAS whether to use the SAS sort, to use the host sort, or to determine which sort is best for the data set.

See Also

- “SORTCUTP System Option” on page 404
- “SORTNAME System Option” on page 406
- “SORTSIZE System Option” on page 408

SORTSIZE System Option

Specifies the amount of memory available to the SORT procedure.

Default: depends on your operating environment

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Sort: Procedure options

System administration: Memory

PROC OPTIONS GROUP= SORT

MEMORY

UNIX specifics: value of MAX

See: SORTSIZE= System Option in *SAS Language Reference: Dictionary*

Syntax

–SORTSIZE *n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

SORTSIZE=*n* | *nK* | *nM* | *nG* | *hexX* | MIN | MAX

n* | *nK* | *nM* | *nG

specifies the number of bytes in multiples of 1 (bytes); 1,024 (kilobytes); 1,048,576 (megabytes); or 1,073,741,824 (gigabytes). You can specify decimal values for the number of kilobytes, megabytes, or gigabytes. For example, a value of **8** specifies 8 bytes, a value of **.782k** specifies 801 bytes, and a value of **3m** specifies 3,145,728 bytes.

hexX

specifies the amount of memory as a hexadecimal value. You must specify the value beginning with a number (0–9), followed by hexadecimal characters (0–9, A–F), and then followed by an X. For example, the value **2dx** sets the amount of memory to 45 bytes.

MIN

specifies 0 bytes, which indicates that there is no limit except the limitation specified by the MEMSIZE system option.

MAX

specifies the maximum addressable memory for the operating environment.

Details

The SORT procedure uses the SORTSIZE system option to specify how much virtual memory the procedure can allocate and use without inducing excessive page swapping. The amount of memory that SAS uses for the SORT procedure also depends on the values of the MEMSIZE and REALMEMSIZE system options. By contrast with the SORTSIZE option, the MEMSIZE system option places a limit on the total amount of virtual memory that SAS dynamically allocates at any time. This virtual memory is supported by a combination of real memory and paging space. The operating environment begins paging when the amount of virtual memory that is required exceeds the real memory that is available. To prevent paging and the associated performance problems, the SORTSIZE and the MEMSIZE system options should be set to a subset of real memory. In most cases, you can set SORTSIZE=MAX because this value will limit the amount of memory that is used by the SORT procedure.

See Also

- “REALMEMSIZE System Option” on page 394
- “MEMSIZE System Option” on page 384
- “SORT Procedure” on page 305

STDIO System Option

Specifies whether SAS should use stdin, stdout, and stderr.

Default: NOSTDIO

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Input control: Data processing

PROC OPTIONS GROUP= INPUTCONTROL

UNIX specifics: all

Syntax

–STDIO | –NOSTDIO

Details

This option tells SAS to take its input from standard input (stdin), to write its log to standard error (stderr), and to write its output to standard output (stdout).

This option is designed for running SAS in batch mode or from a shell script. If you specify this option interactively, SAS starts a line mode session. The STDIO option overrides the DMS, DMSEXP, and EXPLORER system options.

The STDIO option does not affect the assignment of the Stdio, Stdin, and Stderr filerefs. See “Filerefs Assigned by SAS in UNIX Environments” on page 77 for more information.

For example, in the following SAS command, the file **myinput** is used as the source program, and files **myoutput** and **mylog** are used for the procedure output and log respectively.

```
sas -stdio < myinput > myoutput 2> mylog
```

If you are using the C shell, you should use parentheses:

```
(sas -stdio < myinput > myoutput ) >& output_log
```

See Also

- Chapter 4, “Printing and Routing Output,” on page 91

STIMEFMT System Option

Specifies the format that is used to display the time on FULLSTIMER and STIMER output.

Default: M

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log

PROC OPTIONS GROUP= LOGCONTROL

UNIX specifics: all

Syntax

`-STIMEFMT value`

`STIMEFMT=value`

Details

The STIMEFMT system option controls the format of output produced by the FULLSTIMER and STIMER system options.

STIMEFMT takes the following values:

HOURS, H, or Z

prints time statistics in **hh:mm:ss.ss** format.

MINUTES or M

prints time statistics in **mm:ss.ss** format.

SECONDS or S

prints time statistics in **ss.ss** format.

NORMAL or N

prints time statistics in **ss.ss** format if the time is less than one minute, in **mm:ss.ss** format if the time is less than one hour, or in **hh:mm:ss.ss** format otherwise.

STIMER System Option

Specifies whether to write a subset of system performance statistics to the SAS log.

Default: STIMER

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log

PROC OPTIONS GROUP= LOGCONTROL

UNIX specifics: all

Syntax

–STIMER | –NOSTIMER

STIMER | NOSTIMER

STIMER

writes only real time and CPU time to the SAS log.

NOSTIMER

does not write any statistics to the SAS log.

Details

The STIMER system option specifies whether a subset of all the performance statistics of your system that are available to SAS are written to the SAS log. The following is an example of STIMER output:

Output 18.3 STIMER Output

```
real time  1.34 seconds
cpu time   0.04 seconds
```

STIMER displays the following statistics:

Table 18.3 Description of STIMER Statistics

Statistic	Description
real time	the amount of time spent to process the SAS job. Real time is also referred to as elapsed time.
CPU time	the total time spent to execute your SAS code and to perform system overhead tasks on behalf of the SAS process. This value is the combination of the user CPU and system CPU statistics from FULLSTIMER.

If both STIMER and FULLSTIMER are set, the FULLSTIMER statistics are printed.

Note: Starting in SAS 9, some procedures use multiple threads. On computers with multiple CPUs, the operating system can run more than one thread simultaneously. Consequently, CPU time might exceed real time in your STIMER output.

For example, a SAS procedure could use two threads that run on two separate CPUs simultaneously. The value of CPU time would be calculated as the following:

```
CPU1 time + CPU2 time = total CPU time
1 second + 1 second = 2 seconds
```

Because CPU1 can run a thread at the same time that CPU2 runs a separate thread, you can theoretically consume 2 CPU seconds in 1 second of real time. Δ

See Also

- “FULLSTIMER System Option” on page 372

SYSIN System Option

Specifies the default location of SAS source code when running in batch mode.

Default: none

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

`-SYSIN filename | -NOSYSIN`

-SYSIN *filename*

specifies an external file. The value for *filename* must be a valid UNIX filename.

-NOSYSIN

invokes SAS, processes the autoexec file, and then terminates SAS, returning you to the command prompt.

Details

This option applies only when you are using batch mode. It is not necessary to precede the filename with the SYSIN option if the filename immediately follows the keyword **SAS**. For example, the following two SAS commands are equivalent:

```
sas saspjms/report1.sas
sas -sysin saspjms/report1.sas
```

The syntax of the SYSIN system option also enables you to specify NOSYSIN. If you specify NOSYSIN, SAS is invoked, the autoexec file is processed, and then SAS terminates, returning you to the command prompt. The following example shows the syntax:

```
sas -nosysin -autoexec mysas.sas
```

This option is useful if you want to test an autoexec file without actually running a complete SAS session.

See Also

- “Starting SAS Sessions in UNIX Environments” on page 4

SYSPRINT System Option

Specifies the destination for printed output.

Default: default system printer

Valid in: configuration file, SAS invocation, OPTIONS statement, SASV9_OPTIONS environment variable

Category: Log and procedure output control: Procedure output

PROC OPTIONS GROUP= LISTCONTROL and ODSPRINT

UNIX specifics: all

Syntax

–SYSPRINT *destination* | '*destination option-list*'

SYSPRINT=*destination* | '*destination option-list*'

destination

is the name of a hard-copy device at your site. Consult your system administrator for a list of available destinations.

option-list

is the list of options to pass to the **lp** (or **lpr**) command.

Details

The SYSPRINT option specifies a destination for printed output other than the default system printer. You can use the option list to pass options to the **lp** (or **lpr**) command.

Note: When a fileref is assigned, the SYSPRINT option is queried. If the value of the SYSPRINT option is later changed, the fileref does not pick up this change. △

For information, see “Changing the Default Print Command in UNIX Environments” on page 103.

See Also

- “PRINTCMD System Option” on page 393
- Chapter 4, “Printing and Routing Output,” on page 91

USER System Option

Specifies the name of the default permanent SAS library.

Default: none

Valid in: configuration file, SAS invocation, OPTIONS statement, SAS System Options window, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: *pathname* must be a valid UNIX pathname

See: USER= System Option in *SAS Language Reference: Dictionary*

Syntax

-USER *pathname*

USER='pathname' | libref

pathname

identifies the directory containing your default permanent SAS library. It must be a directory name.

libref

is the libref associated with the directory containing your default permanent SAS library. It must already be assigned.

See Also

- “Using One-Level Names to Access Permanent Files (User Library)” on page 59

VERBOSE System Option

Specifies whether SAS writes the system option settings to the SAS log.

Default: NOVERBOSE

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Log and procedure output control: SAS log

PROC OPTIONS GROUP= LOGCONTROL

UNIX specifics: all

Syntax

-VERBOSE | -NOVERBOSE

-VERBOSE

writes the settings of SAS system options from the configuration file, the SAS command, and the SASV9_OPTIONS environment variable to the SAS log. For the CONFIG option, VERBOSE lists the names of the configuration files.

-NOVERBOSE

does not write the settings of the system options to the SAS log.

Details

SAS sends the system option information to standard output. If the standard output is a terminal, the list is displayed with the **more** command. You can also use the **more** command to scroll the file. The ENTER key scrolls one line; the space bar scrolls the entire display.

See Also

- “Customizing Your SAS Session by Using System Options” on page 17
- “OPLIST System Option” on page 390

WORK System Option

Specifies the location of the Work library.

Default: set in the installed `!SASROOT/sasv9.cfg` file

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

See: WORK= System Option in *SAS Language Reference: Dictionary*

Syntax

`-WORK pathname`

pathname

specifies the location of the SAS Work library if *pathname* is a directory. If *pathname* describes a file, then SAS chooses a directory from that file as the location for the Work library for the current SAS session.

Details

The Basics If the value of *pathname* is a directory, then SAS continues its initialization using the directory as the location for the Work library. If the value of *pathname* is a file, then SAS opens the file and selects one of the paths to use as the location for the Work library. SAS either selects a path at random or selects a path based on available space. Once a work location is picked at start-up, all of the work files for the session are sent to the single work directory.

Making the Allocation of Work Libraries More Dynamic

If *pathname* refers to a file, then that file contains a list of locations that can be used for the Work library. Individual SAS Work libraries still reside in a single directory. You use `METHOD=RANDOM` to specify that the Work directory is randomly chosen from the list of directories. This selection enables you to balance the I/O load across multiple hardware systems. You use `METHOD=SPACE` to specify the directory that has the most available space. If `METHOD` is not specified, SAS defaults to choosing a directory at random.

Examples

Example 1: Spreading a Processing Load across Multiple Volumes of Different Disks

The following example shows how to spread an I/O processing load across multiple volumes of different disks. In this case, you use METHOD=RANDOM. A file named /sasinfo/workfiles contains the following information:

```
/disk1/sastempfiles
/disk2/sastempfiles
/disk3/sastempfiles
method=random
```

The Work library for a particular SAS session will be placed on either disk1, disk2, or disk3. The configuration file or command line would include the following:

```
-work /sasinfo/workfiles
```

Example 2: Choosing the Directory That Has the Most Free Space When you process your data, you can choose the directory that has the most free space. In this case, you use METHOD=SPACE. In the following example, /sasinfo/workfiles contains the following directories:

```
/disk1/sastempfiles
/disk2/sastempfiles
/disk3/sastempfiles
method=space
```

The Work library is placed on the disk with the most free space.

See Also

- “WORKINIT System Option” on page 416

WORKINIT System Option

Initializes the Work library.

Default: WORKINIT

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: WORKINIT does not erase files from previous sessions

See: WORKINIT System Option in *SAS Language Reference: Dictionary*

Syntax

```
-WORKINIT | -NOWORKINIT
```

-WORKINIT

specifies that a new subdirectory is to be created in the directory specified in the WORK option.

-NOWORKINIT

specifies that the system is to use the directory specified by the WORK option.

- If the system does not find any old subdirectories, it creates a new one.
- If the system finds more than one old subdirectory, it uses the latest one.
- If file locking is in effect (see the “FILELOCKS System Option” on page 368 option), the system looks for the latest unlocked directory. If it finds none, it creates a new one.

Details

The WORKINIT option controls whether the Work library is initialized at SAS invocation.

See Also

- “FILELOCKS System Option” on page 368
- “WORK System Option” on page 415

WORKPERMS System Option

Sets the permissions of the SAS Work library when it is initially created.

Default: 700

Valid in: configuration file, SAS invocation

Category: Environment control: Files

PROC OPTIONS GROUP= ENVFILES

UNIX specifics: all

Syntax

–WORKPERMS *permission-value*

permission-value

specifies the octal value representing the permissions for the SAS Work directory. Values can be any octal value setting the permission of a UNIX directory. Examples of values include umask, 700, 755, 770, 775, and 777.

Details

The WORKPERMS system option enables you to change or remove the current file mode creation mask value when you initially create a SAS Work library. This means that you can change the value of *permission-value* to change file permissions for a new Work library.

XCMD System Option

Specifies whether the X command is valid in the SAS session.

Default: XCMD

Valid in: configuration file, SAS invocation, SASV9_OPTIONS environment variable

Category: Environment control: Display

PROC OPTIONS GROUP= ENVDISPLAY

UNIX specifics: all

Syntax

-XCMD | -NOXCMD

-XCMD

specifies that the X command is valid in the current SAS session.

-NOXCMD

specifies that the X command is not valid in the current SAS session.

Details


The XCMD system option specifies whether the X command is valid in the current SAS session.

You cannot use several SAS statements, objects, and facilities if you use the NOXCMD system option. Examples of these statements, objects, and facilities include the following:

- the PIPE device type in the FILENAME statement
- the CALL SYSTEM routine
- the %SYSEXEC macro
- any facility that SAS uses to execute a shell-level command


See Also

- “Executing Operating System Commands from Your SAS Session” on page 14
- “X Command” on page 239
- “FILENAME Statement” on page 318
- “CALL SYSTEM Routine” on page 261
- %SYSEXEC macro in “Macro Statements in UNIX Environments” on page 287



Appendixes

- Appendix 1* **The !SASROOT Directory** 421
- Appendix 2* **Tools for the System Administrator** 423
- Appendix 3* **Commands That Are Not Specific to UNIX** 429
- Appendix 4* **Recommended Reading** 475



APPENDIX

1

The !SASROOT Directory

Introduction to the !SASROOT Directory 421

Contents of the !SASROOT Directory 421

Introduction to the !SASROOT Directory

When SAS is installed, its entire directory structure is placed on a node in your file system. This node, which forms the root of SAS, is called the **!SASROOT** directory. The **!SASROOT** directory can be located anywhere in your file system. The default location is `/usr/local/SAS`.

Contents of the !SASROOT Directory

The **!SASROOT** directory contains the files required to use SAS. This directory includes invocation points, configuration files, sample programs, catalogs, data sets, and executable files. You do not need to know the organization of these directories to use SAS.


If all available SAS products are installed on your system, the **!SASROOT** directory contains the files and directories listed in the following tables:

Table A1.1 SAS Files in the !SASROOT Directory

SAS File	Description of Contents
<code>sas</code>	is the default invocation point for SAS.
<code>sassetup</code>	is the invocation point for SAS Setup, the installation program for SAS.
<code>setinit.sas</code>	is the SAS file that was used to update the license information.
<code>sasv9.cfg</code>	is the system configuration file for SAS.
<code>sasv9_local.cfg</code>	is the file that contains user-specified options. This file overrides the options in the default configuration file, and prevents the options from being lost when you reinstall or upgrade SAS.

Table A1.2 SAS Directories in the !SASROOT/utilities/bin Directory

SAS Directory	Description of Contents
bin	contains the invocation scripts for each language listed in the nl s directory. This directory also contains the SASENV script that sets the environment variables that are required by SAS.
dbcs	contains the subdirectories for a DBCS installation.
install	contains the SAS Setup program files.
maps	is a SAS library that contains SAS data sets used by SAS/GRAPH software to produce maps. You receive some maps with SAS/GRAPH software. Additional maps are available in the SAS Map Data Library Series.
misc	contains miscellaneous components, such as Java applets, fonts, and textures. This directory also contains components for various SAS products, such as script files for SAS/CONNECT and thin client interfaces for SAS/SHARE.
nl s	contains subdirectories for national language support. These directories include: en (English), ja (Japanese), ko (Korean), and zh (Simplified Chinese). Each language directory contains a sascfg subdirectory that contains the SAS data files created during installation.
samples	contains sample programs for different SAS products. These programs are organized by product subdirectory. Installing sample programs is optional. Consequently, this directory might not include samples for every SAS product, or the directory might be empty.
sasautos	contains predefined SAS macros. See “Using Autocall Libraries in UNIX Environments” on page 288.
sasexe	contains executable files for different SAS products.
sashelp	is a SAS library that contains online Help files, menus, descriptions of graphics devices, and other catalogs used by SAS procedures that support windows.
sasmsg	contains files that contain all of the messages and notes that are used by SAS.
saspgm	contains various components of SAS products.
sastest	contains files that are used by the Feature Testing Tool.
sasumgmt	obtains and transcodes or decodes the user name and password into Unicode. It then calls the SAS authorization service to authenticate the user. It then exits with an exit status that indicates the success or failure of the authentication.
utilities	contains man pages and utility programs. See “The Utilities Directory in UNIX Environments” on page 423 for more information.
x11	contains the files needed to run SAS with the X Window System. These files include bitmap files, online Help files, and resource files.



APPENDIX

2

Tools for the System Administrator

The Utilities Directory in UNIX Environments 423

Installing Manual Pages 423

The UNIX Authentication API 424

Utilities in the /utilities/bin Directory 424

The Utilities Directory in UNIX Environments

The `!SASROOT/utilities` directory contains two important subdirectories:

man	contains the online manual pages for SAS. “Installing Manual Pages” on page 423 describes how to make these pages accessible to users through the UNIX <code>man</code> command.
bin	contains the executable files for administrative tools. “Utilities in the /utilities/bin Directory” on page 424 describes some of the tools in this directory.
src/auth	contains source files and documentation for the UNIX Authentication API. The API enables administrators to add custom authentication methods to SAS authentication in UNIX environments. For more information, see “The UNIX Authentication API” on page 424.

Installing Manual Pages

To be able to read the manual pages in the `utilities/man` directory, move the files to the `man1` subdirectory of the location of the other man files for your system. This location is usually `/usr/man` or `/usr/local/man`. Execute the UNIX `man man` command to determine the appropriate pathname for your system. When you have found the correct pathname, use the following command to move the SAS man files:

```
cp -r sasroot/utilities/man/* pathname
```

`pathname` is the directory location of your system man files.

For example, the following command enables you to access online Help by moving the SAS man files from the `!SASROOT` directory to the `man1` file in your system’s `man` directory:

```
cp /usr/local/sas91/utilities/man/* /usr/local/man/man1
```

After you have issued this command, you can access online Help with the `man sas` command.

You can also add the directory to your system's MANPATH environment variable if it has been previously defined.

The UNIX Authentication API

The UNIX Authentication Application Programming Interface (API) is a set of predefined routines that provide user authentication, identification, and permissions verification for SAS when running in UNIX environments. The source files provide the ability to add site-specific behavior to the authentication/identification/permissions validations process.

The `!SASROOT/utilities/src/auth/docs.pdf` file describes how to implement custom authentication implementations, and documents the API itself. The document also includes an explanation of how SAS user authentication and identification integrates with authentication features provided by the operating environment. Administrators that need to implement custom behaviors should read the file and follow the instructions.

Utilities in the /utilities/bin Directory

The following table briefly describes some of the tools in the `/utilities/bin` directory. You can use the UNIX `man` command for information about these utilities. You will need ROOT permissions to execute these commands.

Table A2.1 Tools for the System Administrator

Tool Name	Description
<code>authcustom.so</code>	<code>sasauth</code> module for site-specific authentication
<code>authldap.so</code>	<code>sasauth</code> module for LDAP authentication
<code>authpam.so</code>	<code>sasauth</code> module for PAM authentication
<code>bdm</code>	batch driver monitor
<code>cfgpeh</code>	stand-alone scramble command
<code>cleanwork</code>	tool to delete any leftover Work and utility directories whose associated SAS process has ended
<code>docsetup</code>	documentation setup utility invoked by the installer
<code>elsconf</code>	tool to check ELS configuration
<code>elssrv</code>	ELS server, tool to launch subprocesses
<code>inetcfg.pl</code>	application server configuration tool
<code>inetcfg.pl.template</code>	application server configuration tool
<code>inetcfg7.pl</code>	application server configuration tool
<code>inetcfg7.pl.template</code>	application server configuration tool
<code>jproxy</code>	tool used to launch the Java facilities within SAS
<code>ke2j</code>	double-byte input utility
<code>ke2s</code>	double-byte input utility

Tool Name	Description
kj2e	double-byte input utility
ks2e	double-byte input utility
loadmgr	application lead manager
motifxsassm	Motif X session manager
objspawn	object spawner
patchname	resets the name of the SASROOT directory in the specified executable file.
rbrowser	remote browser server for the platform (only supported on Linux, but works on all other Motif platforms)
reshelper	resource helper for X Windows
sasauth	user identification and authentication utility
sasauth.conf	configuration file for sasauth; specifies authentication module used and other options
sasm.elm.mime	script for supporting e-mail from SAS
sasmailer	script for supporting e-mail from SAS
sasperm	user permissions utility
sastcpd	TCP/IP access daemon
saswujms	Japanese input server
setuid	directory
setuid.sh	script to set some commands to run as root

cleanwork Command

Deletes any leftover Work and Utility directories whose associated SAS process has ended.

Syntax

cleanwork *directory* *<-n, -hostmatch>*

directory

names the directory containing the Work and Utility directories. The name must match the value specified in the WORK system option (which is typically **/usr/tmp** and contained in the **!SASROOT/sasv9.cfg** file), or the value specified in the UTILLOC system option.

Note: Unless **cleanwork** is run by root, user permissions might prevent you from deleting a directory. Δ

-n

specifies that SAS list the entries in a directory that can be removed.

-hostmatch

specifies the name of a host from which you can remove Work directories that might still be active in a Network File System (NFS).

Details

The **cleanwork** command deletes any directories that were assigned to the Work data library or directories assigned by the UTILLOC system option. **cleanwork** deletes only the SAS processes that are on the UNIX system where the SAS session is running. Each SAS process has one of the following formats:

SAS_workcode_nodename

SAS_utilcode_nodename

code

is a 12-character code. The first four characters are randomly generated numbers. The next eight characters are based on the hexadecimal process ID of the SAS session. Processes that are active are not deleted.

nodename

specifies the name of the UNIX system where the SAS process is running.

For example, if you are working on nodename *jupiter*, then the **cleanwork** command deletes all directories with inactive processes on *jupiter*. **cleanwork** does not delete a directory that is associated with an orphaned process if that process is still active. In this case, you need to manually kill the process, and then rerun **cleanwork**.

See Also

- “Work Library” on page 58

patchname Command

Resets the name of the !SASROOT directory in the specified executable file.

Syntax

patchname *filepath sasroot-directory-pathname*

filepath

specifies the absolute pathname of the file in which to set the !SASROOT directory.

sasroot-directory-pathname


specifies the absolute pathname of the new !SASROOT directory.

Details

The **patchname** command resets the name of the !SASROOT directory in the specified executable file to the specified directory.

When you install SAS, the installation program uses **patchname** to write the name of the **!SASROOT** directory to the file that needs this information: the executable file that contains the **sas** command.

If you want to relocate SAS to a different directory, follow the steps in “Regenerating SAS Invocation Scripts” on page 5.



APPENDIX

3

**Commands That Are Not Specific
to UNIX**

Commands That Are Not Specific to UNIX 430

- AUTOADD Command* 430
- AUTOFLOW Command* 431
- AUTOSPLIT Command* 432
- AUTOWRAP Command* 433
- BOUNDS Command* 434
- C Command* 435
- CC Command* 436
- CAPS Command* 437
- CL Command* 438
- CCL Command* 439
- CU Command* 440
- CCU Command* 440
- CURSOR Command* 441
- D Command* 442
- DD Command* 442
- DICT Command* 443
- FILL Command* 444
- I Command* 445
- INDENT Command* 446
- JC Command* 447
- JJC Command* 448
- JL Command* 449
- JJL Command* 450
- JR Command* 451
- JJR Command* 452
- KEYS Command* 453
- M Command* 454
- MM Command* 454
- MASK Command* 455
- NUMBERS Command* 456
- R Command* 457
- RR Command* 458
- RESET Command* 459
- SPELL Command* 459
- TC Command* 461
- TF Command* 462
- TS Command* 463
- UNDO Command* 465
- (Command* 465
- ((Command* 466

) Command	467
) Command	468
Shift Left Command	469
Shift Left Block Command	470
Shift Right Command	471
Shift Right Block Command	472

Commands That Are Not Specific to UNIX

Commands that are specific to the text editor are called text-editing commands because they perform editing functions in windows. Text-editing commands can be one of two types:

- line commands
- command-line commands

Most line commands rearrange or reformat text. They perform tasks such as moving, deleting, copying, and aligning lines or blocks of text. Command-line commands, in addition to rearranging and reformatting text, perform other tasks such as reversing the effects of commands or changing the default case of text.

This section describes commands that you can use in the UNIX environment, but that are not specific to UNIX. You can use these commands in any operating environment that supports text-editing commands.

AUTOADD Command

Controls automatic line addition.

Category: Text editing, command-line command

Syntax

AUTOADD <ON | OFF>

Without Arguments

The AUTOADD command is toggled ON or OFF. Issue the command once to reverse the current setting. If the current setting is ON, then issuing the AUTOADD command changes the setting to OFF. If the current setting is OFF, then issuing the AUTOADD command changes the setting to ON.

Arguments

ON

turns on the AUTOADD command in the window so that lines are added automatically.

OFF

turns off the AUTOADD command in the window so that lines are not added automatically.

Details

The AUTOADD command controls whether blank lines are added as you scroll past existing text. The number of lines that are added is determined by the setting of the VSCROLL command, which determines the default scroll amount forward or backward.

See Also

Commands:

“AUTOFLOW Command” on page 431

“AUTOSPLIT Command” on page 432

“AUTOWRAP Command” on page 433

“I Command” on page 445

AUTOFLOW Command

Controls whether text is flowed when it is included, copied, or pasted.

Category: Text editing, command-line command

Syntax

AUTOFLOW <ON | OFF>

Without Arguments

The AUTOFLOW command is toggled on and off. Issue the command once to reverse the current setting. If the current setting is ON, then issuing the AUTOFLOW command changes the setting to OFF. If the current setting is OFF, then issuing the AUTOFLOW command changes the setting to ON.

Arguments**ON**

turns on the AUTOFLOW command in the window so that text flows when it is inserted in the window.

OFF

turns off the AUTOFLOW command in the window so that text retains its previous position when it is inserted in the window.

Details

The AUTOFLOW command controls whether text inserted with the INCLUDE, PASTE, or COPY command automatically flows. When text is flowed, the left and right

boundaries are determined by the settings that were specified with previous executions of the INDENT and BOUNDS commands. The AUTOFLOW command controls all text that is inserted in the window. It does not stop at paragraph boundaries.

Comparisons

The AUTOFLOW command controls whether text inserted in the window flows, while the TF command flows text that is already in the window.

See Also

Commands:

“AUTOSPLIT Command” on page 432

“AUTOWRAP Command” on page 433

“BOUNDS Command” on page 434

“INDENT Command” on page 446

“TF Command” on page 462

AUTOSPLIT Command

Controls whether text is split at the cursor when you press ENTER or RETURN, or when you are at a carriage return.

Category: Text editing, command-line command

Syntax

AUTOSPLIT <ON | OFF>

Without Arguments

The AUTOSPLIT command acts is toggled on and off. The first time you issue the AUTOSPLIT command, it reverses the current setting. If the current setting is ON, then issuing the AUTOSPLIT command changes the setting to OFF. If the current setting is OFF, then issuing the AUTOSPLIT command changes the setting to ON.

Arguments

ON

turns on the AUTOSPLIT command in the window so that when you press ENTER or RETURN, or when you are at a carriage return, text automatically splits at the cursor.

OFF

turns off the AUTOSPLIT command in the window so that when you press ENTER or RETURN, or when you are at a carriage return, text does not automatically split at the cursor.

Details

The AUTOSPLIT command controls whether text is split at the cursor when you press ENTER or RETURN, or when you are at a carriage return. All text on the line, starting with the character on which the cursor is resting, moves to the left margin of the next line. The cursor is repositioned so that it rests on the first character of the new line.

Comparisons

Entering a carriage return with the AUTOSPLIT command turned on is identical to issuing the TS command with the default numeric argument of 1. The results of a carriage return with the AUTOSPLIT command turned on can be reversed by the TC command or undone with the UNDO command.

See Also

Commands:

“AUTOFLOW Command” on page 431

“AUTOWRAP Command” on page 433

“TF Command” on page 462

“TS Command” on page 463

AUTOWRAP Command

Controls whether text is wrapped when it is included, copied, or filed.

Category: Text editing, command-line command

Syntax

AUTOWRAP <ON | OFF>

Without Arguments

The AUTOWRAP command is toggled on or off. The first time you issue the AUTOWRAP command, it reverses the current setting. If the current setting is ON, then issuing the AUTOWRAP command changes the setting to OFF. If the current setting is OFF, then issuing the AUTOWRAP command changes the setting to ON.

Arguments

ON

turns on the AUTOWRAP command in the window so that text is wrapped when it is inserted in the window, or when it is moved to an external file.

OFF

turns off the AUTOWRAP command in the window. Depending on the line length, text can be truncated as it is inserted in the window, or when it is moved to an external file.

Details

When the AUTOWRAP command is turned on, you can use the INCLUDE or COPY commands. These commands can insert a file, which has a line length that exceeds the boundaries of the window, in a window. The text in the file is not truncated. Instead, lines in the file are split at word boundaries. Conversely, the AUTOWRAP command enables you to use the FILE command to send text, which has a line length that exceeds the boundaries of a file, to an external file. The text in the file is not truncated. Lines are split at word boundaries. When the AUTOWRAP command is turned off, text can be truncated depending on the line length of the text and of the window or file.

See Also

Commands:

“AUTOFLOW Command” on page 431

“AUTOSPLIT Command” on page 432

BOUNDS Command

Sets left and right boundaries when text is flowed.

Category: Text editing, command-line command

Syntax

BOUNDS <*left right*>

Without Arguments

The BOUNDS command displays a message identifying the current boundary settings.

Arguments

left

sets the left boundary by column position.

right

sets the right boundary by column position.

Details

The BOUNDS command resets the left and right boundaries for text. Text is reset by column position and must already be in the window and flowed with the TF command. The BOUNDS command sets the left and right boundaries for text inserted in the window with the INCLUDE, COPY, and PASTE commands when the AUTOFLOW command is turned on. When the AUTOFLOW command is turned on, the left boundary setting is maintained when text is split with the TS command.

For example, specify the following command if you want the text flowed between columns 10 and 60:

```
bounds 10 60
```

Each time text is flowed after this BOUNDS command is issued, the text is flowed between spaces 10 and 60. The text is flowed until you issue another BOUNDS command, or the INDENT command is set to ON.

Setting the INDENT command to ON always overrides the current left boundary setting. To ensure that the left boundary setting is used, set the INDENT command to OFF.

Comparisons

The BOUNDS command affects the behaviors of the TF and TS commands. The BOUNDS command is similar to the INDENT command because both can set the left boundary. However, the BOUNDS command can set the right boundary. When text is flowed, setting the INDENT command to ON always sets the left boundary, which overrides the left boundary that is set by the BOUNDS command.

See Also

Commands:

“AUTOFLOW Command” on page 431

“INDENT Command” on page 446

“TF Command” on page 462

“TS Command” on page 463

C Command

Copies one line of text.

Category: Text editing, line command

Syntax

C

intervening text

A | B

Without Arguments

The C command copies one line of text to a new position anywhere in a window.

Arguments

A marks the target position of the line of text to be copied; in this case, after the position where the A argument is typed. You can place the A argument either before or after the line to be copied.

B marks the target position of the line of text to be copied; in this case, before the position where the B argument is typed. You can place the B argument either before or after the line to be copied.

Comparisons

The C and CC commands enable you to specify a target position for the line of text anywhere in the window. The R and RR commands repeat the line or block of text immediately after it first appears.

See Also

Commands:

“CC Command” on page 436

“R Command” on page 457

“RR Command” on page 458

CC Command

Copies a block of lines of text.

Category: Text editing, line command

Syntax

CC

block of text

CC

intervening text

A | B

Without Arguments

The CC command copies a block of lines of text to a new position anywhere in a window.

Arguments

A

marks the target position of the lines of text to be copied; in this case, after the position where the A argument is typed. You can place the A argument either before or after the lines to be copied.

B

marks the target position of the lines of text to be copied; in this case, before the position where the B argument is typed. You can place the B argument either before or after the lines to be copied.

Comparisons

The C and CC commands enable you to specify a target position for the lines of text anywhere in the window. The R and RR commands repeat the block of lines of text immediately after it first appears.

See Also

Command:

“C Command” on page 435

“R Command” on page 457

“RR Command” on page 458

CAPS Command

Changes the default case of text.

Category: Text editing, command-line command

Syntax

CAPS <ON | OFF>

Without Arguments

The CAPS command is toggled on and off. The first time you issue the CAPS command, it reverses the current setting. If the current setting is ON, then issuing the CAPS command changes the setting to OFF. If the current setting is OFF, then issuing the CAPS command changes the setting to ON.

Arguments

ON

turns on the CAPS command. The case of characters that you enter after you turn on the CAPS command is uppercase. Character strings for the FIND and CHANGE

commands are also translated into uppercase unless they are enclosed in quotation marks.

OFF

turns off the CAPS command. The case of characters that you enter after you turn off the CAPS command is unchanged.

Details

The CAPS command changes the case for text not yet entered, or for text that is modified in a window. If you specify CAPS ON and enter text, the text is changed to uppercase as soon as you press ENTER or RETURN. The setting remains in effect for a window until the SAS session ends, or until the setting is changed by another CAPS command. You can use the WSAVE command to save the setting of the CAPS command beyond your current SAS session.

Comparisons

The CAPS ON command is similar to the CU and CCU commands, and to the CL and CCL commands, which change the case of existing text. However, the CAPS command changes the default case of text, not the case of existing text.

See Also

Command:

“CL Command” on page 438 and “CCL Command” on page 439

“CU Command” on page 440 and “CCU Command” on page 440

CL Command

Changes all characters in a designated line of text to lowercase.

Category: Text editing, line command

Syntax

CL <*n*>

Without Arguments

The CL command changes to lowercase all characters in a designated line of text.

Arguments

n

specifies the number of lines of text to be changed to lowercase. Follow the *n* argument with a space.

Comparisons

The CL and CCL commands change existing text to lowercase, while the CAPS OFF command makes the default case of text lowercase. The case of new, inserted text is changed. The CU and CCU commands, which change existing text to uppercase, accomplish the opposite of the CL and CCL commands.

See Also

Commands:

“CCL Command” on page 439

“CAPS Command” on page 437

“CU Command” on page 440 and “CCU Command” on page 440

CCL Command

Changes all characters in designated lines of text to lowercase.

Category: Text editing, line command

Syntax

CCL

block of text

CCL

Without Arguments

The CCL command changes to lowercase all characters in a block of lines of text.

Comparisons

The CL and CCL commands change existing text to lowercase, while the CAPS OFF command makes the default case of text lowercase, which changes the case of new, inserted text. The CU and CCU commands, which change existing text to uppercase, accomplish the opposite of the CL and CCL commands.

See Also

Commands:

“CL Command” on page 438

“CAPS Command” on page 437

“CU Command” on page 440 and “CCU Command” on page 440

CU Command

Changes all characters in a designated line of text to uppercase.

Category: Text editing, line command

Syntax

CU <*n*>

Without Arguments

The CU command changes to uppercase all characters in a designated line of text.

Arguments

n

specifies the number of lines of text to be changed to uppercase. Follow the *n* argument with a space.

Comparisons

The CU and CCU commands are similar to the CAPS ON command. The CU and CCU commands change existing text to uppercase, while the CAPS ON command makes the default case of text uppercase, which changes the case of new, inserted text. The CL and CCL commands, which change existing text to lowercase, accomplish the opposite of the CU and CCU commands.

See Also

Commands:

“CAPS Command” on page 437

“CCU Command” on page 440

“CL Command” on page 438 and “CCL Command” on page 439

CCU Command

Changes all characters in a designated block of lines of text to uppercase.

Category: Text editing, line command

Syntax

CCU

block of text

CCU

Without Arguments

The CCU command changes to uppercase all characters in a block of designated lines of text.

Comparisons

The CU and CCU commands are similar to the CAPS ON command. The CU and CCU commands change existing text to uppercase, while the CAPS ON command makes the default case of text uppercase, which changes the case of new, inserted text. The CL and CCL commands, which change existing text to lowercase, accomplish the opposite of the CU and CCU commands.

See Also

Commands:

“CU Command” on page 440

“CAPS Command” on page 437

“CL Command” on page 438 and “CCL Command” on page 439

CURSOR Command

Moves the cursor to the command line.

Category: Text editing, command-line command

Syntax

CURSOR

Without Arguments

The CURSOR command moves the cursor to the command line. The CURSOR command is designed to be executed with a function key.

Details

The CURSOR command can be used interchangeably with the HOME key.

Comparisons

The CURSOR command has the same results as pressing the HOME key.

See Also

Commands:

“HOME Command” on page 230

D Command

Deletes a designated line.

Category: Text editing, line command

Syntax

D <*n*>

Without Arguments

The D command deletes only the designated line.

Arguments

n

specifies the number of lines to delete. Follow the *n* argument with a space.

See Also

Commands:

“DD Command” on page 442

DD Command

Deletes a designated block of lines.

Category: Text editing, line command

Syntax

DD

block of lines

DD

Without Arguments

The DD command deletes a block of lines of text.

See Also

Commands:

“D Command” on page 442

DICT Command

Includes, releases, and creates an auxiliary dictionary.

Category: Text editing, command-line command

Syntax

DICT INCLUDE *dictionary-name* | **FREE** *dictionary-name* | **CREATE** *dictionary-name*
 <*size*>

Arguments

INCLUDE *dictionary-name*

makes the auxiliary dictionary that is specified available in the current SAS session. Only a one-level name is accepted. The SASUSER.PROFILE catalog is checked first for the dictionary. Then, the SASHELP.BASE catalog is checked. If the auxiliary dictionary is not found, SAS issues an error message. If the auxiliary dictionary is made available from the SASHELP.BASE catalog, no changes to it are saved. If it is made available from the SASUSER.PROFILE catalog, changes to it are saved.

FREE *dictionary-name*

releases the auxiliary dictionary that is specified. A newly created dictionary is not saved in the SASUSER.PROFILE catalog until you issue the DICT command with the FREE argument, or you end the current interactive windowing task. If the auxiliary dictionary has been modified, the changes are saved when you issue the DICT command with the FREE argument. These changes are saved unless the auxiliary dictionary was made available from the SASHELP.BASE catalog.

CREATE *dictionary-name*

creates a new auxiliary dictionary as specified. The dictionary is initially empty. When the dictionary is released, it is saved in the SASUSER.PROFILE catalog. Only a one-level name is accepted.

size

specifies the size in bytes of the auxiliary dictionary. The default is 9,808 bytes.

Comparisons

The DICT command includes, releases, and creates an auxiliary dictionary. The SPELL command checks spelling and flags unrecognized words. In addition, the SPELL command can create and update dictionaries.

See Also

Commands:

“SPELL Command” on page 459

FILL Command

Places fill characters beginning at the current cursor position.

Category: Text editing, command-line command

Syntax

FILL <'fill-character'> <n>

Without Arguments

The FILL command displays a message identifying the fill character and the number of its repetitions.

Arguments

'fill-character'

specifies a customized character that must be enclosed in single quotation marks. The fill character remains in effect until you change it.

n

specifies the exact number of fill characters. The number remains in effect until you change it.

Details

The FILL command places fill characters beginning at the current cursor position. The fill characters extend to the end of a line, or to the space before the next non-blank character, whichever occurs first. By default, the fill character is usually an underscore or dash. If you use the FILL arguments, you can change the fill character and the number of repetitions.

The FILL command is most easily issued with a function key. To place the fill characters at the cursor position, set one of your function keys to issue the FILL command. Move the cursor to a Program Editor field, and then press the function key. The fill characters are displayed.

The following example shows how you can change the default. Issuing the following command makes the default become 10 question marks:

```
fill '?' 10
```

The changed fill character is in effect for the duration of your SAS session, or until you change it. You can use the WSAVE command to permanently save the setting.

I Command

Inserts one or more blank lines.

Category: Text editing, line command

Syntax

I <A | B> <*n*>

Without Arguments

The I command inserts one or more blank lines immediately after the line on which you issued the command.

Arguments

A

inserts one or more blank lines immediately after the line on which you issued the command. You cannot have any characters between the I command and the A argument.

B

inserts one or more the blank lines immediately before the line on which you issued the command. You cannot have any characters between the I command and the B argument.

n

specifies the number of blank lines to insert. Follow the *n* argument with a space. If you use the A or B argument, the *n* argument is specified last. For example, if line 00009 contains a PROC PRINT statement, the following I command specifies that you want to insert three blank lines before the line of text:

```
ib3 9 proc print data=final.educ;
```

Details

The I command inserts one or more blank lines. By default, the lines are blank. You can define content with the MASK command. The I command is most easily issued with a function key.

Comparisons

You can use the MASK command with the I command. The I command inserts one or more blank lines, which can include content set by the MASK command.

See Also

Commands:

“MASK Command” on page 455

INDENT Command

Retains left margin indentation when text is flowed.

Category: Text editing, command-line command

Syntax

INDENT <ON | OFF>

Without Arguments

The INDENT command is toggled on or off. The first time you issue the INDENT command, it reverses the current setting. If the current setting is ON, then issuing the INDENT command changes the setting to OFF. If the current setting is OFF, then issuing the INDENT command changes the setting to ON. When you reissue the INDENT command, it returns to the previous setting.

Arguments

ON

turns on the INDENT command in the window.

Tip: The INDENT ON command indents all lines.

Tip: When the INDENT command is turned on, and you issue the TF command, all of the lines in the paragraph are indented the same as the first line in the paragraph.

OFF

turns off the INDENT command in the window.

Details

The INDENT command specifies that the current left margin indentation is used under the following conditions:

- when existing text in a window is flowed with the TF command
- when text is inserted in a window when the AUTOFLOW command is turned on
- when existing text in a window is split with the TS command

Comparisons

The left boundary can be set by both the INDENT and BOUNDS commands. However, when text is flowed, turning the INDENT command on always determines the left boundary, and overrides the left boundary set by the BOUNDS command.

Examples

This example shows four lines of text. The TF command is typed in the number field of the first line. The first line of the paragraph is indented. The INDENT command is set to ON, and the default boundaries are 1 and 50:

```
tf 01      The purpose of Monday's meeting is to review
00002 the documentation plan and gather your responses. Please
```

```
00003 send a representative
00004 if you are unable to attend.
```

The following example shows the result of pressing ENTER or RETURN to issue the TF command. The indentation is used for all of the lines and the right boundary is 50:

```
tf 01      The purpose of Monday's meeting is to review
00002      the documentation plan and gather your responses. Please
00003      send a representative
00004      if you are unable to attend.
```

See Also

Commands:

“AUTOFLOW Command” on page 431

“BOUNDS Command” on page 434

“TF Command” on page 462

“TS Command” on page 463

JC Command

Centers a designated line of text.

Category: Text editing, line command

Syntax

JC <*n*>

Without Arguments

The JC command centers the designated line of text that contains the line command based on the left and right boundary settings.

Arguments

n
specifies the column position on which to center the designated line of text. Follow the *n* argument with a space.

Details

The JC command centers a designated line of text. Unless you specify a numeric argument, centering is based on the current boundary settings set by the BOUNDS command. A numeric argument overrides these boundary settings.

Comparisons

Like the JL, JLL, JR, and JJR commands, the JC and JJC commands align text.

See Also

Commands:

“JJC Command” on page 448

“BOUNDS Command” on page 434

“JL Command” on page 449 and “JLL Command” on page 450

“JR Command” on page 451 and “JJR Command” on page 452

JJC Command

Centers each line of text in a designated block of text independently.

Category: Text editing, line command

Syntax

JJC

block of text

JJC

Without Arguments

The JJC command centers a block of designated text that contains the line command based on the left and right boundary settings.

Details

The JJC command centers a designated block of text. Each line in the block is centered independently. Centering is based on the current boundary settings set by the BOUNDS command.

Comparisons

Like the JL, JLL, JR, and JJR commands, the JC and JJC commands align text.

See Also

Commands:

“JC Command” on page 447

“BOUNDS Command” on page 434

“JL Command” on page 449 and “JLL Command” on page 450

“JR Command” on page 451 and “JJR Command” on page 452

JL Command

Left-aligns a designated line of text.

Category: Text editing, line command

Syntax

JL <*n*>

Without Arguments

The JL command left-aligns the designated line of text. Alignment is based on the left and right boundary settings.

Arguments

n

specifies the column position on which to left-align the designated line of text. By default, the *n* argument is the left boundary setting. Follow the *n* argument with a space.

Details

The JL command left-aligns a designated line of text. Unless you specify a numeric argument, left-alignment is based on the current boundary settings set by the BOUNDS command. A numeric argument overrides those boundary settings.

Comparisons

Like the JC, JJC, JR, and JJR commands, the JL and J JL commands align text.

See Also

Commands:

“J JL Command” on page 450

“BOUNDS Command” on page 434

“JC Command” on page 447 and “JJC Command” on page 448

“JR Command” on page 451 and “JJR Command” on page 452

JJJ Command

Left-aligns a designated block of text.

Category: Text editing, line command

Syntax

JJJ <*n*>

block of text

JJJ <*n*>

Without Arguments

The JJJ command left-aligns a designated block of text. Alignment is based on the left and right boundary settings.

Arguments

n specifies the column position on which to left-align the designated block of text. By default, the *n* argument is the left boundary setting. Follow the *n* argument with a space. You can specify the numeric argument in the beginning or ending line of the block command or in both. If it is specified in both, then SAS uses the first numeric argument.

Details

The JJJ command left-aligns a designated block of text. Unless you specify a numeric argument, left-alignment is based on the current boundary settings set by the BOUNDS command. A numeric argument overrides these boundary settings.

Comparisons

Like the JC, JJC, JR, and JJR commands, the JL and JJJ commands align text.

See Also

Commands:

“JL Command” on page 449

“BOUNDS Command” on page 434

“JC Command” on page 447 and “JJC Command” on page 448

“JR Command” on page 451 and “JJR Command” on page 452

JR Command

Right-aligns a designated line of text.

Category: Text editing, line command

Syntax

JR <*n*>

Without Arguments

The JR command right-aligns the designated line of text. Alignment is based on the left and right boundary settings.

Arguments

n

specifies the column position on which to right-align the designated line of text. By default, the *n* argument is the right boundary setting. Follow the *n* argument with a space.

Details

The JR command right-aligns a designated line of text. Unless you specify a numeric argument, right-alignment is based on the current boundary settings set by the BOUNDS command. A numeric argument overrides these boundary settings.

Comparisons

Like the JC, JJC, JL, and JJJL commands, the JR and JJR commands align text.

See Also

Commands:

“JJR Command” on page 452

“BOUNDS Command” on page 434

“JC Command” on page 447 and “JJC Command” on page 448

“JL Command” on page 449 and “JJL Command” on page 450

JJR Command

Right-aligns a designated block of text.

Category: Text editing, line command

Syntax

JJR <*n*>

block of text

JJR <*n*>

Without Arguments

The JJR command right-aligns a designated block of text. Alignment is based on the left and right boundary settings.

Arguments

n

specifies the column position on which to right-align the designated block of text. By default, the *n* argument is the right boundary setting. Follow the *n* argument with a space. You can specify the numeric argument in the beginning or ending line of the block command or in both. If it is specified in both, then SAS uses the first numeric argument.

Details

The JJR command right-aligns a designated block of text. Unless you specify a numeric argument, right-alignment is based on the current boundary settings set by the BOUNDS command. A numeric argument overrides these boundary settings.

Comparisons

Like the JC, JJC, JL, and JJJL commands, the JR and JJR commands align text.

See Also

Commands:

“JR Command” on page 451

“BOUNDS Command” on page 434

“JC Command” on page 447 and “JJC Command” on page 448

“JL Command” on page 449 and “JJL Command” on page 450

KEYS Command

Enables you to assign function keys to tasks.

Category: Text editing, command-line command

Syntax

KEYS

Without Arguments

The KEYS command enables you to access the KEYS window so that you can assign function keys to tasks.

M Command

Moves one line of text.

Category: Text editing, line command

Syntax

M <*n*>

intervening text

A | B

Arguments

n

specifies the number of lines of text to move. Follow the *n* argument with a space. Without the *n* argument, only the line of text that contains the line command is moved.

A

marks the target position of the line of text to be moved; in this case, after the line where the A argument is typed. You can place the A argument either before or after the line to be moved.

B

marks the target position of the line of text to be moved; in this case, before the line where the B argument is typed. You can place the B argument either before or after the line to be moved.

Details

The M command moves a designated line of text to a new position anywhere in a window.

See Also

Commands:

“MM Command” on page 454

MM Command

Moves a block of lines of text.

Category: Text editing, line command

Syntax

MM

block of text

MM

intervening text

A | B

Arguments

A

marks the target position of the lines of text to be moved; in this case, after the line where the A argument is typed. You can place the A argument either before or after the lines to be moved.

B

marks the target position of the lines of text to be moved; in this case, before the line where the B argument is typed. You can place the B argument either before or after the lines to be moved.

See Also

Commands:

“M Command” on page 454

MASK Command

Defines the contents of one or more new lines.

Category: Text editing, line command

Syntax

MASK

Without Arguments

The MASK command defines, displays, and enables you to edit the contents of one or more new lines that are created by the I command.

Details

The MASK command defines, displays, and enables you to edit the contents of one or more new lines that are created by the I command. The default setting for the new lines are blank lines. To display or edit a new line, type MASK in the number field of the line, and then press ENTER or RETURN. The line with the contents defined by the

MASK command is inserted. You can then edit the line. A line with the contents defined by the MASK command is inserted each time you issue the I command.

The contents defined by the MASK command remain in effect for that window throughout your current SAS session unless you change them. To change the contents, type over the text. If you want to return to the default (a blank line), do one of the following tasks:

- blank any characters in the text field of the MASK line.
- issue the CLEAR command with MASK as an argument:

```
clear mask
```

The contents of the MASK line are cleared, and a note appears in the log indicating that the MASK line has been cleared.

You can use the RESET command or the D command to not display the contents of the MASK command. The MASK command remains in effect even when it is not displayed. For example, the MASK command remains in effect under the following conditions:

- when you scroll past the MASK line and it is not displayed
- when you issue the D or RESET command without blanking any characters in the text of the MASK line

In some windows, such as the Program Editor window, you can use the WSAVE command to permanently save the contents of the MASK command.

See Also

Commands:

“D Command” on page 442

“RESET Command” on page 459

NUMBERS Command

Adds or removes line numbers.

Category: Text editing, command-line command

Syntax

NUMBERS <ON | OFF>

Without Arguments

The NUMBERS command is toggled on or off. The first time you issue the NUMBERS command, it reverses the current setting. If the current setting is ON, then issuing the NUMBERS command changes the setting to OFF. If the current setting is OFF, then issuing the NUMBERS command changes the setting to ON.

Arguments

ON

turns on the NUMBERS command in the window, so that the lines in the Program Editor are numbered.

OFF

turns off the NUMBERS command in the window, so that the lines in the Program Editor are not numbered.

Details

The NUMBERS command adds or removes line numbers for data lines in windows that allow text editing. When you issue the NUMBERS command to remove line numbers, the line numbers disappear and all text appears to shift left. When you issue the NUMBERS command to add line numbers, the numbers are displayed on the left, and all of the text appears to shift right. The alias for the NUMBERS command is NUMS.

R Command

Repeats a designated line.

Category: Text editing, line command

Syntax

R <*n*>

Without Arguments

The R command repeats a designated line one time.

Arguments

n

specifies the number of times to repeat the designated line. Follow the *n* argument with a space.

Details

The R command repeats a designated line immediately after the designated line. The default is one time.

Comparisons

The R and RR commands repeat the line or block of lines immediately after the line with the R or RR command. The C and CC commands enable you to copy one or more lines anywhere in a window.

See Also

Commands:

“RR Command” on page 458

“C Command” on page 435 and “CC Command” on page 436

RR Command

Repeats a block of lines.

Category: Text editing, line command

Syntax

RR *<n>*

block of text

RR *<n>*

Without Arguments

The RR command repeats the block of lines one time.

Arguments

n

specifies the number of times to repeat the designated block of lines. Follow the *n* argument with a space. You can specify the numeric argument in the beginning or ending line of the block command or in both. If it is specified in both, SAS uses the first numeric argument.

Details

The RR command repeats a designated block of lines immediately after the designated block of lines. The default is one time.

Comparisons

The R and RR commands repeat the line or block of lines immediately after the line with the R or RR command. The C and CC commands enable you to copy one or more lines anywhere in a window.

See Also

Commands:

“R Command” on page 457

“C Command” on page 435 and “CC Command” on page 436

RESET Command

Removes any pending line commands.

Category: Text editing, command-line command

Syntax

RESET

Without Arguments

The RESET command removes any pending line commands.

Details

The RESET command removes any pending line commands. It also removes any MASK lines that were created when the MASK command was issued. The display of the MASK line, not the setting of the MASK command, is removed. If you do not want to complete a command on a block of lines (such as the MM or CC command), you can issue the RESET command to remove the pending command.

Comparisons

The RESET command has the same result as the D command. The RESET command also removes the display of MASK lines from the MASK command.

See Also

Commands:

“D Command” on page 442

SPELL Command

Checks spelling and flags unrecognized words.

Category: Text editing, command-line command

Syntax

SPELL <ALL <SUGGEST>>

SPELL <NEXT | PREV | SUGGEST>

SPELL <REMEMBER <dictionary-name>>

Without Arguments

The SPELL command checks the first word if the cursor is positioned on the command line. Otherwise, the SPELL command checks the word on which the cursor is

positioned. Assign a function key to the SPELL command by using the KEYS command. Use the function key to check the spelling of the word on which the cursor is positioned. If the word is recognized, the message “OK” appears. Otherwise, a message appears indicating that the word is unrecognized.

Arguments

ALL

checks the spelling of all words. If all words are recognized, a message indicates that no unrecognized words have been found.

If a word is unrecognized, the SPELL: Unrecognized Words window appears, listing the unrecognized words and their corresponding line numbers. The window initially displays a blank field for the dictionary that you want to specify. Enter a new dictionary name, or the name of an existing dictionary. Select **Tools ► Remember** to add the unrecognized words to the dictionary.

Tip: Specify the SPELL ALL SUGGEST command to display the SPELL: Suggestions window for each unrecognized word that is found.

SUGGEST

invokes the SPELL: Suggestions window, which displays the last unrecognized word and its line number, and suggestions for changing the unrecognized word. In this window, you can select **Tools ► Remember** to add the unrecognized word to a dictionary. The window will then close, you are returned to the previous window, and a message indicates that the word is now recognized.

You can change the unrecognized word by positioning your cursor on a suggestion, and then pressing ENTER. The suggested word is highlighted. Select **Tools ► Replace**. When you return to the Program Editor window, you will see that the unrecognized word has been changed. If you want to replace all occurrences of the recognized word, first, position the cursor on the phrase **ALL OCCURRENCES**, and then press ENTER. The phrase ALL OCCURRENCES is highlighted. When you return to the Program Editor window, you will see that all occurrences of the unrecognized word have been changed.

Alias: ?

NEXT

finds the next unrecognized word, based on the current cursor position. If all words from the current cursor position to the end of the file are recognized, a message indicates that the end of the file was reached. Otherwise, a message indicates that a word is unrecognized, and the cursor is positioned on the unrecognized word.

PREV

finds the previous unrecognized word, based on the current cursor position. If all words from the current cursor position to the beginning of the file are recognized, a message indicates that the beginning of the file was reached. Otherwise, a message indicates that a word is unrecognized, and the cursor is positioned on the unrecognized word.

REMEMBER *dictionary-name*

adds the last unrecognized word to an auxiliary dictionary, where *dictionary-name* is the name of an auxiliary dictionary. A message indicates that the word has been added to an auxiliary dictionary. If you are using only one auxiliary dictionary, you can omit *dictionary-name*. If no auxiliary dictionary is specified, and *dictionary-name* is omitted, the unrecognized word is saved in a temporary dictionary in the current SAS session only.

You can highlight a word from the SPELL: Unrecognized Words window or from the SPELL: Suggestions window, and then select **Tools ► Remember**. The word that you added is now recognized.

Alias: ADD

Details

The SPELL command checks spelling and flags unrecognized words. You can use the SPELL command to do the following tasks:

- See suggestions for unrecognized words.
- Add unrecognized words to an auxiliary dictionary.
- Replace unrecognized words with suggestions.

The SPELL command checks words with a default dictionary. However, you can specify one or more auxiliary dictionaries to use in addition to the default dictionary.

Any dictionary that you create is stored in your SASUSER.PROFILE catalog. If you update a dictionary using the SPELL REMEMBER command, updates are saved to a temporary dictionary in the current SAS session. The temporary dictionary is created if you do not specify a dictionary name. If you specify a dictionary from the SASHELP.BASE catalog, updates are saved in that dictionary.

Comparisons

The SPELL command checks spelling and flags unrecognized words, and the DICT command includes or creates an auxiliary dictionary. The SPELL command can also create and update auxiliary dictionaries. Use the SPELL command to create a permanent auxiliary dictionary. The word list that is used by the SPELL command acts as a record of the words that are contained in the auxiliary dictionary.

See Also

Commands:

“DICT Command” on page 443

TC Command

Connects two lines of text.

Category: Text editing, line command

Syntax

TC

Without Arguments

The TC command connects two lines of text.

Details

The TC command connects two lines of text. To connect two lines of text, type TC in the number field of a line, and press ENTER or RETURN. The text from the second line moves to the first line. No space appears between text on the first line and text on the second line. To create a space between the last word of the first line and the first word of the second line, start the text of the second line in the second column.

The command does not truncate text.

Comparisons

The TC command is the opposite of the TS command, which splits text at the cursor. It is similar to the TF command, except that it breaks text at boundaries instead of flowing text in a paragraph by removing trailing blanks.

See Also

Commands:

“TF Command” on page 462

“TS Command” on page 463

TF Command

Flows text to a blank line or to the end of the text.

Category: Text editing, line command

Syntax

TF <A> <*n*>

Without Arguments

The TF command flows text to the first blank line, or to the end of the text, whichever comes first, based on left and right boundary settings.

Arguments

A

flows text in a paragraph to the end of the text by removing trailing blanks, continuing over, but not deleting, blank lines. This argument, like the numeric argument, must be specified on the same line as the TF command. You cannot have any characters between the TF command and the A argument.

n

specifies a right boundary to temporarily override the right boundary set by the BOUNDS command. Follow the *n* argument with a space.

Details

The TF command removes trailing blanks and flows text in one of the following ways:

- to the end of the text in a paragraph, as signified by a blank line
- to the end of the text

You can use the TF command to move text into blank space at the ends of lines. You should do this after you delete words and move words to new lines. The TF command never divides a word. The TF command is affected by the INDENT command. If the INDENT command is ON, or an equivalent has been used, the left boundary remains intact when text is flowed.

Comparisons

The TF command is similar to the TC command, which connects two lines of text. The TF command is different from the TS command, which splits text. While the TF command flows text in a window, the AUTOFLOW command determines whether inserted text is flowed in a window. Text flowed as a result of setting AUTOFLOW to ON does not stop at paragraph boundaries. It stops at the end of the new text.

Examples

The following example shows the results of flowing text that has extra spaces. This example shows the text after you type the TF command on line 00001, and before you press ENTER or RETURN:

```
tf001  The TF command
00002           flows a paragraph
00003  by removing
00004  blanks.
```

When you press ENTER or RETURN, the following result is displayed:

```
00001  The TF command flows a paragraph by removing blanks.
```

See Also

Commands:

“AUTOFLOW Command” on page 431

“BOUNDS Command” on page 434

“INDENT Command” on page 446

“TC Command” on page 461

“TS Command” on page 463

TS Command

Splits text at the cursor.

Category: Text editing, line command

Syntax

TS $\langle n \rangle$

Without Arguments

The TS command splits the line of text at the cursor, and moves the remaining text to a new line.

Arguments

n

specifies how many lines down to move the remaining text. The default is one line. Follow the *n* argument with a space.

Details

The TS command splits the line of text at the cursor, and moves the remaining text to a new line starting at the left margin. If you specify a numeric argument, the TS command moves the text down the number of lines specified. With the AUTOFLOW command turned on, the TS command uses the left boundary that is specified by the BOUNDS command. If the INDENT command is turned on, the TS command uses the current indentation at the left margin. With the AUTOFLOW command turned off, the left boundary and the current indentation at the left margin are reset.

Comparisons

The TS command, with the default numeric argument of 1, is the same as entering a carriage return or pressing ENTER or RETURN with the AUTOSPLIT command turned on. The TS command contrasts with the TC command, which connects two lines of text, and the TF command, which flows text to a blank line or to the end of the text. With the AUTOFLOW command turned on, the TS command is affected by both the BOUNDS and INDENT commands.

Examples

This example shows the effect of splitting two statements in a SAS program and placing each statement on a separate line. This example shows the text after you type the TS command on line 0001 and position the cursor after the first statement, and before you press ENTER or RETURN:

```
ts 01 proc print data=temp; run;
```

When you press ENTER or RETURN, the following result is displayed:

```
00001 proc print data=temp;
00002 run;
```

See Also

Commands:

“AUTOSPLIT Command” on page 432

“I Command” on page 445

“TC Command” on page 461

“TF Command” on page 462

UNDO Command

Cancels an action.

Category: Text editing, command-line command

Syntax

UNDO

Without Arguments

The UNDO command cancels the most recent action in an active window that allows text editing.

Details

The UNDO command cancels the most recent action in an active window that allows text editing. The action must be a command that enters or modifies text. If you want to undo more than one action, you must continue to issue the UNDO command. Actions are undone one at a time, starting with the most recent action and moving backward.

Note: The UNDO command cannot undo the SUBMIT command. It cannot reverse the effects of submitted SAS statements. Δ

Comparisons

Although you cannot undo the SUBMIT command, you can use the RECALL command to recall submitted statements back to the Program Editor window.

If you use the CC command to copy and paste a block of text, and then you issue the UNDO command, the block of text that you copied and pasted is deleted. If you use the DD command to delete a block of text, and then you issue the UNDO command, the block of text that you deleted is restored.

(Command

Shifts to the left one designated line of text.

Category: Text editing, line command

Syntax

(<n>

Without Arguments

The (command shifts a designated line of text one space to the left.

Arguments

n

specifies the number of spaces that the designated line of text shifts. The default is one space. Follow the *n* argument with a space.

Details

The (command shifts a designated line of text one or more spaces to the left. If the shift extends past the beginning of the current line, characters are lost.

Comparisons

The)and))commands shift text in the opposite direction from the (and ((commands. The <, <<, >, and >> commands are similar text-shift commands, but they do not lose characters when shifting.

See Also

Commands:

“((Command” on page 466

“) Command” on page 467 and “)) Command” on page 468

“Shift Left Command” on page 469 and “Shift Left Block Command” on page 470

“Shift Right Command” on page 471 and “Shift Right Block Command” on page 472

((Command

Shifts to the left a designated block of lines of text.

Category: Text editing, line command

Syntax

((<*n*>

block of text

((<*n*>

Without Arguments

The ((command shifts a designated block of lines of text one space to the left.

Arguments

n

specifies the number of spaces that the designated block of lines of text shifts. The default is one space. Follow the *n* argument with a space. You can specify the numeric argument in the beginning or ending line of the block command, or in both. If it is specified in both, SAS uses the first numeric argument.

Details

The ((command shifts a designated block of lines of text one or more spaces to the left. If the shift extends past the beginning of the current line, characters are lost.

Comparisons

The)and))commands shift text in the opposite direction from the (and ((commands. The <, <<, >, and >> commands are similar text-shift commands, but they do not lose characters when shifting.

See Also

Commands:

“(Command” on page 465

“) Command” on page 467 and “)) Command” on page 468

“Shift Left Command” on page 469 and “Shift Left Block Command” on page 470

“Shift Right Command” on page 471 and “Shift Right Block Command” on page 472

) Command

Shifts to the right one designated line of text.

Category: Text editing, line command

Syntax

) <*n*>

Without Arguments

The)command shifts a designated line of text one space to the right.

Arguments

n

specifies the number of spaces that the designated line of text shifts. The default is one space. Follow the *n* argument with a space.

Details

The))command shifts a designated line of text one or more spaces to the right. If the shift extends past the end of the current line, characters are lost.

Comparisons

The (and ((commands shift text in the opposite direction from the)and))commands. The <, <<, >, and >> commands are similar text-shift commands, but they do not lose characters when shifting.

See Also

Commands:

“)) Command” on page 468

“(Command” on page 465 and “((Command” on page 466

“Shift Left Command” on page 469 and “Shift Left Block Command” on page 470

“Shift Right Command” on page 471 and “Shift Right Block Command” on page 472

)) Command

Shifts to the right a designated block of lines of text.

Category: Text editing, line command

Syntax

)) <*n*>

block of text

)) <*n*>

Without Arguments

The))command shifts a designated block of lines of text one space to the right.

Arguments

n

specifies the number of spaces that the designated block of lines of text shifts. The default is one space. Follow the n argument with a space. You can specify the numeric argument in the beginning or ending line of the block command, or in both. If it is specified in both, SAS uses the first numeric argument.

Details

The `)` command shifts a designated block of lines of text one or more spaces to the right. If the shift extends past the end of the current line, characters are lost.

Comparisons

The `(` and `((` commands shift text in the opposite direction from the `)` and `)` commands. The `<`, `<<`, `>`, and `>>` commands are similar text-shift commands, but they do not lose characters when shifting.

See Also

Commands:

“`)` Command” on page 467

“`(` Command” on page 465 and “`((` Command” on page 466

“Shift Left Command” on page 469 and “Shift Left Block Command” on page 470

“Shift Right Command” on page 471 and “Shift Right Block Command” on page 472

Shift Left Command

Shifts to the left a designated line of text.

Category: Text editing, line command

Syntax

`< < n >`

Without Arguments

The `<` command shifts a designated line of text one space to the left.

Arguments

n

specifies the number of spaces that the designated line of text shifts. Follow the n argument with a space.

Details

The `<` command shifts a designated line of text one or more spaces to the left. The line shifts the number of spaces that you specify with the n argument, or the line shifts at

the left window border, whichever is less. This text-shift command does not lose characters when shifting.

Comparisons

The > and >> commands shift text in the opposite direction from the < and << commands. The),), (, and ((commands are similar text-shift commands, which, depending on the extent of the shift, can lose characters.

See Also

Commands:

“Shift Left Block Command” on page 470

“Shift Right Command” on page 471 and “Shift Right Block Command” on page 472

“(Command” on page 465 and “((Command” on page 466

) Command” on page 467 and “)) Command” on page 468

Shift Left Block Command

Shifts to the left a designated block of lines of text.

Category: Text editing, line command

Syntax

```
<< <n>
```

```
block of text
```

```
<< <n>
```

Without Arguments

The << command shifts a designated block of lines of text one space to the left.

Arguments

n

specifies the number of spaces that the designated block of lines of text shifts. Follow the *n* argument with a space. You can specify the numeric argument in the beginning or ending line of the block command, or in both. If it is specified in both, SAS uses the first numeric argument.

Details

The << command shifts a designated block of lines of text one or more spaces to the left. The block of lines of text shifts the number of spaces you specify with the *n* argument,

or the block shifts at the left window border, whichever is less. This text-shift command does not lose characters when shifting.

Comparisons

The > and >> commands shift text in the opposite direction from the < and << commands. The),)), (, and ((commands are similar text-shift commands, which, depending on the extent of the shift, can lose characters.

See Also

Commands:

“Shift Left Command” on page 469

“Shift Right Command” on page 471 and “Shift Right Block Command” on page 472

“(Command” on page 465 and “((Command” on page 466

“) Command” on page 467 and “)) Command” on page 468

Shift Right Command

Shifts to the right a designated line of text.

Category: Text editing, line command

Syntax

> <*n*>

Without Arguments

The > command shifts a designated line of text one space to the right.

Arguments

n

specifies the number of spaces that the designated line of text shifts. Follow the *n* argument with a space.

Details

The > command shifts a designated line of text one or more spaces to the right. The line shifts the number of spaces you specify with the *n* argument, or the line shifts at the left window border, whichever is less. This text-shift command does not lose characters when shifting.

Comparisons

The < and << commands shift text in the opposite direction from the > and >> commands. The),), (, and ((commands are similar text-shift commands, which, depending on the extent of the shift, can lose characters.

See Also

Commands:

“Shift Right Block Command” on page 472

“Shift Left Command” on page 469 and “Shift Right Block Command” on page 472

“(Command” on page 465 and “((Command” on page 466

) Command” on page 467 and “)) Command” on page 468

Shift Right Block Command

Shifts to the right a designated block of text.

Category: Text editing, line command

Syntax

```
>> <n>
```

block of text

```
>> <n>
```

Without Arguments

The >> command shifts a designated block of lines of text one space to the right.

Arguments

n

specifies the number of spaces that the designated block of lines of text shifts. Follow the *n* argument with a space. You can specify the numeric argument in the beginning or ending line of the block command, or in both. If it is specified in both, SAS uses the first numeric argument.

Details

The >> command shifts a designated block of lines of text one or more spaces to the right. The block of lines of text shifts the number of spaces you specify with the *n* argument, or the block shifts at the left window border, whichever is less. This text-shift command does not lose characters when shifting.

Comparisons

The < and << commands shift text in the opposite direction from the > and >> commands. The),), (, and ((commands are similar text-shift commands, which, depending on the extent of the shift, can lose characters.

See Also


Commands:

“Shift Right Command” on page 471

“Shift Left Command” on page 469 and “Shift Left Block Command” on page 470

“(Command” on page 465 and “((Command” on page 466

“) Command” on page 467 and “)) Command” on page 468



APPENDIX

4

Recommended Reading

Recommended Reading 475

Recommended Reading

Here is the recommended reading list for this title:

- SAS Language Reference: Concepts*
- SAS Language Reference: Dictionary*
- Base SAS Procedures Guide*
- Moving and Accessing SAS Files*
- SAS Macro Language: Reference*
- SAS National Language Support (NLS): Reference Guide*
- Configuration Guide - SAS 9.2 Foundation for UNIX Environments

For a complete list of SAS publications, go to support.sas.com/bookstore. If you have questions about which titles you need, please contact a SAS Publishing Sales Representative at:

SAS Publishing Sales
SAS Campus Drive
Cary, NC 27513
Telephone: 1-800-727-3228
Fax: 1-919-531-9439
E-mail: sasbook@sas.com
Web address: support.sas.com/bookstore

Customers outside the United States and Canada, please contact your local SAS office for assistance.

Glossary

active window

a window that is open and displayed, and to which keyboard input is directed. Only one window can be active at a time.

aggregate syntax

a convenient way of referring to individual files in a single directory or folder. Instead of assigning a unique fileref to each file, you assign a fileref to the directory or folder. Then, to refer to a specific file in that folder, you enclose the filename in parentheses following the fileref. In UNIX operating environments, aggregate syntax is used in the FILE, INFILE, and %INCLUDE statements.

American Standard Code for Information Interchange

a 7-bit character encoding that is the U.S. national variant of the ISO 646 standard. The ASCII encoding includes the upper- and lowercase letters A-Z, digits, symbols (such as &, #, and mathematical symbols), punctuation marks, and control characters. This set of 128 characters is also included in most other encodings. Short form: ASCII. See also Extended Binary Coded Decimal Interchange Code (EBCDIC) and encoding.

ASCII

See American Standard Code for Information Interchange.

ASCII collating sequence

the rules that are used by a specific ASCII encoding for sorting textual data. Sort order is determined by the location of each code point in the code page of an ASCII encoding. In the Windows Latin1 code page, the sort order of precedence is punctuation characters, numbers, uppercase characters, and lowercase characters. Because the uppercase A (code point 41) precedes the lowercase g (code point 67), A is sorted before g. See also American Standard Code for Information Interchange and EBCDIC collating sequence.

background process

in UNIX environments, a process that executes independently of the shell. When a command is executing in a background process, you can enter other commands or start other background processes without waiting for your initial command to finish executing.

batch mode

a method of executing SAS programs in which a file that contains SAS statements plus any necessary operating environment commands is submitted to the computer's

batch queue. After you submit the program, control returns to your terminal or workstation, and you can perform other tasks. Batch mode is sometimes referred to as running in the background. The program output can be written to files or printed on an output device. Under UNIX, place statements that you want to execute in a file. Then specify that file when you run SAS in the background.

buffer

an area of computer memory that is reserved for use in performing input/output (I/O) operations.

cat

a UNIX command that means concatenate. This command is commonly used to list file contents and to concatenate files.

catalog

See SAS catalog.

class name

a name that provides a way to group individual X resources together. For example, DMSboldFont and DMSFont are two separate X resources, but they are both part of the Font class.

client

(1) a computer or application that requests services, data, or other resources from a server. (2) in the X Window System, an application program that interacts with the X server and can perform tasks such as terminal emulation or window management. For example, SAS is a client because it requests windows to be created, results to be displayed, and so on.

command line

the location in any SAS windowing environment window designated with Command ==>.

command prompt

the symbol after which you enter operating system commands. In UNIX environments, different shells use different command prompts. The default command prompt for the Bourne shell and the Korn shell is \$, and the default prompt for the C shell is %.

container window

any SAS window that contains interior windows.

current directory

the directory that you are working in at any given time. When you log on, your current directory is the starting point for relative pathnames. See also working directory.

device driver

a program that controls the interaction between a computer and an external device such as a printer or a disk drive.

directory

a special type of file in UNIX operating environments that contains a group of files or other directories.

download

to copy a file from a remote host to a local host.

drag

in a graphical user interface, to move an object such as an icon or a window around on a display screen. To drag the object, you usually use a mouse button to select the object, and then move the mouse while keeping the mouse button pressed down.

EBCDIC

See Extended Binary Coded Decimal Interchange Code.

encoding

a set of characters (letters, logograms, digits, punctuation, symbols, control characters, and so on) that have been mapped to numeric values (called code points) that can be used by computers. The code points are assigned to the characters in the character set by applying an encoding method.

environment

under UNIX operating systems, the set of shell variables and their settings that you can access with a shell and with any program that the shell executes.

environment variable

in UNIX environments, a shell variable whose value or values can be accessed by any program that is executed from that shell. The shell assigns default values to some environment variables. For example, the type of terminal and the type of command prompt are specified by the default values of two environment variables.

error message

a message in the SAS log or Message window that indicates that SAS was not able to continue processing the program.

Extended Binary Coded Decimal Interchange Code

a group of 8-bit character encodings that each include up to 256 characters. EBCDIC is used on IBM mainframes and on most IBM mid-range computers, and it includes both graphic (printable) codes and control (nonprintable) codes. Short form: EBCDIC. See also American Standard Code for Information Interchange and encoding.

external file

a file that is created and maintained by a host operating system or by another vendor's software application. SAS can read data from and route output to external files. External files can contain raw data, SAS programming statements, procedure output, or output that was created by the PUT statement. A SAS data set is not an external file. See also fileref.

file descriptor

under UNIX operating systems, a nonnegative integer identifier used to refer to a file opened for reading or writing or both.

file extension

the classification of a file in a directory that identifies what type of information is stored in the file. For example, .sas7bcat is the file extension for UNIX, and .pdf is the file extension for Adobe Acrobat.

fileref

a name that is temporarily assigned to an external file or to an aggregate storage location such as a directory or a folder. The fileref identifies the file or the storage location to SAS. Under the UNIX operating system and its derivatives, you can assign a fileref with a FILENAME statement, or you can define it as an environment variable.

font

a complete set of all the characters of the same design and style. The characters in a font can be figures or symbols as well as alphanumeric characters.

foreground process

in UNIX environments, a process that executes while you wait for the command prompt to reappear. You cannot execute additional commands while the initial command is being executed in a foreground process.

home directory

in UNIX operating environments, the directory in which a user is placed after logging on. The home directory is also called the logon directory.

I/O time

an abbreviation for input/output time. I/O time is the time the computer spends on moving data from storage areas, such as disk or tape, into memory for work (input time) and moving the result out of memory to storage or to a display device, such as a terminal or a printer (output time).

icon

in windowing environments, a pictorial representation of an object. An icon usually represents an application window or is associated with an action such as printing or filing.

index

a component of a SAS data set that enables SAS to access observations in the SAS data set quickly and efficiently. The purpose of SAS indexes is to optimize WHERE-clause processing and to facilitate BY-group processing.

interactive line mode

a method of running SAS programs in which you enter one line of a SAS program at a time at the SAS session prompt. SAS processes each line immediately after you press the ENTER key. Procedure output and informative messages are returned directly to your display device.

libref

a name that is temporarily associated with a SAS library. The complete name of a SAS file consists of two words, separated by a period. The libref, which is the first word, indicates the library. The second word is the name of the specific SAS file. For example, in VLIB.NEWBDAY, the libref VLIB tells SAS which library contains the file NEWBDAY. You assign a libref with a LIBNAME statement or with an operating system command.

local host

the computer on which you use a SAS session to initiate a link with (log on to) a remote host. See also remote host.

lp

under UNIX, a line-printer command, commonly used to direct output to a printer destination via the line printer daemon.

member

a SAS file in a SAS library.

memory

the size of the work area that the central processing unit (CPU) must devote to the operations in a program.

Motif

an X Window System graphical user interface (GUI) that is used in the UNIX environment.

network

an interconnected group of computers.

path

the route through a hierarchical file system that leads to a particular file or directory.

pathname

in UNIX operating systems, a filename that specifies all of the directories that lead to a particular file in the file hierarchy.

PCL

See Printer Command Language.

PID

See process ID.

pipe

under UNIX operating systems and derivatives, the facility that links one command to another so that the standard output of one becomes the standard input of the other.

Printer Command Language

a command language that was developed by Hewlett-Packard for controlling Hewlett-Packard printers. Each PCL command consists of an escape key followed by a series of code numbers. Different versions of PCL have been developed for use with different models or types of Hewlett-Packard printers. Short form: PCL.

process ID

a unique number that is assigned to each process by the operating system. Short form: PID.

protocol

a set of rules that govern data communications between computers and peripheral devices.

menu

the list of menu items or choices that appears when you choose an item from a menu bar or from another menu. See also pop-up menu.

random access

the ability to retrieve records in a file without reading all records sequentially.

redirect

to direct output to a destination other than standard output or to read input from a source other than standard input.

remote browser server

a software agent that runs on your desktop and sends URLs to the browser to display.

remote browsing

a mechanism that is used by SAS to display HTML information (for example, help text and ODS HTML output) using a browser on your desktop.

remote host

a computer that is in a different location than your computer but which you can log on to from your computer. See also local host.

SAS catalog

a SAS file that stores many different kinds of information in smaller units called catalog entries. A single SAS catalog can contain different types of catalog entries.

SAS library

a collection of one or more files that are recognized by SAS and that are referenced and stored as a unit. Each file is a member of the library.

SAS session

See session.

SAS windowing environment

an interactive windowing interface to SAS software. In this environment you can issue commands by typing them on the command line, by pressing function keys, or by selecting items from menus or menu bars. Within one session, you can perform many different tasks, including preparing and submitting programs, viewing and printing results, and debugging and resubmitting programs.

sasauth

a SAS subprocess that performs user authentication and identification functions. The sasauth process is located in the !SASROOT/utilities/bin directory.

sasperm

a SAS subprocess that determines resource access privileges for users.

sasroot

a term that represents the name of the directory or folder in which SAS is installed at your site or on your computer.

Sasuser.Profile catalog

a SAS catalog in which SAS stores information about the attributes of your SAS windowing environment. For example, this catalog contains function-key definitions, fonts for graphics applications, window attributes, and other information that is used by interactive SAS procedures. See also SAS catalog.

sequential access

a method of file access in which the records are read or written one after the other from the beginning of the file to the end.

server

in a network, a computer that is reserved for servicing other computers in the network. Servers can provide different types of services, such as file services and communication services. Servers can also enable users to access shared resources such as disks, data, and modems. See also client.

session

a single period during which a software application is in use, from the time the application is invoked until its execution is terminated.

session gravity

in the X window interface to SAS, the resource that controls the region of the workstation display in which SAS attempts to place its windows.

shell

a UNIX command interpreter. Sample shells are sh, csh, and ksh.

shell script

a file containing commands that can be read and executed by the shell. A shell script is also called a shell procedure or a shell program.

special file

under UNIX operating systems, an interface to an input or output device. Writing to or reading from the file activates the device.

standard error

under UNIX operating systems, the destination of a program's error messages.

standard input

the primary source of data going into a command. Standard input comes from the keyboard unless it is being redirected from a file or piped from another command.

standard output

the primary destination of data coming from a command. Standard output goes to the display unless it is being redirected to a file or piped to another command.

swap

to move data or program code from a computer system's main memory to a storage device such as a hard disk, or vice versa.

swapping

See swap.

threaded kernel

the memory-resident part of a UNIX operating system that manages the computer's resources. The threaded kernel allocates memory, schedules programs for execution, monitors devices, and so on.

toggle

an option, parameter, or other mechanism that enables you to turn on or turn off a processing feature.

toolbox

a part of the SAS windowing environment in which you can place icons that you can associate with SAS commands or macros. Selecting an icon executes its associated command or string of commands.

toolset

a set of predefined tools that is associated with an application. Toolsets make it easier for individual users to customize their application toolboxes.

Universal Printing

a feature of SAS software that enables you to send SAS output to PDF, Postscript, GIF, PNG, SVG, and PCL files, as well as directly to printers. The Universal Printing system also provides many options that enable you to customize your output, and it is available in all of the operating environments that SAS supports.

working directory

the directory in which a software application is invoked.

X resource

a characteristic of a window interface, such as font type, font size, color, gravity, and window size.

X server

in an X Window System, the program that mediates access to the display, mouse, and keyboard from one or more application client programs.

X Window System

a graphical windowing system that was developed at the Massachusetts Institute of Technology.

Index

< command 469
 << command 470
 (command 466
 ((command 466
) command 467
)) command 468

Numbers

32-bit shared libraries 114

A

ABEND option
 ABORT statement 316
 ABORT statement 316
 About SAS dialog box 222
 access descriptor files 36
 aggregate syntax 75
 aliases
 font aliases 195
 ALL option
 FILENAME command 322
 LIBNAME statement 329
 WAITFOR statement 338
 ALTER= data set option 245
 alter passwords 245
 alternate SAS log
 destination for 102
 alternative configuration file 363
 ALTLOG system option 102, 355
 ALTPRINT system option 102, 356
 ANY option
 WAITFOR statement 338
 APPEND system option 356
 application workspace (AWS) 141
 ARG statement 112
 ASCII values
 position of character in ASCII collating sequence 276
 returning characters based on 258
 returning string of 262
 asynchronous tasks
 executing 14, 334
 ATTACH= email option 83
 ATTRIB statement 316
 attribute table 109
 Authentication API 424
 AUTOADD command 430

autocall libraries 288
 setting up and testing macros in 289
 specifying 399
 autoexec files 20
 configuration files versus 21
 specifying 358
 AUTOEXEC system option 358
 AUTOFLOW command 431
 automatic macro variables 285
 automatic paste buffer 155
 disabling 156
 autosave
 location of autosave file 359
 turning on and off 233
 AUTOSAVELOC system option 359
 AUTOSCROLL command 220
 AUTOSPLIT command 432
 AUTOWRAP command 433
 AWS (application workspace) 141

B

background color definitions 197
 background color resources 198
 background process 6
 batch mode 10
 destination for output 393
 executing X statements 17
 Log window destination 381
 source code, default location of 412
 BCC= email option 83
 /bin directory 424
 binary data 216
 binary values
 fixed-point 280
 positive 253, 281
 reading and writing 216
 BLK= option
 FILE command 228, 317
 FILENAME command 320
 INCLUDE command 232, 325
 INFILE command 326
 BLKSIZE= option
 FILE command 228, 317
 FILENAME command 320
 INCLUDE command 232, 325
 INFILE command 326
 blocks
 marking 153

- BMDP engine 62
 - BMDP files 62
 - accessing save files 62
 - converting system files to data sets 296
 - BOUNDS command 434
 - Bourne shell
 - defining environment variables 23
 - file descriptors 78
 - browsers 237
 - buffers
 - allocating for data set processing 246, 360
 - automatic paste buffer 155
 - command recall buffer 208
 - copying marked window contents into 237
 - paste buffers 203
 - permanent page size for output data set 246, 360
 - X synchronization 240
 - BUFNO= data set option 246
 - BUFNO system option 360
 - BUFSIZE= data set option 246
 - BUFSIZE system option 360
 - BY variables
 - dummy 311
 - BYADDR option
 - ARG statement 113
 - BYTE function 258
 - bytes
 - number sorted 405
 - BYVALUE option
 - ARG statement 113
 - \$BYVALw. format
 - MODULE arguments with 123
- C**
- C command 435
 - C language formats 120
 - C shell
 - defining environment variables 23
 - CALL MODULE routine 258
 - CALL routines 257
 - CALL SLEEP routine 260
 - CALL SYSTEM routine 15, 16, 261
 - CALLSEQ= option
 - ROUTINE statement 110
 - CAPS command 221, 437
 - case
 - for notes, warnings, and messages 387
 - in data set names 33
 - mixed case or uppercase filenames 71
 - translating to uppercase 221
 - catalog entries 291
 - CATALOG procedure 291
 - catalogs 35
 - managing entries 291
 - number to keep open 361
 - Sasuser.Profile 56
 - Sasuser.Registry 57
 - writing into transport files 299
 - CATCACHE system option 361
 - CC command 436
 - CC= email option 83
 - CCL command 439
 - CCU command 441
 - CEDA
 - accessing Version 8 or later files 47
 - CENTER system option 105
 - Change dialog box 225
 - Change Working Directory dialog box 152, 222
 - CHAR option
 - ARG statement 112
 - character expressions
 - replacing specific characters in 277
 - character strings
 - marking 153
 - character values
 - converting to/from hexadecimal 252, 280
 - CIMPORT procedure 292
 - CL command 438
 - CLEANUP option
 - SYSTASK statement 336
 - CLEANUP system option 362
 - cleanwork command 425
 - CLEAR option
 - FILENAME command 322
 - LIBNAME statement 329
 - COBOL language formats 121
 - COLLATE function 262
 - collating sequences
 - ASCII 276
 - creating 309
 - COLOR command 197, 221
 - color names 198
 - color resources 197
 - color settings 176
 - Color window 174
 - colors
 - customizing 196
 - window elements 198
 - windows 221
 - command line
 - parentheses in 343
 - toggling cursor to 230
 - command-line commands 430
 - COMMAND option
 - SYSTASK statement 334
 - command recall buffer 208
 - command window
 - configuration for 147
 - invoking with SAS sessions 207
 - opening and closing 148
 - commands 220
 - command-line 430
 - default print command 103
 - executing as asynchronous tasks 334
 - executing several 16
 - executing singly 15
 - executing statements automatically 358
 - for printing tasks 97, 382, 394
 - issuing from SAS sessions 240, 339
 - line commands 430
 - not specific to UNIX 430
 - piping to/from 79
 - piping to UNIX commands 101
 - sending output to UNIX commands 98
 - submitting for execution 261
 - synchronous versus asynchronous 14
 - text-editing 430
 - compatible computer types 42
 - characteristics of 42
 - determining in SAS 9.2 43
 - reading Version 8 or later files from 46

- Releases 6.12 through SAS 9.2 42
 - completion status of jobs 24
 - concatenating directories 51
 - concatenating filenames 75
 - CONFIG system option 363
 - configuration files 5, 20
 - autoexec files versus 21
 - creating 21
 - order of precedence 21
 - overriding system option default values 18
 - specifying 22, 363
 - connecting to X server, preventing SAS from 12
 - console log 26
 - constants
 - as MODULE function arguments 119
 - container windows 142
 - CONTENTS procedure 293
 - contrast 202
 - control keys
 - terminating SAS sessions 25
 - CONVERT procedure 296
 - copying
 - between SAS and other X clients 156
 - copying, cutting, and pasting text 155
 - external files into windows 232
 - text 155
 - window contents to buffer 237
 - CPARMS resources 198, 199
 - CPORT procedure 299
 - \$CSTRw. format
 - MODULE arguments with 122
 - CU command 440
 - CURSOR command 441
 - cursor position 230
 - cut-and-paste 155, 202, 238, 239
 - between SAS and other X clients 156
 - preserving text and attributes 204
- D**
- D command 442
 - data access
 - over NFS mounts 41
 - DATA member type 35
 - data representation 215
 - binary data 216
 - for computer types in SAS 9.2 44
 - missing values 216
 - numeric variables 215
 - data set options 241
 - summary of 241
 - data sets 34
 - allocating buffers for processing 246, 360
 - case sensitivity in names of 33
 - converting BMDP and OSIRIS system files to 296
 - converting SPSS export files to 296
 - descriptor information and data values 34
 - library of map data sets 383
 - number of bytes sorted 405
 - SAS data files 35
 - SAS views 35
 - with same name 52
 - writing into transport files 299
 - DATA step
 - sending e-mail 82, 85
 - sending UNIX command output to 79
 - stopping execution of 316
 - data values 34
 - DATALINES fileref 78
 - DATASETS procedure 300
 - date and time data 410
 - DATE system option 105
 - DBCS system options 363
 - DBMS processes
 - interrupting 29
 - DD command 442
 - decimal data
 - packed 253, 281
 - zoned 254, 283
 - DEFAULT= option
 - LENGTH statement 328
 - descriptor information 34
 - device drivers
 - for graphics output 364
 - listing all available 364
 - DEVICE system option 364
 - devices
 - assigning and deassigning filerefs 73, 268, 319
 - debugging code with DUMMY devices 73
 - DICT command 443
 - DINFO function 264
 - direct I/O 37
 - turning on 38
 - directories
 - assigning and deassigning filerefs 75, 268
 - assigning librefs to several 51
 - changing working directory 222
 - concatenating 51
 - deleting when empty 266
 - of fonts 372
 - opening 264
 - retrieving information on 264, 265
 - !SASROOT directory 421, 426
 - specifying with SAS resources 151
 - staging directory 37
 - unused Work and Utility directories 425
 - utilities directory 423
 - working directory 152
 - DISK files 73
 - disk-format libraries 60
 - disk space
 - for SORT procedure 308
 - out-of-resource conditions 362
 - display option, X command line 13
 - DLGABOUT command 222
 - DLGCDIR command 222
 - DLGENDR command 222
 - DLGFIND command 223
 - DLGFONT command 223
 - DLGOPEN command 224
 - DLGPREF command 225
 - DLGREPLACE command 225
 - DLGSAVE command 226
 - DLGSCRDUMP command 226
 - DLGSMAIL command 227
 - DMLIBASSIGN command 48
 - DOPEN function 264
 - DOPTNAME function 265
 - DOPTNUM function 265
 - double-byte character sets (DBCS) 363
 - drag and drop 156

dummy BY variables
 creating views 311
 DUMMY devices
 debugging code with 73

E

e-mail
 default protocol 158
 FILENAME statement for sending 82
 pipes for sending 81
 Send Mail dialog box 227
 sending from within SAS 158
 system for sending 366

e-mail directives
 specifying in PUT statement 84

ECHO system option 365

echoing messages to computer 365

Edit menu
 selecting text with 155

EDITCMD system option 365

editor
 starting in current window 231

EMAILSYS system option 366

ENCODING= option
 FILE command 228, 317, 318
 FILENAME command 319
 INCLUDE command 232, 325
 INFILE command 326, 327

ENCODING system option 366

ENGINE= system option 331, 367

engines 34
 multiple for a library 53

environment variables 23
 as librefs 53
 assigning filerefs 76
 defining 23, 402
 returning value of 23, 276
 specifying multiple in OPTIONS statement 399

errors
 displaying messages in uppercase 387
 library of SAS error messages 386
 print server errors 96, 97
 SAS console log 26
 specifying stdin, stdout, and stderr 409

executable modules and programs
 renaming !SASROOT directory 426
 specifying search path for 392

executing SAS statements
 autoexec file 20, 21, 358

Exit dialog box
 displaying 207
 invoking 222

exit status
 SAS jobs 24

exiting SAS
 in windowing environment 9
 preferred methods 24

Explorer window 7
 assigning librefs 49

Export as Image dialog box 226

EXPORT option
 DLGSAVE command 226

expressions
 as MODULE function arguments 119
 regular expressions in filenames 151

external files 34, 68
 assigning and deassigning filerefs 73, 268
 associating filerefs with 319
 concatenating filenames 75
 copying window contents 100, 227, 232
 deleting 266
 information items for 270, 272
 opening 150, 274
 reading with INPUT statement 326
 returning names of 274
 routing output into 102
 specifying pathnames 70
 verifying existence of 267
 verifying fileref for current SAS session 269
 wildcards in pathnames 71
 writing data from, with pipes 11

F

FDELETE function 266

FDSTART option
 ARG statement 113

FEXIST function 267

FILE command 227
 copying window contents to external files 100
 routing output with 97

file descriptors 78

file locking 38
 FILELOCKS statement option 39
 FILELOCKS system option 39
 FILELOCKS=CONTINUE 40
 FILELOCKS=NONE 39
 options for 39
 waiting to use locked files 39

file migration
See migrating files

file permissions
 changing for SAS sessions 16
 Work data library 417

file sharing
See sharing files

FILE statement 317

FILECLOSE= data set option 247

FILEEXIST function 267

FILELOCKS= option
 LIBNAME statement 332

FILELOCKS statement 39

FILELOCKS= system option 39
 set to CONTINUE 40
 set to NONE 39

FILELOCKWAITMAX= system option 370

filename extensions 36

FILENAME function 268

FILENAME statement 319
 assigning filerefs to directories 75
 assigning filerefs to external files or devices 73
 assigning filerefs to pipes 80
 concatenating filenames 75
 sending electronic mail 82
 sending output directly to printer 98
 sending output to UNIX commands 98
 specifying pathnames 70

filenames
 concatenating 75
 interpreting log messages 71
 mixed case or uppercase 71

- omitting quotation marks in 70
- regular expressions in 151
- FILEREF function 269
- filerefs 69
 - assigned by SAS 77
 - assigning and deassigning 73, 76, 268
 - assigning and deassigning, directories 75
 - assigning and deassigning, pipes 79
 - associating with files or devices 319
 - PRTFILE and PRINT commands with 98
 - reserved 78
 - verifying external files by 267
 - verifying for current SAS session 269
- files, opening 150
- fill character 229
- FILL command 229, 444
- FILTERS= option
 - DLGOPEN command 224
 - DLGSAVE command 226
- Find dialog box 157, 223
- FINFO function 270
- firewalls, and remote browsing 136
- fixed-point values 252
 - binary 280
 - positive 253
- floating-point values 254, 282
 - converting to/from hexadecimal 251, 279
- FMTSEARCH system option 371
- font aliases 195
- Font dialog box 223
- font resources 193
- FONTLIST command 230
- fonts
 - customizing 192
 - listing all available 230
 - specifying directory containing 372
 - specifying for current session 234
 - windowing environment fonts 192, 193
- Fonts dialog box 193
- FONTSLLOC system option 372
- FOOTNOTE statement 324
- footnotes 324
- FOPTNAME function 270
- FOPTNUM function 272
- foreground color definition 197
- foreground color resources 198
- foreground process 6
- format catalogs
 - order of search 371
- FORMAT= option
 - ARG statement 113
- formats 251
 - associating with variables 316
 - for binary data 216
 - for MODULE arguments 120
- forms printing 96
- FORTTRAN language formats 121
- FTP access method 74
- FULLSTIMER system option 104, 372, 410
- function key definitions 146
- functions 257

G

- GDEVICE procedure 364
- Getting Started Tutorial dialog box 208

- GRAPH windows
 - printing from 96
 - saving contents as image file 226
- graphical user interface (GUI) 140
- graphics output
 - device driver for 364
- gravity, in SAS sessions 141
- grouping SAS variables
 - as structure arguments 117, 129
- GSUBMIT command 230
- GUI (graphical user interface) 140

H

- halting execution
 - ABORT statement for 316
 - DBMS processes 29
 - SAS processes 29
 - SAS sessions 338
- hard links 66
- Help 145, 161
 - customized index files 375
 - installing manual pages 423
 - locations of Sashelp libraries 400
 - table of contents files 377
 - text and index files 376
- HELPHOST system option 374
- HELPINDEX system option 375
- HELPLLOC system option 376
- HELPPORT system option 134
- HELPTOC system option 377
- hexadecimal representation
 - converting to/from character values 252, 280
 - converting to/from real binary 251, 279
- HEXw. format 251
- \$HEXw. informat 280
- HEXw. informat 279
- highlighting windows 221
- HOME command 230
- host computer name 134
- host editor 145, 160
 - configuring SAS for support 160
 - requirements for 160
 - specifying 365
- host sort utility 310, 311
 - passing options to 403
 - passing parameters to 407
 - specifying if used 407
 - temporary files used by 406
 - when to use, based on quantity of observations 404
- HOSTEDIT command 231
 - host editor for 365

I

- I command 445
- I/O, direct 37
 - turning on 38
- IBw.d format 252
- IBw.d informat 280
- iconizing windows 143
- icons
 - toggleing ToolTip text for 236
 - user-defined 206
- IEEE Not-a-Number values 216

- image files
 - GRAPH window contents as 226
- images
 - e-mailing non-text window contents 159
- IML procedure
 - invoking shared library routines 131
- IMPORT option
 - DLGOPEN command 224
- Importing Image dialog box 224
- INCLUDE command 232
- %INCLUDE statement 325
 - concatenating filenames 75
 - specifying pathnames 70
- INDENT command 446
- index files
 - customized index files 375
 - text and index files 376
- indexes 34, 35
- INFILE statement 326
 - concatenating filenames 75
 - specifying pathnames 70
- informats 279
 - associating with variables 316
 - for binary data 216
 - for MODULE arguments 120
- INPUT option
 - ARG statement 112
- INPUT statement
 - specifying external file to read 326
- INSERT system option 378
- installing manual pages 423
- interactive line mode 9
- interface 140
- interface data files 35
- interface library engines 330
- interface SAS views 35
- interior windows 142
- interrupt menu
 - SQL procedure 28
- interrupting SAS 25
 - control keys 25
 - kill command 26
 - messages in console log 26
 - SAS processes 27
 - SAS Session Manager 25
 - SAS sessions 143
- invocation scripts 5
 - regenerating 5
- invoking SAS 4
- invoking SAS sessions
 - as foreground or background process 6
 - batch mode 10
 - in windowing environment 8
 - interactive line mode 9
 - remote host 11

J

- Java Runtime Environment options 379
- JC command 447
- JJC command 448
- JJL command 450
- JJR command 452
- JL command 449
- jobs
 - completion status of 24

- stopping execution of 316
- JR command 451
- JRE (Java Runtime Environment) options 379
- JREOPTIONS system option 379

K

- key definitions
 - creating 185
 - customizing 184
 - defining with Resource Helper 173, 191
 - function keys 146
- key translations 185
 - defining 185
 - keyboard action names 188
- keyboard action names 188
- KEYS command 453
- keysyms 186
- kill command 26
- KILL option
 - SYSTASK statement 335
- Korn shell
 - defining environment variables 23
 - file descriptors 78

L

- labels
 - associating with variables 316
- large files
 - printing with PIPE device type 102
- length of numeric variables 215, 316
 - number of bytes used 328
- LENGTH= option
 - ATTRIB statement 316
 - LIBNAME statement 329
- LENGTH statement 328
- LIBNAME function
 - assigning librefs 48
- LIBNAME statement 328
 - assigning librefs 48
 - named pipes 61
 - omitting engine names 331
 - tape access 61
- LIBNAME window
 - assigning librefs 49
- libraries 34
 - accessing disk-format libraries 60
 - accessing sequential-format libraries 60
 - accessing with librefs 50
 - assigning and deassigning librefs 328
 - default access method 367
 - default permanent name 414
 - listing characteristics of 328
 - migrating 45
 - multiple engines for 53
 - of error messages 386
 - of map data sets 383
 - printing file content descriptions 293
 - returning names of 274
 - Sasuser library 55, 396, 402
 - Work library 58, 415, 417, 425
- library engines 330
- librefs 34, 48
 - accessing permanent SAS libraries 50
 - assigned by SAS 54

- assigning 48
 - assigning and deassigning 328
 - assigning to several directories 51
 - assigning with DMLIBASSIGN command 48
 - assigning with Explorer window 49
 - assigning with LIBNAME function 48
 - assigning with LIBNAME statement 48
 - assigning with LIBNAME window 49
 - environment variables as 53
 - permanently assigning 49
 - referring to SAS files 47
 - Sashelp libref 54
 - Sasuser libref 54
 - User libref 59
 - Work libref 54
 - line commands 430
 - line size 380
 - lines per page 391
 - LINESIZE= system option 104, 105, 380
 - links 66
 - LIST option
 - FILENAME command 322
 - SYSTASK statement 334
 - locked files
 - maximum wait time for 370
 - locking files
 - See file locking
 - log
 - changing default routings 93
 - console log 26
 - content and appearance 104
 - default routings 93
 - defining destinations for 304
 - destination for 102, 355, 381
 - interpreting messages about filenames 71
 - messages to be written to 389
 - MODULE log messages 124
 - routing output from 100
 - writing all system performance statistics to 372, 410
 - writing some system performance statistics to 410, 411
 - writing system option settings to 391
 - LOG fileref 78
 - LOG= option
 - PROC PRINTTO statement 100
 - LOG system option 102, 381
 - Log window
 - controlling display of lines 220
 - echoing messages to 365
 - line size 380
 - lp command 102, 382, 394
 - changing default print command 103
 - lpr command 382, 394
 - changing default print command 103
 - LPTYPE system option 382
 - LRECL= option
 - FILE command 228, 317
 - FILENAME command 320
 - INCLUDE command 232, 325
 - INFILE command 327
- M**
- M command 454
 - macro facility 285
 - autocall libraries 399
 - memory for in-memory macro variables 388
 - memory for macro variable symbol tables 387
 - system options used in 288
 - macro files, naming 288
 - macro functions 287
 - macro statements 287
 - macros
 - setting up and testing in autocall library 289
 - mainframes
 - tapes created on 88
 - manual pages
 - installing 423
 - map data sets
 - data library containing 383
 - mapping windows 143
 - MAPS system option 383
 - MARK command
 - selecting text 155
 - marking text 153
 - MASK command 455
 - MAXARG= option
 - ROUTINE statement 110
 - MAXMEMQUERY system option 384
 - member types 36
 - memory
 - allocating for certain procedures 384
 - allocating for data set processing 246, 360
 - allocating for in-memory macro variables 388
 - allocating for macro variable symbol tables 387
 - allocating for SAS sessions 384
 - amount of real memory available 395
 - available for SORT procedure 408
 - bytes used to store variables 328
 - out-of-resource conditions 362
 - shared libraries 114
 - storing contents of memory addresses 275
 - MEMSIZE system option 384
 - menu facilities
 - defining 303
 - menus 202
 - messages file
 - writing to log 389
 - migrating files 45
 - 32-bit to 64-bit 45
 - benefits of 45
 - SAS libraries 45
 - MINARG= option
 - ROUTINE statement 110
 - missing values 216
 - CONVERT procedure 297
 - mixed case filenames 71
 - MM command 455
 - MNAME= option
 - SYSTASK statement 335
 - MOD option
 - FILE command 318
 - FILENAME command 320
 - MODEXIST function 273
 - MODULE function 109
 - constants and expressions as arguments 119
 - log messages 124
 - shared libraries, accessing efficiently 117
 - MODULE= option
 - ROUTINE statement 110
 - modules
 - calling from shared executable libraries 258
 - search path for executable modules 392

MOPEN function 274
 mouse, selecting text with 154
 MSG system option 386
 MSGCASE system option 387
 MSYMTABMAX system option 288, 387
 multivolume tape libraries 88
 MVARSIZE system option 288, 388

N

-name option, X command line 13
 named pipes
 writing sequential data sets to 61
 names
 macro files 288
 shared libraries 115
 native data files 35
 native library engines 330
 native SAS views 35
 networks
 accessing files on different networks 41
 sharing files in 40
 NEW option
 FILE command 228, 318
 FILENAME command 320
 NEWS system option 104, 389
 NFS mounts
 data access over 41
 NOSUBMIT option
 DLGOPEN command 224
 Not-a-Number values 216
 -noterminal option, X command line 12
 notes, displaying in uppercase 387
 NOTES system option 104
 NOTREQD option
 ARG statement 112
 NOWAIT option
 SYSTASK statement 335
 NUM option
 ARG statement 112
 NUMBER system option 105
 NUMBERS command 456
 numeric variables 215
 length and precision 215, 316
 number of bytes for storing 328
 storing contents of memory address in 275

O

OBS system option 390
 observations
 which to process last 390
 observations, sorting
 host sort utility, passing options to 403
 host sort utility, passing parameters to 407
 host sort versus SAS sort 407
 location of temporary files for 406
 number of observations and 404
 ODS output
 remote browsing with 134
 OLD option
 FILE command 228, 318
 FILENAME command 320
 one-level names 59
 Open dialog box 150, 224
 OPLIST system option 391

OPTIONS procedure 303
 OPTIONS statement
 overriding system option default values 19
 OSIRIS engine 63
 OSIRIS files 62
 accessing 64
 converting system files to data sets 296
 data dictionary files 63
 out-of-resource conditions 362
 output
 content and appearance 104
 default routings 93
 destination for 102, 393
 echoing messages to computer 365
 number of lines per page 391
 previewing 92
 sending directly to printer 98
 title lines for 337
 output devices
 assigning and deassigning filerefs 268
 associating filerefs with 319
 OUTPUT option
 ARG statement 112
 Output window
 controlling display of lines 220
 line size 380
 overriding system option default values 18

P

packed decimal data 253, 281
 page size
 for output SAS data set buffer 246, 360
 PAGENO= system option 105
 pages, number of lines 391
 PAGESIZE= system option 391, 104, 105
 parentheses
 in command line 343
 passwords
 assigning to SAS files 245, 248
 paste buffers 203
 manipulating text 203
 selecting 203
 submitting code from 230
 pasting text 155, 202
 patchname command 426
 PATH system option 392
 PATHNAME function 274
 pathnames 70
 character substitutions in 51
 specifying 50
 pattern resources 208
 pausing execution
 DBMS processes 29
 SAS processes 29
 PDw.d format 253
 PDw.d informat 281
 PEEKCLONG function 116, 275
 accessing returned pointer 128
 PEEKLONG function 116, 275
 performance
 amount of real memory available 395
 memory for SORT procedure 408
 out-of-resource conditions 362
 shared libraries 117

- regular expressions
 - filename selection 151
 - remote browser
 - setting up 135
 - setting up at SAS invocation 135
 - setting up during SAS session 135
 - remote browser server
 - installing 134
 - remote browsing 133
 - firewalls and 136
 - name of computer for 374
 - system options for 134
 - with ODS output 134
 - remote host
 - running SAS on 11
 - Replace dialog box 157
 - opening 157
 - options 157
 - replacing text strings 157, 277
 - REQUIRED option
 - ARG statement 112
 - reserved filerefs 78
 - RESET command 459
 - resource database 165
 - Resource Helper 172
 - modifying window colors 174
 - searching for resource definitions 176
 - setting X resources 172
 - starting 172
 - restoring windows 143
 - RETURN option
 - ABORT statement 316
 - RETURNS= option
 - ROUTINE statement 111
 - RGB values 198
 - ROUTINE statement 109
 - routines
 - calling from shared executable libraries 258
 - routing output
 - default routings 93
 - log and procedure output 93
 - piping to/from commands 79, 101
 - PRINTTO procedure 100
 - SAS logging facility messages to SYSLOGD 95
 - sending directly to printer 98
 - sending to UNIX commands 98
 - system options for 102
 - to printer 101
 - to terminal 102
 - to Universal Printer 101
 - RR command 458
 - RSASUSER system option 396
 - RSUBMIT statement 336
 - RTRACE system option 397
 - RTRACELOC system option 398
 - running SAS
 - as foreground or background process 6
 - batch mode 10
 - in windowing environment 8
 - interactive line mode 9
 - remote host 11
 - interrupting 25
 - invoking 4
 - running in background process 6
 - running in foreground process 6
 - running on remote host 11
 - terminating 25
 - SAS 9.2
 - compatibility of existing files 46
 - compatible computer types 42, 43
 - data representation 44
 - SAS/AF applications
 - previewing output from 92
 - SAS command
 - overriding system option default values 18
 - syntax 5
 - SAS/CONNECT
 - asynchronous processes 336
 - storage locations for script files 401
 - SAS console log 26
 - SAS data files 35
 - SAS files 33
 - accessing for input and update 52
 - accessing for output 52
 - accessing Version 8 or later, with CEDA 47
 - assigning passwords 245, 248
 - common types of 34
 - compatibility with SAS 9.2 46
 - concatenating filenames 75
 - creating for earlier releases 45
 - index files for SAS Help and Documentation 375
 - member types and filename extensions 36
 - migrating 45
 - one-level names for accessing permanent files 59
 - printing content descriptions 293
 - reading and writing on tape 60
 - reading from previous releases or other hosts 46
 - reading Version 6 files 46
 - reading Version 8 or later, from compatible computers 46
 - reading Version 8 or later, from incompatible computers 46
 - referring to, techniques for 47
 - referring to with librefs 47
 - resources read during execution 397, 398
 - sharing 38
 - specifying pathnames 70
 - table of contents files 377
 - SAS/GRAPH
 - device driver for graphics output 364
 - library for map data sets 383
 - SAS/GRAPH drivers
 - printing output 97
 - SAS Help and Documentation
 - customized index files 375
 - locations of Sashelp libraries 400
 - table of contents files 377
 - text and index files 376
 - SAS invocation scripts 5
 - regenerating 5
 - SAS jobs
 - completion status 24
 - SAS logging facility
 - routing messages to SYSLOGD 95
 - SAS processes
 - interrupting 27
 - terminating 29
- S**
- SAS
 - if SAS does not start 6

- SAS programs
 - submitting to batch queue 10
- SAS servers
 - ending processes for running 27
- SAS Session Manager 143
 - closing 145
 - disabling 145
 - features 143
 - interrupting SAS 25
 - starting automatically 208
 - terminating SAS 25
 - terminating SAS sessions 144
- SAS sessions
 - aborting execution 316
 - allocating memory for 384
 - batch mode 10
 - customizing 20
 - customizing color settings 196
 - file permissions for 16
 - font for current session 234
 - if SAS does not start 6
 - interactive line mode 9
 - interrupting 143
 - invocation scripts 5
 - invocation scripts, regenerating 5
 - invoking in windowing environment 8
 - invoking SAS 4
 - issuing commands from 240, 339
 - remote host 11
 - sending e-mail from 158
 - starting 4
 - starting Resource Helper from 172
 - suspending execution 338
 - terminating 144
 - verifying fileref assignment for 269
 - X command, specifying if valid 418
- SAS statements 315
 - autoexec file 358
 - including 325
 - submitting 230
- SAS ToolBox 147
- SAS views 35
- SAS window session ID 141
- SAS.altVisualID resource 207
- SASAUTOS system option 288, 289, 399
- SAS.autoSaveInterval resource 207
- SAS.autoSaveOn resource 207
- SASCBTBL attribute table 109
- SASCOLOR window 196
- SAS.confirmSASExit resource 207
- SAS.defaultCommandWindow resource 207
- SAS.directory resource 207
- Sashelp libraries 400
- Sashelp libref 54
- SASHELP system option 400
- SAS.helpBrowser resource 208
- SAS.htmlUsePassword resource 208
- SAS.insertModeOn resource 208
- SAS.keyboardTranslations resource 185, 187
- SAS.keysWindowLabels resource 188
- SAS.noDoCommandRecall resource 208
- SAS.pattern resource 208
- !SASROOT directory 421
 - renaming 426
 - utilities directory 423
- SASSCRIPT system option 401
- SAS.selectTimeout resource 208
- SAS.startSessionManager resource 208
- SAS.startupLogo resource 208
- SAS.suppressMenuIcons resource 208
- SAS.suppressTutorialDialog resource 208
- SAS.useNativeXmTextTranslations resource 208
- Sasuser data library
 - specifying if updatable 396
 - specifying name of 402
- Sasuser library 55
 - contents of 55
 - Sasuser.Prefs catalog 58
 - Sasuser.Profile catalog 56
 - Sasuser.Registry catalog 57
- Sasuser libref 54
- SASUSER system option 402
- Sasuser.Prefs file 58
- Sasuser.Profile catalog 56
 - checking for uncorrupted 56
 - how SAS accesses 56
 - when catalog does not exist 56
 - when locked or corrupted 57
- Sasuser.Registry catalog 57
 - how SAS accesses 57
- SASV9_OPTIONS environment variable
 - overriding system option default values 18
- SAS.wsaveAllExit resource 209
- Save As dialog box 226
- %SCAN macro function 287
- SCL code
 - sending e-mail 86
- search and replace 157
- searching
 - for resource definitions 176
 - for text strings (Change dialog) 225
 - for text strings (Find dialog) 223
 - format catalogs 371
 - replacing specific characters in expression 277
 - search path for executable modules 392
- security
 - assigning passwords to SAS files 245, 248
 - permissions on Work data library 417
- selecting text 153
 - techniques for 154
 - with Edit menu 155
 - with MARK command 155
 - with mouse 154
- Send Mail dialog box 158, 227
- sequential data sets
 - writing to named pipes 61
- sequential engines 60
- sequential-format libraries
 - accessing 60
 - reading and writing SAS files on tape 60
- session gravity 141
 - customizing 204
- session IDs 141
- session manager 25
 - starting automatically 208
- session workspace
 - customizing 204
- SET system option 402
- SETAUTOSAVE command 233
- SETDMSFONT command 234
- shared executable libraries
 - calling modules and routines from 258

- shared libraries 108
 - 32-bit and 64-bit considerations 114
 - efficient access 117
 - examples of accessing 126
 - formats and informats 120
 - SASCBTBL attribute table 109
- sharing files 38
 - across workstations 40
 - compatible computer types 42
 - file locking 39
 - in networks 40
- SHELL= option
 - SYSTASK statement 335
- shells
 - defining environment variables 23
 - starting remote shells 81
 - starting Resource Helper from 172
- SOCKET access method 74
- software fonts
 - listing all available 230
- SORT procedure 305
 - creating collating sequences 309
 - disk space for 308
 - host sort utility 310, 311
 - performance tuning for 308
- sort utility
 - See host sort utility
- SORTANOM system option 403
- SORTCUT system option 404
- SORTCUTP system option 405
- SORTDEV system option 406
- sorting
 - number of bytes sorted 405
- sorting observations
 - host sort utility, passing options to 403
 - host sort utility, passing parameters to 407
 - host sort versus SAS sort 407
 - location of temporary files for 406
 - number of observations and 404
- SORTPARM system option 407
- SORTPGM system option 407
- SORTSEQ= option
 - PROC SORT statement 311
- SORTSIZE= option
 - PROC SORT statement 306
- SORTSIZE system option 408
- SOURCE system option 104
- SOURCE2 system option 105
- special files 68
- SPELL command 459
- SPSS engine 64
- SPSS files 62
 - converting export files to data sets 296
 - export files 64
 - reformatting 65
- SQL procedure
 - interrupt menu for 28
- staging directory 37, 60
- standard error
 - filerefs for 77
- standard input/output
 - filerefs for 77
 - reading input from standard input 81
 - specifying if SAS should use 409
- starting SAS sessions 4
 - as foreground or background process 6
 - batch mode 10
 - in windowing environment 8
 - interactive line mode 9
 - remote host 11
 - startup logo 208
- startup logo 208
- STATUS= option
 - SYSTASK statement 335
- stderr 409
- stdin 409
- STDIO system option 409
- stdout 409
- STIMEFMT system option 410
- STIMER system option 105, 411
 - controlling format of 410
- stored program files 36
- SUBJECT= email option 83
- SUBMIT option
 - DLGOPEN command 224
- submitting SAS code 230
- submitting SAS statements 230
- suspending execution 260, 338
 - SAS sessions 338
- symbolic links 66, 115
- synchronous tasks 14
- syncsort utility
 - location of temporary files used by 406
 - passing options to 403
 - passing parameters to 407
 - specifying if used 407
 - when to use, based on quantity of observations 404
- SYSCC macro variable 285
- SYSDEVIC macro variable 286
- SYSENV macro variable 286
- %SYSEXEC macro statement 15, 16, 287
- SYSGET function 276
- %SYSGET macro function 287
- SYSIN system option 412
- SYSJOBID macro variable 286
- SYSMAXLONG macro variable 286
- SYSPRINT system option 413
- SYSRC macro variable 286
 - SYSTASK statement return code 337
 - WAITFOR statement return code 339
- SYSTASK statement 334
- system administration tools 424
- system fonts 192
- system options 342
 - adding pathnames to values for 356, 378
 - behavior or syntax specific to UNIX 342
 - customizing SAS sessions 17
 - determining how option was set 343
 - for remote browsing 134
 - listing current settings 303
 - overriding default values 18
 - parentheses in command line 343
 - routing output with 102
 - set in one place 19
 - specifying 18
 - summary of 343
 - used in macro facility 288
 - writing settings to log 391
 - writing settings to terminal 414

T

TAGSORT option
 PROC SORT statement 307

tape
 reading and writing SAS files on 60

tape storage
 multivolume 88
 positioning after file closes 247

TASKNAME= option
 SYSTASK statement 335

TC command 461

TEMP devices 74

temporary files 74

TERMINAL devices
 accessing 74
 routing output to 102

terminals
 routing output to 102

terminating execution
 DBMS processes 29
 SAS processes 29

terminating SAS 25
 control keys 25
 kill command 26
 messages in console log 26
 SAS Session Manager 25

terminating SAS sessions
 with session manager 144

TERMSTR= option
 FILE command 317
 FILENAME command 320
 INFILE command 326

text
 copying 155
 cut-and-paste 155
 selecting 153

text attributes
 transferring 161

text-editing commands 430

text editor windows
 specifying host editor 365

text windows
 e-mailing contents of 159
 printing contents of 95

TF command 462

TIMEOUT= option
 WAITFOR statement 338

title lines 337

-title option, X command line 13

TITLE statement 337

TO= email option 83

Tool Editor 179
 invoking 179
 invoking on specified toolbox 235

toolbars
 default configuration 147
 opening and closing 148

toolboxes
 adding tools 182
 button size 235
 changing appearance of 180
 closing 234
 creating 183
 customizing 147, 177, 183
 deleting tools 182
 invoking Tool Editor on 235

loading 236
 SAS ToolBox 147
 saving changes 182
 toggling ToolTip text for icons in 236
 X resources for 178

TOOLCLOSE command 234

TOOLEDIT command 235

TOOLLARGE command 235

TOOLLOAD command 236

tools
 changing attributes of 180
 deleting 182

toolsets 179
 creating 183
 customizing 183
 saving changes 182

ToolTips
 toggling for icons in toolbox 236

TOOLTIPS command 236

top-level windows 142

trace information 397, 398

TRANSLATE function 277

transport files
 restoring 292
 writing data sets and catalogs into 299

TRANPOSE= option
 ROUTINE statement 110

troubleshooting
 connection problems 12
 data access over NFS mounts 41
 key definitions 174
 out-of-resource conditions 362
 print server errors 96, 97
 starting SAS sessions 6
 transfer of text attributes 161

TS command 464

U

UNBUF option
 FILE command 318
 FILENAME command 321

undo 239

UNDO command 465

Universal Printer
 routing output to 101

Universal Printing
 as default printing mode 96
 previewing output with 92
 printing from GRAPH windows 96

UNIX Authentication API 424

UNIX editor
 starting in current window 231

UPDATE option
 ARG statement 112

updating Sasuser data library 396

uppercase filenames 71

UPRINTER device type 101

user-defined icons 206
 locating 206
 X resources for specifying 206

User libref 59
 assigning 59

USER system option 414

utilities directory 423
 deleting when unused 425

V

variable names
 CONVERT procedure 297

variables
 grouping as structure arguments 117, 129
 number of bytes for storing 328
 numeric 215

VERBOSE system option 414

VERIFY option
 DLGOPEN command 224
 DLGSAVE command 226

Version 6
 reading files 46

Version 8
 accessing files with CEDA 47
 reading files 46
 reading files, from incompatible computer types 46

view engines 329

VIEW member type 35

views
 creating with dummy BY variables 311

virtual keysyms 186

W

WAIT option
 SYSTASK statement 335

warnings
 displaying in uppercase 387

WBROWSE command 237

WCOPY command 237

WCUT command 238

WDEF command 238

Web browsers
 invoking 237

wildcards
 in pathnames 71

window colors
 modifying with Resource Helper 174

window element definitions 198

window managers 141

window sizes 204

window types 142

windowing environment 7
 command window configuration 147
 copying text 155
 customizing 141
 cut-and-paste 155
 drag and drop 156
 e-mail 158
 fonts 192
 fonts, default 193
 Help 161
 interface 140
 invoking SAS in 8
 opening files 150
 SAS Session Manager 25, 143, 208
 SAS ToolBox 147
 search and replace 157
 selecting text 153
 toolbar configuration 147
 types of windows 142
 working directory 152

X resources for customizing 164, 209

X Window System and 141

windows

 color and highlighting of 221

 container windows 142

 copying contents into external files 227

 copying contents to external files 100

 copying external file contents into 232

 copying marked contents to buffer 237

 e-mailing contents of 159

 iconizing 143

 interior windows 142

 line size 380

 mapping 143

 pasting buffer contents into 239

 positioning 142

 printing contents of 95, 98

 resizing 142, 238

 restoring 143

 specifying host editor 365

 top-level 142

 types of 142

Work library 58
 deleting unused Work directories 425
 initializing 417
 multiple directories 59
 name of 415
 setting permissions when created 417

Work libref 54

WORK system option 415

working directory 152
 changing 152, 222

WORKINIT system option 417

WORKPERMS system option 417

workstations
 sharing files 40

WPASTE command 239

write passwords
 assigning to SAS files 248

writing binary data 216

writing SAS files on tape 60

writing sequential data sets
 to named pipes 61

WSAVE ALL command 209

WUNDO command 239

X

X command 240
 executing a single command 15
 executing several commands 16
 specifying if valid 418

X command line options
 unsupported 14

X resources 165
 controlling toolbox behavior 178
 customizing 165
 modifying with Preferences dialog box 167
 searching for resource definitions 176
 setting with Resource Helper 172
 specifying user-defined icons 206
 summary of 209
 syntax 165

X server
 preventing SAS connection to 12

X statement 339
 executing a single command 15
 executing in batch mode 17

- executing several commands 16
- X synchronization 240
- X window managers 141
- X Window System 141
 - interface 140
 - SAS window session ID 141
 - session gravity 141
 - specifying options 13
 - window managers 141
 - window types 142

- XCMD system option 418
- xrm option, X command line 13
- XSYNC command 240

Z

- ZDw.d format 254
- ZDw.d informat 283
- zoned decimal data 254, 283

Your Turn

We welcome your feedback.

- If you have comments about this book, please send them to **yourturn@sas.com**. Include the full title and page numbers (if applicable).
- If you have comments about the software, please send them to **suggest@sas.com**.

SAS® Publishing Delivers!

Whether you are new to the work force or an experienced professional, you need to distinguish yourself in this rapidly changing and competitive job market. SAS® Publishing provides you with a wide range of resources to help you set yourself apart. Visit us online at support.sas.com/bookstore.

SAS® Press

Need to learn the basics? Struggling with a programming problem? You'll find the expert answers that you need in example-rich books from SAS Press. Written by experienced SAS professionals from around the world, SAS Press books deliver real-world insights on a broad range of topics for all skill levels.

support.sas.com/saspress

SAS® Documentation

To successfully implement applications using SAS software, companies in every industry and on every continent all turn to the one source for accurate, timely, and reliable information: SAS documentation. We currently produce the following types of reference documentation to improve your work experience:

- Online help that is built into the software.
- Tutorials that are integrated into the product.
- Reference documentation delivered in HTML and PDF – free on the Web.
- Hard-copy books.

support.sas.com/publishing

SAS® Publishing News

Subscribe to SAS Publishing News to receive up-to-date information about all new SAS titles, author podcasts, and new Web site features via e-mail. Complete instructions on how to subscribe, as well as access to past issues, are available at our Web site.

support.sas.com/spn



sas

THE
POWER
TO KNOW®

