

## CHAPTER

## 1

# Getting Started with SAS in UNIX Environments

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

<i>Changing the File Permissions for Your SAS Session</i>	15
<i>Executing X Statements in Batch Mode</i>	15
<i>Customizing Your SAS Registry Files</i>	16
<i>Customizing Your SAS Session Using Configuration and Autoexec Files</i>	16
<i>Introduction to Configuration and Autoexec Files</i>	16
<i>Differences between Configuration and Autoexec Files</i>	16
<i>Creating a Configuration File</i>	17
<i>Order of Precedence for SAS Configuration Files</i>	17
<i>Specifying a Configuration File for SAS to Use</i>	18
<i>Customizing Your SAS Session Using System Options</i>	18
<i>Ways to Specify a SAS System Option</i>	18
<i>Overriding the Default Value for a System Option</i>	19
<i>How SAS Processes System Options Set in One Place</i>	20
<i>How SAS Processes System Options Set in Multiple Places</i>	20
<i>Precedence for Processing System Options</i>	20
<i>Defining Environment Variables in UNIX Environments</i>	21
<i>What Is an Environment Variable?</i>	21
<i>How to Define an Environment Variable for Your Shell</i>	21
<i>Bourne and Korn Shells</i>	21
<i>C Shell</i>	22
<i>Displaying the Value of an Environment Variable</i>	22
<i>Determining the Completion Status of a SAS Job in UNIX Environments</i>	22
<i>Interrupting or Terminating Your SAS Session in UNIX Environments</i>	22
<i>Preferred Methods of Exiting SAS</i>	22
<i>Additional Methods for Interrupting or Terminating SAS</i>	23
<i>Using Control Keys</i>	23
<i>Using the SAS Session Manager</i>	23
<i>Using the UNIX kill Command</i>	24
<i>Messages in the SAS Console Log</i>	24
<i>Ending a Process That Is Running as a SAS Server</i>	24
<i>Ending a SAS Process on a Relational Database</i>	25
<i>How to Interrupt a SAS Process</i>	25
<i>Example: Interrupt Menu for PROC SQL</i>	26
<i>How to Terminate a SAS Process</i>	26
<i>What Happens When You Interrupt a SAS Process and the Underlying DBMS Process</i>	27

---

## Starting SAS Sessions in UNIX Environments

---

### Invoking SAS

The command that you use to invoke your SAS session is defined during the SAS installation process and is added to the list of commands that are recognized by the operating environment. Ask your system administrator what the command is that invokes SAS at your site. At many sites, the command to invoke SAS is simply **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 “Interrupting or Terminating Your SAS Session in UNIX Environments” on page 22). Also, if you cannot stop your session, contact your system administrator; do not turn off your machine, especially if your machine is part of a network.  $\triangle$

---

## 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 17, “System Options under UNIX,” on page 311 and “X Command Line Options” on page 11 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: Invoking 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 6.

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

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

To run a SAS program and pass parameters to it, enter

```
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.

---

## 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 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. 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 tasks without waiting for your initial command to execute. You can run SAS in the foreground or in the background.

*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 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 your keyboard, mouse, pull-down menus, pop-up menus, and icons. The windowing environment includes, but is not limited to, the Explorer, Program Editor, Output, Log, and Results windows.

Your SAS session may 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).

For more information about using the windowing environment, see Chapter 2, “Working in the SAS Windowing Environment,” on page 29.

*Note:* If you are not using an X display, then you need to invoke SAS in interactive line mode using the NODMS system option. For more information, see “Interactive Line Mode in UNIX Environments” on page 7.  $\triangle$

### 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 SAS Help and Documentation.

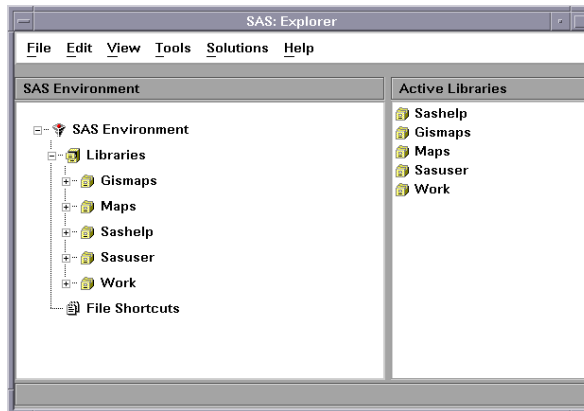
---

## Invoking SAS in the Windowing Environment

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

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

```
sas -explorer
```



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

```
sas -dms
```

- You can use the DMSEXP system option to open the Program Editor, Output, Log, and Results windows and the Explorer:

```
sas -dmsexp
```

SAS also opens the toolbox from which you can open additional SAS windows. For more information about the toolbox, see to Chapter 2, “Working in the SAS Windowing Environment,” on page 29.

---

## Exiting SAS in the Windowing Environment

To end your SAS session, enter the BYE or ENDSAS command on the command line or select



from the pull-down menu of the session that you want to end.

---

# Interactive Line Mode in UNIX Environments

---

## Introduction to Interactive Line Mode

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

- 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.

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 23) or by issuing the ENDSAS statement:

```
endsas;
```

---

# Batch Mode in UNIX Environments

---

## Introduction to Running SAS in Batch Mode

To run SAS in batch mode, you specify your SAS application name in the SAS 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.rpt is the file containing the SAS statements to be executed, and you want to specify the NODATE and LINESIZE system options, you would enter

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

You do not need to specify the SYSIN option as with some other platforms.

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.rpt from your shell prompt as follows:

```
$ at 2am
sas weekly.rpt
<control-D>
warning: commands will be executed using /usr/bin/sh
job 8400.a at Wed Jun 11 02:00:00 2003
$
```

If you create a file that contains the SAS command necessary to run your program, for example CmdFile.sas, then you can enter the following command at your shell prompt:

```
at 2am < cmdfile.sas
```

SAS sends the output to a file that has the same name as the program and an extension of .lst, and the log goes to a file with an extension of .log. Both of these files are written to your current directory. Refer to the man pages for these commands for more information on submitting jobs to the batch queue. For more details on routing output, see Chapter 6, “Printing and Routing Output,” on page 153.

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

---

## Writing Data from an External File Using Pipes

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

```
INFILE STDIN;
```

Issue this command to have MyProg.sas read data from MyData:

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

For details about using external files, see Chapter 5, “Using External Files and Devices,” on page 131. See also “File Descriptors in the Bourne and Korn Shells” on page 140 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 machine's name. For example, if your machine 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. Most systems control this by using the **xhost** client. Other systems control access through a session manager. 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 +
```

To run this command automatically each time you log in, enter this command in a file named **.xhost**.

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

- 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 do this 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 **setenv** command:

```
setenv DISPLAY green:0.0
```

The DISPLAY variable will be used by all Xclients on the system.

*Note:* To determine the shell for your system, type **ps** at the command prompt or check the value of the SHELL environment variable.  $\triangle$

- b Use the **-display** 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
```

*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.  $\triangle$



---

## 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** option on the command line.

---

## 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: Cannot open Xdisplay. Check display name/server access authorization.
```

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 might need to contact your 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, you might enter

```
-display wizard:0.0
```

or

```
-display 10.22.1.1:0
```

**-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
```

**-noterminal**

specifies that you do not want to display the SAS session. You must specify this option to generate a graph in batch mode. For more information, see “Running SAS/GRAPH Programs” in *SAS/GRAPH Reference, Volumes 1 and 2*.

*Note:* To prevent SAS from attempting to connect to the X server, unset the `DISPLAY` environment variable and use the `-noterminal` option on the command line. △

**-title *string***

specifies the title up to six characters long for your SAS session window. 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 55 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 88 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 88 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 305 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 239 and “Macro Statements in UNIX Environments” on page 265 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 on quoting functions. △

### Example 1: Executing a UNIX Command 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 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 239.

## 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. 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 15.

---

## 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. △

## 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:

---

\* The shell used depends on the SHELL environment variable.

```
%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 14.

---

## Starting a Shell

If you are not running in the SAS windowing environment, you can start a shell by not specifying any UNIX commands in the X statement:

```
X;
```

SAS responds with

```
Enter 'exit' to return to your SAS session.
```

SAS then starts a shell.

Enter any UNIX commands. When you are ready to return to the SAS session, enter the `exit` command.

Even if you changed directories while in the shell, you will be in the same directory as when you started the shell.

---

## 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 the file permissions from within SAS, issue the `umask mask` command on the X statement. The `umask mask` command sets the permissions for any new file that you create. The value of `mask` can be either numeric or symbolic. For more information about this command, see the man page for `umask`.

When SAS executes the `umask mask` command, it changes the file permissions of the current SAS session, but it does not change the permissions in the parent shell. Any subsequent file that you create during this SAS session will inherit the permissions that you specified.

---

## 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. Some of the data stored in the registry includes

- the libraries and file shortcuts that SAS assigns at startup. 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 `regstry.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.  $\triangle$

---

## Customizing Your SAS Session Using Configuration and Autoexec Files

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 3, “Customizing the SAS Windowing Environment,” on page 53.

---

### Introduction to Configuration and Autoexec Files

You can customize your SAS session by defining configuration and/or 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 data libraries. Refer to *SAS Language Reference: Dictionary* for a complete description of system options.)

The configuration file (for SAS 9.1) 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 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 SAS Configuration Files” on page 17.

### 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 data 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 SAS Configuration Files” on page 17 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 – except 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 Close the configuration file.

---

## Order of Precedence for SAS Configuration Files

SAS is shipped with a default configuration file in the **!SASROOT** directory. Your SAS Installation Representative 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 397.)
- 2 `.sasv9.cfg` in your home directory. (Notice the leading period.)
- 3 `sasv9.cfg` in your home directory.
- 4 `sasv9.cfg` in your current directory.
- 5 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 *SAS System Configuration Guide for UNIX*.

SAS uses the last value it encounters for a system option. 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.

---

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

---

## 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 SAS Configuration Files” on page 17.

*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 21. 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 21. For example, in the Korn shell, you would use:

```
export SASV9_CONFIG=filename
```

*filename* is the name of a file containing 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.

---

## Customizing Your SAS Session Using System Options

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 3, “Customizing the SAS Windowing Environment,” on page 53.

---

### 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 17.3 on page 385 shows where each SAS system option can be specified.

Any options that do not affect the initialization of the 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 SAS Configuration Files” on page 17) 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 21.

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='-xwait -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 on 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 17.3 on page 385.

OPTIONS statement in an autoexec file

Specify SAS system options in an OPTIONS statement in a 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 4.6 on page 115 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. See “APPEND System Option” on page 315 and “INSERT System Option” on page 337 for more information.

## 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 will override only the setting for the same system option in your configuration file.

### 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 SAS Configuration Files” on page 17.

For example, if a configuration file specifies NOSTIMER, you can override the setting in the SAS command.

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 315 and “INSERT System Option” on page 337 for more information.

## Defining Environment Variables in UNIX Environments

### What Is an Environment Variable?

*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.

*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. △

### 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 into one command:

```
$ export scname=phonelist
```

If you change the value of **scname**, 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, 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 sname 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.

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 warning(s)	1
SAS System issued error(s)	2
User issued ABORT statement	3
User issued ABORT RETURN statement	4
User issued ABORT ABEND statement	5
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 may not print what the user expects because the code is interpreted as a signed byte.

---

## Interrupting or Terminating Your SAS Session in UNIX Environments

---

### Preferred Methods of Exiting SAS

The preferred methods of exiting a SAS session are the following:

- select



if you are using SAS in the windowing environment

- use **endsas**;
- enter **BYE** in the command line
- use CTRL+D if you are using SAS in interactive line mode.

## Additional Methods for Interrupting or Terminating SAS

In addition to the preferred methods, you can terminate SAS in the following ways:

- Press the interrupt or quit control key
- Use the session manager
- Enter the UNIX **kill** command.

Although you can terminate SAS using these techniques, you should try one of the four preferred techniques listed first.

### Using Control Keys

Control keys enable you to interrupt or terminate your session by simply 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 **stty** UNIX man page before using it. 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 (^) stands for the CTRL key. In this example, CTRL+C is the interrupt key and CTRL+\ is the quit key.

### Using the SAS Session Manager

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

*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. △

## Using the UNIX kill Command

*Note:* Only use the **kill** command after you have tried all other methods to exit your SAS session.  $\triangle$

The **kill** command sends an interrupt or quit 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 to determine the process identification number (*pid*) of the SAS session that you want to interrupt or terminate.

For example, suppose you want to stop a SAS job running in the background. First, issue the **ps** command to determine the PID of the SAS job.

```
> ps
  PID TTY      TIME COMMAND
  2103 ttyu0    0:00 motifxsa
  2111 ttyu0    0:01 sas
  2116 ttyu0    0:00 ps
  3856 ttyu2    0:03 ksh
```

Four PIDs appear, but only one is for a SAS program. (**motifxsa** is the SAS session manager. See “The SAS Session Manager (motifxsassm) in UNIX” on page 33 for more information.) Therefore, to send the interrupt signal to that SAS program, you would issue this command:

```
kill -SIGINT 2111
```

SAS replies with a prompt:

```
Press Y to cancel submitted statements,
N to continue.
```

For more information, refer to 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.

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

---

## 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 if 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 24. △

---

## 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 either case, however, 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.

*Note:* This control key sequence will not interrupt a SAS process if you issue it in a SAS window, such as the Program Editor or SAS log. △

- Use the **-SIGINT** option in the **kill** command. For more information, see “Using the UNIX kill Command” on page 24
- If you are running the SAS windowing environment in the foreground, then click **Interrupt** in the SAS Session Manager.
- If you are running an interactive SAS process in the background, then you must click **Interrupt** in the SAS Session Manager. You cannot use a control key sequence to interrupt the SAS process.

The interrupt signal is sent to the DATA step or procedure that is currently executing for interpretation. The interrupt signal is interpreted differently by DATA steps and procedures. The actions that you can take appear in the interrupt menu. Since 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 options available for PROC SORT differ from those available for PROC SQL.
- 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. If the DATA step or procedure does not respond to the interrupt in a given amount of time, then SAS does not terminate the current DATA step or procedure. △

## 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 1.2** 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. 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"> <li><input type="checkbox"/> If you select <b>C</b>, then the menu will disappear, but since 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.</li> <li><input type="checkbox"/> If you select <b>Q</b>, then SAS cancels the current query even if it is on a relational database. SAS continues processing the next statement.</li> <li><input type="checkbox"/> If you select <b>S</b>, then all of the PROC SQL statements that you submitted are cancelled.</li> <li><input type="checkbox"/> If you select <b>X</b>, SAS exits the current SQL procedure and starts processing the next statement in the submit block.</li> </ul>
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 <b>C</b> will end the current query. If you are running a short query and SAS has the control, then selecting <b>C</b> will cause the interrupt menu to disappear and the current query will continue. To determine whether or not 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 the SAS System	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 in interactive line mode or in batch mode using a foreground process, use the **-SIGTERM** option in the **kill** command. For more information, see “Using the UNIX kill Command” on page 24:
- If you are running the SAS windowing environment in the foreground, then click **Terminate** in the SAS Session Manager.
- If you are running an interactive SAS process in the background, then you must click **Terminate** in the SAS Session Manager. 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. △

*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. △

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.

