



THE
POWER
TO KNOW.

SAS[®] 9.2

Language Reference

Concepts

Second Edition



The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2010. *SAS® 9.2 Language Reference: Concepts, Second Edition*. Cary, NC: SAS Institute Inc.

SAS® 9.2 Language Reference: Concepts, Second Edition

Copyright © 2010, SAS Institute Inc., Cary, NC, USA

ISBN 978-1-60764-448-4

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, May 2010

1st printing, May 2010

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

PART 1 **SAS System Concepts** **1**

Chapter 1 △ **Essential Concepts of Base SAS Software** **3**

- What Is SAS? **3**
- Overview of Base SAS Software **4**
- Components of the SAS Language **4**
- Ways to Run Your SAS Session **7**
- Customizing Your SAS Session **9**
- Conceptual Information about Base SAS Software **10**

Chapter 2 △ **SAS Processing** **13**

- Definition of SAS Processing **13**
- Types of Input to a SAS Program **14**
- The DATA Step **15**
- The PROC Step **16**

Chapter 3 △ **Rules for Words and Names in the SAS Language** **17**

- Words in the SAS Language **17**
- Names in the SAS Language **20**

Chapter 4 △ **SAS Variables** **25**

- Definition of SAS Variables **26**
- SAS Variable Attributes **26**
- Ways to Create Variables **28**
- Variable Type Conversions **32**
- Aligning Variable Values **32**
- Automatic Variables **33**
- SAS Variable Lists **34**
- Dropping, Keeping, and Renaming Variables **36**
- Numeric Precision in SAS Software **38**

Chapter 5 △ **Missing Values** **49**

- Definition of Missing Values **49**
- Creating Special Missing Values **50**
- Order of Missing Values **51**
- When Variable Values Are Automatically Set to Missing by SAS **52**
- When Missing Values Are Generated by SAS **53**
- Working with Missing Values **55**

Chapter 6 △ **Expressions** **57**

- Definitions for SAS Expressions **58**
- Examples of SAS Expressions **58**

SAS Constants in Expressions	58
SAS Variables in Expressions	64
SAS Functions in Expressions	65
SAS Operators in Expressions	65
Chapter 7 △ Dates, Times, and Intervals	77
About SAS Date, Time, and Datetime Values	77
About Date and Time Intervals	88
Chapter 8 △ Error Processing and Debugging	99
Types of Errors in SAS	99
Error Processing in SAS	108
Debugging Logic Errors in the DATA Step	114
Chapter 9 △ SAS Output	115
Definitions for SAS Output	115
Routing SAS Output	117
The SAS Log	117
Traditional SAS Listing Output	126
Changing the Destination of the Log and the Output	128
Output Delivery System	128
Chapter 10 △ BY-Group Processing in SAS Programs	131
Definition of BY-Group Processing	131
References for BY-Group Processing	131
Chapter 11 △ WHERE-Expression Processing	133
Definition of WHERE-Expression Processing	133
Where to Use a WHERE Expression	134
Syntax of WHERE Expression	135
Combining Expressions by Using Logical Operators	144
Improving Performance of WHERE Processing	146
Processing a Segment of Data That Is Conditionally Selected	146
Deciding Whether to Use a WHERE Expression or a Subsetting IF Statement	148
Chapter 12 △ Optimizing System Performance	151
Definitions for Optimizing System Performance	151
Collecting and Interpreting Performance Statistics	152
Techniques for Optimizing I/O	153
Techniques for Optimizing Memory Usage	158
Techniques for Optimizing CPU Performance	158
Calculating Data Set Size	160
Chapter 13 △ Support for Parallel Processing	161
Definition of Parallel Processing	161
Threaded I/O	161
Threaded Application Processing	162

Chapter 14 △ **Monitoring Performance Using Application Response Measurement (ARM)** 163

What Is ARM? 163

Chapter 15 △ **The SAS Registry** 165

Introduction to the SAS Registry 166

Managing the SAS Registry 168

Configuring Your Registry 176

Chapter 16 △ **Printing with SAS** 181

Universal Printing 183

Managing Printing Tasks with the Universal Printing User Interface 185

Configuring Universal Printing with Programming Statements 204

Forms Printing 210

Using TrueType Fonts with Universal Printers and SAS/GRAPH Devices in SAS 9.2 211

Creating Scalable Vector Graphics Using Universal Printing 223

Creating Portable Network Graphics Using Universal Printing 241

PART 2 **Windowing Environment Concepts** 245

Chapter 17 △ **Introduction to the SAS Windowing Environment** 247

What Is the SAS Windowing Environment? 247

Main Windows in the SAS Windowing Environment 248

Navigating in the SAS Windowing Environment 255

Getting Help in SAS 259

List of SAS Windows and Window Commands 260

Chapter 18 △ **Managing Your Data in the SAS Windowing Environment** 263

Introduction to Managing Your Data in the SAS Windowing Environment 263

Managing Data with SAS Explorer 264

Working with VIEWTABLE 268

Subsetting Data by Using the WHERE Expression 273

Exporting a Subset of Data 276

Importing Data into a Table 279

PART 3 **DATA Step Concepts** 283

Chapter 19 △ **DATA Step Processing** 285

Why Use a DATA Step? 285

Overview of DATA Step Processing 286

Processing a DATA Step: A Walkthrough 289

About DATA Step Execution 294

About Creating a SAS Data Set with a DATA Step 299

Writing a Report with a DATA Step 303

The DATA Step and ODS 310

Chapter 20	△ Reading Raw Data	311
	Definition of Reading Raw Data	311
	Ways to Read Raw Data	312
	Types of Data	312
	Sources of Raw Data	315
	Reading Raw Data with the INPUT Statement	316
	How SAS Handles Invalid Data	321
	Reading Missing Values in Raw Data	322
	Reading Binary Data	323
	Reading Column-Binary Data	325
Chapter 21	△ BY-Group Processing in the DATA Step	329
	Definitions for BY-Group Processing	329
	Syntax for BY-Group Processing	330
	Understanding BY Groups	331
	Invoking BY-Group Processing	332
	Determining Whether the Data Requires Preprocessing for BY-Group Processing	332
	Preprocessing Input Data for BY-Group Processing	333
	How the DATA Step Identifies BY Groups	334
	Processing BY-Groups in the DATA Step	336
Chapter 22	△ Reading, Combining, and Modifying SAS Data Sets	341
	Definitions for Reading, Combining, and Modifying SAS Data Sets	343
	Overview of Tools	343
	Reading SAS Data Sets	344
	Combining SAS Data Sets: Basic Concepts	345
	Combining SAS Data Sets: Methods	356
	Error Checking When Using Indexes to Randomly Access or Update Data	382
Chapter 23	△ Using DATA Step Component Objects	391
	Introduction to DATA Step Component Objects	392
	Using the Hash Object	392
	Using the Hash Iterator Object	406
	Using the Java Object	409
Chapter 24	△ Array Processing	429
	Definitions for Array Processing	429
	A Conceptual View of Arrays	430
	Syntax for Defining and Referencing an Array	431
	Processing Simple Arrays	432
	Variations on Basic Array Processing	436
	Multidimensional Arrays: Creating and Processing	437
	Specifying Array Bounds	439
	Examples of Array Processing	441

Chapter 25	△ SAS Libraries	447
Definition of a SAS Library		447
Library Engines		449
Library Names		449
Library Concatenation		452
Permanent and Temporary Libraries		454
SAS System Libraries		455
Sequential Data Libraries		457
Tools for Managing Libraries		458
Chapter 26	△ SAS Data Sets	461
Definition of a SAS Data Set		461
Descriptor Information for a SAS Data Set		462
Data Set Names		463
Data Set Lists		464
Special SAS Data Sets		465
Sorted Data Sets		466
Tools for Managing Data Sets		472
Viewing and Editing SAS Data Sets		473
Chapter 27	△ SAS Data Files	475
Definition of a SAS Data File		477
Differences between Data Files and SAS Views		477
Understanding an Audit Trail		479
Understanding Generation Data Sets		487
Understanding Integrity Constraints		493
Understanding SAS Indexes		505
Compressing Data Files		526
Chapter 28	△ SAS Views	529
Definition of SAS Views		529
Benefits of Using SAS Views		530
When to Use SAS Views		531
DATA Step Views		531
PROC SQL Views		535
Comparing DATA Step and PROC SQL Views		536
SAS/ACCESS Views		536
Chapter 29	△ Stored Compiled DATA Step Programs	539
Definition of a Stored Compiled DATA Step Program		539
Uses for Stored Compiled DATA Step Programs		539
Restrictions and Requirements for Stored Compiled DATA Step Programs		540
How SAS Processes Stored Compiled DATA Step Programs		540
Creating a Stored Compiled DATA Step Program		541
Executing a Stored Compiled DATA Step Program		542
Differences between Stored Compiled DATA Step Programs and DATA Step Views		546

Examples of DATA Step Programs	546
Chapter 30 △ DICTIONARY Tables	549
Definition of a DICTIONARY Table	549
How to View DICTIONARY Tables	549
Chapter 31 △ SAS Catalogs	553
Definition of a SAS Catalog	553
SAS Catalog Names	553
Tools for Managing SAS Catalogs	554
Profile Catalog	555
Catalog Concatenation	557
Chapter 32 △ About SAS/ACCESS Software	561
Definition of SAS/ACCESS Software	561
Dynamic LIBNAME Engine	561
SQL Procedure Pass-Through Facility	563
ACCESS Procedure and Interface View Engine	564
DBLOAD Procedure	565
Interface DATA Step Engine	565
Chapter 33 △ Processing Data Using Cross-Environment Data Access (CEDA)	567
Definition of Cross-Environment Data Access (CEDA)	567
Advantages of CEDA	568
SAS File Processing with CEDA	568
Alternatives to Using CEDA	573
Creating Files in a Foreign Data Representation	574
Examples of Using CEDA	574
Chapter 34 △ SAS 9.2 Compatibility with SAS Files from Earlier Releases	577
Introduction to Version Compatibility	577
Comparing SAS System 9 to Earlier Releases	577
Using SAS Library Engines	578
Chapter 35 △ File Protection	581
Definition of a Password	581
Assigning Passwords	582
Removing or Changing Passwords	584
Using Password-Protected SAS Files in DATA and PROC Steps	584
How SAS Handles Incorrect Passwords	585
Assigning Complete Protection with the PW= Data Set Option	585
Encoded Passwords	586
Using Passwords with Views	586
SAS Data File Encryption	588
Chapter 36 △ SAS Engines	589
Definition of a SAS Engine	589

Specifying an Engine	589
How Engines Work with SAS Files	590
Engine Characteristics	591
About Library Engines	593
Special-Purpose Engines	596

Chapter 37 △ **SAS File Management** 599

Improving Performance of SAS Applications	599
Moving SAS Files between Operating Environments	599
Repairing Damaged SAS Files	600

Chapter 38 △ **External Files** 605

Definition of External Files	605
Referencing External Files Directly	606
Referencing External Files Indirectly	606
Referencing Many External Files Efficiently	608
Referencing External Files with Other Access Methods	608
Working with External Files	609

PART 5 **Industry Protocols Used in SAS** 611

Chapter 39 △ **The SMTP E-Mail Interface** 613

Sending E-Mail through SMTP	613
System Options That Control SMTP E-Mail	613
Statements That Control SMTP E-mail	614

Chapter 40 △ **Universal Unique Identifiers** 617

Universal Unique Identifiers and the Object Spawner	617
Using SAS Language Elements to Assign UUIDs	619

Chapter 41 △ **Internet Protocol Version 6 (IPv6)** 621


Overview of IPv6	621
IPv6 Address Format	621
Examples of IPv6 Addresses	622
Fully Qualified Domain Names (FQDN)	622

PART 6 **Appendix** 625

Appendix 1 △ **Recommended Reading** 627


Recommended Reading	627
---------------------	-----

Index	629
-------	-----



SAS System Concepts

<i>Chapter 1</i>	Essential Concepts of Base SAS Software	3
<i>Chapter 2</i>	SAS Processing	13
<i>Chapter 3</i>	Rules for Words and Names in the SAS Language	17
<i>Chapter 4</i>	SAS Variables	25
<i>Chapter 5</i>	Missing Values	49
<i>Chapter 6</i>	Expressions	57
<i>Chapter 7</i>	Dates, Times, and Intervals	77
<i>Chapter 8</i>	Error Processing and Debugging	99
<i>Chapter 9</i>	SAS Output	115
<i>Chapter 10</i>	BY-Group Processing in SAS Programs	131
<i>Chapter 11</i>	WHERE-Expression Processing	133
<i>Chapter 12</i>	Optimizing System Performance	151
<i>Chapter 13</i>	Support for Parallel Processing	161
<i>Chapter 14</i>	Monitoring Performance Using Application Response Measurement (ARM)	163
<i>Chapter 15</i>	The SAS Registry	165
<i>Chapter 16</i>	Printing with SAS	181



CHAPTER

1

Essential Concepts of Base SAS Software

<i>What Is SAS?</i>	3
<i>Overview of Base SAS Software</i>	4
<i>Components of the SAS Language</i>	4
<i>SAS Files</i>	4
<i>SAS Data Sets</i>	5
<i>External Files</i>	5
<i>Database Management System Files</i>	6
<i>SAS Language Elements</i>	6
<i>SAS Macro Facility</i>	6
<i>Ways to Run Your SAS Session</i>	7
<i>Starting a SAS Session</i>	7
<i>Different Types of SAS Sessions</i>	7
<i>SAS Windowing Environment</i>	7
<i>Interactive Line Mode</i>	8
<i>Noninteractive Mode</i>	8
<i>Batch Mode</i>	9
<i>Objectserver Mode</i>	9
<i>Customizing Your SAS Session</i>	9
<i>Setting Default System Option Settings</i>	9
<i>Executing Statements Automatically</i>	9
<i>Customizing the SAS Windowing Environment</i>	10
<i>Conceptual Information about Base SAS Software</i>	10
<i>SAS System Concepts</i>	10
<i>DATA Step Concepts</i>	10
<i>SAS Files Concepts</i>	10

What Is SAS?

SAS is a set of solutions for enterprise-wide business users and provides a powerful fourth-generation programming language for performing tasks such as these:

- data entry, retrieval, and management
- report writing and graphics
- statistical and mathematical analysis
- business planning, forecasting, and decision support
- operations research and project management
- quality improvement
- applications development

With Base SAS software as the foundation, you can integrate with SAS many SAS business solutions that enable you to perform large scale business functions, such as data warehousing and data mining, human resources management and decision support, financial management and decision support, and others.

Overview of Base SAS Software

The core of the SAS System is Base SAS software, which consists of the following:

<i>DATA step</i>	a programming language that you use to manipulate and manage your data.
<i>SAS procedures</i>	software tools for data analysis and reporting.
<i>macro facility</i>	a tool for extending and customizing SAS software programs and for reducing text in your programs.
<i>DATA step debugger</i>	a programming tool that helps you find logic problems in DATA step programs.
<i>Output Delivery System (ODS)</i>	a system that delivers output in a variety of easy-to-access formats, such as SAS data sets, procedure output files, or Hypertext Markup Language (HTML).
<i>SAS windowing environment</i>	an interactive, graphical user interface that enables you to easily run and test your SAS programs.

This document, along with *SAS Language Reference: Dictionary*, covers only the SAS language. For a complete guide to Base SAS software functionality, also see these documents: *SAS Output Delivery System: User's Guide*, *SAS National Language Support (NLS): Reference Guide*, *Base SAS Procedures Guide*, *SAS Metadata LIBNAME Engine: User's Guide*, *SAS XML LIBNAME Engine: User's Guide*, *SAS Macro Language: Reference*, and the *Getting Started with SAS* online tutorial. The SAS windowing environment is described in the online Help.

Components of the SAS Language

SAS Files

When you work with SAS, you use files that are created and maintained by SAS, as well as files that are created and maintained by your operating environment, and that are not related to SAS. Files with formats or structures known to SAS are referred to as *SAS files*. All SAS files reside in a *SAS library*.

The most commonly used SAS file is a *SAS data set*. A SAS data set is structured in a format that SAS can process. Another common type of SAS file is a *SAS catalog*. Many different types of information that are used in a SAS job are stored in SAS catalogs, such as instructions for reading and printing data values, or function key settings that you use in the SAS windowing environment. A *SAS stored program* is a type of SAS file that contains compiled code that you create and save for repeated use.

Operating Environment Information: In some operating environments, a SAS data library is a physical relationship among files; in others, it is a logical relationship. Refer

to the SAS documentation for your operating environment for details about the characteristics of SAS libraries in your operating environment. △

SAS Data Sets

There are two types of SAS data sets:

- SAS data file
- SAS view


A *SAS data file* both describes and physically stores your data values. A *SAS view*, on the other hand, does not actually store values. Instead, it is a query that creates a logical SAS data set that you can use as if it were a single SAS data set. It enables you to look at data stored in one or more SAS data sets or in other vendors' software files. SAS views enable you to create logical SAS data sets without using the storage space required by SAS data files.

A SAS data set consists of the following:

- descriptor information
- data values

The descriptor information describes the contents of the SAS data set to SAS. The *data values* are data that has been collected or calculated. They are organized into rows, called observations, and columns, called variables. An *observation* is a collection of data values that usually relate to a single object. A *variable* is the set of data values that describe a given characteristic. The following figure represents a SAS data set.

Figure 1.1 Representation of a SAS Data Set



Usually, an observation is the data that is associated with an entity such as an inventory item, a regional sales office, a client, or a patient in a medical clinic. Variables are characteristics of these entities, such as sale price, number in stock, and originating vendor. When data values are incomplete, SAS uses a *missing value* to represent a missing variable within an observation.

External Files

Data files that you use to read and write data, but which are in a structure unknown to SAS, are called *external files*. External files can be used for storing

- raw data that you want to read into a SAS data file
- SAS program statements
- procedure output

Operating Environment Information: Refer to the SAS documentation for your operating environment for details about the characteristics of external files in your operating environment. △

Database Management System Files

SAS software is able to read and write data to and from other vendors' software, such as many common database management system (DBMS) files. In addition to Base SAS software, you must license the SAS/ACCESS software for your DBMS and operating environment.

SAS Language Elements

The SAS language consists of statements, expressions, options, formats, and functions similar to those of many other programming languages. In SAS, you use these elements within one of two groups of SAS statements:

- DATA steps
- PROC steps

A *DATA step* consists of a group of statements in the SAS language that can

- read data from external files
- write data to external files
- read SAS data sets and SAS views
- create SAS data sets and SAS views

Once your data is accessible as a SAS data set, you can analyze the data and write reports by using a set of tools known as SAS procedures.

A group of procedure statements is called a *PROC step*. SAS procedures analyze data in SAS data sets to produce statistics, tables, reports, charts, and plots, to create SQL queries, and to perform other analyses and operations on your data. They also provide ways to manage and print SAS files.

You can also use global SAS statements and options outside of a DATA step or PROC step.

SAS Macro Facility

Base SAS software includes the SAS Macro Facility, a powerful programming tool for extending and customizing your SAS programs, and for reducing the amount of code that you must enter to do common tasks. Macros are SAS files that contain compiled macro program statements and stored text. You can use macros to automatically generate SAS statements and commands, write messages to the SAS log, accept input, or create and change the values of macro variables. For complete documentation, see *SAS Macro Language: Reference*.

Ways to Run Your SAS Session

Starting a SAS Session

You start a SAS session with the SAS command, which follows the rules for other commands in your operating environment. In some operating environments, you include the SAS command in a file of system commands or control statements; in others, you enter the SAS command at a system prompt or select SAS from a menu.

Different Types of SAS Sessions

You can run SAS in any of several ways that might be available for your operating environment:

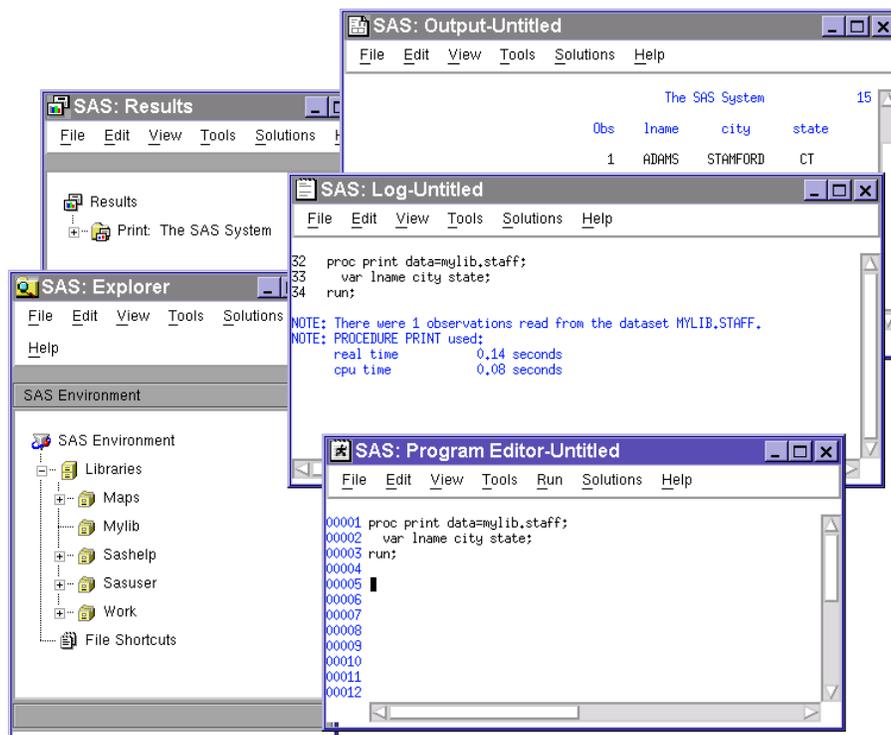
- SAS windowing environment
- interactive line mode
- noninteractive mode
- batch (or background) mode

In addition, SAS/ASSIST software provides a menu-driven system for creating and running your SAS programs. For more information about SAS/ASSIST, see *Getting Started with SAS/ASSIST*.

SAS Windowing Environment

In the *SAS windowing environment*, you can edit and execute programming statements, display the SAS log, procedure output, and online Help, and more. The following figure shows the SAS windowing environment.

Figure 1.2 SAS Windowing Environment



In the Explorer window, you can view and manage your SAS files, which are stored in libraries, and create shortcuts to external files. The Results window helps you navigate and manage output from SAS programs that you submit; you can view, save, and manage individual output items. You use the Program Editor, Log, and Output windows to enter, edit, and submit SAS programs, view messages about your SAS session and programs that you submit, and browse output from programs that you submit. For more detailed information about the SAS windowing environment, see Chapter 17, “Introduction to the SAS Windowing Environment,” on page 247.

Interactive Line Mode

In *interactive line mode*, you enter program statements in sequence in response to prompts from the SAS System. DATA and PROC steps execute when

- a RUN, QUIT, or a semicolon on a line by itself after lines of data are entered
- another DATA or PROC statement is entered
- the ENDSAS statement is encountered

By default, the SAS log and output are displayed immediately following the program statements.

Noninteractive Mode

In *noninteractive mode*, SAS program statements are stored in an external file. The statements in the file execute immediately after you issue a SAS command referencing the file. Depending on your operating environment and the SAS system options that you use, the SAS log and output are either written to separate external files or displayed.

Operating Environment Information: Refer to the SAS documentation for your operating environment for information about how these files are named and where they are stored. △

Batch Mode

You can run SAS jobs in *batch mode* in operating environments that support batch or background execution. Place your SAS statements in a file and submit them for execution along with the control statements and system commands required at your site.

When you submit a SAS job in batch mode, one file is created to contain the SAS log for the job, and another is created to hold output that is produced in a PROC step or, when directed, output that is produced in a DATA step by a PUT statement.

Operating Environment Information: Refer to the SAS documentation for your operating environment for information about executing SAS jobs in batch mode. Also, see the documentation specific to your site for local requirements for running jobs in batch and for viewing output from batch jobs. △

Objectserver Mode

When SAS runs in objectserver mode, SAS runs as an IOM server. Some examples of SAS IOM servers are the SAS Metadata Server, the SAS Workspace Server, the SAS Stored Process Server, and the SAS OLAP Server. For information about running SAS in objectserver mode, see *SAS Intelligence Platform: Application Server Administration Guide*.

Customizing Your SAS Session

Setting Default System Option Settings

You can use a *configuration file* to store system options with the settings that you want. When you invoke SAS, these settings are in effect. SAS system options determine how SAS initializes its interfaces with your computer hardware and the operating environment, how it reads and writes data, how output appears, and other global functions.

By placing SAS system options in a configuration file, you can avoid having to specify the options every time that you invoke SAS. For example, you can specify the NODATE system option in your configuration file to prevent the date from appearing at the top of each page of your output.

Operating Environment Information: See the SAS documentation for your operating environment for more information about the configuration file. In some operating environments, you can use both a system-wide and a user-specific configuration file. △

Executing Statements Automatically

To execute SAS statements automatically each time you invoke SAS, store them in an *autoexec file*. SAS executes the statements automatically after the system is initialized. You can activate this file by specifying the AUTOEXEC= system option.

Any SAS statement can be included in an autoexec file. For example, you can set report titles, footnotes, or create macros or macro variables automatically with an autoexec file.

Operating Environment Information: See the SAS documentation for your operating environment for information on how autoexec files should be set up so that they can be located by SAS. Δ

Customizing the SAS Windowing Environment

You can customize many aspects of the SAS windowing environment and store your settings for use in future sessions. With the SAS windowing environment, you can

- change the appearance and sorting order of items in the Explorer window
- customize the Explorer window by registering member, entry, and file types
- set up favorite folders
- customize the toolbar
- set fonts, colors, and preferences.

See the SAS online Help for more information and for additional ways to customize your SAS windowing environment.

Conceptual Information about Base SAS Software

SAS System Concepts

SAS system-wide concepts include the building blocks of SAS language: rules for words and names, variables, missing values, expressions, dates, times, and intervals, and each of the six SAS language elements — data set options, formats, functions, informats, statements, and system options.

SAS system-wide concepts also include introductory information that helps you begin to use SAS, including information about the SAS log, SAS output, error processing, WHERE processing, and debugging. Information about SAS processing prepares you to write SAS programs. Information on how to optimize system performance as well as how to monitor performance.

DATA Step Concepts

Understanding essential DATA step concepts can help you construct DATA step programs effectively. These concepts include how SAS processes the DATA step, how to read raw data to create a SAS data set, and how to write a report with a DATA step.


More advanced concepts include how to combine and modify information once you have created a SAS data set, how to perform BY-group processing of your data, how to use array processing for more efficient programming, and how to create stored compiled DATA step programs.

SAS Files Concepts

SAS file concepts include advanced topics that are helpful for advanced applications, though not strictly necessary for writing simple SAS programs. These topics include the

elements that comprise the physical file structure that SAS uses, including data libraries, data files, SAS views, catalogs, file protection, engines, and external files.

Advanced topics for data files include the audit trail, generation data sets, integrity constraints, indexes, and file compression. In addition, these topics include compatibility issues with earlier releases and how to process files across operating environments.



CHAPTER

2

SAS Processing


<i>Definition of SAS Processing</i>	13
<i>Types of Input to a SAS Program</i>	14
<i>The DATA Step</i>	15
<i>DATA Step Output</i>	15
<i>The PROC Step</i>	16
<i>PROC Step Output</i>	16

Definition of SAS Processing

SAS processing is the way that the SAS language reads and transforms input data and generates the kind of output that you request. The DATA step and the procedure (PROC) step are the two steps in the SAS language. Generally, the DATA step manipulates data, and the PROC step analyzes data, produces output, or manages SAS files. These two types of steps, used alone or combined, form the basis of SAS programs.

The following figure shows a high-level view of SAS processing using a DATA step and a PROC step. The figure focuses primarily on the DATA step.

Figure 2.1 SAS Processing



You can use different types of data as input to a DATA step. The DATA step includes SAS statements that you write, which contain instructions for processing the data. As each DATA step in a SAS program is compiling or executing, SAS generates a log that contains processing messages and error messages. These messages can help you debug a SAS program.

Types of Input to a SAS Program

You can use different sources of input data in your SAS program:

SAS data sets	can be one of two types:
SAS data files	store actual data values. A SAS data file consists of a descriptor portion that describes the data in the file, and a data portion.
SAS views	contain references to data stored elsewhere. A SAS view uses descriptor information and data from other files. It enables you to dynamically combine data from various sources, without using storage space to create a new data set. SAS views consist of DATA step views, PROC SQL views, and SAS/ACCESS views. In most cases, you can use a SAS view as if it were a SAS data file.
	For more information, see Chapter 27, “SAS Data Files,” on page 475, and Chapter 28, “SAS Views,” on page 529.
Raw data	specifies unprocessed data that have not been read into a SAS data set. You can read raw data from two sources:
External files	contain records comprised of formatted data (data is arranged in columns) or free-formatted data (data that are not arranged in columns).
Instream data	is data included in your program. You use the DATALINES statement at the beginning of your data to identify the instream data.
	For more information about raw data, see Chapter 20, “Reading Raw Data,” on page 311.
Remote access	allows you to read input data from nontraditional sources such as a TCP/IP socket or a URL. SAS treats this data as if it were coming from an external file. SAS allows you to access your input data remotely in the following ways:
SAS catalog	specifies the access method that enables you to reference a SAS catalog as an external file.
FTP	specifies the access method that enables you to use File Transfer Protocol (FTP) to read from or write to a file from any host computer that is connected to a network with an FTP server running.
TCP/IP socket	specifies the access method that enables you to read from or write to a Transmission Control Protocol/Internet Protocol (TCP/IP) socket.

URL specifies the access method that enables you to use the uniform resource locator (URL) to read from and write to a file from any host computer that is connected to a network with a URL server running.

For more information about accessing data remotely, see FILENAME, CATALOG Access Method; FILENAME, FTP Access Method; FILENAME, SOCKET Access Method; and FILENAME, URL Access Method statements in the Statements section of *SAS Language Reference: Dictionary*.

The DATA Step

The DATA step processes input data. In a DATA step, you can create a SAS data set, which can be a SAS data file or a SAS view. The DATA step uses input from raw data, remote access, assignment statements, or SAS data sets. The DATA step can, for example, compute values, select specific input records for processing, and use conditional logic. The output from the DATA step can be of several types, such as a SAS data set or a report. You can also write data to the SAS log or to an external data file. For more information about DATA step processing, see Chapter 19, “DATA Step Processing,” on page 285.

DATA Step Output

The output from the DATA step can be a SAS data set or an external file such as the program log, a report, or an external data file. You can also update an existing file in place, without creating a separate data set. Data must be in the form of a SAS data set to be processed by many SAS procedures. You can create the following types of DATA step output:

SAS log	contains a list of processing messages and program errors. The SAS log is produced by default.		
SAS data file	is a SAS data set that contains two parts: a data portion and a data descriptor portion.		
SAS view	is a SAS data set that uses descriptor information and data from other files. SAS views allow you to dynamically combine data from various sources without using disk space to create a new data set. While a SAS data file actually contains data values, SAS views contain only references to data stored elsewhere. SAS views are of member type VIEW. In most cases, you can use a SAS view as if it were a SAS data file.		
External data file	contains the results of DATA step processing. These files are data or text files. The data can be records that are formatted or free-formatted.		
Report	contains the results of DATA step processing. Although you usually generate a report by using a PROC step, you can generate the following two types of reports from the DATA step: <table> <tr> <td>Procedure output file</td> <td>contains printed results of DATA step processing, and usually contains headers and page breaks.</td> </tr> </table>	Procedure output file	contains printed results of DATA step processing, and usually contains headers and page breaks.
Procedure output file	contains printed results of DATA step processing, and usually contains headers and page breaks.		


HTML file contains results that you can display on the World Wide Web. This type of output is generated through the Output Delivery System (ODS). For complete information about ODS, see *SAS Output Delivery System: User's Guide*.

The PROC Step

The PROC step consists of a group of SAS statements that call and execute a procedure, usually with a SAS data set as input. Use PROCs to analyze the data in a SAS data set, produce formatted reports or other results, or provide ways to manage SAS files. You can modify PROCs with minimal effort to generate the output you need. PROCs can also perform functions such as displaying information about a SAS data set. For more information about SAS procedures, see *Base SAS Procedures Guide*.

PROC Step Output

The output from a PROC step can provide univariate descriptive statistics, frequency tables, crosstabulation tables, tabular reports consisting of descriptive statistics, charts, plots, and so on. Output can also be in the form of an updated data set. For more information about procedure output, see *Base SAS Procedures Guide* and the *SAS Output Delivery System: User's Guide*.



CHAPTER

3

Rules for Words and Names in the SAS Language

<i>Words in the SAS Language</i>	17
<i>Definition of Word</i>	17
<i>Types of Words or Tokens</i>	18
<i>Placement and Spacing of Words in SAS Statements</i>	19
<i>Spacing Requirements</i>	19
<i>Examples</i>	19
<i>Names in the SAS Language</i>	20
<i>Definition of a SAS Name</i>	20
<i>Rules for User-Supplied SAS Names</i>	20
<i>Rules for Most SAS Names</i>	20
<i>Rules for SAS Variable Names</i>	22
<i>SAS Name Literals</i>	23
<i>Definition of SAS Name Literals</i>	23
<i>Important Restrictions</i>	23
<i>Avoiding Errors When Using Name Literals</i>	24
<i>Examples</i>	24

Words in the SAS Language

Definition of Word

A *word* or *token* in the SAS programming language is a collection of characters that communicates a meaning to SAS and which cannot be divided into smaller units that can be used independently. A word can contain a maximum of 32,767 characters.

A word or token ends when SAS encounters one of the following:

- the beginning of a new token
- a blank after a name or a number token
- the ending quotation mark of a literal token

Each word or token in the SAS language belongs to one of four categories:

- names
- literals
- numbers
- special characters

Types of Words or Tokens

There are four basic types of words or tokens:

name

is a series of characters that begin with a letter or an underscore. Later characters can include letters, underscores, and numeric digits. A name token can contain up to 32,767 characters. In most contexts, however, SAS names are limited to a shorter maximum length, such as 32 or 8 characters. See Table 3.1 on page 21. Here are some examples of name tokens:

- data
- _new
- yearcutoff
- year_99
- descending
- _n_

literal

consists of 1 to 32,767 characters enclosed in single or double quotation marks. Here are some examples of literals:

- 'Chicago'
- "1990-91"
- 'Amelia Earhart'
- 'Amelia Earhart' 's plane'
- "Report for the Third Quarter"

Note: The surrounding quotation marks identify the token as a literal, but SAS does not store these marks as part of the literal token. Δ

number

in general, is composed entirely of numeric digits, with an optional decimal point and a leading plus or minus sign. SAS also recognizes numeric values in the following forms as number tokens: scientific (E-) notation, hexadecimal notation, missing value symbols, and date and time literals. Here are some examples of number tokens:

- 5683
- 2.35
- 0b0x
- 5
- 5.4E-1
- '24aug90'd

special character

is usually any single keyboard character other than letters, numbers, the underscore, and the blank. In general, each special character is a single token, although some two-character operators, such as ** and <=, form single tokens. The blank can end a name or a number token, but it is not a token. Here are some examples of special-character tokens:

- =
- ;
- '
- +
- @
- /

Placement and Spacing of Words in SAS Statements

Spacing Requirements

- 1 You can begin SAS statements in any column of a line and write several statements on the same line.
- 2 You can begin a statement on one line and continue it on another line, but you cannot split a word between two lines.
- 3 A blank is not treated as a character in a SAS statement unless it is enclosed in quotation marks as a literal or part of a literal. Therefore, you can put multiple blanks any place in a SAS statement where you can put a single blank, with no effect on the syntax.
- 4 The rules for recognizing the boundaries of words or tokens determine the use of spacing between them in SAS programs. If SAS can determine the beginning of each token due to cues such as operators, you do not need to include blanks. If SAS cannot determine the beginning of each token, you must use blanks. See “Examples” on page 19.

Although SAS does not have rigid spacing requirements, SAS programs are easier to read and maintain if you consistently indent statements. The examples illustrate useful spacing conventions.

Examples

- In this statement, blanks are *not required* because SAS can determine the boundary of every token by examining the beginning of the next token:

```
total=x+y;
```

The first special-character token, the equal sign, marks the end of the name token **total**. The plus sign, another special-character token, marks the end of the name token **x**. The last special-character token, the semicolon, marks the end of the **y** token. Though blanks are not needed to end any tokens in this example, you can add them for readability, as shown here:

```
total = x + y;
```

- This statement *requires blanks* because SAS cannot recognize the individual tokens without them:

```
input group 15 room 20;
```

Without blanks, the entire statement up to the semicolon fits the rules for a name token: it begins with a letter or underscore, contains letters, digits, or

underscores thereafter, and is less than 32,767 characters long. Therefore, this statement requires blanks to distinguish individual name and number tokens.

Names in the SAS Language

Definition of a SAS Name

A *SAS name* is a name token that represents

- variables
- SAS data sets
- formats or informats
- SAS procedures
- options
- arrays
- statement labels
- SAS macros or macro variables
- SAS catalog entries
- librefs or filerefs
- component objects

There are two types of names in SAS:

- names of elements of the SAS language
- names supplied by SAS users

Rules for User-Supplied SAS Names

Rules for Most SAS Names

Note: The rules are more flexible for SAS variable names than for other language elements. See “Rules for SAS Variable Names” on page 22. \triangle

- The length of a SAS name depends on the element it is assigned to. Many SAS names can be 32 characters long; others have a maximum length of 8.
- The first character must be an English letter (A, B, C, . . . , Z) or underscore (_). Subsequent characters can be letters, numeric digits (0, 1, . . . , 9), or underscores.
- You can use upper or lowercase letters.
- Blanks cannot appear in SAS names.
- Special characters, except for the underscore, are not allowed. In filerefs only, you can use the dollar sign (\$), pound sign (#), and at sign (@).
- SAS reserves a few names for automatic variables and variable lists, SAS data sets, and librefs.
 - When creating variables, do not use the names of special SAS automatic variables (for example, `_N_` and `_ERROR_`) or special variable list names (for example, `_CHARACTER_`, `_NUMERIC_`, and `_ALL_`).

- When associating a libref with a SAS library, do not use these:
 - SASHELP
 - SASMSG
 - SASUSER
 - WORK
- When you create SAS data sets, do not use these names:
 - _NULL_
 - _DATA_
 - _LAST_
- When assigning a fileref to an external file, do not use the filename SASCAT.
- When you create a macro variable, do not use names that begin with **sys**.

Table 3.1 Maximum Length of User-Supplied SAS Names

SAS Language Element	Maximum Length
Arrays	32
CALL routines	16
Catalog entries	32
DATA step statement labels	32
DATA step variable labels	256
DATA step variables	32
DATA step windows	32
Engines	8
Filerefs	8
Formats, character	31
Formats, numeric	32
Functions	16
Generation data sets	28
Informats, character	30
Informats, numeric	31
Librefs	8
Macro variables	32
Macro windows	32
Macros	32
Members of SAS libraries (SAS data sets, SAS views, catalogs, indexes) except for generation data sets	32
Passwords	8

SAS Language Element	Maximum Length
Procedure names (first eight characters must be unique, and cannot begin with “SAS”)	16
SCL variables	32

Rules for SAS Variable Names

The rules for SAS variable names have expanded to provide more functionality. The setting of the VALIDVARNAME= system option determines what rules apply to the variables that you can create and process in your SAS session as well as to variables that you want to read from existing data sets. The VALIDVARNAME= option has three settings (V7, UPCASE, and ANY), each with varying degrees of flexibility for variable names:

V7

is the default setting.

Variable name rules are as follows:

- SAS variable names can be up to 32 characters in length.
- The first character must begin with an English letter or an underscore. Subsequent characters can be English letters, numeric digits, or underscores.
- A variable name cannot contain blanks.
- A variable name cannot contain any special characters other than the underscore.
- A variable name can contain mixed case. Mixed case is remembered and used for presentation purposes only. (SAS stores the case used in the first reference to a variable.) When SAS processes variable names, however, it internally uppercases them. You cannot, therefore, use the same letters with different combinations of lowercase and uppercase to represent different variables. For example, **cat**, **Cat**, and **CAT** all represent the same variable.
- You do not assign the names of special SAS automatic variables (such as **_N_** and **_ERROR_**) or variable list names (such as **_NUMERIC_**, **_CHARACTER_**, and **_ALL_**) to variables.

UPCASE

is the same as V7, except that variable names are uppercased, as in earlier versions of SAS.

ANY

- SAS variable names can be up to 32 characters in length.
- The name can contain special and multi-byte characters not to exceed 32 bytes.
- The name cannot contain any null bytes.
- A name can contain blanks. The leading blanks are preserved, but the trailing blanks are ignored.
- A name must contain at least one character. A name with all blanks is not permitted.
- The name can start with or contain any characters, including blanks.

Note: If you use any characters other than the ones that are valid when the VALIDVARNAME system option is set to V7 (letters of the Latin alphabet, numerals, or underscores), then you must express the variable name as a *name literal* and you must set VALIDVARNAME=ANY.

If you use either the percent sign (%) or the ampersand (&), then you must use single quotation marks in the name literal in order to avoid interaction with the SAS Macro Facility. See “SAS Name Literals” on page 23. △

- A variable name can contain mixed-case letters. Mixed-case letters are stored and used for presentation purposes only. (SAS stores the case that is used in the first reference to a variable.) When SAS processes variable names, however, it internally uppercases them. Therefore, you cannot use the same letters with different combinations of lowercase and uppercase to represent different variables. For example, **cat**, **Cat**, and **CAT** all represent the same variable.

Warning: The intent of the VALIDVARNAME=ANY system option is to enable compatibility with other DBMS variable (column) naming conventions, such as allowing embedded blanks and national characters. Throughout SAS, using the name literal syntax with variable names that exceed the 32-byte limit or have excessive embedded quotation marks might cause unexpected results.

SAS Name Literals

Definition of SAS Name Literals

A SAS *name literal* is a name token that is expressed as a string within quotation marks, followed by the upper- or lowercase letter *n*. Name literals enable you to use special characters (including blanks) that are not otherwise allowed in SAS names when you specify a SAS data set or a variable. Blanks between the closing quotation mark and the *n* are not valid when you specify a name literal. When the name literal of a data set or variable contains any characters that are not allowed when VALIDVARNAME=V7, then you must set the VALIDVARNAME= system option to ANY. Note that even if you set the system option to ANY, the V6 engine does not support names that have intervening blanks.

Name literals are especially useful for expressing DBMS column and table names that contain special characters.

The following is an example of a VAR statement and a name literal:

```
var 'a b'n;
```

The following is an example of a VAR statement with variables A and B:

```
var a b;
```

Important Restrictions

- You can use a name literal only for variables, statement labels, and DBMS column and table names.
- When the name literal of a variable or DBMS column contains any characters that are not allowed when VALIDVARNAME=V7, then you must set the system option VALIDVARNAME=ANY.

Note: For more details about the VALIDVARNAME=ANY system option, see *SAS Language Reference: Dictionary*. △

- If you use either the percent sign (%) or the ampersand (&), then you must use single quotation marks in the name literal in order to avoid interaction with the SAS Macro Facility.

- When the name literal of a DBMS table or column contains any characters that are not valid for SAS rules, then you might need to specify a SAS/ACCESS LIBNAME statement option.

Note: For more details and examples about the SAS/ACCESS LIBNAME statement and about using DBMS table and column names that do not conform to SAS naming conventions, see *SAS/ACCESS for Relational Databases: Reference*. \triangle

- In a quoted string, SAS preserves and uses leading blanks, but SAS ignores and trims trailing blanks.


Avoiding Errors When Using Name Literals

For information on how to avoid creating name literals in error, see “Avoiding a Common Error with Constants” on page 63.

Examples

Examples of SAS name literals are

- `input 'Bob's Asset Number'n;`
- `input ''Bob's Asset Number''n;`
- `libname foo SAS/ACCESS-engine-name
SAS/ACCESS-engine-connection-options;
data foo.'My Table'n;`
- `input 'Amount Budgeted'n 'Amount Spent'n
'Amount Difference'n;`



CHAPTER

4

SAS Variables

<i>Definition of SAS Variables</i>	26
<i>SAS Variable Attributes</i>	26
<i>Ways to Create Variables</i>	28
<i>Overview</i>	28
<i>Using an Assignment Statement</i>	29
<i>Reading Data with the INPUT Statement in a DATA Step</i>	30
<i>Specifying a New Variable in a FORMAT or an INFORMAT Statement</i>	30
<i>Specifying a New Variable in a LENGTH Statement</i>	30
<i>Specifying a New Variable in an ATTRIB Statement</i>	31
<i>Using the IN= Data Set Option</i>	31
<i>Variable Type Conversions</i>	32
<i>Aligning Variable Values</i>	32
<i>Automatic Variables</i>	33
<i>SAS Variable Lists</i>	34
<i>Definition</i>	34
<i>Numbered Range Lists</i>	34
<i>Name Range Lists</i>	35
<i>Name Prefix Lists</i>	35
<i>Special SAS Name Lists</i>	35
<i>Dropping, Keeping, and Renaming Variables</i>	36
<i>Using Statements or Data Set Options</i>	36
<i>Using the Input or Output Data Set</i>	36
<i>Order of Application</i>	37
<i>Examples of Dropping, Keeping, and Renaming Variables</i>	37
<i>Numeric Precision in SAS Software</i>	38
<i>How SAS Stores Numeric Values</i>	38
<i>Troubleshooting Problems Regarding Floating-Point Representation</i>	39
<i>Overview</i>	39
<i>Floating-Point Representation on IBM Mainframes</i>	39
<i>Floating-Point Representation on OpenVMS</i>	41
<i>Floating-Point Representation Using the IEEE Standard</i>	42
<i>Precision Versus Magnitude</i>	42
<i>Computational Considerations of Fractions</i>	42
<i>Numeric Comparison Considerations</i>	43
<i>Storing Numbers with Less Precision</i>	43
<i>Truncating Numbers and Making Comparisons</i>	45
<i>Determining How Many Bytes Are Needed to Store a Number Accurately</i>	45
<i>Double-Precision Versus Single-Precision Floating-Point Numbers</i>	46
<i>Transferring Data between Operating Systems</i>	46

Definition of SAS Variables

variables

are containers that you create within a program to store and use character and numeric values. Variables have attributes, such as name and type, that enable you to identify them and that define how they can be used.

character variables

are variables of type *character* that contain alphabetic characters, numeric digits 0 through 9, and other special characters.

numeric variables

are variables of type *numeric* that are stored as floating-point numbers, including dates and times.

numeric precision

refers to the degree of accuracy with which numeric variables are stored in your operating environment.

SAS Variable Attributes

A SAS variable has the attributes that are listed in the following table:

Table 4.1 Variable Attributes

Variable Attribute	Possible Values	Default Value
Name	Any valid SAS name. See Chapter 3, “Rules for Words and Names in the SAS Language,” on page 17.	None
Type ¹	Numeric and character	Numeric
Length ¹	2 to 8 bytes ² 1 to 32,767 bytes for character	8 bytes for numeric and character
Format	See “Formats” in <i>SAS Language Reference: Dictionary</i>	BEST12. for numeric, \$w. for character
Informat	See “Informats” in <i>SAS Language Reference: Dictionary</i>	w.d for numeric, \$w. for character
Label	Up to 256 characters	None
Position in observation	1- <i>n</i>	None
Index type	NONE, SIMPLE, COMPOSITE, or BOTH	None

¹ If they are not explicitly defined, a variable’s type and length are implicitly defined by its first occurrence in a DATA step.

² The minimum length is 2 bytes in some operating environments, 3 bytes in others. See the SAS documentation for your operating environment.

Note: Starting with SAS 9.1, the maximum number of variables can be greater than 32,767. The maximum number depends on your environment and the file’s attributes. For example, the maximum number of variables depends on the total length of all the variables and cannot exceed the maximum page size. \triangle

You can use the CONTENTS procedure, or the functions that are named in the following definitions, to obtain information about a variable's attributes:

name

identifies a variable. A variable name must conform to SAS naming rules. A SAS name can be up to 32 characters long. The first character must be a letter (A, B, C, . . . , Z) or underscore (_). Subsequent characters can be letters, digits (0 to 9), or underscores. Note that blanks are not used. Mixed case variables are used. See Chapter 3, "Rules for Words and Names in the SAS Language," on page 17 for more details on mixed case variables.

The names `_N_`, `_ERROR_`, `_FILE_`, `_INFILE_`, `_MSG_`, `_IORC_`, and `_CMD_` are reserved for the variables that are generated automatically for a DATA step. Note that SAS products use variable names that start and end with an underscore; it is recommended that you do not use names that start and end with an underscore in your own applications. See "Automatic Variables" on page 33 for more information.

To determine the value of this attribute, use the `VNAME` or `VARNAME` function.

Note: The rules for variable names that are described in this section apply when the `VALIDVARNAME=` system option is set to `VALIDVARNAME=V7`, which is the default setting. Other rules apply when this option is set differently. See Chapter 3, "Rules for Words and Names in the SAS Language," on page 17 for more information. Δ

type

identifies a variable as numeric or character. Within a DATA step, a variable is assumed to be numeric unless character is indicated. Numeric values represent numbers, can be read in a variety of ways, and are stored in floating-point format. Character values can contain letters, numbers, and special characters and can be from 1 to 32,767 characters long.

To determine the value of this attribute, use the `VTYPE` or `VARTYPE` function.

length

refers to the number of bytes used to store each of the variable's values in a SAS data set. You can use a `LENGTH` statement to set the length of both numeric and character variables. Variable lengths specified in a `LENGTH` statement affect the length of numeric variables only in the output data set; during processing, all numeric variables have a length of 8. Lengths of character variables specified in a `LENGTH` statement affect both the length during processing and the length in the output data set.

In an `INPUT` statement, you can assign a length other than the default length to character variables. You can also assign a length to a variable in the `ATTRIB` statement. A variable that appears for the first time on the left side of an assignment statement has the same length as the expression on the right side of the assignment statement.

To determine the value of this attribute, use the `VLENGTH` or `VARLEN` function.

format

refers to the instructions that SAS uses when printing variable values. If no format is specified, the default format is `BEST12.` for a numeric variable, and `$w.` for a character variable. You can assign SAS formats to a variable in the `FORMAT` or `ATTRIB` statement. You can use the `FORMAT` procedure to create your own format for a variable.

To determine the value of this attribute, use the `VFORMAT` or `VARFMT` function.

informat

refers to the instructions that SAS uses when reading data values. If no informat is specified, the default informat is *w.d* for a numeric variable, and *\$w.* for a character variable. You can assign SAS informats to a variable in the **INFORMAT** or **ATTRIB** statement. You can use the **FORMAT** procedure to create your own informat for a variable.

To determine the value of this attribute, use the **VINFORMAT** or **VARINFMT** function.

label

refers to a descriptive label up to 256 characters long. A variable label, which can be printed by some SAS procedures, is useful in report writing. You can assign a label to a variable with a **LABEL** or **ATTRIB** statement.

To determine the value of this attribute, use the **VLABEL** or **VARLABEL** function.

position in observation

is determined by the order in which the variables are defined in the **DATA** step. You can find the position of a variable in the observations of a SAS data set by using the **CONTENTS** procedure. This attribute is generally not important within the **DATA** step except in variable lists, such as the following:

```
var rent-phone;
```

See “SAS Variable Lists” on page 34 for more information.

The positions of variables in a SAS data set affect the order in which they appear in the output of SAS procedures, unless you control the order within your program, for example, with a **VAR** statement.

To determine the value of this attribute, use the **VARNUM** function.

index type

indicates whether the variable is part of an index for the data set. See “Understanding SAS Indexes” on page 505 for more information.

To determine the value of this attribute, use the **OUT=** option with the **CONTENTS** procedure to create an output data set. The **IDXUSAGE** variable in the output data set contains one of the following values for each variable:

Table 4.2 Index Type Attribute Values

Value	Definition
NONE	The variable is not indexed
SIMPLE	The variable is part of a simple index
COMPOSITE	The variable is part of one or more composite indexes
BOTH	The variable is part of both simple and composite indexes

Ways to Create Variables

Overview

These are some of the most common ways that you can create variables in a **DATA** step:

- use an assignment statement
- read data with the INPUT statement in a DATA step
- specify a new variable in a FORMAT or INFORMAT statement
- specify a new variable in a LENGTH statement
- specify a new variable in an ATTRIB statement.

Note: This list is not exhaustive. For example, the SET, MERGE, MODIFY, and UPDATE statements can also create variables. Δ

Using an Assignment Statement

In a DATA step, you can create a new variable and assign it a value by using it for the first time on the left side of an assignment statement. SAS determines the length of a variable from its first occurrence in the DATA step. The new variable gets the same type and length as the expression on the right side of the assignment statement.

When the type and length of a variable are not explicitly set, SAS gives the variable a default type and length, as shown in the examples in the following table.

Table 4.3 Resulting Variable Types and Lengths Produced When They Are Not Explicitly Set

Expression	Example	Resulting Type of X	Resulting Length of X	Explanation
Numeric variable	<code>length a 4;</code> <code>x=a;</code>	Numeric variable	8	Default numeric length (8 bytes unless otherwise specified)
Character variable	<code>length a \$ 4;</code> <code>x=a;</code>	Character variable	4	Length of source variable
Character literal	<code>x='ABC';</code> <code>x='ABCDE';</code>	Character variable	3	Length of first literal encountered
Concatenation of variables	<code>length a \$ 4</code> <code>b \$ 6</code> <code>c \$ 2;</code> <code>x=a b c;</code>	Character variable	12	Sum of the lengths of all variables
Concatenation of variables and literal	<code>length a \$ 4;</code> <code>x=a 'CAT';</code> <code>x=a 'CATNIP';</code>	Character variable	7	Sum of the lengths of variables and literals encountered in first assignment statement

If a variable appears for the first time on the right side of an assignment statement, SAS assumes that it is a numeric variable and that its value is missing. If no later statement gives it a value, SAS prints a note in the log that the variable is uninitialized.

Note: A RETAIN statement initializes a variable and can assign it an initial value, even if the RETAIN statement appears after the assignment statement. Δ

Reading Data with the INPUT Statement in a DATA Step

When you read raw data in SAS by using an INPUT statement, you define variables based on positions in the raw data. You can use one of the following methods with the INPUT statement to provide information to SAS about how the raw data is organized:

- column input
- list input (simple or modified)
- formatted input
- named input.

See *SAS Language Reference: Dictionary* for more information about using each method.

The following example uses simple list input to create a SAS data set named GEMS and defines four variables based on the data provided:

```
data gems;
  input Name $ Color $ Carats Owner $;
  datalines;
emerald green 1 smith
sapphire blue 2 johnson
ruby red 1 clark
;
```

Specifying a New Variable in a FORMAT or an INFORMAT Statement

You can create a variable and specify its format or informat with a FORMAT or an INFORMAT statement. For example, the following FORMAT statement creates a variable named SALE_PRICE with a format of 6.2 in a new data set named SALES:

```
data sales;
  Sale_Price=49.99;
  format Sale_Price 6.2;
run;
```

SAS creates a numeric variable with the name SALE_PRICE and a length of 8.

See *SAS Language Reference: Dictionary* for more information about using the FORMAT and INFORMAT statements.

Specifying a New Variable in a LENGTH Statement

You can use the LENGTH statement to create a variable and set the length of the variable, as in the following example:

```
data sales;
  length Salesperson $20;
run;
```

For character variables, you must use the longest possible value in the first statement that uses the variable, because you cannot change the length with a subsequent LENGTH statement within the same DATA step. The maximum length of any character variable in SAS is 32,767 bytes. For numeric variables, you can change the length of the variable by using a subsequent LENGTH statement.

When SAS assigns a value to a character variable, it pads the value with blanks or truncates the value on the right side, if necessary, to make it match the length of the target variable. Consider the following statements:

```
length address1 address2 address3 $ 200;
address3=address1||address2;
```

Because the length of ADDRESS3 is 200 bytes, only the first 200 bytes of the concatenation (the value of ADDRESS1) are assigned to ADDRESS3. You might be able to avoid this problem by using the TRIM function to remove trailing blanks from ADDRESS1 before performing the concatenation, as follows:

```
address3=trim(address1)||address2;
```

See *SAS Language Reference: Dictionary* for more information about using the LENGTH statement.

Specifying a New Variable in an ATTRIB Statement

The ATTRIB statement enables you to specify one or more of the following variable attributes for an existing variable:

- FORMAT=
- INFORMAT=
- LABEL=
- LENGTH=.

If the variable does not already exist, one or more of the FORMAT=, INFORMAT=, and LENGTH= attributes can be used to create a new variable. For example, the following DATA step creates a variable named FLAVOR in a data set named LOLLIPOPS:

```
data lollipops;
  Flavor="Cherry";
  attrib Flavor format=$10.;
run;
```

Note: You cannot create a new variable by using a LABEL statement or the ATTRIB statement's LABEL= attribute by itself; labels can be applied only to existing variables. Δ

See *SAS Language Reference: Dictionary* for more information about using the ATTRIB statement.

Using the IN= Data Set Option

The IN= data set option creates a special Boolean variable that indicates whether the data set contributed data to the current observation. The variable has a value of 1 when true, and a value of 0 when false. You can use IN= on the SET, MERGE, and UPDATE statements in a DATA step.

The following example shows a merge of the OLD and NEW data sets where the IN= option is used to create a variable named X that indicates whether the NEW data set contributed data to the observation:

```
data master missing;
  merge old new(in=x);
  by id;
  if x=0 then output missing;
  else output master;
run;
```

Variable Type Conversions

If you define a numeric variable and assign the result of a character expression to it, SAS tries to convert the character result of the expression to a numeric value and to execute the statement. If the conversion is not possible, SAS prints a note to the log, assigns the numeric variable a value of missing, and sets the automatic variable `_ERROR_` to 1. For a listing of the rules by which SAS automatically converts character variables to numeric variables and vice-versa, see “Automatic Numeric-Character Conversion” on page 64.

If you define a character variable and assign the result of a numeric expression to it, SAS tries to convert the numeric result of the expression to a character value using the `BESTw.` format, where w is the width of the character variable and has a maximum value of 32. SAS then tries to execute the statement. If the character variable you use is not long enough to contain a character representation of the number, SAS prints a note to the log and assigns the character variable asterisks. If the value is too small, SAS provides no error message and assigns the character variable the character zero (0).

Output 4.1 Automatic Variable Type Conversions (partial SAS log)

```

1  data _null_;
2  x= 3626885;
3  length y $ 4;
4  y=x;
5  put y;
NOTE: Numeric values have been converted to character
      values at the places given by: (Line):(Column).
      4:6
36E5

6  data _null_;
7  x1= 3626885;
8  length y1 $ 1;
9  y1=x1;
10 xs=0.000005;
11 length ys $ 1;
12 ys=xs;
13 put y1= ys=;
14 run;
NOTE: Numeric values have been converted to character
      values at the places given by: (Line):(Column).
      9:7 12:7
NOTE: Invalid character data, x1=3626885.00 , at line 9 column 7.
y1=* ys=0
x1=3626885 y1=* xs=5E-6 ys=0 _ERROR_=1 _N_=1
NOTE: At least one W.D format was too small for the number to be printed.
      The decimal might be shifted by the "BEST" format.
```

In the first DATA step of the example, SAS is able to fit the value of Y into a 4-byte field by representing its value in scientific notation. In the second DATA step, SAS cannot fit the value of Y1 into a 1-byte field and displays an asterisk (*) instead.

Aligning Variable Values

In SAS, numeric variables are right-aligned and character values are left-aligned. You can further control their alignment by using a format.

However, when you assign a character *value* in an assignment statement, SAS stores the value as it appears in the statement and does not perform any alignment. Output 4.2 illustrates the character value alignment produced by the following program:

```
data aircode;
  input city $1-13;
  length airport $ 10;
  if city='San Francisco' then airport='SFO';
  else if city='Honolulu' then airport='HNL';
  else if city='New York' then airport='JFK or EWR';
  else if city='Miami' then airport='  MIA  ';
  datalines;
San Francisco
Honolulu
New York
Miami
;

proc print data=aircode;
run;
```

This example produces the following output:

Output 4.2 Output from the PRINT Procedure

The SAS System		
OBS	CITY	AIRPORT
1	San Francisco	SFO
2	Honolulu	HNL
3	New York	JFK or EWR
4	Miami	MIA

Automatic Variables

Automatic variables are created automatically by the DATA step or by DATA step statements. These variables are added to the program data vector but are not output to the data set being created. The values of automatic variables are retained from one iteration of the DATA step to the next, rather than set to missing.

Automatic variables that are created by specific statements are documented with those statements. For examples, see the BY statement, the MODIFY statement, and the WINDOW statement in *SAS Language Reference: Dictionary*.

Two automatic variables are created by every DATA step: `_N_` and `_ERROR_`.

`_N_`

is initially set to 1. Each time the DATA step loops past the DATA statement, the variable `_N_` increments by 1. The value of `_N_` represents the number of times the DATA step has iterated.

`_ERROR_`

is 0 by default but is set to 1 whenever an error is encountered, such as an input data error, a conversion error, or a math error, as in division by 0 or a floating point overflow. You can use the value of this variable to help locate errors in data records and to print an error message to the SAS log.

For example, either of the two following statements writes to the SAS log, during each iteration of the DATA step, the contents of an input record in which an input error is encountered:

```
if _error_=1 then put _infile_;

if _error_ then put _infile_;
```

SAS Variable Lists

Definition

A SAS *variable list* is an abbreviated method of referring to a list of variable names. SAS enables you to use the following variable lists:

- numbered range lists
- name range lists
- name prefix lists
- special SAS name lists.

With the exception of the numbered range list, you refer to the variables in a variable list in the same order that SAS uses to keep track of the variables. SAS keeps track of active variables in the order that the compiler encounters them within a DATA step, whether they are read from existing data sets, an external file, or created in the step. In a numbered range list, you can refer to variables that were created in any order, provided that their names have the same prefix.

You can use variable lists in many SAS statements and data set options, including those that define variables. However, they are especially useful *after* you define all of the variables in your SAS program because they provide a quick way to reference existing groups of data.

Note: Only the numbered range list is used in the RENAME= option. \triangle

Numbered Range Lists

Numbered range lists require you to have a series of variables with the same name, except for the last character or characters, which are consecutive numbers. For example, the following two lists refer to the same variables:

```
x1,x2,x3,...,xn
```

```
x1-xn
```

In a numbered range list, you can begin with any number and end with any number as long as you do not violate the rules for user-supplied variable names and the numbers are consecutive.

For example, suppose you decide to give some of your numeric variables sequential names, as in VAR1, VAR2, and so on. Then, you can write an INPUT statement as follows:

```
input idnum name $ var1-var3;
```

Note that the character variable NAME is not included in the abbreviated list.

Name Range Lists

Name range lists rely on the order of variable definition, as shown in the following table:

Table 4.4 Name Range Lists

Variable List	Included Variables
x--a	all variables in order of variable definition, from X to A inclusive
x-numeric-a	all numeric variables from X to A inclusive
x-character-a	all character variables from X to A inclusive

You can use the VARNUM option in PROC CONTENTS to print the variables in the order of definition.

For example, consider the following INPUT statement:

```
input idnum name $ weight pulse chins;
```

In later statements you can use these variable lists:

```
/* keeps only the numeric variables idnum, weight, and pulse */
```

```
keep idnum-numeric-pulse;
```

```
/* keeps the consecutive variables name, weight, and pulse */
```

```
keep name--pulse;
```

Name Prefix Lists

Some SAS functions and statements enable you to use a name prefix list to refer to all variables that begin with a specified character string:

```
sum(of SALES:)
```

tells SAS to calculate the sum of all the variables that begin with “SALES,” such as SALES_JAN, SALES_FEB, and SALES_MAR.

Special SAS Name Lists

Special SAS name lists include

NUMERIC

specifies all numeric variables that are already defined in the current DATA step.

CHARACTER

specifies all character variables that are currently defined in the current DATA step.

ALL

specifies all variables that are currently defined in the current DATA step.

Dropping, Keeping, and Renaming Variables

Using Statements or Data Set Options

The DROP, KEEP, and RENAME statements or the DROP=, KEEP=, and RENAME= data set options control which variables are processed or output during the DATA step. You can use one or a combination of these statements and data set options to achieve the results you want. The action taken by SAS depends largely on whether you

- use a statement or data set option or both
- specify the data set options on an input or an output data set.

The following table summarizes the general differences between the DROP, KEEP, and RENAME statements and the DROP=, KEEP=, and RENAME= data set options.

Table 4.5 Statements versus Data Set Options for Dropping, Keeping, and Renaming Variables

Statements	Data Set Options
apply to output data sets only.	apply to output or input data sets.
affect all output data sets.	affect individual data sets.
can be used in DATA steps only.	can be used in DATA steps and PROC steps.
can appear anywhere in DATA steps.	must immediately follow the name of each data set to which they apply.

Using the Input or Output Data Set

You must also consider whether you want to drop, keep, or rename the variable before it is read into the program data vector or as it is written to the new SAS data set. If you use the DROP, KEEP, or RENAME statement, the action always occurs as the variables are written to the output data set. With SAS data set options, where you use the option determines when the action occurs. If the option is used on an input data set, the variable is dropped, kept, or renamed before it is read into the program data vector. If used on an output data set, the data set option is applied as the variable is written to the new SAS data set. (In the DATA step, an input data set is one that is specified in a SET, MERGE, or UPDATE statement. An output data set is one that is specified in the DATA statement.) Consider the following facts when you make your decision:

- If variables are not written to the output data set and they do not require any processing, using an input data set option to exclude them from the DATA step is more efficient.
- If you want to rename a variable before processing it in a DATA step, you must use the RENAME= data set option in the input data set.
- If the action applies to output data sets, you can use either a statement or a data set option in the output data set.

The following table summarizes the action of data set options and statements when they are specified for input and output data sets. The last column of the table tells whether the variable is available for processing in the DATA step. If you want to rename the variable, use the information in the last column.

Table 4.6 Status of Variables and Variable Names When Dropping, Keeping, and Renaming Variables

Where Specified	Data Set Option or Statement	Purpose	Status of Variable or Variable Name
Input data set	DROP= KEEP=	includes or excludes variables from processing	if excluded, variables are not available for use in DATA step
	RENAME=	changes name of variable before processing	use new name in program statements and output data set options; use old name in other input data set options
Output data set	DROP, KEEP	specifies which variables are written to all output data sets	all variables available for processing
	RENAME	changes name of variables in all output data sets	use old name in program statements; use new name in output data set options
	DROP= KEEP=	specifies which variables are written to individual output data sets	all variables are available for processing
	RENAME=	changes name of variables in individual output data sets	use old name in program statements and other output data set options

Order of Application

If your program requires that you use more than one data set option or a combination of data set options and statements, it is helpful to know that SAS drops, keeps, and renames variables in the following order:

- First, options on input data sets are evaluated left to right within SET, MERGE, and UPDATE statements. DROP= and KEEP= options are applied before the RENAME= option.
- Next, DROP and KEEP statements are applied, followed by the RENAME statement.
- Finally, options on output data sets are evaluated left to right within the DATA statement. DROP= and KEEP= options are applied before the RENAME= option.

Examples of Dropping, Keeping, and Renaming Variables

The following examples show specific ways to handle dropping, keeping, and renaming variables:

- This example uses the DROP= and RENAME= data set options and the INPUT function to convert the variable POPRANK from character to numeric. The name POPRANK is changed to TEMPVAR before processing so that a new variable POPRANK can be written to the output data set. Note that the variable TEMPVAR is dropped from the output data set and that the new name TEMPVAR is used in the program statements.

```

data newstate(drop=tempvar);
  length poprank 8;
  set state(rename=(poprank=tempvar));
  poprank=input(tempvar,8.);
run;

```

- This example uses the DROP statement and the DROP= data set option to control the output of variables to two new SAS data sets. The DROP statement applies to both data sets, CORN and BEAN. You must use the RENAME= data set option to rename the output variables BEANWT and CORNWT in each data set.

```

data corn(rename=(cornwt=yield) drop=beanwt)
  bean(rename=(beanwt=yield) drop=cornwt);
  set harvest;
  if crop='corn' then output corn;
  else if crop='bean' then output bean;
  drop crop;
run;

```

- This example shows how to use data set options in the DATA statement and the RENAME statement together. Note that the new name QTRTOT is used in the DROP= data set option.

```

data qtr1 qtr2 ytd(drop=qtrtot);
  set ytdsales;
  if qtr=1 then output qtr1;
  else if qtr=2 then output qtr2;
  else output ytd;
  rename total=qtrtot;
run;

```

Numeric Precision in SAS Software

How SAS Stores Numeric Values

To store numbers of large magnitude and to perform computations that require many digits of precision to the right of the decimal point, SAS stores all numeric values using *floating-point*, or real binary, representation. Floating-point representation is an implementation of what is generally known as scientific notation, in which values are represented as numbers between 0 and 1 times a power of 10. The following is an example of a number in scientific notation:

$$.1234 \times 10^4$$

Numbers in scientific notation consist of the following parts:

- The *base* is the number of significant digits, including zero, that a positional numeral system uses to represent the number; in this example, the base is 10.
- The *mantissa* are the digits that define the number's magnitude; in this example, the mantissa is .1234.
- The *exponent* indicates how many times the base is to be multiplied; in this example, the exponent is 4.

Floating-point representation is a form of scientific notation, except that on most operating systems the base is not 10, but is either 2 or 16. The following table summarizes various representations of floating-point numbers that are stored in 8 bytes.

Table 4.7 Summary of Floating-Point Numbers Stored in 8 Bytes

Representation	Base	Exponent Bits	Maximum Mantissa Bits
IBM mainframe	16	7	56
IEEE	2	11	52

SAS uses truncated floating-point numbers via the LENGTH statement, which reduces the number of mantissa bits. For more information on the effects of truncated lengths, see “Storing Numbers with Less Precision” on page 43.

Troubleshooting Problems Regarding Floating-Point Representation

Overview

In most situations, the way that SAS stores numeric values does not affect you as a user. However, floating-point representation can account for anomalies you might notice in SAS program behavior. The following sections identify the types of problems that can occur in various operating environments and how you can anticipate and avoid them.

Floating-Point Representation on IBM Mainframes

SAS for z/OS uses the traditional IBM mainframe floating-point representation as follows:

```
SEEEEEEE MMMMMMMM MMMMMMMM MMMMMMMM
byte 1   byte 2   byte 3   byte 4
```

```
MMMMMMMM MMMMMMMM MMMMMMMM MMMMMMMM
byte 5   byte 6   byte 7   byte 8
```

This representation corresponds to bytes of data with each character being 1 bit, as follows:

- The S in byte 1 is the *sign bit* of the number. A value of 0 in the sign bit is used to represent positive numbers.
- The seven E characters in byte 1 represent a binary integer known as the *characteristic*. The characteristic represents a signed exponent and is obtained by adding the bias to the actual exponent. The *bias* is an offset used to enable both negative and positive exponents with the bias representing 0. If a bias is not used, an additional sign bit for the exponent must be allocated. For example, if a system uses a bias of 64, a characteristic with the value 66 represents an exponent of +2, while a characteristic of 61 represents an exponent of -3.
- The remaining M characters in bytes 2 through 8 represent the bits of the mantissa. There is an implied *radix point* before the leftmost bit of the mantissa; therefore, the mantissa is always less than 1. The term radix point is used instead of decimal point because decimal point implies that you are working with decimal

(base 10) numbers, which might not be the case. The radix point can be thought of as the generic form of decimal point.

The exponent has a base associated with it. Do not confuse this with the base in which the exponent is represented; the exponent is always represented in binary, but the exponent is used to determine how many times the base should be multiplied by the mantissa. In the case of the IBM mainframes, the exponent's base is 16. For other computers, it is commonly either 2 or 16.

Each bit in the mantissa represents a fraction whose numerator is 1 and whose denominator is a power of 2. For example, the leftmost bit in byte 2 represents $(\frac{1}{2})^1$, the next bit represents $(\frac{1}{2})^2$, and so on. In other words, the mantissa is the sum of a series of fractions such as $\frac{1}{2}$, $\frac{1}{4}$, $\frac{1}{8}$, and so on. Therefore, for any floating-point number to be represented exactly, you must be able to express it as the previously mentioned sum. For example, 100 is represented as the following expression:

$$\left(\frac{1}{4} + \frac{1}{8} + \frac{1}{64}\right) \times 16^2$$

To illustrate how the above expression is obtained, two examples follow. The first example is in base 10. The value 100 is represented as follows:

100.

The period in this number is the radix point. The mantissa must be less than 1; therefore, you normalize this value by shifting the radix point three places to the left, which produces the following value:

.100

Because the radix point is shifted three places to the left, 3 is the exponent:

$$.100 \times 10^3 = 100$$

The second example is in base 16. In hexadecimal notation, 100 (base 10) is written as follows:

64.

Shifting the radix point two places to the left produces the following value:

.64

Shifting the radix point also produces an exponent of 2, as in:

$$.64 \times 16^2$$

The binary value of this number is **.01100100**, which can be represented in the following expression:

$$\left(\frac{1}{2}\right)^2 + \left(\frac{1}{2}\right)^3 + \left(\frac{1}{2}\right)^6 = \frac{1}{4} + \frac{1}{8} + \frac{1}{64}$$

In this example, the exponent is 2. To represent the exponent, you add the bias of 64 to the exponent. The hexadecimal representation of the resulting value, 66, is 42. The binary representation is as follows:

```
01000010 01100100 00000000 00000000
00000000 00000000 00000000 00000000
```

Floating-Point Representation on OpenVMS

On OpenVMS, SAS stores numeric values in the D-floating format, which has the following scheme:

```
MMMMMMMM MMMMMMMM MMMMMMMM MMMMMMMM
byte 8   byte 7   byte 6   byte 5

MMMMMMMM MMMMMMMM SEEEEEEE EMMMMMMM
byte 4   byte 3   byte 2   byte 1
```

In D-floating format, the exponent is 8 bits instead of 7, but uses base 2 instead of base 16 and a bias of 128, which means the magnitude of the D-floating format is not as great as the magnitude of the IBM representation. The mantissa of the D-floating format is, physically, 55 bits. However, all floating-point values under OpenVMS are normalized, which means it is guaranteed that the high-order bit will always be 1. Because of this guarantee, there is no need to physically represent the high-order bit in the mantissa; therefore, the high-order bit is hidden.

For example, the decimal value 100 represented in binary is as follows:

```
01100100.
```

This value can be normalized by shifting the radix point as follows:

```
0.1100100
```

Because the radix was shifted to the left seven places, the exponent, 7 plus the bias of 128, is 135. Represented in binary, the number is as follows:

```
10000111
```

To represent the mantissa, subtract the hidden bit from the fraction field:

```
.100100
```

You can combine the sign (0), the exponent, and the mantissa to produce the D-floating format:

```
MMMMMMMM MMMMMMMM MMMMMMMM MMMMMMMM
00000000 00000000 00000000 00000000

MMMMMMMM MMMMMMMM SEEEEEEE EMMMMMMM
00000000 00000000 01000011 11001000
```

Floating-Point Representation Using the IEEE Standard

The Institute of Electrical and Electronic Engineers (IEEE) representation is used by many operating systems, including Windows and UNIX. The IEEE representation uses an 11-bit exponent with a base of 2 and bias of 1023, which means that it has much greater magnitude than the IBM mainframe representation, but sometimes at the expense of 3 bits less in the mantissa. The value of 1 represented by the IEEE standard is as follows:

```
3F F0 00 00 00 00 00 00
```

Precision Versus Magnitude

As discussed in previous sections, floating-point representation enables numbers of very large magnitude (numbers such as 2 to the 30th power) and high degrees of precision (many digits to the right of the decimal place). However, operating systems differ on how much precision and how much magnitude they use.

In “How SAS Stores Numeric Values” on page 38, you can see that the number of exponent bits and mantissa bits varies. The more bits that are reserved for the mantissa, the more precise the number; the more bits that are reserved for the exponent, the greater the magnitude the number can have.

Whether precision or magnitude is more important depends on the characteristics of your data. For example, if you are working with physics applications, very large numbers might be needed, and magnitude is probably more important. However, if you are working with banking applications, where every digit is important but the number of digits is not great, then precision is more important. Most often, SAS applications need a moderate amount of both precision and magnitude, which is sufficiently provided by floating-point representation.

Computational Considerations of Fractions

Regardless of how much precision is available, there is still the problem that some numbers cannot be represented exactly. In the decimal number system, the fraction 1/3 cannot be represented exactly in decimal notation. Likewise, most decimal fractions (for example, .1) cannot be represented exactly in base 2 or base 16 numbering systems. This is the principle reason for difficulty in storing fractional numbers in floating-point representation.

Consider the IBM mainframe representation of .1:

```
40 19 99 99 99 99 99 99
```

Notice the trailing 9 digit, similar to the trailing 3 digit in the attempted decimal representation of 1/3 (.3333 ...). This lack of precision is aggravated by arithmetic operations. Consider what would happen if you added the decimal representation of 1/3 several times. When you add .33333 ... to .99999 ... , the theoretical answer is 1.33333 ... 2, but in practice, this answer is not possible. The sums become imprecise as the values continue.

Likewise, the same process happens when the following DATA step is executed:

```
data _null_;
  do i=-1 to 1 by .1;
    if i=0 then put 'AT ZERO';
  end;
run;
```

The AT ZERO message in the DATA step is never printed because the accumulation of the imprecise number introduces enough error that the exact value of 0 is never

encountered. The number is close, but never exactly 0. This problem is easily resolved by explicitly rounding with each iteration, as the following statements illustrate:

```
data _null_;
  i=-1;
  do while(i<=1);
    i=round(i+.1,.001);
    if i=0 then put 'AT ZERO';
  end;
run;
```

Numeric Comparison Considerations

As discussed in “Computational Considerations of Fractions” on page 42, imprecision can cause problems with computations. Imprecision can also cause problems with comparisons. Consider the following example in which the PUT statement is not executed:

```
data _null_;
  x=1/3;
  if x=.33333 then put 'MATCH';
run;
```

However, if you add the ROUND function, as in the following example, the PUT statement is executed:

```
data _null_;
  x=1/3;
  if round(x,.00001)=.33333 then put 'MATCH';
run;
```

In general, if you are doing comparisons with fractional values, it is good practice to use the ROUND function.

Storing Numbers with Less Precision

As discussed in “How SAS Stores Numeric Values” on page 38, SAS enables numeric values to be stored on disk with less than full precision. Use the LENGTH statement to control the number of bytes that are used to store the floating-point number. Use the LENGTH statement carefully to avoid significant data loss.

For example, the IBM mainframe representation uses 8 bytes for full precision, but you can store as few as 2 bytes on disk. The value 1 is represented as 41 10 00 00 00 00 00 00 in 8 bytes. In 2 bytes, it would be truncated to 41 10. You still have the full range of magnitude because the exponent remains intact; there are simply fewer digits involved. A decrease in the number of digits means either fewer digits to the right of the decimal place or fewer digits to the left of the decimal place before trailing zeros must be used.

For example, consider the number 1234567890, which would be .1234567890 to the 10th power of 10 (in base 10). If you have only five digits of precision, the number becomes 123460000 (rounding up). Note that this is the case regardless of the power of 10 that is used (.12346, 12.346, .0000012346, and so on).

The only reason to truncate length by using the LENGTH statement is to save disk space. All values are expanded to full size to perform computations in DATA and PROC steps. In addition, you must be careful in your choice of lengths, as the previous discussion shows.

Consider a length of 2 bytes on an IBM mainframe system. This value enables 1 byte to store the exponent and sign, and 1 byte for the mantissa. The largest value that can

be stored in 1 byte is 255. Therefore, if the exponent is 0 (meaning 16 to the 0th power, or 1 multiplied by the mantissa), then the largest integer that can be stored with complete certainty is 255. However, some larger integers can be stored because they are multiples of 16. For example, consider the 8-byte representation of the numbers 256 to 272 in the following table:

Table 4.8 Representation of the Numbers 256 to 272 in Eight Bytes

Value	Sign/Exp	Mantissa 1	Mantissa 2-7	Considerations
256	43	10	000000000000	trailing zeros; multiple of 16
257	43	10	100000000000	extra byte needed
258	43	10	200000000000	
259	43	10	300000000000	
.	.	.	.	
271	43	10	F00000000000	
272	43	11	000000000000	trailing zeros; multiple of 16

The numbers from 257 to 271 cannot be stored exactly in the first 2 bytes; a third byte is needed to store the number precisely. As a result, the following code produces misleading results:

```
data temp;
  length x 2;
  x=257;
  y1=x+1;
run;

data _null_;
  set temp;
  if x=257 then put 'FOUND';
  y2=x+1;
run;
```

The PUT statement is never executed because the value of X is actually 256 (the value 257 truncated to 2 bytes). Recall that 256 is stored in 2 bytes as 4310, but 257 is also stored in 2 bytes as 4310, with the third byte of 10 truncated.

You receive no warning that the value of 257 is truncated in the first DATA step. Note, however, that Y1 has the value 258 because the values of X are kept in full, 8-byte floating-point representation in the program data vector. The value is truncated only when stored in a SAS data set. Y2 has the value 257, because X is truncated before the number is read into the program data vector.

CAUTION:

Do not use the LENGTH statement if your variable values are not integers. Fractional numbers lose precision if truncated. Also, use the LENGTH statement to truncate values only when disk space is limited. Refer to the length table in the SAS documentation for your operating environment for maximum values. Δ

Truncating Numbers and Making Comparisons

The TRUNC function truncates a number to a requested length and then expands the number back to full length. The truncation and subsequent expansion duplicate the effect of storing numbers in less than full length and then reading them. For example, if the variable

```
x=1/3;
```

is stored with a length of 3, then the following comparison is not true:

```
if x=1/3 then ...;
```

However, adding the TRUNC function makes the comparison true, as in the following:

```
if x=trunc(1/3,3) then ...;
```

Determining How Many Bytes Are Needed to Store a Number Accurately

To determine the minimum number of bytes needed to store a value accurately, you can use the TRUNC function. For example, the following program finds the minimum length of bytes (MINLEN) needed for numbers stored in a native SAS data set named NUMBERS. The data set NUMBERS contains the variable VALUE. VALUE contains a range of numbers, in this example, from 269 to 272:

```
data numbers;
  input value;
  datalines;
269
270
271
272
;

data temp;
  set numbers;
  x=value;
  do L=8 to 1 by -1;
    if x NE trunc(x,L) then
      do;
        minlen=L+1;
        output;
        return;
      end;
  end;
run;

proc print noobs;
  var value minlen;
run;
```

The following output shows the results from this code.

Output 4.3 Using the TRUNC Function

The SAS System	
VALUE	MINLEN
269	3
270	3
271	3
272	2

Note that the minimum length required for the value 271 is greater than the minimum required for the value 272. This fact illustrates that it is possible for the largest number in a range of numbers to require fewer bytes of storage than a smaller number. If precision is needed for all numbers in a range, you should obtain the minimum length for all the numbers, not just the largest one.

Double-Precision Versus Single-Precision Floating-Point Numbers

You might have data created by an external program that you want to read into a SAS data set. If the data is in floating-point representation, you can use the `RBw.d` informat to read in the data. However, there are exceptions.

The `RBw.d` informat might truncate double-precision floating-point numbers if the `w` value is less than the size of the double-precision floating-point number (8 on all the operating systems discussed in this section). Therefore, the `RB8.` informat corresponds to a full 8-byte floating point. The `RB4.` informat corresponds to an 8-byte floating point truncated to 4 bytes, exactly the same as a `LENGTH 4` in the `DATA` step.

An 8-byte floating point that is truncated to 4 bytes might not be the same as *float* in a C program. In the C language, an 8-byte floating-point number is called a *double*. In FORTRAN, it is a `REAL*8`. In IBM PL/I, it is a `FLOAT BINARY(53)`. A 4-byte floating-point number is called a *float* in the C language, `REAL*4` in FORTRAN, and `FLOAT BINARY(21)` in IBM PL/I.

On the IBM mainframes, a single-precision floating-point number is exactly the same as a double-precision number truncated to 4 bytes. On operating systems that use the IEEE standard, this is not the case; a single-precision floating-point number uses a different number of bits for its exponent and uses a different bias, so that reading in values using the `RB4.` informat does not produce the expected results.

Transferring Data between Operating Systems

The problems of precision and magnitude when you use floating-point numbers are not confined to a single operating system. Additional problems can arise when you move from one operating system to another, unless you use caution. This section discusses factors to consider when you are transporting data sets with very large or very small numeric values by using the `UPLOAD` and `DOWNLOAD` procedures, the `CPORT` and `CIMPORT` procedures, or transport engines.


Table 4.7 on page 39 shows the maximum number of digits of the base, exponent, and mantissa. Because there are differences in the maximum values that can be stored in different operating environments, there might be problems in transferring your floating-point data from one computer to another.

Consider, for example, transporting data between an IBM mainframe and a PC. The IBM mainframe has a range limit of approximately `.54E-78` to `.72E76` (and their negative equivalents and 0) for its floating-point numbers. Other computers, such as the PC, have wider limits (the PC has an upper limit of approximately `1E308`). Therefore, if you are transferring numbers in the magnitude of `1E100` from a PC to a

mainframe, you lose that magnitude. During data transfer, the number is set to the minimum or maximum allowable on that operating system, so 1E100 on a PC is converted to a value that is approximately .72E76 on an IBM mainframe.

CAUTION:

Transfer of data between computers can affect numeric precision. If you are transferring data from an IBM mainframe to a PC, notice that the number of bits for the mantissa is 4 less than that for an IBM mainframe, which means you lose 4 bits when moving to a PC. This precision and magnitude difference is a factor when moving from one operating environment to any other where the floating-point representation is different. Δ



CHAPTER

5

Missing Values

<i>Definition of Missing Values</i>	49
<i>Creating Special Missing Values</i>	50
<i>Definition</i>	50
<i>Tips</i>	50
<i>Example</i>	50
<i>Order of Missing Values</i>	51
<i>Numeric Variables</i>	51
<i>Character Variables</i>	52
<i>When Variable Values Are Automatically Set to Missing by SAS</i>	52
<i>When Reading Raw Data</i>	52
<i>When Reading a SAS Data Set</i>	53
<i>When Missing Values Are Generated by SAS</i>	53
<i>Propagation of Missing Values in Calculations</i>	53
<i>Illegal Operations</i>	53
<i>Illegal Character-to-Numeric Conversions</i>	53
<i>Creating Special Missing Values</i>	54
<i>Preventing Propagation of Missing Values</i>	54
<i>Working with Missing Values</i>	55
<i>How to Represent Missing Values in Raw Data</i>	55
<i>How to Set Variable Values to Missing in a DATA Step</i>	55
<i>How to Check for Missing Values in a DATA Step</i>	56

Definition of Missing Values

missing value

is a value that indicates that no data value is stored for the variable in the current observation. There are three kinds of missing values:

- numeric
- character
- special numeric.

By default, SAS prints a missing numeric value as a single period (.) and a missing character value as a blank space. See “Creating Special Missing Values” on page 50 for more information about special numeric missing values.

Creating Special Missing Values

Definition

special missing value

is a type of numeric missing value that enables you to represent different categories of missing data by using the letters A-Z or an underscore.

Tips

- SAS accepts either uppercase or lowercase letters. Values are displayed and printed as uppercase.
- If you do not begin a special numeric missing value with a period, SAS identifies it as a variable name. Therefore, to use a special numeric missing value in a SAS expression or assignment statement, you must begin the value with a period, followed by the letter or underscore, as in the following example:

```
x=.d;
```

- When SAS prints a special missing value, it prints only the letter or underscore.
 - When data values contain characters in numeric fields that you want SAS to interpret as special missing values, use the MISSING statement to specify those characters. For further information, see the MISSING statement in *SAS Language Reference: Dictionary*.
-

Example

The following example uses data from a marketing research company. Five testers were hired to test five different products for ease of use and effectiveness. If a tester was absent, there is no rating to report, and the value is recorded with an X for “absent.” If the tester was unable to test the product adequately, there is no rating, and the value is recorded with an I for “incomplete test.” The following program reads the data and displays the resulting SAS data set. Note the special missing values in the first and third data lines:

```
data period_a;
  missing X I;
  input Id $4. Foodpr1 Foodpr2 Foodpr3 Coffeem1 Coffeem2;
  datalines;
1001 115 45 65 I 78
1002 86 27 55 72 86
1004 93 52 X 76 88
1015 73 35 43 112 108
1027 101 127 39 76 79
;

proc print data=period_a;
  title 'Results of Test Period A';
  footnote1 'X indicates TESTER ABSENT';
```

```

footnote2 'I indicates TEST WAS INCOMPLETE';
run;

```

The following output is produced:

Output 5.1 Output with Multiple Missing Values

Obs	Id	Results of Test Period A				
		Foodpr1	Foodpr2	Foodpr3	Coffeem1	Coffeem2
1	1001	115	45	65	I	78
2	1002	86	27	55	72	86
3	1004	93	52	X	76	88
4	1015	73	35	43	112	108
5	1027	101	127	39	76	79

X indicates TESTER ABSENT
 I indicates TEST WAS INCOMPLETE

Order of Missing Values

Numeric Variables

Within SAS, a missing value for a numeric variable is smaller than all numbers; if you sort your data set by a numeric variable, observations with missing values for that variable appear first in the sorted data set. For numeric variables, you can compare special missing values with numbers and with each other. Table 5.1 on page 51 shows the sorting order of numeric values.

Table 5.1 Numeric Value Sort Order

Sort Order	Symbol	Description
smallest	._	underscore
	.	period
	.A-Z	special missing values A (smallest) through Z (largest)
	-n	negative numbers
	0	zero
largest	+n	positive numbers

For example, the numeric missing value (.) is sorted before the special numeric missing value .A, and both are sorted before the special missing value .Z. SAS does not distinguish between lowercase and uppercase letters when sorting special numeric missing values.

Note: The numeric missing value sort order is the same regardless of whether your system uses the ASCII or EBCDIC collating sequence. Δ

Character Variables

Missing values of character variables are smaller than any printable character value. Therefore, when you sort a data set by a character variable, observations with missing (blank) values of the BY variable always appear before observations in which values of the BY variable contain only printable characters. However, some usually unprintable characters (for example, machine carriage-control characters and real or binary numeric data that have been read in error as character data) have values less than the blank. Therefore, when your data includes unprintable characters, missing values might not appear first in a sorted data set.

When Variable Values Are Automatically Set to Missing by SAS

When Reading Raw Data

At the beginning of each iteration of the DATA step, SAS sets the value of each variable you create in the DATA step to missing, with the following exceptions:

- variables named in a RETAIN statement
- variables created in a SUM statement
- data elements in a `_TEMPORARY_` array
- variables created with options in the FILE or INFILE statements
- variables created by the FGET function
- data elements which are initialized in an ARRAY statement
- automatic variables.

SAS replaces the missing values as it encounters values that you assign to the variables. Thus, if you use program statements to create new variables, their values in each observation are missing until you assign the values in an assignment statement, as shown in the following DATA step:

```
data new;
  input x;
  if x=1 then y=2;
  datalines;
4
1
3
1
;
```

This DATA step produces a SAS data set with the following variable values:

OBS	X	Y
1	4	.
2	1	2
3	3	.
4	1	2

When X equals 1, the value of Y is set to 2. Since no other statements set Y's value when X is not equal to 1, Y remains missing (.) for those observations.

When Reading a SAS Data Set

When variables are read with a SET, MERGE, or UPDATE statement, SAS sets the values to missing only before the first iteration of the DATA step. (If you use a BY statement, the variable values are also set to missing when the BY group changes.) The variables retain their values until new values become available; for example, through an assignment statement or through the next execution of the SET, MERGE, or UPDATE statement. Variables created with options in the SET, MERGE, and UPDATE statements also retain their values from one iteration to the next.

When all of the rows in a data set in a match-merge operation (with a BY statement) have been processed, the variables in the output data set retain their values as described earlier. That is, as long as there is no change in the BY value in effect when all of the rows in the data set have been processed, the variables in the output data set retain their values from the final observation. `FIRST.variable` and `LAST.variable`, the automatic variables that are generated by the BY statement, both retain their values. Their initial value is 1.

When the BY value changes, the variables are set to missing and remain missing because the data set contains no additional observations to provide replacement values. When all of the rows in a data set in a one-to-one merge operation (without a BY statement) have been processed, the variables in the output data set are set to missing and remain missing.

When Missing Values Are Generated by SAS

Propagation of Missing Values in Calculations

SAS assigns missing values to prevent problems from arising. If you use a missing value in an arithmetic calculation, SAS sets the result of that calculation to missing. Then, if you use that result in another calculation, the next result is also missing. This action is called *propagation of missing values*. SAS prints notes in the log to notify you which arithmetic expressions have missing values and when they were created; however, processing continues.

Illegal Operations

SAS prints a note in the log and assigns a missing value to the result if you try to perform an illegal operation, such as the following:

- dividing by zero
- taking the logarithm of zero
- using an expression to produce a number too large to be represented as a floating-point number (known as overflow).

Illegal Character-to-Numeric Conversions

SAS automatically converts character values to numeric values if a character variable is used in an arithmetic expression. If a character value contains nonnumerical information and SAS tries to convert it to a numeric value, a note is

printed in the log, the result of the conversion is set to missing, and the `_ERROR_` automatic variable is set to 1.

Creating Special Missing Values

The result of any numeric missing value in a SAS expression is a period. Thus, both special missing values and ordinary numeric missing values propagate as a period.

```
data a;
    x=.d;
    y=x+1;
    put y=;
run;
```

This DATA step results in the following log:

Output 5.2 SAS Log Results for a Missing Value

```
1 data a;
2 x= .d;
3 y=x+1;
4 put y=;
5 run;

y=.
NOTE: Missing values were generated as a result of performing an
operation on missing values.
Each place is given by:
(Number of times) at (Line):(Column).
1 at 3:6
NOTE: The data set WORK.A has 1 observations and 2 variables.
NOTE: DATA statement used:
real time          0.58 seconds
cpu time           0.05 seconds
```

Preventing Propagation of Missing Values

If you do not want missing values to propagate in your arithmetic expressions, you can omit missing values from computations by using the sample statistic functions. For a list of these functions, see the descriptive statistics category in “Functions and CALL Routines by Category” in the *SAS Language Reference: Dictionary*. For example, consider the following DATA step:

```
data test;
    x=.;
    y=5;
    a=x+y;
    b=sum(x,y);
    c=5;
    c+x;
    put a= b= c=;
run;
```

Output 5.3 SAS Log Results for a Missing Value in a Statistic Function

```

1  data test;
2  x=.;
3  y=5;
4  a=x+y;
5  b=sum(x,y);
6  c=5;
7  c+x;
8  put a= b= c=;
9  run;

a=. b=5 c=5
NOTE: Missing values were generated as a result of performing
      an operation on missing values.
      Each place is given by:
      (Number of times) at (Line):(Column).
      1 at 4:6
NOTE: The data set WORK.TEST has 1 observations and 5 variables.
NOTE: DATA statement used:
      real time          0.11 seconds
      cpu time           0.03 seconds

```

Adding X and Y together in an expression produces a missing result because the value of X is missing. The value of A, therefore, is missing. However, since the SUM function ignores missing values, adding X to Y produces the value 5, not a missing value.

Note: The SUM statement also ignores missing values, so the value of C is also 5. Δ

Working with Missing Values

How to Represent Missing Values in Raw Data

Table 5.2 on page 55 shows how to represent each type of missing value in raw data so that SAS will read and store the value appropriately.

Table 5.2 Representing Missing Values

Missing Values	Representation in Data
Numeric	. (a single decimal point)
Character	' ' (a blank enclosed in quotes)
Special	. <i>letter</i> (a decimal point followed by a letter, for example, .B)
Special	._ (a decimal point followed by an underscore)

How to Set Variable Values to Missing in a DATA Step

You can set values to missing within your DATA step by using program statements such as this one:

```
if age<0 then age=.;
```

This statement sets the stored value of AGE to a numeric missing value if AGE has a value less than 0.

Note: You can display a missing numeric value with a character other than a period by using the DATA step's MISSING statement or the MISSING= system option. Δ

The following example sets the stored value of NAME to a missing character value if NAME has a value of "none":

```
if name="none" then name=' ';
```

Alternatively, if you want to set to a missing value for one or more variable values, you can use the CALL MISSING routine. For example,

```
call missing(sales, name);
```

sets both variable values to a missing value.

Note: You can mix character and numeric variables in the CALL MISSING routine argument list. Δ

How to Check for Missing Values in a DATA Step

You can use the N and NMISS functions to return the number of nonmissing and missing values, respectively, from a list of numeric arguments.

When you check for ordinary missing numeric values, you can use code that is similar to the following:

```
if numvar=. then do;
```

If your data contains special missing values, you can check for either an ordinary or special missing value with a statement that is similar to the following:

```
if numvar<=.z then do;
```

To check for a missing character value, you can use a statement that is similar to the following:


```
if charvar=' ' then do;
```

The MISSING function enables you to check for either a character or numeric missing value, as in:

```
if missing(var) then do;
```

In each case, SAS checks whether the value of the variable in the current observation satisfies the condition specified. If it does, SAS executes the DO group.

Note: Missing values have a value of **false** when you use them with logical operators such as AND or OR. Δ



CHAPTER

6

Expressions

<i>Definitions for SAS Expressions</i>	58
<i>Examples of SAS Expressions</i>	58
<i>SAS Constants in Expressions</i>	58
<i>Definition</i>	58
<i>Character Constants</i>	59
<i>Using Quotation Marks with Character Constants</i>	59
<i>Comparing Character Constants and Character Variables</i>	59
<i>Character Constants Expressed in Hexadecimal Notation</i>	60
<i>Numeric Constants</i>	60
<i>Numeric Constants Expressed in Standard Notation</i>	60
<i>Numeric Constants Expressed in Scientific Notation</i>	61
<i>Numeric Constants Expressed in Hexadecimal Notation</i>	61
<i>Date, Time, and Datetime Constants</i>	61
<i>Bit Testing Constants</i>	62
<i>Avoiding a Common Error with Constants</i>	63
<i>SAS Variables in Expressions</i>	64
<i>Definition</i>	64
<i>Automatic Numeric-Character Conversion</i>	64
<i>SAS Functions in Expressions</i>	65
<i>SAS Operators in Expressions</i>	65
<i>Definitions</i>	65
<i>Arithmetic Operators</i>	66
<i>Comparison Operators</i>	66
<i>Numeric Comparisons</i>	67
<i>The IN Operator in Numeric Comparisons</i>	68
<i>Character Comparisons</i>	69
<i>The IN Operator in Character Comparisons</i>	69
<i>Logical (Boolean) Operators and Expressions</i>	70
<i>The AND Operator</i>	71
<i>The OR Operator</i>	71
<i>The NOT Operator</i>	71
<i>Boolean Numeric Expressions</i>	72
<i>The MIN and MAX Operators</i>	72
<i>The Concatenation Operator</i>	72
<i>Order of Evaluation in Compound Expressions</i>	73

Definitions for SAS Expressions

expression

is generally a sequence of operands and operators that form a set of instructions that are performed to produce a resulting value. You use expressions in SAS program statements to create variables, assign values, calculate new values, transform variables, and perform conditional processing. SAS expressions can resolve to numeric values, character values, or Boolean values.

operands

are constants or variables that can be numeric or character.

operators

are symbols that represent a comparison, arithmetic calculation, or logical operation; a SAS function; or grouping parentheses.

simple expression

is an expression with no more than one operator. A simple expression can consist of a single

- constant
- variable
- function.

compound expression

is an expression that includes several operators. When SAS encounters a compound expression, it follows rules to determine the order in which to evaluate each part of the expression.

WHERE expressions

is a type of SAS expression that is used within a WHERE statement or WHERE= data set option to specify a condition for selecting observations for processing in a DATA or PROC step. For syntax and further information on WHERE expressions, see Chapter 11, “WHERE-Expression Processing,” on page 133 and *SAS Language Reference: Dictionary*

Examples of SAS Expressions

The following are examples of SAS expressions:

- `3`
- `x`
- `x+1`
- `age<100`
- `trim(last)||', '||first`

SAS Constants in Expressions

Definition

A SAS *constant* is a number or a character string that indicates a fixed value. Constants can be used as expressions in many SAS statements, including variable

assignment and IF-THEN statements. They can also be used as values for certain options. Constants are also called *literals*.

The following are types of SAS constants:

- character
- numeric
- date, time, and datetime
- bit testing.

Character Constants

A *character constant* consists of 1 to 32,767 characters and must be enclosed in quotation marks. Character constants can also be represented in hexadecimal form.

Using Quotation Marks with Character Constants

In the following SAS statement, **Tom** is a character constant:

```
if name='Tom' then do;
```

If a character constant includes a single quotation mark, surround it with double quotation marks. For example, to specify the character value **Tom's** as a constant, enter

```
name="Tom's"
```

Another way to write the same string is to enclose the string in single quotation marks and to express the apostrophe as two consecutive quotation marks. SAS treats the two consecutive quotation marks as one quotation mark:

```
name='Tom''s'
```

The same principle holds true for double quotation marks:

```
name="Tom""s"
```

CAUTION:

Matching quotation marks correctly is important. Missing or extraneous quotation marks cause SAS to misread both the erroneous statement and the statements that follow it. For example, in `name='O'Brien'`, **O** is the character value of NAME, **Brien** is extraneous, and `';` begins another quoted string. Δ

Comparing Character Constants and Character Variables

It is important to remember that character constants are enclosed in quotation marks, but names of character variables are not. This distinction applies wherever you can use a character constant, such as in titles, footnotes, labels, and other descriptive strings; in option values; and in operating environment-specific strings, such as file specifications and commands.

The following statements use character constants:

- `x='abc';`
- `if name='Smith' then do;`

The following statements use character variables:

- `x=abc;`
- `if name=Smith then do;`

In the second set of examples, SAS searches for variables named ABC and SMITH, instead of constants.

Note: SAS distinguishes between uppercase and lowercase when comparing quoted values. For example, the character values 'Smith' and 'SMITH' are not equivalent. Δ

Character Constants Expressed in Hexadecimal Notation

SAS character constants can be expressed in hexadecimal notation. A character hexadecimal constant is a string of an even number of hexadecimal characters enclosed in single or double quotation marks, followed immediately by an X, as in this example:

```
'534153'x
```

A comma can be used to make the string more readable, but it is not part of and does not alter the hexadecimal value. If the string contains a comma, the comma must separate an even number of hexadecimal characters within the string, as in this example:

```
if value='3132,3334'x then do;
```

Note: Any trailing blanks or leading blanks within the quotation marks cause an error message to be written to the log. Δ

Numeric Constants

A *numeric constant* is a number that appears in a SAS statement. Numeric constants can be presented in many forms, including

- standard notation
- scientific (E) notation
- hexadecimal notation.

Numeric Constants Expressed in Standard Notation

Most numeric constants are written just as numeric data values are. The numeric constant in the following expression is 100:

```
part/all*100
```

Numeric constants can be expressed in standard notation in the following ways:

Table 6.1 Standard Notation for Numeric Constants

Numeric Constant	Description
1	is an unsigned integer
-5	contains a minus sign
+49	contains a plus sign

Numeric Constant	Description
1.23	contains decimal places
01	contains a leading zero which is not significant

Numeric Constants Expressed in Scientific Notation

In scientific notation, the number before the E is multiplied by the power of ten that is indicated by the number after the E. For example, 2E4 is the same as 2×10^4 or 20,000. For numeric constants larger than $(10^{32})-1$, you must use scientific notation. Additional examples follow:

- 1.2e23
- 0.5e-10

Numeric Constants Expressed in Hexadecimal Notation

A numeric constant that is expressed as a hexadecimal value starts with a numeric digit (usually 0), can be followed by more hexadecimal characters, and ends with the letter X. The constant can contain up to 16 valid hexadecimal characters (0 to 9, A to F). The following are numeric hexadecimal constants:

- 0c1x
- 9x

You can use numeric hexadecimal constants in a DATA step, as follows:

```
data test;
  input abend pib2.;
  if abend=0c1x or abend=0b0ax then do;
    ... more SAS statements ...
  run;
```

Date, Time, and Datetime Constants

You can create a *date constant*, *time constant*, or *datetime constant* by specifying the date or time in single or double quotation marks, followed by a D (date), T (time), or DT (datetime) to indicate the type of value.

Any trailing blanks or leading blanks included within the quotation marks will not affect the processing of the *date constant*, *time constant*, or *datetime constant*.

Use the following patterns to create date and time constants:

'ddmmm<yy>yy'D or "ddmmm<yy>yy"D represents a SAS date value:

- date='1jan2006'd;
- date='01jan04'd;

'hh:mm<:ss.s>'T or "hh:mm<:ss.s>"T represents a SAS time value:

- time='9:25't;
- time='9:25:19pm't;

'ddmmm<yy>yy:hh:mm<:ss.s>'DT or "ddmmm<yy>yy:hh:mm<:ss.s>"DT represents a SAS datetime value:

- if begin='01may04:9:30:00'dt then end='31dec90:5:00:00'dt;
- dttime='18jan2003:9:27:05am'dt;

For more information on SAS dates, refer to Chapter 7, “Dates, Times, and Intervals,” on page 77.

Bit Testing Constants

Bit masks are used in bit testing to compare internal bits in a value’s representation. You can perform bit testing on both character and numeric variables. The general form of the operation is:

expression comparison-operator bit-mask

The following are the components of the bit-testing operation:

expression

can be any valid SAS expression. Both character and numeric variables can be bit tested. When SAS tests a character value, it aligns the left-most bit of the mask with the left-most bit of the string; the test proceeds through the corresponding bits, moving to the right. When SAS tests a numeric value, the value is truncated from a floating-point number to a 32-bit integer. The right-most bit of the mask is aligned with the right-most bit of the number, and the test proceeds through the corresponding bits, moving to the left.

comparison-operator

compares an expression with the bit mask. Refer to “Comparison Operators” on page 66 for a discussion of these operators.

bit-mask

is a string of 0s, 1s, and periods in quotation marks that is immediately followed by a B. Zeros test whether the bit is off; ones test whether the bit is on; and periods ignore the bit. Commas and blanks can be inserted in the bit mask for readability without affecting its meaning.

CAUTION:

Truncation can occur when SAS uses a bit mask. If the expression is longer than the bit mask, SAS truncates the expression before it compares it with the bit mask. A false comparison might result. An expression’s length (in bits) must be less than or equal to the length of the bit mask. If the bit mask is longer than a character expression, SAS prints a warning in the log, stating that the bit mask is truncated on the left, and continues processing. Δ

The following example tests a character variable:

```
if a='..1.0000'b then do;
```

If the third bit of A (counting from the left) is on, and the fifth through eighth bits are off, the comparison is true and the expression result is 1. Otherwise, the comparison is false and the expression result is 0. The following is a more detailed example:

```
data test;
  input @88 bits $char1.;
  if bits='10000000'b
    then category='a';
  else if bits='01000000'b
    then category='b';
  else if bits='00100000'b
    then category='c';
run;
```

Note: Bit masks cannot be used as bit literals in assignment statements. For example, the following statement is not valid:

```
x='0101'b;      /* incorrect */
```

Δ

The \$BINARY w . and BINARY w . formats and the \$BINARY w ., BINARY $w.d$., and BITS $w.d$ informats can be useful for bit testing. You can use them to convert character and numeric values to their binary values, and vice versa, and to extract specified bits from input data. See *SAS Language Reference: Dictionary* for complete descriptions of these formats and informats.

Avoiding a Common Error with Constants

When you use a string in quotation marks followed by a variable name, always put a blank space between the closing quotation mark and the variable name. Otherwise, SAS might interpret a character constant followed by a variable name as a special SAS constant as illustrated in this table.

Table 6.2 Characters That Cause Misinterpretation When Following a Character Constant

Characters That Follow a Character Constant	Possible Interpretation	Examples
b	bit testing constant	'00100000'b
d	date constant	'01jan04'd
dt	datetime constant	'18jan2005:9:27:05am'dt
n	name literal	'My Table'n
t	time constant	'9:25:19pm't
x	hexadecimal notation	'534153'x

In the following example, '821't is evaluated as a time constant. For more information about SAS time constants, see "Date, Time, and Datetime Constants" on page 61.

```
data work.europe;
  set ia.europe;
  if flight='821'then
    flight='230';
run;
```

The program writes the following lines to the SAS log:

Output 6.1 Log Results from an Error Caused by a Time Literal Misinterpretation

```
ERROR: Invalid date/time/datetime constant '821't.
ERROR 77--185: Invalid number conversion on '821't.

ERROR 388--185: Expecting an arithmetic operator.
```

Inserting a blank space between the ending quotation mark and the succeeding character in the IF statement eliminates this misinterpretation. No error message is

generated and all observations with a FLIGHT value of 821 are replaced with a value of 230.

```
if flight='821' then
  flight='230';
```

SAS Variables in Expressions

Definition

variable

is a set of data values that describe a given characteristic. A variable can be used in an expression.

Automatic Numeric-Character Conversion

If you specify a variable in an expression, but the variable value does not match the type called for, SAS attempts to convert the value to the expected type. SAS automatically converts character variables to numeric variables and numeric variables to character variables, according to the following rules:

- If you use a character variable with an operator that requires numeric operands, such as the plus sign, SAS converts the character variable to numeric.
- If you use a comparison operator, such as the equal sign, to compare a character variable and a numeric variable, the character variable is converted to numeric.
- If you use a numeric variable with an operator that requires a character value, such as the concatenation operator, the numeric value is converted to character using the BEST12. format. Because SAS stores the results of the conversion beginning with the right-most byte, you must store the converted values in a variable of sufficient length to accommodate the BEST12. format. You can use the LEFT function to left-justify a result.
- If you use a numeric variable on the left side of an assignment statement and a character variable on the right, the character variable is converted to numeric. In the opposite situation, where the character variable is on the left and the numeric is on the right, SAS converts the numeric variable to character using the BEST n . format, where n is the length of the variable on the left.

When SAS performs an automatic conversion, it prints a note in the SAS log informing you that the conversion took place. If converting a character variable to numeric produces invalid numeric values, SAS assigns a missing value to the result, prints an error message in the log, and sets the value of the automatic variable `_ERROR_` to 1.

Note: You can also use the PUT and INPUT functions to convert data values. These functions can be more efficient than automatic conversion. See “The Concatenation Operator” on page 72 for an example of the PUT function. See *SAS Language Reference: Dictionary* for more details on these functions. Δ

For more information on SAS variables, see Chapter 4, “SAS Variables,” on page 25 or *SAS Language Reference: Dictionary*.

SAS Functions in Expressions

A SAS *function* is a keyword that you use to perform a specific computation or system manipulation. Functions return a value, might require one or more arguments, and can be used in expressions. For further information on SAS functions, see *SAS Language Reference: Dictionary*.

SAS Operators in Expressions

Definitions

A SAS *operator* is a symbol that represents a comparison, arithmetic calculation, or logical operation; a SAS function; or grouping parentheses. SAS uses two major types of operators:

- prefix operators
- infix operators.

A *prefix operator* is an operator that is applied to the variable, constant, function, or parenthetical expression that immediately follows it. The plus sign (+) and minus sign (–) can be used as prefix operators. The word NOT and its equivalent symbols are also prefix operators. The following are examples of prefix operators used with variables, constants, functions, and parenthetical expressions:

- +y**
- 25**
- cos(angle1)**
- +(x*y)**

An *infix operator* applies to the operands on each side of it, for example, 6<8. Infix operators include the following:

- arithmetic
- comparison
- logical, or Boolean
- minimum
- maximum
- concatenation.

When used to perform arithmetic operations, the plus and minus signs are infix operators.

SAS also provides several other operators that are used only with certain SAS statements. The WHERE statement uses a special group of SAS operators, valid only when used with WHERE expressions. For a discussion of these operators, see Chapter 11, “WHERE-Expression Processing,” on page 133. The _NEW_ operator is used to create an instance of a DATA step component object. For more information, see Chapter 23, “Using DATA Step Component Objects,” on page 391.

Arithmetic Operators

Arithmetic operators indicate that an arithmetic calculation is performed, as shown in the following table:

Table 6.3 Arithmetic Operators

Symbol	Definition	Example	Result
**	exponentiation	a**3	raise A to the third power
*	multiplication ¹	2*y	multiply 2 by the value of Y
/	division	var/5	divide the value of VAR by 5
+	addition	num+3	add 3 to the value of NUM
-	subtraction	sale-discount	subtract the value of DISCOUNT from the value of SALE

¹ The asterisk (*) is always necessary to indicate multiplication; **2Y** and **2(Y)** are not valid expressions.

If a missing value is an operand for an arithmetic operator, the result is a missing value. See Chapter 5, “Missing Values,” on page 49 for a discussion of how to prevent the propagation of missing values.

See “Order of Evaluation in Compound Expressions” on page 73 for the order in which SAS evaluates these operators.

Comparison Operators

Comparison operators set up a comparison, operation, or calculation with two variables, constants, or expressions. If the comparison is true, the result is 1. If the comparison is false, the result is 0.

Comparison operators can be expressed as symbols or with their mnemonic equivalents, which are shown in the following table:

Table 6.4 Comparison Operators

Symbol	Mnemonic Equivalent	Definition	Example
=	EQ	equal to	a=3
^=	NE	not equal to ¹	a ne 3
≠	NE	not equal to	
~=	NE	not equal to	
>	GT	greater than	num>5
<	LT	less than	num<8
>=	GE	greater than or equal to ²	sales>=300

Symbol	Mnemonic Equivalent	Definition	Example
<=	LE	less than or equal to ³	<code>sales<=100</code>
	IN	equal to one of a list	<code>num in (3, 4, 5)</code>

- 1 The symbol you use for NE depends on your personal computer.
- 2 The symbol => is also accepted for compatibility with previous releases of SAS. It is not supported in WHERE clauses or in PROC SQL.
- 3 The symbol =< is also accepted for compatibility with previous releases of SAS. It is not supported in WHERE clauses or in PROC SQL.

See “Order of Evaluation in Compound Expressions” on page 73 for the order in which SAS evaluates these operators.

Note: You can add a colon (:) modifier to any of the operators to compare only a specified prefix of a character string. See “Character Comparisons” on page 69 for details. Δ

Note: You can use the IN operator to compare a value that is produced by an expression on the left of the operator to a list of values that are given on the right. The form of the comparison is:

```
expression IN (value-1<... ,value-n>)
```

The components of the comparison are as follows:

expression can be any valid SAS expression, but is usually a variable name when it is used with the IN operator.

value must be a constant.

For examples of using the IN operator, see “The IN Operator in Numeric Comparisons” on page 68. Δ

Numeric Comparisons

SAS makes numeric comparisons that are based on values. In the expression $A \leq B$, if A has the value 4 and B has the value 3, then $A \leq B$ has the value 0, or false. If A is 5 and B is 9, then the expression has the value 1, or true. If A and B each have the value 47, then the expression is true and has the value 1.

Comparison operators appear frequently in IF-THEN statements, as in this example:

```
if x<y then c=5;
   else c=12;
```

You can also use comparisons in expressions in assignment statements. For example, the preceding statements can be recoded as follows:

```
c=5*(x<y)+12*(x>=y);
```

Since SAS evaluates quantities inside parentheses before performing any operations, the expressions $(x < y)$ and $(x \geq y)$ are evaluated first and the result (1 or 0) is substituted for the expressions in parentheses. Therefore, if $X=6$ and $Y=8$, the expression evaluates as follows:

```
c=5*(1)+12*(0)
```

The result of this statement is $C=5$.

You might get an incorrect result when you compare numeric values of different lengths because values less than 8 bytes have less precision than those longer than 8

bytes. Rounding also affects the outcome of numeric comparisons. See Chapter 4, “SAS Variables,” on page 25 for a complete discussion of numeric precision.

A missing numeric value is smaller than any other numeric value, and missing numeric values have their own sort order. See Chapter 5, “Missing Values,” on page 49 for more information.

The IN Operator in Numeric Comparisons

You can use a shorthand notation to specify a range of sequential integers to search. The range is specified by using the syntax `M:N` as a value in the list to search, where `M` is the lower bound and `N` is the upper bound. `M` and `N` must be integers, and `M`, `N`, and all the integers between `M` and `N` are included in the range. For example, the following statements are equivalent.

- `y = x in (1, 2, 3, 4, 5, 6, 7, 8, 9, 10);`
- `y = x in (1:10);`

You can use multiple ranges in the same `IN` list, and you can use ranges with other constants in an `IN` list. The following example shows a range that is used with other constants to test if `X` is 0, 1, 2, 3, 4, 5, or 9.

```
if x in (0,9,1:5);
```

You can also use the `IN` operator to search an array of numeric values. For example, the following code creates an array `a`, defines a constant `x`, and then uses the `IN` operator to search for `x` in array `a`. Note that the array initialization syntax of `array a{10} (2*1:5)` creates an array that contains the initial values of 1, 2, 3, 4, 5, 1, 2, 3, 4, 5.

```
data _null_;
  array a{10} (2*1:5);
  x=99;
  y = x in a;
  put y=;
  a{5} = 99;
  y = x in a;
  put y=;
run;
```

Output 6.2 Results From Using the `IN` Operator to Search an Array of Numeric Values (partial output)

```
1 data _null_;
2   array a[10] (2*1:5);
3   x=99;
4   y = x in a;
5   put y=;
6   a[5] = 99;
7   y = x in a;
8   put y=;
9   run;
y=0
y=1
```

Note: PROC SQL does not support this syntax. \triangle

Character Comparisons

You can perform comparisons on character operands, but the comparison always yields a numeric result (1 or 0). Character operands are compared character by character from left to right. Character order depends on the *collating sequence*, usually ASCII or EBCDIC, used by your computer.

For example, in the EBCDIC and ASCII collating sequences, **G** is greater than **A**; therefore, this expression is true:

```
'Gray' > 'Adams'
```

Two character values of unequal length are compared as if blanks were attached to the end of the shorter value before the comparison is made. A blank, or missing character value, is smaller than any other printable character value. For example, because **.** is less than **h**, this expression is true:

```
'C. Jones' < 'Charles Jones'
```

Since trailing blanks are ignored in a comparison, **'fox '** is equivalent to **'fox'**. However, because blanks at the beginning and in the middle of a character value are significant to SAS, **' fox'** is not equivalent to **'fox'**.

You can compare only a specified prefix of a character expression by using a colon (:) after the comparison operator. SAS truncates the longer value to the length of the shorter value during the comparison. In the following example, the colon modifier after the equal sign tells SAS to look at only the first character of values of the variable LASTNAME and to select the observations with names beginning with the letter **S**:

```
if lastname='S';
```

Because printable characters are greater than blanks, both of the following statements select observations with values of LASTNAME that are greater than or equal to the letter **S**:

```
□ if lastname >= 'S';
□ if lastname >= 'S';
```

Note: If you compare a zero-length character value with any other character value in either an IN: comparison or an EQ: comparison, the two character values are not considered equal. The result always evaluates to 0, or false. Δ

The operations that are discussed in this section show you how to compare entire character strings and the beginnings of character strings. Several SAS character functions enable you to search for and extract values from within character strings. See *SAS Language Reference: Dictionary* for complete descriptions of all SAS functions.

The IN Operator in Character Comparisons

You can use the IN operator with character strings to determine whether a variable's value is among a list of character values. The following statements produce the same results:

```
□ if state in ('NY', 'NJ', 'PA') then region+1;
□ if state='NY' or state='NJ' or state='PA' then region+1;
```

You can also use the IN operator to search an array of character values. For example, the following code creates an array **a**, defines a constant **x**, and then uses the IN operator to search for **x** in array **a**.

```
data _null_;
  array a{5} $ (5*' ');
```

```

x='b1';
y = x in a;
put y=;
a{5} = 'b1';
y = x in a;
put y=;
run;

```

Output 6.3 Results From Using the IN Operator to Search an Array of Character Values (partial output)

```

1  data _null_;
2  array a{5} $ (5*' ');
3  x='b1';
4  y = x in a;
5  put y=;
6  a{5} = 'b1';
7  y = x in a;
8  put y=;
9  run;
y=0
y=1
NOTE: DATA statement used (Total process time):
      real time           0.01 seconds
      cpu time            0.01 seconds

```

Logical (Boolean) Operators and Expressions

Logical operators, also called *Boolean operators*, are usually used in expressions to link sequences of comparisons. The logical operators are shown in the following table:

Table 6.5 Logical Operators

Symbol	Mnemonic Equivalent	Example
&	AND	(a>b & c>d)
	OR ¹	(a>b or c>d)
!	OR	
∣	OR	
¬	NOT ²	not (a>b)
^	NOT	
~	NOT	

¹ The symbol you use for OR depends on your operating environment.

² The symbol you use for NOT depends on your operating environment.

See “Order of Evaluation in Compound Expressions” on page 73 for the order in which SAS evaluates these operators.

In addition, a numeric expression without any logical operators can serve as a Boolean expression. For an example of Boolean numeric expressions, see “Boolean Numeric Expressions” on page 72.

The AND Operator

If *both* of the quantities linked by AND are 1 (true), then the result of the AND operation is 1; otherwise, the result is 0. For example, in the following comparison:

```
a<b & c>0
```

the result is true (has a value of 1) only when both $A < B$ and $C > 0$ are 1 (true): that is, when A is less than B and C is positive.

Two comparisons with a common variable linked by AND can be condensed with an implied AND. For example, the following two subsetting IF statements produce the same result:

```
□ if 16<=age and age<=65;
□ if 16<=age<=65;
```

The OR Operator

If *either* of the quantities linked by an OR is 1 (true), then the result of the OR operation is 1 (true); otherwise, the OR operation produces a 0. For example, consider the following comparison:

```
a<b|c>0
```

The result is true (with a value of 1) when $A < B$ is 1 (true) regardless of the value of C . It is also true when the value of $C > 0$ is 1 (true), regardless of the values of A and B . Therefore, it is true when either or both of those relationships hold.

Be careful when using the OR operator with a series of comparisons (in an IF, SELECT, or WHERE statement, for example). Remember that only one comparison in a series of OR comparisons must be true to make a condition true, and any nonzero, nonmissing constant is always evaluated as true (see “Boolean Numeric Expressions” on page 72). Therefore, the following subsetting IF statement is always true:

```
if x=1 or 2;
```

SAS first evaluates $X=1$, and the result can be either true or false; however, since the 2 is evaluated as nonzero and nonmissing (true), the entire expression is true. In this statement, however, the condition is not necessarily true because either comparison can evaluate as true or false:

```
if x=1 or x=2;
```

The NOT Operator

The prefix operator NOT is also a logical operator. The result of putting NOT in front of a quantity whose value is 0 (false) is 1 (true). That is, the result of negating a false statement is 1 (true). For example, if $X=Y$ is 0 (false) then $\text{NOT}(X=Y)$ is 1 (true). The result of NOT in front of a quantity whose value is missing is also 1 (true). The result of NOT in front of a quantity with a nonzero, nonmissing value is 0 (false). That is, the result of negating a true statement is 0 (false).

For example, the following two expressions are equivalent:

```
□ not (name='SMITH' )
□ name ne 'SMITH'
```

Furthermore, NOT(A&B) is equivalent to NOT A|NOT B, and NOT(A|B) is the same as NOT A & NOT B. For example, the following two expressions are equivalent:

- `not(a=b & c>d)`
- `a ne b | c le d`

Boolean Numeric Expressions

In computing terms, a value of true is a 1 and a value of false is a 0. In SAS, any numeric value other than 0 or missing is true, and a value of 0 or missing is false. Therefore, a numeric variable or expression can stand alone in a condition. If its value is a number other than 0 or missing, the condition is true; if its value is 0 or missing, the condition is false.

```
0 | . = False
1 = True
```

For example, suppose that you want to fill in variable REMARKS depending on whether the value of COST is present for a given observation. You can write the IF-THEN statement as follows:

```
if cost then remarks='Ready to budget';
```

This statement is equivalent to:

```
if cost ne . and cost ne 0
  then remarks='Ready to budget';
```

A numeric expression can be simply a numeric constant, as follows:

```
if 5 then do;
```

The numeric value that is returned by a function is also a valid numeric expression:

```
if index(address,'Avenue') then do;
```

The MIN and MAX Operators

The MIN and MAX operators are used to find the minimum or maximum value of two quantities. Surround the operators with the two quantities whose minimum or maximum value you want to know. The MIN (><) operator returns the lower of the two values. The MAX (<>) operator returns the higher of the two values. For example, if A<B, then A><B returns the value of A.

If missing values are part of the comparison, SAS uses the sorting order for missing values that is described in “Order of Missing Values” on page 51. For example, the maximum value that is returned by .A<>.Z is the value .Z.

Note: In a WHERE statement or clause, the <> operator is equivalent to NE. Δ

The Concatenation Operator

The concatenation operator (||) concatenates character values. The results of a concatenation operation are usually stored in a variable with an assignment statement, as in `level='grade' || 'A'`. The length of the resulting variable is the sum of the lengths of each variable or constant in the concatenation operation, unless you use a LENGTH or ATTRIB statement to specify a different length for the new variable.

The concatenation operator does not trim leading or trailing blanks. If variables are padded with trailing blanks, check the lengths of the variables and use the TRIM

function to trim trailing blanks from values before concatenating them. See *SAS Language Reference: Dictionary* for descriptions and examples of additional character functions.

For example, in this DATA step, the value that results from the concatenation contains blanks because the length of the COLOR variable is eight:

```
data namegame;
  length color name $8 game $12;
  color='black';
  name='jack';
  game=color||name;
  put game=;
run;
```

The value of GAME is '**black jack**'. To correct this problem, use the TRIM function in the concatenation operation as follows:

```
game=trim(color)||name;
```

This statement produces a value of '**blackjack**' for the variable GAME. The following additional examples demonstrate uses of the concatenation operator:

- If A has the value '**fortune**', B has the value '**five**', and C has the value '**hundred**', then the following statement produces the value '**fortunefivehundred**' for the variable D:

```
d=a||b||c;
```

- This example concatenates the value of a variable with a character constant.

```
newname='Mr. or Ms. '||oldname;
```

If the value of OLDNAME is '**Jones**', then NEWNAME will have the value '**Mr. or Ms. Jones**'.

- Because the concatenation operation does not trim blanks, the following expression produces the value '**JOHN SMITH**':

```
name='JOHN '||'SMITH';
```

- This example uses the PUT function to convert a numeric value to a character value. The TRIM function is used to trim blanks.

```
month='sep ';
year=99;
date=trim(month) || left(put(year,8.));
```

The value of DATE is the character value '**sep99**'.

Order of Evaluation in Compound Expressions

Table 6.6 on page 74 shows the order of evaluation in compound expressions. The table contains the following columns:

Priority

lists the priority of evaluation. In compound expressions, SAS evaluates the part of the expression containing operators in Group I first, then each group in order.

Order of Evaluation

lists the rules governing which part of the expression SAS evaluates first.

Parentheses are often used in compound expressions to group operands; expressions within parentheses are evaluated before those outside of them. The

rules also list how a compound expression that contains more than one operator from the same group is evaluated.

Symbols

lists the symbols that you use to request the comparisons, operations, and calculations.

Mnemonic Equivalent

lists alternate forms of the symbol. In some cases, such as when your keyboard does not support special symbols, you should use the alternate form.

Definition

defines the symbol.

Example


provides an example of how to use the symbol or mnemonic equivalent in a SAS expression.

Table 6.6 Order of Evaluation in Compound Expressions

Priority	Order of Evaluation	Symbols	Mnemonic Equivalent	Definition	Example
Group I	right to left	**		exponentiation ¹	y=a**2;
		+		positive prefix ²	y+=(a*b);
		-		negative prefix ³	z=-(a+b);
		^ ~	NOT	logical not ⁴	if not z then put x;
		><	MIN	minimum ⁵	x=(a<b);
		<>	MAX	maximum	x=(a>b);
Group II	left to right	*		multiplication	c=a*b;
		/		division	f=g/h;
Group III	left to right	+		addition	c=a+b;
		-		subtraction	f=g-h;
Group IV	left to right	!!		concatenate character values ⁶	name= 'J' 'SMITH';
Group V ⁷	left to right ⁸	<	LT	less than	if x<y then c=5;
		<=	LE	less than or equal to	if x le y then a=0;
		=	EQ	equal to	if y eq (x+a) then output;
		≠	NE	not equal to	if x ne z then output;
		>=	GE	greater than or equal to	if y>=a then output;

Priority	Order of Evaluation	Symbols	Mnemonic Equivalent	Definition	Example
		>	GT	greater than	if z>a then output;
			IN	equal to one of a list	if state in ('NY', 'NJ', 'PA') then region='NE'; y = x in (1:10);
Group VI	left to right	&	AND	logical and	if a=b & c=d then x=1;
Group VII	left to right	; !	OR	logical or ⁹	if y=2 or x=3 then a=d;

- 1 Because Group I operators are evaluated from right to left, the expression $x=2**3**4$ is evaluated as $x=(2**(3**4))$.
- 2 The plus (+) sign can be either a prefix or arithmetic operator. A plus sign is a prefix operator only when it appears at the beginning of an expression or when it is immediately preceded by an open parenthesis or another operator.
- 3 The minus (-) sign can be either a prefix or arithmetic operator. A minus sign is a prefix operator only when it appears at the beginning of an expression or when it is immediately preceded by an open parenthesis or another operator.
- 4 Depending on the characters available on your keyboard, the symbol can be the not sign (\neg), tilde (\sim), or caret (^). The SAS system option CHARCODE allows various other substitutions for unavailable special characters.
- 5 For example, the SAS System evaluates $-3><-3$ as $-(3><-3)$, which is equal to $-(-3)$, which equals $+3$. This is because Group I operators are evaluated from right to left.
- 6 Depending on the characters available on your keyboard, the symbol you use as the concatenation operator can be a double vertical bar (||), broken vertical bar (;|), or exclamation mark (!!).
- 7 Group V operators are comparison operators. The result of a comparison operation is 1 if the comparison is true and 0 if it is false. Missing values are the lowest in any comparison operation. The symbols =< (less than or equal to) are also allowed for compatibility with previous versions of the SAS System. When making character comparisons, you can use a colon (:) after any of the comparison operators to compare only the first character or characters of the value. SAS truncates the longer value to the length of the shorter value during the comparison. For example, if **name=: 'P'** compares the value of the first character of NAME to the letter P.
- 8 An exception to this rule occurs when two comparison operators surround a quantity. For example, the expression $x<y<z$ is evaluated as $(x<y)$ and $(y<z)$.
- 9 Depending on the characters available on your keyboard, the symbol you use for the logical or can be a single vertical bar (|), broken vertical bar (;|), or exclamation mark (!). You can also use the mnemonic equivalent OR.



CHAPTER

7

Dates, Times, and Intervals

<i>About SAS Date, Time, and Datetime Values</i>	77
<i>Definitions</i>	77
<i>Two-Digit and Four-Digit Years</i>	78
<i>The Year 2000</i>	78
<i>Using the YEARCUTOFF= System Option</i>	78
<i>Example: How YEARCUTOFF= Affects Two and Four-Digit Years</i>	79
<i>Practices That Help Ensure Date Integrity</i>	80
<i>Working with SAS Dates and Times</i>	80
<i>Informats and Formats</i>	80
<i>Date and Time Tools by Task</i>	81
<i>Examples</i>	86
<i>Example 1: Displaying Date, Time, and Datetime Values as Recognizable Dates and Times</i>	86
<i>Example 2: Reading, Writing, and Calculating Date Values</i>	87
<i>About Date and Time Intervals</i>	88
<i>Definitions</i>	88
<i>Syntax</i>	88
<i>Intervals By Category</i>	89
<i>Example: Calculating a Duration</i>	91
<i>Boundaries of Intervals</i>	91
<i>Single-Unit Intervals</i>	92
<i>Multi-Unit Intervals</i>	93
<i>Multi-Unit Intervals Other Than Multi-Week Intervals</i>	93
<i>Multi-Week Intervals</i>	93
<i>Shifted Intervals</i>	94
<i>How to Use Shifted Intervals</i>	94
<i>How SAS Creates Shifted Intervals</i>	95
<i>Custom Intervals</i>	95
<i>Retail Calendar Intervals: ISO 8601 Compliant</i>	95

About SAS Date, Time, and Datetime Values

Definitions

SAS date value

is a value that represents the number of days between January 1, 1960, and a specified date. SAS can perform calculations on dates ranging from A.D. 1582 to A.D. 19,900. Dates before January 1, 1960, are negative numbers; dates after are positive numbers.

YEARCUTOFF= option to work with date ranges for your data, or switch to four-digit years. The following program changes the YEARCUTOFF= value to 1950. This change means that all two digit dates are now assumed to fall in the 100-year span from 1950 to 2049.

```
options yearcutoff=1950;
data _null_;
  a='26oct02'd;
  put 'SAS date='a;
  put 'formatted date='a date9.;
run;
```

The PUT statement writes the following lines to the SAS log:

```
SAS date=15639
formatted date=26OCT2002
```

Note: Whenever possible, specify a year using all four digits. Most SAS date and time language elements support four digit year values. Δ

Example: How YEARCUTOFF= Affects Two and Four-Digit Years

The following example shows what happens with data that contains both two and four-digit years. Note how the YEARCUTOFF= option is set to 1920.

```
options yearcutoff=1920 nodate pageno=1 linesize=80 pagesize=60;

data schedule;
  input @1 jobid $ @6 projdate mmddy10.;
  datalines;
A100 01/15/25
A110 03/15/2025
A200 01/30/96
B100 02/05/00
B200 06/15/2000
;

proc print data=schedule;
  format projdate mmddy10.;
run;
```

The resulting output from the PROC PRINT statement looks like this:

Output 7.1 Output from The Previous DATA Step Showing Four-Digit Years That Result from Setting YEARCUTOFF= to 1920

	The SAS System		1
Obs	jobid	projdate	
1	A100	01/15/1925	
2	A110	03/15/2025	
3	A200	01/30/1996	
4	B100	02/05/2000	
5	B200	06/15/2000	

Here are some facts to note in this example:

- In the data lines in the DATA step, the first record contains a two-digit year of 25, and the second record contains a four-digit year of 2025. Because the YEARCUTOFF= system option is set to 1920, the two-digit year defaults to a year in the 1900s in observation number 1. The four-digit year in observation number 2 is unaffected by the YEARCUTOFF= option.
- The third record is similar to the first and defaults to a year in the 1900s based on the value of YEARCUTOFF=.
- The output from records 4 and 5 shows results that are similar to records 1 and 2. The fourth record specifies a two-digit year of 00, and the fifth one specifies a four-digit year of 2000. Because of the value of the YEARCUTOFF= option, the years in the two resulting observations are the same.

As you can see, specifying a two-digit year might or might not result in the intended century prefix. The optimal value of the YEARCUTOFF= option depends on the range of the dates that you are processing.

In Releases 6.06 through 6.12 of SAS, the default value for the YEARCUTOFF= system option is 1900; starting with SAS 7, the default value is 1920.

For more information on how SAS handles dates, see the section on dates, times, and datetime values.

Practices That Help Ensure Date Integrity

The following practices help ensure that your date values are correct during all the conversions that occur during processing:

- Store dates as SAS date values, not as simple numeric or character values.
- Use the YEARCUTOFF= system option when converting two-digit dates to SAS date values.
- Examine sets of raw data coming into your SAS process to make sure that any dates containing two-digit years will be correctly interpreted by the YEARCUTOFF= system option. Look out for
 - two-digit years that are distributed over more than a 100-year period. For dates covering more than a 100-year span, you must either use four digit years in the data, or use conditional logic in a DATA step to interpret them correctly.
 - two-digit years that need an adjustment to the default YEARCUTOFF= range. For example, if the default value for YEARCUTOFF= in your operating environment is 1920 and you have a two-digit date in your data that represents 1919, you will have to adjust your YEARCUTOFF= value downward by a year in the SAS program that processes this value.
- Make sure that output SAS data sets represent dates as SAS date values.
- Check your SAS programs to make sure that formats and informats that use two-digit years, such as DATE7., MMDDYY6., or MMDDYY8., are reading and writing data correctly.

Note: The YEARCUTOFF= option has no effect on dates that are already stored as SAS date values. \triangle

Working with SAS Dates and Times

Informats and Formats

SAS converts date, time, and datetime values back and forth between calendar dates and clock times with SAS language elements called *formats* and *informats*.

- Formats present a value, recognized by SAS, such as a time or date value, as a calendar date or clock time in a variety of lengths and notations.
- Informats read notations or a value, such as a clock time or a calendar date, which might be in a variety of lengths, and then convert the data to a SAS date, time, or datetime value.

SAS can read date and time values that are delimited by the following characters:

! # \$ % & () * + - . / : ; < = > ? [\] ^ _ { | } ~

The blank character can also be used.

Only one delimiter can be used for a date. Otherwise, an error message is written to the SAS log. For example, 01/Jan/2007 uses a single delimiter, and can be read by SAS. In the case of 01-Jan/2007, two different delimiters separate the parts of the date, which results in an error message.

Date and Time Tools by Task

The following table correlates tasks with various SAS language elements that are available for working with time and date data.

Table 7.1 Tasks with Dates and Times, Part 1

Task	Type of Language Element	Language Element	Input	Result
Write SAS date values in recognizable forms	Date formats	DATE.	14686	17MAR00
		DATE9.	14686	17MAR2000
		DAY.	14686	17
		DDMMYY.	14686	17/03/00
		DDMMYY10.	14686	17/03/2000
		DDMMYYB.	14686	17 03 00
		DDMMYYB10.	14686	17 03 2000
		DDMMYYC.	14686	17:03:20
		DDMMYYC10.	14686	17:03:2000
		DDMMYYD.	14686	17-03-00
		DDMMYYD10.	14686	17-03-2000
		DDMMYYN.	14686	17032000
		DDMMYYN6.	14686	170300
		DDMMYYP.	14686	17.03.00
		DDMMYYP10.	14686	17.03.2000
		DDMMYYYS.	14686	17/03/00
		DDMMYYYS10.	14686	17/03/2000
DOWNAME.	14686	Friday		
JULDAY.	14686	77		
JULIAN.	14686	00077		
MMDDYY.	14686	03/17/00		

Task	Type of Language Element	Language Element	Input	Result
		MMDDYY10.	14686	03/17/2000
		MMDDYYB.	14686	03 17 00
		MMDDYYB10.	14686	03 17 2000
		MMDDYYC.	14686	03:17:00
		MMDDYYC10.	14686	03:17:2000
		MMDDYYD.	14686	03-17-00
		MMDDYYD10.	14686	03-17-2000
		MMDDYYN.	14686	031700
		MMDDYYN8.	14686	03172000
		MMDDYYP.	14686	03.17.00
		MMDDYYP10.	14686	03.17.2000
		MMDDYYYS.	14686	03/17/00
		MMDDYYYS10.	14686	03/17/2000
		MMYY.	14686	03M2000
		MMYYC.	14686	03:2000
		MMYYD.	14686	03-2000
		MMYYN.	14686	032000
		MMYYP.	14686	03.2000
		MMYYYS.	14686	03/2000
		MONNAME.	14686	March
		MONTH.	14686	3
		MONYY.	14686	MAR2000
		PDJULG.	14686	2000077F
		PDJULI.	14686	0100077F
		WEEKDATE.	14686	Friday, March 17, 2000
		WEEKDAY.	14686	6
		WORDDATE.	14686	March 17, 2000
		WORDDATX.	14686	17 MARCH 2000
	Quarter formats	QTR.	14686	1
		QTRR.	14686	I
	Time formats	TIME.	14686	4:04:46
		TIMEAMP.	14686	4:04:46 AM

Task	Type of Language Element	Language Element	Input	Result
		TOD.	14686	04:04:46
	Year formats	YEAR.	14686	2000
		YYMM.	14686	2000M03
		YYMMC.	14686	2000:03
		YYMMD.	14686	2000-03
		YYMMP.	14686	2000.03
		YYMMS.	14686	2000/03
		YYMMN.	14686	200003
		YYMMDD.	14686	00-03-17
		YYMON.	14686	2000MAR
	Year/Quarter formats	YYQ.	14686	2000Q1
		YYQC.	14686	2000:1
		YYQD.	14686	2000-1
		YYQP.	14686	2000.1
		YYQS.	14686	2000/1
		YYQN.	14686	20001
		YYQR.	14686	2000QI
		YYQRC.	14686	2000:I
		YYQRD.	14686	2000-I
		YYQRP.	14686	2000.I
		YYQRS.	14686	2000/I
		YYQRN.	14686	III

Table 7.2 Tasks with Dates and Times, Part 2

Task	Type of Language Element	Language Element	Input	Result
Date Tasks				
Read calendar dates as SAS date	Date informats	DATE.	17MAR2000	14686
Note:				
YEARCUTOFF=1920		DATE9.	17MAR2000	14686
		DDMMYY.	170300	14686
		DDMMYY8.	17032000	14686
		JULIAN.	0077	14686

Task	Type of Language Element	Language Element	Input	Result
		JULIAN7.	2000077	14686
		MMDDYY.	031700	14686
		MMDDYY10.	03172000	14686
		MONYY.	MAR00	14670
		YYMMDD.	000317	14686
		YYMMDD10.	20000317	14686
		YYQ.	00Q1	14610
Create date values from pieces	Date functions	DATEJUL	2000077	14686
		DATETIME	'17MAR2000'D, 00,00,00	1268870400
		TIME	14,45,32	53132
		MDY	03,17,00	14686
		MDY	03,17,2000	14686
		YYQ	00,1	14610
Extract a date from a datetime value	Date functions	DATEPART	'17MAR00:00:00 'DT	14686
Return today's date as a SAS date	Date functions	DATE() or TODAY() (equivalent)	()	SAS date for today
Extract calendar dates from SAS	Date functions	DAY	14686	17
		HOUR	14686	4
		JULDATE	14686	77
		JULDATE7	14686	2000077
		MINUTE	14686	4
		MONTH	14686	3
		QTR	14686	3
		SECOND	14686	46
		WEEKDAY	14686	6
		YEAR	14686	2000
Write a date as a constant in an expression	SAS date constant	'ddmmyy'd or 'ddmmyyyy'	'17mar00'd '17mar2000'd	14686

Task	Type of Language Element	Language Element	Input	Result	
Write today's date as a string	SYSDATE automatic macro variable	SYSDATE	&SYSDATE	Date at time of SAS initialization in DDMMYY	
	SYSDATE9	SYSDATE9	&SYSDATE9	Date at time of SAS initialization in DDMMYYYY	
Time Tasks					
Write SAS time values as time values	time formats	HHMM.	53132	14:46	
		HOURL.	53132	15	
		MMSS.	53132	885	
		TIME.	53132	14:45:32	
		TIMEAMPM.	53132	2:45:32 PM	
Read time values as SAS time values	Time informats	TOD.	53132	14:45:32	
		TIME.	14:45:32	53132	
		SYSTIME automatic macro variable	SYSTIME	&SYSTIME	Time at moment of execution in HH:MM
		Time functions	TIME()	()	SAS time value at moment of execution in NNNNN.NN
		Time functions	TIMEPART	SAS datetime value in NNNNNNNNN.N	SAS time value part of date value in NNNNN.NN
Datetime Tasks					
Write SAS datetime values as datetime values	Datetime formats	DATEAMPM	1217083532	26JUL98:02:45 PM	
		DATETIME	1268870400	17MAR00:00:00 :00	
Read datetime values as SAS datetime values	Datetime informats	DATETIME	17MAR00:00:00:00	1268870400	

Task	Type of Language Element	Language Element	Input	Result
Return the current date and time of day as a SAS datetime value	Datetime functions	DATETIME()	()	SAS datetime value at moment of execution in NNNNNNNNN.N
Interval Tasks				
Return the number of specified time intervals that lie between the two date or datetime values	Interval functions	INTCK	week 2 01aug60 01jan01	1055
Advances a date, time, or datetime value by a given interval, and returns a date, time, or datetime value	Interval functions	INTNX	day 14086 01jan60	14086

SAS also supports international formats and informats that are equivalent to some of the most commonly used English-language date formats and informats. For details, see the SAS formats and informats in *SAS Language Reference: Dictionary*.

Examples

Example 1: Displaying Date, Time, and Datetime Values as Recognizable Dates and Times

The following example demonstrates how a value might be displayed as a date, a time, or a datetime. Remember to select the SAS language element that converts a SAS date, time, or datetime value to the intended date, time or datetime format. See the previous tables for examples.

Note:

- Time formats count the number of seconds within a day, so the values will be between 0 and 86400.
- DATETIME formats count the number of seconds since January 1, 1960, so for datetimes that are greater than 02JAN1960:00:00:01 (integer of 86401), the datetime value will always be greater than the time value.
- When in doubt, look at the contents of your data set for clues as to which type of value you are dealing with.

Δ

This program uses the DATETIME, DATE, and TIMEAMPM formats to display the value 86399 to a date and time, a calendar date, and a time.

```
data test;
  Time1=86399;
  format Time1 datetime.;
  Date1=86399;
```

```

format Datel date.;
Time2=86399;
format Time2 timeampm.;
run;
proc print data=test;
  title 'Same Number, Different SAS Values';
  footnote1 'Time1 is a SAS DATETIME value';
  footnote2 'Datel is a SAS DATE value';
  footnote3 'Time2 is a SAS TIME value'.;
run;
footnote;

```

Output 7.2 Datetime, Date and Time Values for 86399

Same Number, Different SAS Values				1
Obs	Time1	Datel	Time2	
1	01JAN60:23:59:59	20JUL96	11:59:59 PM	
Time1 is a SAS DATETIME value Datel is a SAS DATE value Time2 is a SAS TIME value.				

Example 2: Reading, Writing, and Calculating Date Values

This program reads four regional meeting dates and calculates the dates on which announcements should be mailed.

```

data meeting;

  input region $ mtg : mmddyy8.;
  sendmail=mtg-45;
  datalines;
N 11-24-99
S 12-28-99
E 12-03-99
W 10-04-99
;

proc print data=meeting;
  format mtg sendmail date9.;
  title 'When To Send Announcements';
run;

```

Output 7.3 Calculated Date Values: When to Send Mail

When To Send Announcements				
Obs	region	mtg	sendmail	
1	N	24NOV1999	10OCT1999	
2	S	28DEC1999	13NOV1999	
3	E	03DEC1999	19OCT1999	
4	W	04OCT1999	20AUG1999	

About Date and Time Intervals

Definitions

duration

is an integer representing the difference between any two dates or times or datetimes. Date durations are integer values representing the difference, in the number of days, between two SAS dates. Time durations are decimal values representing the number of seconds between two times or datetimes.

Tip: Date and datetimes durations can be easily calculated by subtracting the smaller date or datetime from the larger. When dealing with SAS times, special care must be taken if the beginning and the end of a duration are on different calendar days. Whenever possible, the simplest solution is to use datetimes rather than times.

interval

is a unit of measurement that SAS can count within an elapsed period of time, such as DAYS, MONTHS, or HOURS. SAS determines date and time intervals based on fixed points on the calendar, the clock, or both. The starting point of an interval calculation defaults to the beginning of the period in which the beginning value falls, which might not be the actual beginning value specified. For example, if you are using the INTCK function to count the months between two dates, regardless of the actual day of the month specified by the date in the beginning value, SAS treats it as the first of that month.

Syntax

SAS provides date, time, and datetime intervals for counting different periods of elapsed time. You can create multiples of the intervals and shift their starting point. Use them with the INTCK and INTNX functions and with procedures that support numbered lists (such as the PLOT procedure). The form of an interval is

name<*multiple*><.starting-point>

The terms in an interval have the following definitions:

name

is the name of the interval. See the following table for a list of intervals and their definitions.

multiple

creates a multiple of the interval. *Multiple* can be any positive number. The default is 1. For example, YEAR2 indicates a two-year interval.

.starting-point

is the starting point of the interval. By default, the starting point is 1. A value greater than 1 shifts the start to a later point within the interval. The unit for shifting depends on the interval, as shown in the following table. For example, YEAR.3 specifies a yearly period from the first of March through the end of February of the following year.

Intervals By Category

Table 7.3 Intervals Used with Date and Time Functions

Category	Interval	Definition	Default Starting Point	Shift Period	Example	Description
Date	DAY	Daily intervals	Each day	Days	DAY3	Three-day intervals starting on Sunday
	WEEK	Weekly intervals of seven days	Each Sunday	Days (1=Sunday ... 7=Saturday)	WEEK.7	Weekly with Saturday as the first day of the week
	WEEKDAY <daysW>	Daily intervals with Friday-Saturday-Sunday counted as the same day (five-day work week with a Saturday-Sunday weekend). <i>Days</i> identifies the weekend days by number (1=Sunday ... 7=Saturday). By default, <i>days</i> =17.	Each day	Days	WEEKDAY1W WEEKDAY35W	Six-day week with Sunday as a weekend day Five-day week with Tuesday and Thursday as weekend days (W indicates that day 3 and day 5 are weekend days)

Category	Interval	Definition	Default Starting Point	Shift Period	Example	Description
	TENDAY	Ten-day intervals (a U.S. automobile industry convention)	First, eleventh, and twenty-first of each month	Ten-day periods	TENDAY4.2	Four ten-day periods starting at the second TENDAY period
	SEMIMONTH	Half-month intervals	First and sixteenth of each month	Semi-monthly periods	SEMIMONTH2.2	Intervals from the sixteenth of one month through the fifteenth of the next month
	MONTH	Monthly intervals	First of each month	Months	MONTH2.2	February-March, April-May, June-July, August-September, October-November, and December-January of the following year
	QTR	Quarterly (three-month) intervals	January 1 April 1 July 1 October 1	Months	QTR3.2	Three-month intervals starting on April 1, July 1, October 1, and January 1
	SEMIYEAR	Semiannual (six-month) intervals	January 1 July 1	Months	SEMIYEAR.3	Six-month intervals, March-August and September-February
	YEAR	Yearly intervals	January 1	Months		
Datetime	Add DT to any of the date intervals	Interval corresponding to the associated date interval	Midnight of January 1, 1960		DTMONTH DTWEEKDAY	

Category	Interval	Definition	Default Starting Point	Shift Period	Example	Description
Time	SECOND	Second intervals	Start of the day (midnight)	Seconds		
	MINUTE	Minute intervals	Start of the day (midnight)	Minutes		
	HOUR	Hourly intervals	Start of the day (midnight)	Hours		

Example: Calculating a Duration

This program reads the project start and end dates and calculates the duration between them.

```

data projects;
options nodate pageno=1 linesize=80 pagesize=60;
  input Projid startdate date9. enddate date9.;
  Duration=enddate-startdate;
  datalines;
398 17oct1997 02nov1997
942 22jan1998 10mar1998
167 15dec1999 15feb2000
250 04jan2001 11jan2001
;

proc print data=projects;
  format startdate enddate date9.;
  title 'Days Between Project Start and Project End';
run;

```

Output 7.4 Output from the PRINT Procedure

Days Between Project Start and Project End run						8
Obs	Projid	Startdate	Enddate	Duration		
1	398	17OCT1997	02NOV1997	16		
2	942	22JAN1998	10MAR1998	47		
3	167	15DEC1999	15FEB2000	62		
4	250	04JAN2001	11JAN2001	7		

Boundaries of Intervals

SAS associates date and time intervals with fixed points on the calendar. For example, the MONTH interval represents the time from the beginning of one calendar

month to the next, not a period of 30 or 31 days. When you use date and time intervals (for example, with the INTCK or INTNX functions), SAS bases its calculations on the calendar divisions that are present. Consider the following examples:

Table 7.4 Using INTCK And INTNX

Example	Results	Explanation
<code>mnthnum1=intck('month', '25aug2000'd, '05sep2000'd);</code>	mnthnum1=1	The number of MONTH intervals the INTCK function counts depends on whether the first day of a month falls within the period.
<code>mnthnum2=intck('month', '01aug2000'd, '31aug2000'd);</code>	mnthnum2=0	
<code>next=intnx('month', '25aug2000'd,1);</code>	next represents 01sep2000	The INTNX function produces the SAS date value that corresponds to the beginning of the next interval.

Note: The only intervals that do not begin on the same date in each year are WEEK and WEEKDAY. A Sunday can occur on any date because the year is not divided evenly into weeks. \triangle

Single-Unit Intervals

Single-unit intervals begin at the following points on the calendar:

Table 7.5 Single-Unit Intervals

Single-Unit Interval	Beginning Point on the Calendar
DAY	each day
WEEKDAY	for a standard weekday: <i>Start day–End day</i> Monday–Monday Tuesday–Tuesday Wednesday–Wednesday Thursday–Thursday Friday–Sunday
WEEK	each Sunday
TENDAY	the first, eleventh, and twenty-first of each month
SEMIMONTH	the first and sixteenth of each month
MONTH	the first of each month
QTR	the first of January, April, July, and October
SEMIYEAR	the first of January and July
YEAR	the first of January

Single-unit time intervals begin as follows:

Table 7.6 Single-Unit Time Intervals

Single-Unit Time Intervals	Beginning Point
SECOND	each second
MINUTE	each minute
HOUR	each hour

Multi-Unit Intervals

Multi-Unit Intervals Other Than Multi-Week Intervals

Multi-unit intervals, such as MONTH2 or DAY50, also depend on calendar measures, but they introduce a new problem: SAS can find the beginning of a unit (for example, the first of a month), but where does that unit fall in the interval? For example, does the first of October mark the first or the second month in a two-month interval?

For all multi-unit intervals except multi-week intervals, SAS creates an interval beginning on January 1, 1960, and counts forward from that date to determine where individual intervals begin on the calendar. As a practical matter, when a year can be divided evenly by an interval, think of the intervals as beginning with the current year. Thus, MONTH2 intervals begin with January, March, May, July, September, and November. Consider this example:

Table 7.7 Month2 Intervals

SAS statements	Results
<code>howmany1=intck('month2','15feb2000'd, '15mar2000'd);</code>	howmany1=1
<code>count=intck('day50','01oct1998'd, '01jan1999'd);</code>	count=1

In the above example, SAS counts 50 days beginning with January 1, 1960; then another 50 days; and so on. As part of this count, SAS counts one DAY50 interval between October 1, 1998 and January 1, 1999. As an example, to determine the date on which the next DAY50 interval begins, use the INTNX function, as follows:

Table 7.8 Using the INTNX Function

SAS statements	Results
<code>start=intnx('day50','01oct98'd,1);</code>	SAS date value 14200, or Nov 17, 1998

The next interval begins on November 17, 1998.

Time intervals (those that represent divisions of a day) are aligned with the start of the day, that is, midnight. For example, HOUR8 intervals divide the day into the periods 00:00 to 08:00, 8:00 to 16:00, and 16:00 to 24:00 (the next midnight).

Multi-Week Intervals

Multi-week intervals, such as WEEK2, present a special case. In general, weekly intervals begin on Sunday, and SAS counts a week whenever it passes a Sunday.

However, SAS cannot calculate multi-week intervals based on January 1, 1960, because that date fell on a Friday, as shown:

Figure 7.2 Calculating Multi-Week Intervals

Dec	Su	Mo	Tu	We	Th	Fr	Sa	Jan
1959	27	28	29	30	31	1	2	1960

Therefore, SAS begins the first interval on Sunday of the week containing January 1, 1960—that is, on Sunday, December 27, 1959. SAS counts multi-week intervals from that point. The following example counts the number of two-week intervals in the month of August 1998:

Table 7.9 Counting Two-Week Intervals

SAS Statements	Results
<code>count=intck('week2', '01aug98'd, '31aug98'd);</code>	count=3

To see the beginning date of the next interval, use the INTNX function, as shown here:

Table 7.10 Using INTNX to See the Beginning Date of an Interval

SAS Statements	Results
<code>begin=intnx('week2', '01aug1998'd, 1);</code>	“Begin” represents SAS date 14093 or August 02, 1998

The next interval begins on August 16.

Shifted Intervals

Shifting the beginning point of an interval is useful when you want to make the interval represent a period in your data. For example, if your company’s fiscal year begins on July 1, you can create a year beginning in July by specifying the YEAR.7 interval. Similarly, you can create a period matching U.S. presidential elections by specifying the YEAR4.11 interval. This section discusses how to use shifted intervals and how SAS creates them.

How to Use Shifted Intervals

When you shift a time interval by a subperiod, the shift value must be less than or equal to the number of subperiods in the interval. For example, YEAR.12 is valid (yearly periods beginning in December), but YEAR.13 is not. Similarly, YEAR2.25 is not valid because there is no twenty-fifth month in the two-year period.

In addition, you cannot shift an interval by itself. For example, you cannot shift the interval MONTH because the shifting subperiod for MONTH is one month and MONTH contains only one monthly subperiod. However, you can shift multi-unit intervals by the subperiod. For example, MONTH2.2 specifies bimonthly periods starting on the first day of the second month.

How SAS Creates Shifted Intervals

For all intervals except those based on weeks, SAS creates shifted intervals by creating the interval based on January 1, 1960, by moving forward the required number of subperiods, and by counting shifted intervals from that point. For example, suppose you create a shifted interval called DAY50.5. SAS creates a 50-day interval in which January 1, 1960 is day 1. SAS then moves forward to day 5. (Note that the *difference*, or amount of movement, is four days.) SAS begins counting shifted intervals from that point. The INTNX function demonstrates that the next interval begins on January 5, 1960:

Table 7.11 Using INTNX to Determine When an Interval Begins

SAS Statements	Results
<code>start=intnx('day50.5','01jan1960'd,1);</code>	SAS date value 4, or Jan 5, 1960

For shifted intervals based on weeks, SAS first creates an interval based on Sunday of the week containing January 1, 1960 (that is, December 27, 1959), then moves forward the required number of days. For example, suppose you want to create the interval WEEK2.8 (biweekly periods beginning on the second Sunday of the period). SAS measures a two-week interval based on Sunday of the week containing January 1, 1960, and begins counting shifted intervals on the eighth day of that. The INTNX function shows the beginning of the next interval:

Table 7.12 Using the INTNX Function to Show the Beginning of the Next Interval

SAS Statements	Results
<code>start=intnx('week2.8','01jan1960'd,1);</code>	SAS date value 2, or Jan 3, 1960

You can also shift time intervals. For example, HOUR8.7 intervals divide the day into the periods 06:00 to 14:00, 14:00 to 22:00, and 22:00 to 06:00.

Custom Intervals

You can define custom intervals and associate interval data sets with new interval names when you use the INTERVALDS system option. The dates for these intervals are located in a SAS data set that you create. The data set must contain two variables, BEGIN and END. Each observation represents one interval with the BEGIN variable that contains the start of the interval and the END variable that contains the end of the interval. After the custom intervals have been defined, you can use them with the INTCK and INTNX functions just as you would use standard intervals.

The INTERVALDS system option enables you to increase the number of allowable intervals. In addition to the standard list of intervals (DAY, WEEKDAY, and so on), the names that are listed in INTERVALDS are valid as well.

Retail Calendar Intervals: ISO 8601 Compliant

The retail industry often accounts for its data by dividing the yearly calendar into four 13-week periods, based on one of the following formats: 4-4-5, 4-5-4, and 5-4-4. The first, second, and third numbers specify the number of weeks in the first, second, and

third month of each period, respectively. Retail calendar intervals facilitate comparisons across years, because week definitions remain consistent from year to year.


The intervals that are created from the formats can be used in any of the following functions: INTCINDEX, INTCK, INTCYCLE, INTFIT, INTFMT, INTGET, INTINDEX, INTNX, INTSEAS, INTSHIFT, and INTTEST.

The following table lists calendar intervals that are used in the retail industry and that are ISO 8601 compliant.

Table 7.13 Calendar Intervals Used in the Retail Industry

Interval	Description
YEARV	specifies ISO 8601 yearly intervals. The ISO 8601 year begins on the Monday on or immediately preceding January 4. Note that it is possible for the ISO 8601 year to begin in December of the preceding year. Also, some ISO 8601 years contain a leap week. The beginning subperiod s is written in ISO 8601 weeks (WEEKV).
R445YR	is the same as YEARV except that in the retail industry the beginning subperiod s is 4-4-5 months (R445MON).
R454YR	is the same as YEARV except that in the retail industry the beginning subperiod s is 4-5-4 months (R454MON).
R544YR	is the same as YEARV except that in the retail industry the beginning subperiod s is 5-4-4 months (R544MON).
R445QTR	specifies retail 4-4-5 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week. The beginning subperiod s is 4-4-5 months (R445MON).
R454QTR	specifies retail 4-5-4 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week. The beginning subperiod s is 4-5-4 months (R454MON).
R544QTR	specifies retail 5-4-4 quarterly intervals (every 13 ISO 8601 weeks). Some fourth quarters will contain a leap week. The beginning subperiod s is 5-4-4 months (R544MON).
R445MON	specifies retail 4-4-5 monthly intervals. The 3rd, 6th, 9th, and 12th months are five ISO 8601 weeks long with the exception that some 12th months contain leap weeks. All other months are four ISO 8601 weeks long. R445MON intervals begin with the 1st, 5th, 9th, 14th, 18th, 22nd, 27th, 31st, 35th, 40th, 44th, and 48th weeks of the ISO year. The beginning subperiod s is 4-4-5 months (R445MON).
R454MON	specifies retail 4-5-4 monthly intervals. The 2nd, 5th, 8th, and 11th months are five ISO 8601 weeks long with the exception that some 12th months contain leap weeks. R454MON intervals begin with the 1st, 5th, 10th, 14th, 18th, 23rd, 27th, 31st, 36th, 40th, 44th, and 49th weeks of the ISO year. The beginning subperiod s is 4-5-4 months (R454MON).

Interval	Description
R544MON	specifies retail 5-4-4 monthly intervals. The 1st, 4th, 7th, and 10th months are five ISO 8601 weeks long. All other months are four ISO 8601 weeks long with the exception that some 12th months contain leap weeks. R544MON intervals begin with the 1st, 6th, 10th, 14th, 19th, 23rd, 27th, 32nd, 36th, 40th, 45th, and 49th weeks of the ISO year. The beginning subperiod s is 5-4-4 months (R544MON).
WEEKV	specifies ISO 8601 weekly intervals of seven days. Each week begins on Monday. The beginning subperiod s is calculated in days (DAY). Note that WEEKV differs from WEEK in that WEEKV.1 begins on Monday, WEEKV.2 begins on Tuesday, and so on.



CHAPTER

8

Error Processing and Debugging

<i>Types of Errors in SAS</i>	99
<i>Summary of Types of Errors That SAS Recognizes</i>	99
<i>Syntax Errors</i>	100
<i>Semantic Errors</i>	102
<i>Execution-Time Errors</i>	103
<i>Definition</i>	103
<i>Out-of-Resources Condition</i>	104
<i>Examples</i>	104
<i>Data Errors</i>	106
<i>Format Modifiers for Error Reporting</i>	108
<i>Macro-related Errors</i>	108
<i>Error Processing in SAS</i>	108
<i>Syntax Check Mode</i>	108
<i>Overview of Syntax Check Mode</i>	108
<i>Enabling Syntax Check Mode</i>	109
<i>Processing Multiple Errors</i>	109
<i>Checkpoint Mode and Restart Mode</i>	110
<i>Overview of Checkpoint Mode and Restart Mode</i>	110
<i>Setting Up and Executing Checkpoint Mode and Restart Mode</i>	111
<i>Restarting Batch Programs</i>	112
<i>Using System Options to Control Error Handling</i>	113
<i>Using Return Codes</i>	114
<i>Other Error-Checking Options</i>	114
<i>Debugging Logic Errors in the DATA Step</i>	114

Types of Errors in SAS

Summary of Types of Errors That SAS Recognizes

SAS performs *error processing* during both the compilation and the execution phases of SAS processing. You can *debug* SAS programs by understanding processing messages in the SAS log and then fixing your code. You can use the DATA Step Debugger to detect logic errors in a DATA step during execution.

SAS recognizes five types of errors.

Type of error	When this error occurs	When the error is detected
syntax	when programming statements do not conform to the rules of the SAS language	compile time
semantic	when the language element is correct, but the element might not be valid for a particular usage	compile time
execution-time	when SAS attempts to execute a program and execution fails	execution time
data	when data values are invalid	execution time
macro-related	when you use the macro facility incorrectly	macro compile time or execution time, DATA or PROC step compile time or execution time

Syntax Errors

Syntax errors occur when program statements do not conform to the rules of the SAS language. Here are some examples of syntax errors:

- misspelled SAS keyword
- unmatched quotation marks
- missing a semicolon
- invalid statement option
- invalid data set option.

When SAS encounters a syntax error, it first attempts to correct the error by attempting to interpret what you mean. Then SAS continues processing your program based on its assumptions. If SAS cannot correct the error, it prints an error message to the log.

In the following example, the DATA statement is misspelled, and SAS prints a warning message to the log. Because SAS could interpret the misspelled word, the program runs and produces output.

```

date temp;
  x=1;
run;

proc print data=temp;
run;

```

Output 8.1 SAS Log: Syntax Error (misspelled key word)

```

1   date temp;
      ----
      14
WARNING 14-169: Assuming the symbol DATA was misspelled as date.

2       x=1;
3   run;

NOTE: The data set WORK.TEMP has 1 observations and 1 variables.
NOTE: DATA statement used:
      real time           0.17 seconds
      cpu time            0.04 seconds

4
5   proc print data=temp;
6   run;

NOTE: PROCEDURE PRINT used:
      real time           0.14 seconds
      cpu time            0.03 seconds

```

Some errors are explained fully by the message that SAS prints in the log; other error messages are not as easy to interpret because SAS is not always able to detect exactly where the error occurred. For example, when you fail to end a SAS statement with a semicolon, SAS does not always detect the error at the point where it occurs because SAS statements are free-format (they can begin and end anywhere). In the following example, the semicolon at the end of the DATA statement is missing. SAS prints the word ERROR in the log, identifies the possible location of the error, prints an explanation of the error, and stops processing the DATA step.

```

data temp
  x=1;
run;

proc print data=temp;
run;

```

Output 8.2 SAS Log: Syntax Error (missing semicolon)

```

1  data temp
2  x=1;
   -
   76
ERROR 76-322: Syntax error, statement will be ignored.

3  run;

NOTE: The SAS System stopped processing this step because of errors.
NOTE: DATA statement used:
      real time          0.11 seconds
      cpu time           0.02 seconds

4
5  proc print data=temp;
ERROR: File WORK.TEMP.DATA does not exist.
6  run;

NOTE: The SAS System stopped processing this step because of errors.
NOTE: PROCEDURE PRINT used:
      real time          0.06 seconds
      cpu time           0.01 seconds

```

Whether subsequent steps are executed depends on which method of running SAS you use, as well as on your operating environment.

Note: You can add these lines to your code to fix unmatched comment tags, unmatched quotation marks, and missing semicolons.

```

/* '; * "; */;
quit;
run;

```

\triangle

Semantic Errors

Semantic errors occur when the form of the elements in a SAS statement is correct, but the elements are not valid for that usage. Semantic errors are detected at compile time and can cause SAS to enter syntax check mode. (For a description of syntax check mode, see “Syntax Check Mode” on page 108.)

Examples of semantic errors include the following:

- specifying the wrong number of arguments for a function
- using a numeric variable name where only a character variable is valid
- using illegal references to an array.

In the following example, SAS detects an illegal reference to the array ALL.

```

data _null_;
  array all{*} x1-x5;
  all=3;
  datalines;
1 1.5
. 3
2 4.5
3 2 7
3 . .

```

```

;

run;

```

Output 8.3 SAS Log: Semantic Error (illegal reference to an array)

```

          cpu time          0.02 seconds

1  data _null_;
2  array all{*} x1-x5;
ERROR: Illegal reference to the array all.
3  all=3;
4  datalines;
NOTE: The SAS System stopped processing this step because of errors.
NOTE: DATA statement used:
      real time          2.28 seconds
      cpu time           0.06 seconds

10 ;
11

```

The following is another example of a semantic error. In this DATA step, the libref SOMELIB has not been previously assigned in a LIBNAME statement.

```

data test;
  set somelib.old;
run;

```

Output 8.4 SAS Log: Semantic Error (libref not previously assigned)

```

          cpu time          0.00 seconds

1  data test;
ERROR: Libname SOMELIB is not assigned.
2  set somelib.old;
3  run;
NOTE: The SAS System stopped processing this step because of errors.
WARNING: The data set WORK.TEST might be incomplete.  When this step was stopped
        there were 0 observations and 0 variables.
NOTE: DATA statement used:
      real time          0.17 seconds

```

Execution-Time Errors

Definition

Execution-time errors are errors that occur when SAS executes a program that processes data values. Most execution-time errors produce warning messages or notes in the SAS log but allow the program to continue executing. *The location of an

* When you run SAS in noninteractive mode, more serious errors can cause SAS to enter syntax check mode and stop processing the program.

execution-time error is usually given as line and column numbers in a note or error message.

Common execution-time errors include the following:

- illegal arguments to functions
- illegal mathematical operations (for example, division by 0)
- observations in the wrong order for BY-group processing
- reference to a nonexistent member of an array (occurs when the array's subscript is out of range)
- open and close errors on SAS data sets and other files in INFILE and FILE statements
- INPUT statements that do not match the data lines (for example, an INPUT statement in which you list the wrong columns for a variable or fail to indicate that the variable is a character variable).

Out-of-Resources Condition

An execution-time error can also occur when you encounter an out-of-resources condition, such as a full disk, or insufficient memory for a SAS procedure to complete. When these conditions occur, SAS attempts to find resources for current use. For example, SAS might ask the user for permission to delete temporary data sets that might no longer be needed, or to free the memory in which macro variables are stored.

When an out-of-resources condition occurs in a windowing environment, you can use the SAS CLEANUP system option to display a requestor panel that enables you to choose how to resolve the error. When you run SAS in batch, noninteractive, or interactive line mode, the operation of CLEANUP depends on your operating environment. For more information, see the CLEANUP system option in *SAS Language Reference: Dictionary*, and in the SAS documentation for your operating environment.

Examples

In the following example, an execution-time error occurs when SAS uses data values from the second observation to perform the division operation in the assignment statement. Division by 0 is an illegal mathematical operation and causes an execution-time error.

```
options linesize=64 nodate pageno=1 pagesize=25;

data inventory;
  input Item $ 1-14 TotalCost 15-20
        UnitsOnHand 21-23;
  UnitCost=TotalCost/UnitsOnHand;
  datalines;
Hammers      440    55
Nylon cord   35     0
Ceiling fans 1155   30
;

proc print data=inventory;
  format TotalCost dollar8.2 UnitCost dollar8.2;
run;
```

Output 8.5 SAS Log: Execution-Time Error (division by 0)

```

      cpu time          0.02 seconds

1
2  options linesize=64 nodate pageno=1 pagesize=25;
3
4  data inventory;
5      input Item $ 1-14 TotalCost 15-20
6          UnitsOnHand 21-23;
7      UnitCost=TotalCost/UnitsOnHand;
8      datalines;
NOTE: Division by zero detected at line 12 column 22.
RULE:----+----1----+----2----+----3----+----4----+----5----+----
10  Nylon cord      35      0
Item=Nylon cord TotalCost=35 UnitsOnHand=0 UnitCost=._ERROR_=1
   _N_=2
NOTE: Mathematical operations could not be performed at the
      following places. The results of the operations have been
      set to missing values.
      Each place is given by:
      (Number of times) at (Line):(Column).
      1 at 12:22
NOTE: The data set WORK.INVENTORY has 3 observations and 4
      variables.
NOTE: DATA statement used:
      real time          2.78 seconds
      cpu time           0.08 seconds

12 ;
13
14  proc print data=inventory;
15      format TotalCost dollar8.2 UnitCost dollar8.2;
16  run;
NOTE: There were 3 observations read from the dataset
      WORK.INVENTORY.
NOTE: PROCEDURE PRINT used:
      real time          2.62 seconds

```

Output 8.6 SAS Listing Output: Execution-Time Error (division by 0)

The SAS System					1
Obs	Item	Total Cost	Units OnHand	UnitCost	
1	Hammers	\$440.00	55	\$8.00	
2	Nylon cord	\$35.00	0	.	
3	Ceiling fans	\$1155.00	30	\$38.50	

SAS executes the entire step, assigns a missing value for the variable UnitCost in the output, and writes the following to the SAS log:

- a note that describes the error
- the values that are stored in the input buffer
- the contents of the program data vector at the time the error occurred
- a note explaining the error.

Note that the values that are listed in the program data vector include the `_N_` and `_ERROR_` automatic variables. These automatic variables are assigned temporarily to each observation and are not stored with the data set.

In the following example of an execution-time error, the program processes an array and SAS encounters a value of the array's subscript that is out of range. SAS prints an error message to the log and stops processing.

```
options linesize=64 nodate pageno=1 pagesize=25;

data test;
  array all{*} x1-x3;
  input I measure;
  if measure > 0 then
    all{I} = measure;
  datalines;
1 1.5
. 3
2 4.5
;

proc print data=test;
run;
```

Output 8.7 SAS Log: Execution-Time Error (subscript out of range)

```

      cpu time          0.02 seconds

1  options linesize=64 nodate pageno=1 pagesize=25;
2
3  data test;
4      array all{*} x1-x3;
5      input I measure;
6      if measure > 0 then
7          all{I} = measure;
8      datalines;
ERROR: Array subscript out of range at line 12 column 7.
RULE:-----+-----1-----+-----2-----+-----3-----+-----4-----+-----5-----+-----
10   . 3
x1=. x2=. x3=. I=. measure=3 _ERROR_=1 _N_=2
NOTE: The SAS System stopped processing this step because of
      errors.
WARNING: The data set WORK.TEST might be incomplete.  When this
         step was stopped there were 1 observations and 5
         variables.
NOTE: DATA statement used:
      real time          0.90 seconds
      cpu time           0.09 seconds

12 ;
13
14 proc print data=test;
15 run;
NOTE: There were 1 observations read from the dataset WORK.TEST.
NOTE: PROCEDURE PRINT used:
      real time          0.81 seconds
```

Data Errors

Data errors occur when some data values are not appropriate for the SAS statements that you have specified in the program. For example, if you define a variable as numeric, but the data value is actually character, SAS generates a data error. SAS

detects data errors during program execution and continues to execute the program, and does the following:

- writes an invalid data note to the SAS log.
- prints the input line and column numbers that contain the invalid value in the SAS log. Unprintable characters appear in hexadecimal. To help determine column numbers, SAS prints a rule line above the input line.
- prints the observation under the rule line.
- sets the automatic variable `_ERROR_` to 1 for the current observation.

In this example, a character value in the Number variable results in a data error during program execution:

```
options linesize=64 nodate pageno=1 pagesize=25;

data age;
  input Name $ Number;
  datalines;
Sue 35
Joe xx
Steve 22
;

proc print data=age;
run;
```

The SAS log shows that there is an error in line 8, position 5–6 of the program.

Output 8.8 SAS Log: Data Error

```

      cpu time          0.01 seconds

1
2  options linesize=64 nodate pageno=1 pagesize=25;
3
4  data age;
5      input Name $ Number;
6      datalines;
NOTE: Invalid data for Number in line 8 5-6.
RULE:----+----1----+----2----+----3----+----4----+----5----+----6
8  Joe xx
Name=Joe Number=. _ERROR_=1 _N_=2
NOTE: The data set WORK.AGE has 3 observations and 2 variables.
NOTE: DATA statement used:
      real time          0.06 seconds
      cpu time           0.02 seconds

10 ;
11
12 proc print data=age;
13 run;
NOTE: There were 3 observations read from the dataset WORK.AGE.
NOTE: PROCEDURE PRINT used:
      real time          0.01 seconds
```

Output 8.9 SAS Listing Output: Data Error

The SAS System			1
Obs	Name	Number	
1	Sue	35	
2	Joe	.	
3	Steve	22	

You can also use the `INVALIDDATA=` system option to assign a value to a variable when your program encounters invalid data. For more information, see the `INVALIDDATA=` system option in *SAS Language Reference: Dictionary*.

Format Modifiers for Error Reporting

The `INPUT` statement uses the `?` and the `??` format modifiers for error reporting. The format modifiers control the amount of information that is written to the SAS log. Both the `?` and the `??` modifiers suppress the invalid data message. However, the `??` modifier also sets the automatic variable `_ERROR_` to 0. For example, these two sets of statements are equivalent:

- `input x ?? 10-12;`
- `input x ? 10-12;`
`_error_=0;`

In either case, SAS sets the invalid values of `X` to missing values.

Macro-related Errors

Several types of macro-related errors exist:

- macro compile time and macro execution-time errors, generated when you use the macro facility itself
- errors in the SAS code produced by the macro facility.

For more information about macros, see *SAS Macro Language: Reference*.

Error Processing in SAS

Syntax Check Mode

Overview of Syntax Check Mode

If you want processing to stop when a statement in a `DATA` step has a syntax error, you can enable SAS to enter syntax check mode. You do this by setting the `SYNTAXCHECK` system option in batch or non-interactive mode, or by setting the `DMSSYNCHK` system option in the windowing environment.

SAS can enter syntax check mode only if your program creates a data set. If you use the `DATA _NULL_` statement, then SAS cannot enter syntax check mode because no data set is created. In this case, using the `SYNTAXCHECK` or `DMSSYNCHK` system option has no effect.

In syntax check mode, SAS internally sets the `OBS=` option to 0 and the `REPLACE/NOREPLACE` option to `NOREPLACE`. When these options are in effect, SAS acts as follows:

- reads the remaining statements in the `DATA` step or `PROC` step
- checks that statements are valid SAS statements
- executes global statements
- writes errors to the SAS log
- creates the descriptor portion of any output data sets that are specified in program statements
- does not write any observations to new data sets that SAS creates
- does not execute most of the subsequent `DATA` steps or procedures in the program (exceptions include `PROC DATASETS` and `PROC CONTENTS`).

Note: Any data sets that are created after SAS has entered syntax check mode do not replace existing data sets with the same name. Δ

When syntax checking is enabled, if SAS encounters a syntax or semantic error in a `DATA` step, SAS underlines the point where it detects the error and identifies the error by number. SAS then enters syntax check mode and remains in this mode until the program finishes executing. When SAS enters syntax check mode, all `DATA` step statements and `PROC` step statements are validated.

Enabling Syntax Check Mode

You use the `SYNTAXCHECK` system option to enable syntax check mode when you run SAS in non-interactive or batch mode. You use the `DMSSYNCHK` system option to enable syntax check mode when you run SAS in the windowing environment. You can use these system options only if your program creates a data set. If you use the `DATA _NULL_` statement, then these options are ignored.

To disable syntax check mode, use the `NOSYNTAXCHECK` and `NODMSSYNCHK` system options.

In an `OPTIONS` statement, place the `OPTIONS` statement that enables `SYNTAXCHECK` or `DMSSYNCHK` before the step for which you want it to apply. If you place the `OPTIONS` statement inside a step, then `SYNTAXCHECK` or `DMSSYNCHK` will not take effect until the beginning of the next step.

For more information about these system options, see the `DMSSYNCHK` system option and the `SYNTAXCHECK` system option, in *SAS Language Reference: Dictionary*.

Processing Multiple Errors

Depending on the type and severity of the error, the method you use to run SAS, and your operating environment, SAS either stops program processing or flags errors and continues processing. SAS continues to check individual statements in procedures after it finds certain types of errors. Thus, in some cases SAS can detect multiple errors in a single statement and might issue more error messages for a given situation, particularly if the statement containing the error creates an output SAS data set.

The following example illustrates a statement with two errors:

```
data temporary;
  Item1=4;
```

```

run;

proc print data=temporary;
  var Item1 Item2 Item3;
run;

```

Output 8.10 SAS Log: Multiple Program Errors

```

          cpu time          0.00 seconds

1 data temporary;
2   Item1=4;
3 run;
NOTE: The data set WORK.TEMPORARY has 1 observations and 1
      variables.
NOTE: DATA statement used:
      real time          0.10 seconds
      cpu time           0.01 seconds

4
5 proc print data=temporary;
ERROR: Variable ITEM2 not found.
ERROR: Variable ITEM3 not found.
6   var Item1 Item2 Item3;
7 run;
NOTE: The SAS System stopped processing this step because of
      errors.
NOTE: PROCEDURE PRINT used:
      real time          0.53 seconds
      cpu time           0.01 seconds

```

SAS displays two error messages, one for the variable Item2 and one for the variable Item3.

When you are running debugged production programs that are unlikely to encounter errors, you might want to force SAS to abend after a single error occurs. You can use the ERRORABEND system option to do this.

Checkpoint Mode and Restart Mode

Overview of Checkpoint Mode and Restart Mode

Used together, checkpoint mode and restart mode enable batch programs that terminate before completing to be resubmitted, resuming execution with the DATA or PROC step that was executing when the failure occurred. DATA and PROC steps that already completed will not be re-executed.

When checkpoint mode is enabled, SAS records information about DATA and PROC steps in a checkpoint library. When a batch program terminates prematurely, you can resubmit the program in restart mode to complete execution. In restart mode, global statements and macros are re-executed and SAS reads the data in the checkpoint library to determine which steps completed. Program execution resumes with the step that was executing when the failure occurred.

The checkpoint-restart data contains only information about the DATA and PROC steps that completed and the step that did not complete. It does not contain information about macro variables, macro definitions, SAS data sets, or any other information that might have been processed in the step that did not complete.

If a step must be reexecuted, you can add the global statement `CHECKPOINT EXECUTE_ALWAYS` immediately before the step. This statement tells SAS to always execute the following step without considering the checkpoint-restart data.

Checkpoint mode is not valid for batch programs that contain the `DM` statement to submit commands to SAS. If checkpoint mode is enabled and SAS encounters a `DM` statement, checkpoint mode is disabled and the checkpoint catalog entry is deleted.

You enable checkpoint mode and restart mode by using system options when you start the batch program in SAS. You can also use one or more `CHECKPOINT` statements in the batch program:

- `STEPCHKPT` system option enables checkpoint mode, which indicates to SAS to record checkpoint-restart data
- `STEPCHKPTLIB` system option identifies a user-specified checkpoint-restart library
- `STEPRESTART` system option enables restart mode, ensuring that execution resumes with the `DATA` or `PROC` step indicated by the checkpoint-restart library.
- `CHECKPOINT EXECUTE_ALWAYS` is a global statement that can be placed before a `DATA` or `PROC` step to indicate that the step that immediately follows must be executed, even if the checkpoint-restart data shows that the step already completed. It is applicable only to the step that follows the statement.

In order for checkpoint mode and restart mode to work successfully, the number and order of the `DATA` and `PROC` steps in the batch program must not change between SAS invocations. By specifying the `ERRORABEND` and `ERRORCHECK` system options when SAS starts, SAS will terminate for most error conditions in order to maintain valid checkpoint-restart data.

The checkpoint-restart library can be a user-specified library or, if no library is specified, the checkpoint-restart data is saved to the `WORK` library. If the `WORK` library is used, the `NOWORKTERM` and `NOWORKINIT` system options must be specified when SAS starts to ensure that the checkpoint-restart data is not lost. SAS writes the name of the `WORK` library to the SAS log.

For information about the system options and the statement, see the `STEPCHKPT` system option, `STEPCHKPTLIB` system option, `STEPRESTART` system option, and `CHECKPOINT EXECUTE_ALWAYS` statement in *SAS Language Reference: Dictionary*

Setting Up and Executing Checkpoint Mode and Restart Mode

To set up checkpoint mode and restart mode, make the following modification to your batch program:

- Add the `CHECKPOINT EXECUTE_ALWAYS` statement before any `DATA` and `PROC` steps that you want to execute each time the batch program is submitted.
- If your checkpoint-restart library is a user-defined library, you must add the `LIBNAME` statement that defines the checkpoint-restart libref as the first statement in the batch program. If you use the `WORK` library as your checkpoint library, no `LIBNAME` statement is necessary.

Once the batch program has been modified, you start the program using the appropriate system options:

- For checkpoint-restart data that is saved in the `WORK` library, start a batch SAS session that specifies these system options:
 - `SYSIN`, if required in your operating environment, names the batch program
 - `STEPCHKPT` enables checkpoint mode
 - `NOWORKTERM` saves the `WORK` library when SAS ends
 - `NOWORKINIT` does not initialize the `WORK` library when SAS starts

- `ERRORCHECK STRICT` puts SAS in syntax-check mode when an error occurs in the `LIBNAME`, `FILENAME`, `%INCLUDE`, and `LOCK` statements
- `ERRORABEND` specifies whether SAS terminates for most errors

In the Windows operating environment, the following SAS command starts a batch program in checkpoint mode using the `WORK` library as the checkpoint-restart library:

```
sas -sysin 'c:\mysas\myprogram.sas' -stepchkpt -noworkterm -noworkinit
      -errorcheck strict -errorabend
```

- For checkpoint-restart data that is saved in a user-specified library, start a batch SAS session that includes these system options:
 - `SYSIN`, if required in your operating environment, names the batch program
 - `STEPCHKPT` enables checkpoint mode
 - `STEPCHKPTLIB` specifies the libref of the library where SAS saves the checkpoint-restart data
 - `ERRORCHECK STRICT` puts SAS in syntax-check mode when an error occurs in the `LIBNAME`, `FILENAME`, `%INCLUDE`, and `LOCK` statements
 - `ERRORABEND` specifies whether SAS terminates for most errors

In the Windows operating environment, the following SAS command starts a batch program in checkpoint mode using a user-specified checkpoint-restart library:

```
sas -sysin 'c:\mysas\myprogram.sas' -stepchkpt -stepchkptlib mylibref
      -errorcheck strict -errorabend
```

In this case, the first statement in `myprogram.sas` is the `LIBNAME` statement that defines the `mylibref` libref.

Restarting Batch Programs

To resubmit a batch SAS session using the checkpoint-restart data that is saved in the `WORK` library, include these system options when SAS starts:

- `SYSIN`, if required in your operating environment, names the batch program
- `STEPCHKPT` continues checkpoint mode
- `STEPRESTART` enables restart mode, indicating to SAS to use the checkpoint-restart data
- `NOWORKINIT` starts SAS using the `WORK` library from the previous SAS session
- `NOWORKTERM` saves the `WORK` library when SAS ends
- `ERRORCHECK STRICT` puts SAS in syntax-check mode when an error occurs in the `LIBNAME`, `FILENAME`, `%INCLUDE`, and `LOCK` statements
- `ERRORABEND` specifies whether SAS terminates for most errors

In the Windows operating environment, the following SAS command resubmits a batch program whose checkpoint-restart data was saved to the `WORK` library:

```
sas -sysin 'c:\mysas\mysasprogram.sas' -stepchkpt -steprestart -noworkinit
      -noworkterm -errorcheck strict -errorabend
```

By specifying the `STEPCHKPT` and `NOWORKTERM` system options, checkpoint mode continues to be enabled once the batch program restarts.

To resubmit a batch SAS session using the checkpoint-restart data that is saved in a user-specified library, include these system options when SAS starts:

- `SYSIN`, if required in your operating environment, names the batch program

- STEPCHKPT continues checkpoint mode
- STEPRESTART enables restart mode, indicating to SAS to use the checkpoint-restart data
- STEPCHKPTLIB specifies the libref of the checkpoint-restart library
- ERRORCHECK STRICT puts SAS in syntax-check mode when an error occurs in the LIBNAME, FILENAME, %INCLUDE, and LOCK statements
- ERRORABEND specifies whether SAS terminates for most errors

In the Windows operating environment, the following SAS command resubmits a batch program whose checkpoint-restart data was saved to a user-specified library:

```
sas -sysin 'c:\mysas\mysasprogram.sas' -stepchkpt -steprestart -stepchklib
mylibref -errorcheck strict -errorabend
```

Using System Options to Control Error Handling

You can use the following system options to control error handling (resolve errors) in your program:

BYERR	specifies whether SAS produces errors when the SORT procedure attempts to process a <code>_NULL_</code> data set.
DKRCOND=	specifies the level of error detection to report when a variable is missing from an input data set during the processing of a <code>DROP=</code> , <code>KEEP=</code> , and <code>RENAME=</code> data set option.
DKROCOND=	specifies the level of error detection to report when a variable is missing from an output data set during the processing of a <code>DROP=</code> , <code>KEEP=</code> , and <code>RENAME=</code> data set option.
DSNFERR	when a SAS data set cannot be found, specifies whether SAS issues an error message.
ERRORABEND	specifies whether SAS responds to errors by terminating.
ERRORCHECK=	specifies whether SAS enters syntax-check mode when errors are found in the <code>LIBNAME</code> , <code>FILENAME</code> , <code>%INCLUDE</code> , and <code>LOCK</code> statements.
ERRORS=	specifies the maximum number of observations for which SAS issues complete error messages.
FMTERR	when a variable format cannot be found, specifies whether SAS generates an error or continues processing.
INVALIDDATA=	specifies the value that SAS assigns to a variable when invalid numeric data is encountered.
MERROR	specifies whether SAS issues a warning message when a macro-like name does not match a macro keyword.
QUOTELENMAX	if a quoted string exceeds the maximum length allowed, specifies whether SAS writes a warning message to the SAS log.
SERROR	specifies whether SAS issues a warning message when a macro variable reference does not match a macro variable.
STEPCHKPT	specifies whether checkpoint-restart data is to be recorded for a batch program.
STEPCHKPTLIB	specifies the libref of the library where checkpoint-restart data is saved.

STEPRESTART	specifies whether to execute a batch program by using checkpoint-restart data.
VNFERR	specifies whether SAS issues an error or warning when a BY variable exists in one data set but not another data set when processing the SET, MERGE, UPDATE, or MODIFY statements.

For more information about SAS system options, see *SAS Language Reference: Dictionary*.

Using Return Codes

In some operating environments, SAS passes a return code to the system, but the way in which return codes are accessed is specific to your operating environment.

Operating Environment Information: For more information about return codes, see the SAS documentation for your operating environment. Δ

Other Error-Checking Options


To help determine your programming errors, you can use the following methods:

- the `_IORC_` automatic variable that SAS creates (and the associated `IORCMSG` function) when you use the `MODIFY` statement or the `KEY=` data set option in the `SET` statement
- the `ERRORS=` system option to limit the number of identical errors that SAS writes to the log
- the `SYSRC` and `SYSMSG` functions to return information when a data set or external-files access function encounters an error condition
- the `SYSRC` automatic macro variable to receive return codes
- the `SYSERR` automatic macro variable to detect major system errors, such as out of memory or failure of the component system
- log control options:

<code>MSGLEVEL=</code>	controls the level of detail in messages that are written to the SAS log.
<code>PRINTMSGLIST</code>	controls the printing of extended lists of messages to the SAS log.
<code>SOURCE</code>	controls whether SAS writes source statements to the SAS log.
<code>SOURCE2</code>	controls whether SAS writes source statements included by <code>%INCLUDE</code> to the SAS log.

Debugging Logic Errors in the DATA Step

To debug logic errors in a `DATA` step, you can use the `DATA` step debugger. This tool enables you to issue commands to execute `DATA` step statements one by one and then pause to display the resulting variables' values in a window. By observing the results that are displayed, you can determine where the logic error lies. Because the debugger is interactive, you can repeat the process of issuing commands and observing results as many times as needed in a single debugging session. To invoke the debugger, add the `DEBUG` option to the `DATA` statement and execute the program. For detailed information about how to use the `DATA` step debugger, see *SAS Language Reference: Dictionary*.



CHAPTER

9

SAS Output

<i>Definitions for SAS Output</i>	115
<i>Routing SAS Output</i>	117
<i>The SAS Log</i>	117
<i>Structure of the Log</i>	117
<i>The SAS Log in Interactive Mode</i>	119
<i>The SAS Log in Batch, Line, or Objectserver Modes</i>	119
<i>Overview of the SAS Log in Batch, Line, or Objectserver Modes</i>	119
<i>Appending to or Replacing the SAS Log</i>	120
<i>Specifying When to Write to the SAS Log</i>	120
<i>Rolling Over the SAS Log</i>	120
<i>Writing to the Log in All Modes</i>	122
<i>Customizing the Log</i>	122
<i>Altering the Contents of the Log</i>	122
<i>Customizing the Appearance of the Log</i>	124
<i>Other System Options That Affect the SAS Log</i>	125
<i>Traditional SAS Listing Output</i>	126
<i>Example of Traditional Listing Output</i>	126
<i>Making Output Descriptive</i>	127
<i>Reformatting Values</i>	128
<i>Printing Missing Values</i>	128
<i>Changing the Destination of the Log and the Output</i>	128
<i>Output Delivery System</i>	128

Definitions for SAS Output

SAS output is the result of executing SAS programs. Most SAS procedures and some DATA step applications produce output.

There are three types of SAS output:

SAS log

contains a description of the SAS session and lists the lines of source code that were executed. Depending on the setting of SAS system options, the method of running SAS, and the program statements that you specify, the log can include the following types of information:

- program statements
- names of data sets created by the program
- notes, warnings, or error messages encountered during program execution
- the number of variables and observations each data set contains
- processing time required for each step.

You can write specific information to the SAS log (such as variable values or text strings) by using the SAS statements that are described in “Writing to the Log in All Modes” on page 122.

The log is also used by some of the SAS procedures that perform utility functions, for example the DATASETS and OPTIONS procedures. See the *Base SAS Procedures Guide* for more information.

Because the SAS log provides a journal of program processing, it is an essential debugging tool. However, certain system options must be in effect to make the log effective for debugging your SAS programs. “Customizing the Log” on page 122 describes several SAS system options that you can use.

program results

contain the results of most SAS procedures and some DATA step applications. Results can be routed to a file, and printed as a listing. If you use the Output Delivery System (ODS), you can produce results for a high resolution printer or create HTML output for use with a web browser. You can customize your output by modifying or creating your own table definitions, which are descriptions of how you want to format your output. For more information about the flexibility of ODS, see the *SAS Output Delivery System: User’s Guide*.

SAS console log

When the SAS log is not active, it contains information, warning, and error messages. When the SAS log is active, the SAS console log is used only for fatal system initialization errors or late termination messages.

Operating Environment Information: See the SAS documentation for your operating environment for specific information on the destination of the SAS console log. \triangle

SAS logging facility logs

contain log messages that you create using the SAS logging facility. Logging facility messages can be created within SAS programs or they can be created by SAS for logging SAS server messages. The logging facility log messages are based on message categories such as authentication, administration, performance, or customized message categories in SAS programs. In SAS programs, you use logging facility functions, autocall macros, or DATA step component objects to create the logging facility environment.

The logging facility environment consists of loggers, appenders, and log events. A logger defines a message category, references one or more appenders, and specifies the logger’s message level threshold. The message level threshold can be one of the following, from lowest to highest: trace, debug, info, warn, error, or fatal. An appender defines the physical location to write log messages and the format of the message. A log event consists of a log message, a message threshold, and a logger. Log events are initiated by SAS servers and SAS programs.

When SAS processes a logging facility log event, it compares the message level in the log event to the message threshold that is defined for the logger that is named in the log event. If the log event message threshold is the same or higher than the logger’s message threshold, the message is written to the locations that are specified by the appenders that are referenced in the logger definition. If the log event is not accepted by the logger, the message is discarded.

Appenders are defined for the duration of a macro program or a DATA step. Loggers are defined for the duration of the SAS session.

For more information, see SAS Logging Facility in *SAS Logging: Configuration and Programming Reference*.

Routing SAS Output

The destination of your output depends on

- the operating environment
- the setting of SAS system options
- whether you use ODS
- the method of running SAS

There are several ways to route SAS output to a destination other than the default destination. You can route the output to your personal computer, to an external file, or directly to a printer. If you use ODS, you can route output to a data set for most procedures.

Operating Environment Information: See the SAS documentation for your operating environment for specific information. Δ

The following table shows the default destination of SAS output for each method of operation.

Table 9.1 Default Destinations of SAS Output

Method of Running SAS	Destination of SAS Output
windowing mode (Explorer window)	the Log window or the output window
interactive line mode	the personal computer display (as statements are entered)
noninteractive mode	depends on the operating environment
batch mode	depends on the operating environment

Operating Environment Information: The default destination for SAS output is specific to your operating environment. For specific information, see the SAS documentation for your operating environment. Δ

The SAS Log

Structure of the Log

The SAS log is a record of everything that you do in your SAS session or with your SAS program. Original program statements are identified by line numbers. Interspersed with SAS statements are messages from SAS. These messages might begin with the words NOTE, INFO, WARNING, ERROR, or an error number, and they might refer to a SAS statement by its line number in the log.

For example, in the following output, the number 1 prints to the left of the OPTIONS statement. This means that it is the first line in the program. In interactive mode, SAS continues with the sequence of line numbering until you end your session. If you submit the program again (or submit other programs in your current SAS session), the first program line number will be the next consecutive number.

Operating Environment Information: The SAS log appears differently depending on your operating environment. See the SAS documentation for your operating environment. Δ

Output 9.1 Sample SAS Log

```

NOTE: Copyright (c) 2007--2008 by SAS Institute Inc., Cary, NC, USA. ①
NOTE: SAS (r) 9.2 (TS1B0) ②
      Licensed to SAS Institute Inc., Site 0000000001. ③
NOTE: This session is executing on the XP_PRO platform. ④

NOTE: SAS initialization used:
      real time          4.20 seconds
      cpu time           1.18 seconds

1  options pagesize=24 linesize=64 nodate; ⑤
2
3  data logsample; ⑥
4      infile '/u/abcdef/testdir/sampleddata.dat';
5      input LastName $ 1-12 ID $ Gender $ Birth : date7. ⑦
6  ! score1 score2 score3
6                                     score
6  ! 4 score5 score6;
7      format Birth mmddyy8.;
8  run;

NOTE: The infile '/u/abcdef/testdir/sampleddata.dat' is: ⑧
      Filename=/u/abcdef/testdir/sampleddata.dat,
      Owner Name=abcdef,Group Name=pubs,
      Access Permission=rw-r--r--,
      File Size (bytes)=296

NOTE: 5 records were read from the infile ⑨
      '/u/abcdef/testdir/sampleddata.dat'.
      The minimum record length was 58.
      The maximum record length was 59.
NOTE: The data set WORK.LOGSAMPLE has 5 observations and 10
      variables. ⑩
NOTE: DATA statement used:
      real time          0.08 seconds ⑪
      cpu time           0.02 seconds

9
10 proc sort data=logsample; ⑫
11     by LastName;
12

NOTE: There were 5 observations read from the data set
      WORK.LOGSAMPLE.
NOTE: The data set WORK.LOGSAMPLE has 5 observations and 10
      variables. ⑬
NOTE: PROCEDURE SORT used:
      real time          0.03 seconds
      cpu time           0.01 seconds

13 proc print data=logsample; ⑭
14     by LastName;
15 run;

NOTE: There were 5 observations read from the data set
      WORK.LOGSAMPLE.
NOTE: PROCEDURE PRINT used:
      real time          0.09 seconds
      cpu time           0.05 seconds

```

The following list corresponds to the circled numbers in the SAS log shown above:

- ① copyright information.
- ② SAS system release used to run this program.
- ③ name and site number of the computer installation where the program ran.

- ④ platform used to run the program.
- ⑤ OPTIONS statement. This statement uses SAS system options to set a page size of 24 and a line size of 64, and to suppress the date in the output.
- ⑥ SAS statements that make up the program (if the SAS system option SOURCE is enabled).
- ⑦ long statement continued to the next line. Note that the continuation line is preceded by an exclamation point (!), and that the line number does not change.
- ⑧ input file information-notes or warning messages about the raw data and where they were obtained (if the SAS system option NOTES is enabled).
- ⑨ the number and record length of records read from the input file (if the SAS system option NOTES is enabled).
- ⑩ SAS data set that your program created; notes that contain the number of observations and variables for each data set created (if the SAS system option NOTES is enabled).
- ⑪ reported performance statistics when the STIMER option or the FULLSTIMER option is set
- ⑫ procedure that sorts your data set
- ⑬ note about the sorted SAS data set
- ⑭ procedure that prints your data set

The SAS Log in Interactive Mode

In interactive mode, the SAS log is opened when SAS starts. The SAS log is not named until you save it in the active window. The name that you specify must follow the file naming conventions for your operating environment. The SAS log cannot be saved automatically in interactive mode. However, SAS can create a second copy of the SAS log if the ALTLOG= system option is set either at SAS invocation or in a configuration file.

The SAS Log in Batch, Line, or Objectserver Modes

Overview of the SAS Log in Batch, Line, or Objectserver Modes

In batch, line, and objectserver modes, you can configure the SAS log by using the LOG= system option and the LOGPARM= system option if the LOGCONFIGLOC= system option is not specified when SAS starts. If the LOGCONFIGLOC= system option is specified, logging is performed by the SAS logging facility. The LOGPARM= option is ignored. The LOG= option is honored only when the %S{App.Log} conversion character is specified in the logging configuration file.

The following sections discuss the log options that you can configure using the LOGPARM= system option and how you would name the SAS log for those options when the logging facility has not been initiated.

The LOG= system option names the SAS log. The LOGPARM= system option enables you to perform the following tasks:

- append or replace an existing SAS log
- determine when to write to the SAS log
- start a new SAS log under certain conditions

For information about these log system options, see the LOGPARM= system option in *SAS Language Reference: Dictionary* and the LOG= system option in the documentation for your operating environment.

For information about the SAS logging facility, see SAS Logging Facility in *SAS Logging: Configuration and Programming Reference*.

Appending to or Replacing the SAS Log

If you specify a destination for the SAS log in the LOG= system option, SAS verifies if a SAS log already exists. If the log does exist, you can specify how content is written to the SAS log by using the OPEN= option of the LOGPARM= system option:

OPEN=APPEND

appends the SAS log content to the existing SAS log

OPEN=REPLACE

replaces the existing SAS log

OPEN=REPLACEOLD

replace the existing SAS log if it is older than 24 hours

In the following SAS command, both the LOG= and LOGPARM= system options are specified in order to replace an existing SAS log that is more than one day old:

```
sas -sysin "my-batch-program" -log "c:\sas\SASlogs\mylog"
-logparm open=replaceold
```

The OPEN= option is ignored when the ROLLOVER= option of the LOGPARM= system option is set to a specific size, *n*.

Specifying When to Write to the SAS Log

Content can be written to the SAS log either as the log content is produced or it can be buffered and written when the buffer is full. By default, SAS writes to the log when the log buffer is full. By buffering the log content, SAS performs more efficiently by writing to the log file periodically instead of writing one line at a time.

Operating Environment Information: Under Windows, the buffered log contents are written periodically, using a SAS-specified interval. Δ

You use the WRITE= option of the LOGPARM= system option to configure when the SAS log contents are written. Set LOGPARM="WRITE=IMMEDIATE" for the log content to be written as it is produced and set LOGPARM="WRITE=BUFFERED" for the log content to be written when the buffer is full.

Rolling Over the SAS Log

Overview of Rolling Over the SAS Log: The SAS log can get very large for long running servers and for batch jobs. By using the LOGPARM= and LOG= system options together, you can specify to roll over the SAS log to a new SAS log. When SAS rolls over the log, it closes the log and opens a new log.

The LOGPARM= system option controls when log files are opened and closed and the LOG= system option names the SAS log file. Logs can be rolled over automatically, when a SAS session starts, when the log has reached a specific size, or not at all. By using formatting directives in the SAS log name, each SAS log can be named with unique identifiers.

Using Directives to Name the SAS Log: For the SAS log, a directive is a processing instruction that is used to uniquely name the SAS log. By using directives, you can add information to the SAS log name such as the day, the hour, the system node name, or a

unique identifier. You can include one or more directives in the name of the SAS log when you specify the log name in the LOG= system option. For example, if you want the SAS log name to include the year, the month, and the day, the LOG= system option might look like this:

```
-log='c:\saslog\#Y#b#dsas.log'
```

When the SAS log is created on February 2, 2009, the name of the log is 2009Feb02sas.log.

Directives resolve only when the value of the ROLLOVER= option of the LOGPARM= system option is set to AUTO or SESSION. If directives are specified in the log name and the value of the ROLLOVER option is NONE or a specific size, *n*, the directive characters, such as #b or #Y, become part of the log name. Using the example above for the LOG= system option, if the LOGPARM= system option specifies ROLLOVER=NONE, the name of the SAS log is #Y%b#dsas.log.

For a complete list of directives, see the LOGPARM=System Option in *SAS Language Reference: Dictionary*.

Automatically Rolling Over the SAS Log When Directives Change: When the SAS log name contains one or more directives and the ROLLOVER= option of the LOGPARM= system option is set to AUTO, SAS closes the log and opens a new log when the directive values change. The new SAS log name contains the new directive values.

The follow table shows some of the log names that are created when SAS is started on the second of the month at 6:15 AM, using this SAS command:

```
sas -objectserver -log "london#n#d#%H.log" -logparm "rollover=auto"
```

The directive #n inserts the system node name into the log name. #d adds the day of the month to the log name. #H adds the hour to the log name. The node name for this example is Thames. The log for this SAS session rolls over when the hour changes and when the day changes.

Table 9.2 Log Names for Rolled Over Logs

Rollover Time	Log Name
SAS initialization	londonThames0206.log
First rollover	londonThames0207.log
Last log of the day	londonThames0223.log
First log past midnight	londonThames0300.log

Rolling Over the SAS Log by SAS Session: To roll over the log at the start of a SAS session, specify the LOGPARM="ROLLOVER=SESSION" option when SAS starts. SAS resolves the system-specific directives by using the system information obtained when SAS starts. No roll over occurs during the SAS session and the log file is closed at the end of the SAS session.

Rolling Over the SAS Log by the Log Size: To roll over the log when the log reaches a specific size, specify the LOGPARM="ROLLOVER=*n*" option when SAS starts. *n* is the maximum size of the log, in bytes, and it cannot be smaller than 10K (10,240) bytes. When the log reaches the specified size, SAS closes the log and appends the text "old" to the filename. For example, londonold.log. SAS opens a new log using the value of the LOG= option for the log name, ignoring the OPEN= option of the LOGPARM= system option so that it never writes over an existing log file. Directives in log names are ignored for logs that roll over based on log size.

To ensure unique log filenames between servers, SAS creates a lock file that is based on the log filename. The lock filename is *logname.lck*, where *logname* is the value of the LOG= option. If a lock file exists for a server log and another server specifies the same

log name, the log and lock filenames for the second server have a number appended to the names. The numbers begin with 2 and increment by 1 for subsequent requests for the same log filename. For example, if a lock exists for the log file london.log, the second server log would be london2.log and the lock file would be london2.lck.

No SAS Log Roll Over: To not roll over the log at all, specify the LOGPARM="ROLLOVER=NONE" option when SAS starts. Directives are not resolved and no rollover occurs. For example, if LOG="March#b.log", the directive #b does not resolve and the log name is March#b.log.

Writing to the Log in All Modes

In all modes, you can instruct SAS to write additional information to the log by using the following statements:

PUT statement

writes selected lines (including text strings and DATA step variable values) to the SAS log in the current iteration of a DATA step. If a FILE statement with the LOG destination executes before a PUT statement, the PUT statement output is directed to a destination that is specified by the FILE statement.

%PUT statement

enables you to write a text string or macro variable values to the SAS log. %PUT is a SAS macro program statement that is independent of the DATA step and can be used anywhere.

PUTLOG statement

writes a user-specified message to the SAS log. Use the PUTLOG statement in a DATA step.

LIST statement

writes to the SAS log the input data records for the data line that is being processed. The LIST statement operates only on data that are read with an INPUT statement. It has no effect on data that are read with a SET, MERGE, MODIFY, or UPDATE statement. Use the LIST statement in a DATA step.

DATA statement with /NESTING option

writes to the SAS log a note for the beginning and end for each nesting level of DO-END and SELECT-END statements. This enables you to debug mismatched DO-END and SELECT-END statements.

ERROR statement

sets the automatic _ERROR_ variable to 1 and (OPTIONAL) writes to the log a message that you specify. Use the ERROR statement in a DATA step.

Use the PUT, PUTLOG, LIST, DATA, and ERROR statements in combination with conditional processing to debug DATA steps by writing selected information to the log.

Customizing the Log

Altering the Contents of the Log

When you have large SAS production programs or an application that you run on a regular basis without changes, you might want to suppress part of the log. SAS system options enable you to suppress SAS statements and system messages, as well as to limit the number of error messages. Note that all SAS system options remain in effect

for the duration of your session or until you change the options. You should not suppress log messages until you have successfully executed the program without errors.

The following list describes some of the SAS system options that you can use to alter the contents of the log:

CPUID | NOCPUID

specifies whether hardware information is written to the SAS log.

ECHO

specifies a message to be written to the SAS log while SAS initializes. The ECHO system option is valid only under the Windows and UNIX operating environments.

ECHOAUTO | NOECHOAUTO

specifies whether autoexec code in an input file is written to the log.

ERRORS=*n*

specifies the maximum number of observations for which data error messages are printed.

FULLSTATS

writes expanded statistics to the SAS log. The FULLSTATS system option is valid only under z/OS.

FULLSTIMER

writes a subset of system performance statistics to the SAS log.

ISPNOTES

specifies whether ISPF error messages are written to the SAS log. The ISPNOTES system option is valid only under the z/OS operating environment.

LOGPARM "OPEN=APPEND | REPLACE | REPLACEOLD"

when a log file already exists and SAS is opening the log, the LOGPARM option specifies whether to append to the existing log or to replace the existing log. The REPLACEOLD option specifies to replace logs that are more than one day old.

MEMRPT

specifies whether memory usage statistics are written to the SAS log for each step. The MEMRPT system option is valid only under the z/OS operating environment.

MLOGIC

writes macro execution trace information to the SAS log.

MLOGICNEST

writes macro nesting execution trace information to the SAS log.

MPRINT | NOMPRINT

specifies whether SAS statements that are generated by macro execution are written to the SAS log.

MSGLEVEL=N | I

specifies the level of detail in messages that are written to the SAS log. If the MSGLEVEL system option is set to N, the log displays notes, warnings, and error messages only. If MSGLEVEL is set to I, the log displays additional notes pertaining to index usage, merge processing, and sort utilities, along with standard notes, warnings, and error messages.

NEWS=*external-file*

specifies whether news information that is maintained at your site is written to the SAS log.

NOTES | NONOTES

specifies whether notes (messages beginning with NOTE) are written to the SAS log. NONOTES does not suppress error or warning messages.

OPLIST

specifies whether to write to the SAS log the values of all system options that are specified when SAS is invoked.

OVP | NOOVP

specifies whether error messages that are printed by SAS are overprinted.

PAGEBREAKINITIAL

specifies whether the SAS log and the listing file begin on a new page.

PRINTMSGLIST | NOPRINTMSGLIST

specifies whether extended lists of messages are written to the SAS log.

RTRACE

produces a list of resources that are read during SAS execution and writes them to the SAS log if a location is not specified for the RTRACELOC= system option. The RTRACE system option is valid only for the Windows and UNIX operating environments.

SOURCE | NOSOURCE

specifies whether SAS writes source statements to the SAS log.

SOURCE2 | NOSOURCE2

specifies whether SAS writes secondary source statements from files included by %INCLUDE statements to the SAS log.

SYMBOLGEN | NOSYMBOLGEN

specifies whether the results of resolving macro variable references are written to the SAS log.

VERBOSE

specifies whether SAS writes to the batch log or to the computer monitor the values of the system options that are specified in the configuration file.

See *SAS Language Reference: Dictionary* for more information about how to use these and other SAS system options.

Operating Environment Information: See the documentation for your operating environment for other options that affect log output. \triangle

Customizing the Appearance of the Log

The following SAS statements and SAS system options enable you to customize the log. Customizing the log is helpful when you use the log for report writing or for creating a permanent record.

DATE system option

controls whether the date and time that the SAS job began are printed at the top of each page of the SAS log and any output created by SAS.

DETAILS | NODETAILS

specifies whether to include additional information when files are listed in a SAS library.

DMSLOGSIZE= system option

specifies the maximum number of rows to display in the SAS log window.

DTRESET | NODTRESET

specifies whether to update the date and time in the SAS log and in the listing file.

FILE statement

enables you to write the results of PUT statements to an external file. You can use the following two options in the FILE statement to customize the log for that report.

LINESIZE=*value* specifies the maximum number of columns per line for reports and the maximum record length for data files.

PAGESIZE=*value* specifies the maximum number of lines to print on each page of output.

Note: FILE statement options apply only to the output specified in the FILE statement, whereas the LINESIZE= and PAGESIZE= SAS system options apply to all subsequent listings. Δ

LINESIZE= system option

specifies the line size (printer line width) for the SAS log and SAS output that are used by the DATA step and procedures.

MSGCASE

specifies whether to display notes, warning, and error messages in uppercase letters or lowercase letters.

MISSING= system option

specifies the character to be printed for missing numeric variable values.

NUMBER system option

controls whether the page number prints on the first title line of each page of printed output.

PAGE statement

skips to a new page in the SAS log and continues printing from there.

PAGESIZE= system option

specifies the number of lines that you can print per page of SAS output.

SKIP statement

skips a specified number of lines in the SAS log.

STIMEFMT= system option

specifies the format to use for displaying the read and cpu processing times when the STIMER system option is set. The STIMEFMT= system option is valid under Windows, VMS, and UNIX operating environments.

Operating Environment Information: The range of values for the FILE statement and for SAS system options depends on your operating environment. See the SAS documentation for your operating environment for more information. Δ

For more information about how to use these and other SAS system options and statements, see *SAS Language Reference: Dictionary*.

Other System Options That Affect the SAS Log

The following system options pertain to the SAS log other than by the content and appearance of the SAS log:

ALTLOG= system option

specifies the destination for a copy of the SAS log.

LOG= system option

specifies the destination for the SAS log when SAS is run in batch mode.

Traditional SAS Listing Output

Example of Traditional Listing Output

Many SAS procedures process or analyze data and can produce output as one result. You can also generate a listing by the DATA step, using a combination of the FILE and PUT statements.

See the procedure descriptions in *Base SAS Procedures Guide* for examples of output from SAS procedures. For a discussion and examples of DATA step output, see the FILE and PUT statements in *SAS Language Reference: Dictionary*.

This example produces a listing that is generated by the PUT and FILE statements in a DATA step. The input file is the SAS data set GRAIN_PRODUCERS.

```
options pagesize=60 linesize=64 nodate pageno=1;

title 'Leading Grain Producers';
title2 'for 1996';

data _null_;
  set grain_producers;
  file print header=newpage;
  if year=1996;
  format country $entry.;
  label type='Grain';
  put country @25 type @50 kilotons;
  return;
newpage:
  put 'Country' @25 'Grain' @50 'Kilotons';
  put 60*'=';
  return;
run;
```

Leading Grain Producers for 1996			1
Country	Grain	Kilotons	
Brazil	Wheat	3302	
Brazil	Rice	10035	
Brazil	Corn	31975	
China	Wheat	109000	
China	Rice	190100	
China	Corn	119350	
India	Wheat	62620	
India	Rice	120012	
India	Corn	8660	
Indonesia	Wheat	.	
Indonesia	Rice	51165	
Indonesia	Corn	8925	
United States	Wheat	62099	
United States	Rice	7771	
United States	Corn	236064	

Making Output Descriptive

There are several ways to customize SAS procedure output and DATA step output. You can change the look of output by adding informative titles, footnotes, and labels, and by changing the way the information is formatted on the page. The following list describes some of the statements and SAS system options that you can use.

CENTER | NOCENTER system option

controls whether output is centered. By default, SAS centers titles and procedure output on the page and on the personal computer display.

DATE | NODATE system option

controls printing of date and time values. When this option is enabled, SAS prints on the top of each page of output the date and time the SAS job started. When you run SAS in interactive mode, the date and time the job started is the date and time you started your SAS session.

FOOTNOTE statements

print footnotes at the bottom of each output page. You can also use the FOOTNOTES window for this purpose.

FORMCHAR=

specifies the default output formatting characters for some procedures such as CALENDAR, FREQ, REPORT, and TABULATE.

FORMDLIM=

specifies a character that is used to delimit page breaks in SAS output.

LABEL statement

associates descriptive labels with variables. With most procedure output, SAS writes the label rather than the variable name.

The LABEL statement also provides descriptive labels when it is used with certain SAS procedures. See *Base SAS Procedures Guide* for information on using the LABEL statement with a specific procedure (for example, the PRINT procedure).

LINESIZE= and PAGESIZE= system options

can change the default number of lines per page (page size) and characters per line (line size) for printed output. The default depends on the method of running SAS and the settings of certain SAS system options. Specify new page and line sizes in the OPTIONS statement or OPTIONS window. You can also specify line and page size for DATA step output in the FILE statement.

The values you use for the LINESIZE= and PAGESIZE= system options can significantly affect the appearance of the output that is produced by some SAS procedures.

NUMBER | NONUMBER and PAGENO= system options

control page numbering. The NUMBER system option controls whether the page number prints on the first title line of each page of printed output. You can also specify a beginning page number for the next page of output produced by SAS by using the PAGENO= system option.

TITLE statements

print titles at the top of each output page. By default, SAS prints the following title:

```
The SAS System
```

You can use the TITLE statement or TITLES window to replace the default title or specify other descriptive titles for SAS programs. You can use the null title statement (`title;`) to suppress a TITLE statement.

See *SAS Language Reference: Dictionary* for more information about how to use these and other SAS system options and statements.

Reformatting Values

Certain SAS statements, procedures, and options enable you to print values using specified formats. You can apply or change formats with the `FORMAT` and `ATTRIB` statements, or with the `VAR` window in a windowing environment.

The `FORMAT` procedure enables you to design your own formats and informats, giving you added flexibility in displaying values. See the `FORMAT` procedure in *Base SAS Procedures Guide* for more information.

Printing Missing Values

SAS represents ordinary missing numeric values in a SAS listing as a single period, and missing character values as a blank space. If you specified special missing values for numeric variables, SAS writes the letter or the underscore. For character variables, SAS writes a series of blanks equal to the length of the variable.

The `MISSING=` system option enables you to specify a character to print in place of the period for ordinary missing numeric values.

Changing the Destination of the Log and the Output

You can redirect both the SAS log and procedure output to your personal computer display, to a printer, or to an external file. You can redirect output using the following methods:

`PRINTTO` procedure

routes `DATA` step, log, or procedure output from the system default destinations to the destination you choose.

`FILENAME` statement

associates a fileref with an external file or output device and enables you to specify file and device attributes.

`FILE` command

stores the contents of the `LOG` or `OUTPUT` windows in files you specify, when the command is issued from within the windowing environment.

SAS system options

redefine the destination of log and output for an entire SAS program. These system options are specified when you invoke SAS. The system options used to route output are the `ALTLOG=`, `ALTPRINT=`, `LOG=`, and `PRINT=` options.

Operating Environment Information: The way you specify output destinations when you use SAS system options is dependent on your operating environment. See the SAS documentation for your operating environment for details. \triangle

Output Delivery System


The SAS Output Delivery System (ODS) gives you greater flexibility in generating, storing, and reproducing SAS procedure and `DATA` step output, with a wide range of

formatting options. ODS provides formatting functionality that is not available when using individual procedures or the DATA step without ODS. ODS overcomes these limitations and enables you to format your output more easily.

Before Version 7, most SAS procedures generated output that was designed for a traditional line-printer. This type of output has limitations that prevents you from getting the most value from your results:

- Traditional SAS output is limited to monospace fonts. With today's desktop document editors and publishing systems, you need more versatility in printed output.
- Some commonly used procedures produce printed output but do not create an output data set. Many times it would be very convenient to produce not only printed output but also an output data set that you could use as input to another SAS procedure or to a DATA step.

For complete information about the Output Delivery System, see the *SAS Output Delivery System: User's Guide*.



CHAPTER

10

BY-Group Processing in SAS Programs

Definition of BY-Group Processing 131

References for BY-Group Processing 131


Definition of BY-Group Processing

BY-group processing is a method of processing observations from one or more SAS data sets that are grouped or ordered by values of one or more common variables. You can use BY-group processing in both DATA and PROC steps.

The most common use of BY-group processing in the DATA step is to combine two or more SAS data sets by using the BY statement with a SET, MERGE, MODIFY, or UPDATE statement. When you create reports or summaries with SAS procedures, BY-group processing allows you to group information in the output according to values of one or more variables.

References for BY-Group Processing

- For more information about BY-Group processing, see Chapter 21, “BY-Group Processing in the DATA Step,” on page 329.
- For information about how to use BY-group processing with SAS procedures, see “Fundamental Concepts for Using Base SAS Procedures” and individual procedures in *Base SAS Procedures Guide*.
- For information about using BY-group processing to combine information from multiple SAS data sets, see Chapter 22, “Reading, Combining, and Modifying SAS Data Sets,” on page 341. For even more extensive examples of BY-group processing, see *Combining and Modifying SAS Data Sets: Examples*.
- For information about the BY statement, see Statements in *SAS Language Reference: Dictionary*.
- For information about how to use BY-group processing with other software products, see the SAS documentation for those products.



CHAPTER

11**WHERE-Expression Processing**

<i>Definition of WHERE-Expression Processing</i>	133
<i>Where to Use a WHERE Expression</i>	134
<i>Syntax of WHERE Expression</i>	135
<i>Specifying an Operand</i>	135
Variable	135
SAS Function	136
Constant	137
<i>Specifying an Operator</i>	138
Arithmetic Operators	138
Comparison Operators	138
IN Operator	139
Fully Bounded Range Condition	139
BETWEEN-AND Operator	140
CONTAINS Operator	140
IS NULL or IS MISSING Operator	141
LIKE Operator	141
Sounds-like Operator	142
SAME-AND Operator	143
MIN and MAX Operators	144
Concatenation Operator	144
Prefix Operators	144
<i>Combining Expressions by Using Logical Operators</i>	144
Syntax	144
Processing Compound Expressions	145
Using Parentheses to Control Order of Evaluation	145
<i>Improving Performance of WHERE Processing</i>	146
<i>Processing a Segment of Data That Is Conditionally Selected</i>	146
Applying FIRSTOBS= and OBS= to a Subset of Data	147
Processing a SAS View	147
<i>Deciding Whether to Use a WHERE Expression or a Subsetting IF Statement</i>	148

Definition of WHERE-Expression Processing**WHERE-expression processing**

enables you to conditionally select a subset of observations, so that SAS processes only the observations that meet a set of specified conditions. For example, if you have a SAS data set that contains sales records, you might want to print just the subset of observations for which the sales are greater than \$300,000 but less than \$600,000. In addition, WHERE-expression processing can improve efficiency of a request. For example, if a WHERE expression can be optimized with an index, it

is not necessary for SAS to read all observations in the data set in order to perform the request.

WHERE expression

defines a condition that selected observations must satisfy in order to be processed. You can have a single WHERE expression, referred to as a *simple expression*, such as the following:

```
where sales gt 600000;
```

Or you can have multiple WHERE expressions, referred to as a *compound expression*, such as the following:

```
where sales gt 600000 and salary lt 100000;
```

Where to Use a WHERE Expression

In SAS, you can use a WHERE expression in the following situations:

- WHERE statement in both DATA and PROC steps. For example, the following PRINT procedure includes a WHERE statement so that only the observations where the year is greater than 2001 are printed:

```
proc print data=employees;
  where startdate > '01jan2001'd;
run;
```

- WHERE= data set option. The following PRINT procedure includes the WHERE= data set option:

```
proc print data=employees (where=(startdate > '01jan2001'd));
run;
```

- WHERE clause in the SQL procedure, SCL, and SAS/IML software. For example, the following SQL procedure includes a WHERE clause to select only the states where the murder count is greater than seven:

```
proc sql;
  select state from crime
  where murder > 7;
```

- WHERE command in windowing environments like SAS/FSP software:

```
where age > 15
```

- SAS view (DATA step view, SAS/ACCESS view, PROC SQL view), stored with the definition. For example, the following SQL procedure creates an SQL view named STAT from the data file CRIME and defines a WHERE expression for the SQL view definition:

```
proc sql;
  create view stat as
  select * from crime
  where murder > 7;
```

In some cases, you can combine the methods that you use to specify a WHERE expression. That is, you can use a WHERE statement as follows:

- in conjunction with a WHERE= data set option
- along with the WHERE= data set option in windowing procedures, and in conjunction with the WHERE command

- on a SAS view that has a stored WHERE expression.

For example, it might be useful to combine methods when you merge data sets. That is, you might want different criteria to apply to each data set when you create a subset of data. However, when you combine methods to create a subset of data, there are some restrictions. For example, in the DATA step, if a WHERE statement and a WHERE= data set option apply to the same data set, the data set option takes precedence. For details, see the documentation for the method you are using to specify a WHERE expression.

Note: By default, a WHERE expression does not evaluate added and modified observations. To specify whether a WHERE expression should evaluate updates, you can specify the WHEREUP= data set option. See the WHEREUP= data set option in *SAS Language Reference: Dictionary*. Δ

Syntax of WHERE Expression

A WHERE expression is a type of SAS expression that defines a condition for selecting observations. A WHERE expression can be as simple as a single variable name or a constant (which is a fixed value). A WHERE expression can be a SAS function, or it can be a sequence of operands and operators that define a condition for selecting observations. In general, the syntax of a WHERE expression is as follows:

WHERE *operand* <operator> <operand>

<i>operand</i>	something to be operated on. An operand can be a variable, a SAS function, or a constant. See “Specifying an Operand” on page 135.
<i>operator</i>	a symbol that requests a comparison, logical operation, or arithmetic calculation. All SAS expression operators are valid for a WHERE expression, which include arithmetic, comparison, logical, minimum and maximum, concatenation, parentheses to control order of evaluation, and prefix operators. In addition, you can use special WHERE expression operators, which include BETWEEN-AND, CONTAINS, IS NULL or IS MISSING, LIKE, sounds-like, and SAME-AND. See “Specifying an Operator” on page 138.

For more information on SAS expressions, see Chapter 6, “Expressions,” on page 57.

Specifying an Operand

Variable

A variable is a column in a SAS data set. Each SAS variable has attributes like name and type (character or numeric). The variable type determines how you specify the value for which you are searching. For example:

```
where score > 50;
where date >= '01jan2001'd and time >= '9:00't;
where state = 'Texas';
```

In a WHERE expression, you cannot use automatic variables created by the DATA step (for example, FIRST.variable, LAST.variable, _N_, or variables created in assignment statements).

As in other SAS expressions, the names of numeric variables can stand alone. SAS treats numeric values of 0 or missing as false; other values as true. In the following example, the WHERE expression will return all rows where EMPNUM is not missing and not zero *and* ID is not missing and not zero:

```
where empnum and id;
```

The names of character variables can also stand alone. SAS selects observations where the value of the character variable is not blank. For example, the following WHERE expression returns all values not equal to blank:

```
where lastname;
```

SAS Function

A SAS function returns a value from a computation or system manipulation. Most functions use arguments that you supply, but a few obtain their arguments from the operating environment. To use a SAS function in a WHERE expression, type its name and arguments enclosed in parentheses. Some functions you might want to specify include:

- SUBSTR extracts a substring
- TODAY returns the current date
- PUT returns a given value using a given format.

The following DATA step produces a SAS data set that contains only observations from data set CUSTOMER in which the value of NAME begins with **Mac** and the value of variable CITY is **Charleston** or **Atlanta**:

```
data testmacs;
  set customer;
  where substr (name,1,3) = 'Mac' and
    (city='Charleston' or city='Atlanta');
run;
```

The OF syntax is permitted in some SAS functions, but it cannot be used when using those functions that are specified in a WHERE clause. In the following DATA step example, OF can be used with RANGE.

```
data newfile;
  x1=2;
  x2=3;
  x3=4;
  r=range(of x1-x3);
run;
```

When you use the WHERE clause with RANGE and OF, however, the following error is written to the SAS log:

Output 11.1 Output When WHERE Clause Is Used with OF

```

proc print data=abc;
where range(of x1-x3)=6;
      --
      22
      76
ERROR: Syntax error while parsing WHERE clause.
ERROR 22-322: Syntax error, expecting one of the following: !, !!, &, (, *, **,
+, ', ', -, /, <, <=, <>, =, >, >=, ?,
      AND, BETWEEN, CONTAINS, EQ, GE, GT, LE, LIKE, LT, NE, OR, ^=, |,
||, ~=.
ERROR 76-322: Syntax error, statement will be ignored.
run;

```

Below is a table of SAS functions that can use the OF syntax:

Table 11.1 SAS Functions that Use the OF Syntax

CAT	HARMEANZ	RMS
CATS	KURTOSIS	SKEWNESS
CATT	MAX	STD
CATX	MEAN	STDERR
CSS	MIN	SUM
CV	N	USS
GEOMEAN	NMISS	VAR
GEOMEANZ	ORDINAL	
HARMEAN	RANGE	

Note: The SAS functions that are used in a WHERE expression and can be optimized by an index are the SUBSTR function and the TRIM function. Δ

For more information on SAS functions, see "Functions and CALL Routines" in *SAS Language Reference: Dictionary*.

Constant

A constant is a fixed value such as a number or quoted character string, that is, the value for which you are searching. A constant is a value of a variable obtained from the SAS data set, or values created within the WHERE expression itself. Constants are also called literals. For example, a constant could be a flight number or the name of a city. A constant can also be a time, date, or datetime value.

The value is either numeric or character. Note the following rules regarding whether to use quotation marks:

- If the value is numeric, do not use quotation marks.

```
where price > 200;
```

- If the value is character, use quotation marks.

```
where lastname eq 'Martin';
```

- You can use either single or double quotation marks, but do not mix them. Quoted values must be exact matches, including case.

- It might be necessary to use single quotation marks when double quotation marks appear in the value, or use double quotation marks when single quotation marks appear in the value.

```
where item = '6" decorative pot';
where name ? "D'Amico";
```

- A SAS date constant must be enclosed in quotation marks. When you specify date values, case is not important. You can use single or double quotation marks. The following expressions are equivalent:

```
where birthday = '24sep1975'd;
where birthday = "24sep1975"d;
```

Specifying an Operator

Arithmetic Operators

Arithmetic operators enable you to perform a mathematical operation. The arithmetic operators include the following:

Table 11.2 Arithmetic Operators

Symbol	Definition	Example
*	multiplication	where bonus = salary * .10;
/	division	where f = g/h;
+	addition	where c = a+b;
-	subtraction	where f = g-h;
**	exponentiation	where y = a**2;

Comparison Operators

Comparison operators (also called binary operators) compare a variable with a value or with another variable. Comparison operators propose a relationship and ask SAS to determine whether that relationship holds. For example, the following WHERE expression accesses only those observations that have the value 78753 for the numeric variable ZIPCODE:

```
where zipcode eq 78753;
```

The following table lists the comparison operators:

Table 11.3 Comparison Operators

Symbol	Mnemonic Equivalent	Definition	Example
=	EQ	equal to	where empnum eq 3374;
^= or ~= or \neq or <>	NE	not equal to	where status ne full-time;
>	GT	greater than	where hiredate gt '01jun1982'd;

Symbol	Mnemonic Equivalent	Definition	Example
<	LT	less than	where empnum < 2000;
>=	GE	greater than or equal to	where empnum >= 3374;
<=	LE	less than or equal to	where empnum <= 3374;
	IN	equal to one from a list of values	where state in ('NC','TX');

When you do character comparisons, you can use the colon (:) modifier to compare only a specified prefix of a character string. For example, in the following WHERE expression, the colon modifier, used after the equal sign, tells SAS to look at only the first character in the values for variable LASTNAME and to select the observations with names beginning with the letter **s**:

```
where lastname=: 'S';
```

Note that in the SQL procedure, the colon modifier that is used in conjunction with an operator is not supported; you can use the LIKE operator instead.

IN Operator

The IN operator, which is a comparison operator, searches for character and numeric values that are equal to one from a list of values. The list of values must be in parentheses, with each character value in quotation marks and separated by either a comma or blank.

For example, suppose you want all sites that are in North Carolina or Texas. You could specify:

```
where state = 'NC' or state = 'X';
```

However, it is easier to use the IN operator, which selects any state in the list:

```
where state in ('NC','TX');
```

In addition, you can use the NOT logical operator to exclude a list.

```
where state not in ('CA', 'TN', 'MA');
```

You can use a shorthand notation to specify a range of sequential integers to search. The range is specified by using the syntax M:N as a value in the list to search, where M is the lower bound and N is the upper bound. M and N must be integers, and M, N, and all the integers between M and N are included in the range. For example, the following statements are equivalent.

```
□ y = x in (1, 2, 3, 4, 5, 6, 7, 8, 9, 10);
```

```
□ y = x in (1:10);
```

Fully Bounded Range Condition

A fully bounded range condition consists of a variable between two comparison operators, specifying both an upper and lower limit. For example, the following expression returns the employee numbers that fall within the range of 500 to 1000 (inclusive):

```
where 500 <= empnum <= 1000;
```

Note that the previous range condition expression is equivalent to the following:

```
where empnum >= 500 and empnum <= 1000;
```

You can combine the NOT logical operator with a fully bounded range condition to select observations that fall outside the range. Note that parentheses are required:

```
where not (500 <= empnum <= 1000);
```

BETWEEN-AND Operator

The BETWEEN-AND operator is also considered a fully bounded range condition that selects observations in which the value of a variable falls within an inclusive range of values.

You can specify the limits of the range as constants or expressions. Any range you specify is an inclusive range, so that a value equal to one of the limits of the range is within the range. The general syntax for using BETWEEN-AND is as follows:

```
WHERE variable BETWEEN value AND value;
```

For example:

```
where empnum between 500 and 1000;
where taxes between salary*0.30 and salary*0.50;
```

You can combine the NOT logical operator with the BETWEEN-AND operator to select observations that fall outside the range:

```
where empnum not between 500 and 1000;
```

Note: The BETWEEN-AND operator and a fully bounded range condition produce the same results. That is, the following WHERE expressions are equivalent:

```
where 500 <= empnum <= 1000;
where empnum between 500 and 1000;
```

Δ

CONTAINS Operator

The most common usage of the CONTAINS (?) operator is to select observations by searching for a specified set of characters within the values of a character variable. The position of the string within the variable's values does not matter; however, the operator is case sensitive when making comparisons.

The following examples select observations having the values **Mobay** and **Brisbayne** for the variable COMPANY, but they do not select observations containing **Bayview**:

```
where company contains 'bay';
where company ? 'bay';
```

You can combine the NOT logical operator with the CONTAINS operator to select observations that are not included in a specified string:

```
where company not contains 'bay';
```

You can also use the CONTAINS operator with two variables, that is, to determine whether one variable is contained in another. When you specify two variables, keep in mind the possibility of trailing spaces, which can be resolved using the TRIM function.

```
proc sql;
  select *
  from table1 as a, table2 as b
  where a.fullname contains trim(b.lastname) and
        a.fullname contains trim(b.firstname);
```

In addition, the TRIM function is helpful when you search on a macro variable.

```
proc print;
  where fullname contains trim("&lname");
run;
```

IS NULL or IS MISSING Operator

The IS NULL or IS MISSING operator selects observations in which the value of a variable is missing. The operator selects observations with both regular or special missing value characters and can be used for both character and numeric variables.

```
where idnum is missing;
where name is null;
```

The following are equivalent for character data:

```
where name is null;
where name = ' ';
```

And the following is equivalent for numeric data. This statement differentiates missing values with special missing value characters:

```
where idnum <= .Z;
```

You can combine the NOT logical operator with IS NULL or IS MISSING to select nonmissing values, as follows:

```
where salary is not missing;
```

LIKE Operator

The LIKE operator selects observations by comparing the values of a character variable to a specified pattern, which is referred to as pattern matching. The LIKE operator is case sensitive. There are two special characters available for specifying a pattern:

percent sign (%) specifies that any number of characters can occupy that position. The following WHERE expression selects all employees with a name that starts with the letter **N**. The names can be of any length.

```
where lastname like 'N%';
```

underscore (_) matches just one character in the value for each underscore character. You can specify more than one consecutive underscore character in a pattern, and you can specify a percent sign and an underscore in the same pattern. For example, you can use different forms of the LIKE operator to select character values from this list of first names:

```
Diana
Diane
Dianna
Dianthus
Dyan
```

The following table shows which of these names is selected by using various forms of the LIKE operator:

Pattern	Name Selected
like 'D_an'	Dyan
like 'D_an_'	Diana, Diane
like 'D_an__'	Dianna
like 'D_an%'	all names from list

You can use a SAS character expression to specify a pattern, but you cannot use a SAS character expression that uses a SAS function.

You can combine the NOT logical operator with LIKE to select values that do not have the specified pattern, such as the following:

```
where firstname not like 'D_an%';
```

Because the % and _ characters have special meaning for the LIKE operator, you must use an escape character when searching for the % and _ characters in values. An escape character is a single character that, in a sequence of characters, signifies that what follows takes an alternative meaning. For the LIKE operator, an escape character signifies to search for literal instances of the % and _ characters in the variable's values instead of performing the special-character function.

For example, if the variable X contains the values **abc**, **a_b**, and **axb**, the following LIKE operator with an escape character selects only the value **a_b**. The escape character (/) specifies that the pattern searches for a literal '_' that is surrounded by the characters **a** and **b**. The escape character (/) is not part of the search.

```
where x like 'a/_b' escape '/';
```

Without an escape character, the following LIKE operator would select the values **a_b** and **axb**. The special character underscore in the search pattern matches any single **b** character, including the value with the underscore:

```
where x like 'a_b';
```

To specify an escape character, include the character in the pattern-matching expression, and then the keyword ESCAPE followed by the escape-character expression. When you include an escape character, the pattern-matching expression must be enclosed in quotation marks, and it cannot contain a column name. The escape-character expression evaluates to a single character. The operands must be character or string literals. If it is a single character, enclose it in quotation marks.

```
LIKE 'pattern-matching-expression' ESCAPE 'escape-character-expression'
```

Sounds-like Operator

The sounds-like (=*) operator selects observations that contain a spelling variation of a specified word or words. The operator uses the Soundex algorithm to compare the variable value and the operand. For more information, see the SOUNDEX function in *SAS Language Reference: Dictionary*.

Note: Note that the SOUNDEX algorithm is English-biased, and is less useful for languages other than English. Δ

Although the sounds-like operator is useful, it does not always select all possible values. For example, consider that you want to select observations from the following list of names that sound like Smith:

```
Schmitt
Smith
Smithson
Smitt
Smythe
```

The following WHERE expression selects all the names from this list except **Smithson**:

```
where lastname=* 'Smith';
```

You can combine the NOT logical operator with the sounds-like operator to select values that do not contain a spelling variation of a specified word or words, such as:

```
where lastname not =* 'Smith';
```

Note: The sounds-like operator cannot be optimized with an index. △

SAME-AND Operator

Use the SAME-AND operator to add more conditions to an existing WHERE expression later in the program without retyping the original conditions. This capability is useful with the following:

- interactive SAS procedures
- full-screen SAS procedures that enable you to type a WHERE expression on the command line
- any kind of RUN-group processing.

Use the SAME-AND operator when you already have a WHERE expression defined and you want to insert additional conditions. The SAME-AND operator has the following form:

```
where-expression-1;
... SAS statements...
WHERE SAME AND where-expression-2;
... SAS statements...
WHERE SAME AND where-expression-n;
```

SAS selects observations that satisfy the conditions after the SAME-AND operator in addition to any previously defined conditions. SAS treats all of the existing conditions as if they were conditions separated by AND operators in a single WHERE expression.

The following example shows how to use the SAME-AND operator within RUN groups in the GPLOT procedure. The SAS data set YEARS has three variables and contains quarterly data for the 1990–1997 period:

```
proc gplot data=years;
  plot unit*quar=year;
run;

  where year > '01jan1991'd;
run;

  where same and year < '01jan1996'd;
run;
```

The following WHERE expression is equivalent to the preceding code:

```
where year > '01jan1991'd and year < '01jan1996'd;
```

MIN and MAX Operators

Use the MIN or MAX operators to find the minimum or maximum value of two quantities. Surround the operators with the two quantities whose minimum or maximum value you want to know.

- The MIN operator returns the lower of the two values.
- The MAX operator returns the higher of two values.

For example, if A is less than B, then the following would return the value of A:

```
where x = (a min b);
```

Note: The symbol representation $><$ is not supported, and $<>$ is interpreted as “not equal to.” Δ

Concatenation Operator

The concatenation operator concatenates character values. You indicate the concatenation operator as follows:

- || (two OR symbols)
- !! (two exclamation marks)
- ;| (two broken vertical bars).

For example,

```
where name = 'John' || 'Smith';
```

Prefix Operators

The plus sign (+) and minus sign (–) can be either prefix operators or arithmetic operators. They are prefix operators when they appear at the beginning of an expression or immediately preceding an open parentheses. A prefix operator is applied to the variable, constant, SAS function, or parenthetical expression.

```
where z = -(x + y);
```

Note: The NOT operator is also considered a prefix operator. Δ

Combining Expressions by Using Logical Operators

Syntax

You can combine or modify WHERE expressions by using the logical operators (also called Boolean operators) AND, OR, and NOT. The basic syntax of a compound WHERE expression is as follows:

```
WHERE where-expression-1 logical-operator where-expression-n ;
```

AND combines two conditions by finding observations that satisfy both conditions. For example:

	<code>where skill eq 'java' and years eq 4;</code>
OR	combines two conditions by finding observations that satisfy either condition or both. For example: <code>where skill eq 'java' or years eq 4;</code>
NOT	modifies a condition by finding the complement of the specified criteria. You can use the NOT logical operator in combination with any SAS and WHERE expression operator. And you can combine the NOT operator with AND and OR. For example: <code>where skill not eq 'java' or years not eq 4;</code>

The logical operators and their equivalent symbols are shown in the following table:

Table 11.4 Logical (Boolean) Operators

Symbol	Mnemonic Equivalent
&	AND
! or or !	OR
^ or ~ or \neg	NOT

Processing Compound Expressions

When SAS encounters a compound WHERE expression (multiple conditions), the software follows rules to determine the order in which to evaluate each expression. When WHERE expressions are combined, SAS processes the conditions in a specific order:

- 1 The NOT expression is processed first.
- 2 Then the expressions joined by AND are processed.
- 3 Finally, the expressions joined by OR are processed.

For a complete discussion of the rules for evaluating compound expressions, see “Order of Evaluation in Compound Expressions” on page 73.

Using Parentheses to Control Order of Evaluation

Even though SAS evaluates logical operators in a specific order, you can control the order of evaluation by nesting expressions in parentheses. That is, an expression enclosed in parentheses is processed before one not enclosed. The expression within the innermost set of parentheses is processed first, followed by the next deepest, moving outward until all parentheses have been processed.

For example, suppose you want a list of all the Canadian sites that have both SAS/GRAPH and SAS/STAT software, so you issue the following expression:

```
where product='GRAPH' or product='STAT' and country='Canada';
```

The result, however, includes all sites that license SAS/GRAPH software along with the Canadian sites that license SAS/STAT software. To obtain the correct results, you can use parentheses, which causes SAS to evaluate the comparisons within the parentheses first, providing a list of sites with either product licenses, then the result is used for the remaining condition:

```
where (product='GRAPH' or product='STAT') and country='Canada';
```

Improving Performance of WHERE Processing

Indexing a SAS data set can significantly improve the performance of WHERE processing. An index is an optional file that you can create for SAS data files in order to provide direct access to specific observations.

Processing a WHERE expression without an index requires SAS to sequentially read observations in order to find the ones that match the selection criteria. Without an index, SAS first checks for the sort indicator, which is stored with the data file from a previous SORT procedure or SORTEDBY= data set option. If the sort indicator is validated, SAS takes advantage of it and stops reading the file once it is clear there are no more values that satisfy the WHERE expression. For example, consider a data set that is sorted by Age, without an index. To process the expression **where age le 25**, SAS stops reading observations after it finds an observation that is greater than 25. Note that while SAS can determine when to stop reading observations, without an index, there is no indication where to begin, so SAS always begins with the first observation, which can require reading a lot of observations.

Having an index enables SAS to determine which observations satisfy the criteria, which is referred to as optimizing the WHERE expression. However, by default, SAS decides whether to use the index or read the entire data set sequentially. For details on how SAS uses an index to process a WHERE expression, see “Using an Index for WHERE Processing” on page 515.

In addition to creating indexes for the data set, here are some guidelines for writing efficient WHERE expressions:

Table 11.5 Constructing Efficient WHERE Expressions

Guideline	Efficient	Inefficient
Avoid using the LIKE operator that begins with % or _.	where country like 'A%INA';	where country like '%INA';
Avoid using arithmetic expressions.	where salary > 48000;	where salary > 12*4000;

Processing a Segment of Data That Is Conditionally Selected

When you conditionally select a subset of observations with a WHERE expression, you can also segment that subset by applying FIRSTOBS=, OBS= or both processing (as data set options and system options). When used with a WHERE expression,

- FIRSTOBS= specifies the observation number within the subset of data selected by the WHERE expression to begin processing.
- OBS= specifies when to stop processing observations from the subset of data selected by the WHERE expression.

When used with a WHERE expression, the values specified for OBS= and FIRSTOBS= are not the physical observation number in the data set, but a logical number in the subset. For example, **obs=3** does not mean the third observation number in the data set; instead, it means the third observation in the subset of data selected by the WHERE expression.

Applying OBS= and FIRSTOBS= processing to a subset of data is supported for the WHERE statement, WHERE= data set option, and WHERE clause in the SQL procedure.

If you are processing a SAS view that is a view of another view (nested views), applying OBS= and FIRSTOBS= to a subset of data could produce unexpected results. For nested views, OBS= and FIRSTOBS= processing is applied to each SAS view, starting with the root (lowest-level) view, and then filtering observations for each SAS view. The result could be that no observations meet the subset and segment criteria. See “Processing a SAS View” on page 147.

Applying FIRSTOBS= and OBS= to a Subset of Data

The following SAS program illustrates how to specify a condition to subset data, and how to specify a segment of the subset of data to process.

```
data A; ❶
  do I=1 to 100;
    X=I + 1;
    output;
  end;
run;

proc print data=work.a (firstobs=2 ❸ obs=4 ❹ ;
  where I > 90; ❷
run;
```

- ❶ The DATA step creates a data set named WORK.A containing 100 observations and two variables: I and X.
- ❷ The WHERE expression `I > 90` tells SAS to process only the observations that meet the specified condition, which results in the subset of observations 91 through 100.
- ❸ The FIRSTOBS= data set option tells SAS to begin processing with the 2nd observation in the subset of data, which is observation 92.
- ❹ The OBS= data set option tells SAS to stop processing when it reaches the 4th observation in the subset of data, which is observation 94.

The result of PROC PRINT is observations 92, 93, and 94.

Processing a SAS View

The following SAS program creates a data set, a SAS view for the data set, then a second SAS view that subsets data from the first SAS view. Both a WHERE statement and the OBS= system option are used.

```
data a; ❶
  do I=1 to 100;
    X=I + 1;
    output;
  end;
run;

data viewa/view=viewa; ❷
  set a;
  Z = X+1;
run;
```

```

data viewb/view=viewb; ❸
  set viewa;
  where I > 90;
run;

options obs=3; ❹

proc print data=work.viewb; ❺
run;

```

- ❶ The first DATA step creates a data set named WORK.A, which contains 100 observations and two variables: I and X.
- ❷ The second DATA step creates a SAS view named WORK.VIEWA containing 100 observations and three variables: I, X (from data set WORK.A), and Z (assigned in this DATA step).
- ❸ The third DATA step creates a SAS view named WORK.VIEWB and subsets the data with a WHERE statement, which results in the view accessing ten observations.
- ❹ The OBS= system option applies to the previous SET VIEWA statement, which tells SAS to stop processing when it reaches the 3rd observation in the subset of data being processed.
- ❺ When SAS processes the PRINT procedure, the following occurs:
 - 1 First, SAS applies **obs=3** to WORK.VIEWA, which stops processing at the 3rd observation.
 - 2 Next, SAS applies the condition **I > 90** to the three observations being processed. None of the observations meet the criteria.
 - 3 PROC PRINT results in no observations.

To prevent the potential of unexpected results, you can specify **obs=max** when creating WORK.VIEWA to force SAS to read all the observations in the root (lowest-level) view:

```

data viewa/view=viewa;
  set a (obs=max);
  z = x+1;
run;

```

The PRINT procedure processes observations 91, 92, and 93.

Deciding Whether to Use a WHERE Expression or a Subsetting IF Statement

To conditionally select observations from a SAS data set, you can use either a WHERE expression or a subsetting IF statement. While they both test a condition to determine whether SAS should process an observation, they differ as follows:

- The subsetting IF statement can be used only in a DATA step. A subsetting IF statement tests the condition after an observation is read into the Program Data Vector (PDV). If the condition is true, SAS continues processing the current observation. Otherwise, the observation is discarded, and processing continues with the next observation.
- You can use a WHERE expression in both a DATA step and SAS procedures, as well as in a windowing environment, SCL programs, and as a data set option. A


WHERE expression tests the condition before an observation is read into the PDV. If the condition is true, the observation is read into the PDV and processed. If the condition is false, the observation is not read into the PDV, and processing continues with the next observation, which can yield substantial savings when observations contain many variables or very long character variables (up to 32K bytes). Additionally, a WHERE expression can be optimized with an index, and the WHERE expression enables more operators, such as LIKE and CONTAINS.

Note: Although it is generally more efficient to use a WHERE expression and avoid the move to the PDV before processing, if the data set contains observations with very few variables, the move to the PDV could be cheap. However, one variable containing 32K bytes of character data is not cheap, even though it is only one variable. Δ

In most cases, you can use either method. However, the following table provides a list of tasks that require you to use a specific method:

Table 11.6 Tasks Requiring Either WHERE Expression or Subsetting IF Statement

Task	Method
Make the selection in a procedure without using a preceding DATA step	WHERE expression
Take advantage of the efficiency available with an indexed data set	WHERE expression
Use one of a group of special operators, such as BETWEEN-AND, CONTAINS, IS MISSING or IS NULL, LIKE, SAME-AND, and Sounds-Like	WHERE expression
Base the selection on anything other than a variable value that already exists in a SAS data set. For example, you can select a value that is read from raw data, or a value that is calculated or assigned during the course of the DATA step	subsetting IF
Make the selection at some point during a DATA step rather than at the beginning	subsetting IF
Execute the selection conditionally	subsetting IF



CHAPTER

12

Optimizing System Performance

<i>Definitions for Optimizing System Performance</i>	151
<i>Collecting and Interpreting Performance Statistics</i>	152
<i>Using the FULLSTIMER and STIMER System Options</i>	152
<i>Interpreting FULLSTIMER and STIMER Statistics</i>	152
<i>Techniques for Optimizing I/O</i>	153
<i>Overview of Techniques for Optimizing I/O</i>	153
<i>Using WHERE Processing</i>	153
<i>Using DROP and KEEP Statements</i>	154
<i>Using LENGTH Statements</i>	154
<i>Using the OBS= and FIRSTOBS= Data Set Options</i>	154
<i>Creating SAS Data Sets</i>	155
<i>Using Indexes</i>	155
<i>Accessing Data through SAS Views</i>	155
<i>Using Engines Efficiently</i>	156
<i>Setting the BUFNO=, BUFSIZE=, CATCACHE=, and COMPRESS= System Options</i>	156
<i>Using the SASFILE Statement</i>	157
<i>Techniques for Optimizing Memory Usage</i>	158
<i>Techniques for Optimizing CPU Performance</i>	158
<i>Reducing CPU Time by Using More Memory or Reducing I/O</i>	158
<i>Storing a Compiled Program for Computation-Intensive DATA Steps</i>	158
<i>Reducing Search Time for SAS Executable Files</i>	158
<i>Specifying Variable Lengths</i>	159
<i>Using Parallel Processing</i>	159
<i>Reducing CPU Time by Modifying Program Compilation Optimization</i>	159
<i>Calculating Data Set Size</i>	160

Definitions for Optimizing System Performance

performance statistics

are measurements of the total input and output operations (I/O), memory, and CPU time used to process individual DATA or PROC steps. You can obtain these statistics by using SAS system options that can help you measure your job's initial performance and to determine how to improve performance.

system performance

is measured by the overall amount of I/O, memory, and CPU time that your system uses to process SAS programs. By using the techniques discussed in the following sections, you can reduce or reallocate your usage of these three critical resources to improve system performance. While you might not be able to take

advantage of every technique for every situation, you can choose the ones that are best suited for a particular situation.

Collecting and Interpreting Performance Statistics

Using the FULLSTIMER and STIMER System Options

The FULLSTIMER and STIMER system options control the printing of performance statistics in the SAS log. These options produce different results, depending on your operating environment. See the SAS documentation for your operating environment for details about the output that SAS generates for these options.

The following output shows an example of the FULLSTIMER output in the SAS log, as produced in a UNIX operating environment.

Output 12.1 Sample Results of Using the FULLSTIMER Option in a UNIX Operating Environment

```
NOTE: DATA statement used:
      real time           0.19 seconds
      user cpu time       0.06 seconds
      system cpu time     0.01 seconds
      Memory                460k
      Semaphores   exclusive 194 shared 9 contended 0
      SAS Task context switches      1 splits 0
```

The STIMER option reports a subset of the FULLSTIMER statistics. The following output shows an example of the STIMER output in the SAS log in a UNIX operating environment.

Output 12.2 Sample Results of Using the STIMER Option in a UNIX Operating Environment

```
NOTE: DATA statement used:
      real time           1.16 seconds
      cpu time            0.09 seconds
```

Operating Environment Information: See the documentation for your operating environment for information about how STIMER differs from FULLSTIMER in your operating environment. The information that these options display varies depending on your operating environment, so statistics that you see might differ from the ones shown. Δ

Interpreting FULLSTIMER and STIMER Statistics

Several types of resource usage statistics are reported by the STIMER and FULLSTIMER options, including real time (elapsed time) and CPU time. *Real time* represents the clock time it took to execute a job or step; it is heavily dependent on the capacity of the system and the current load. As more users share a particular resource,

less of that resource is available to you. *CPU time* represents the actual processing time required by the CPU to execute the job, exclusive of capacity and load factors. If you must wait longer for a resource, your CPU time will not increase, but your real time will increase. It is not advisable to use real time as the only criterion for the efficiency of your program because you cannot always control the capacity and load demands on your system. A more accurate assessment of system performance is CPU time, which decreases more predictably as you modify your program to become more efficient.

The statistics reported by FULLSTIMER relate to the three critical computer resources: I/O, memory, and CPU time. Under many circumstances, reducing the use of any of these three resources usually results in better throughput of a particular job and a reduction of real time used. However, there are exceptions, as described in the following sections.

Techniques for Optimizing I/O

Overview of Techniques for Optimizing I/O

I/O is one of the most important factors for optimizing performance. Most SAS jobs consist of repeated cycles of reading a particular set of data to perform various data analysis and data manipulation tasks. To improve the performance of a SAS job, you must reduce the number of times SAS accesses disk or tape devices.

To do this, you can modify your SAS programs to process only the necessary variables and observations by:

- using WHERE processing
- using DROP and KEEP statements
- using LENGTH statements
- using the OBS= and FIRSTOBS= data set options.

You can also modify your programs to reduce the number of times it processes the data internally by:

- creating SAS data sets
- using indexes
- accessing data through SAS views
- using engines efficiently.

You can reduce the number of data accesses by processing more data each time a device is accessed by

- setting the BUFNO=, BUFSIZE=, CATCACHE=, and COMPRESS= system options
- using the SASFILE global statement to open a SAS data set and allocate enough buffers to hold the entire data set in memory.

Note: Sometimes you might be able to use more than one method, making your SAS job even more efficient. Δ

Using WHERE Processing

You might be able to use a WHERE statement in a procedure in order to perform the same task as a DATA step with a subsetting IF statement. The WHERE statement can

eliminate extra DATA step processing when performing certain analyses because unneeded observations are not processed.

For example, the following DATA step creates a data set SEATBELT, which contains only those observations from the AUTO.SURVEY data set for which the value of SEATBELT is YES. The new data set is then printed.

```
libname auto '/users/autodata';
data seatbelt;
    set auto.survey;
    if seatbelt='yes';
run;

proc print data=seatbelt;
run;
```

However, you can get the same output from the PROC PRINT step without creating a data set if you use a WHERE statement in the PRINT procedure, as in the following example:

```
proc print data=auto.survey;
    where seatbelt='yes';
run;
```

The WHERE statement can save resources by eliminating the number of times you process the data. In this example, you might be able to use less time and memory by eliminating the DATA step. Also, you use less I/O because there is no intermediate data set. Note that you cannot use a WHERE statement in a DATA step that reads raw data.

The extent of savings that you can achieve depends on many factors, including the size of the data set. It is recommended that you test your programs to determine which is the most efficient solution. See “Deciding Whether to Use a WHERE Expression or a Subsetting IF Statement” on page 148 for more information.

Using DROP and KEEP Statements

Another way to improve efficiency is to use DROP and KEEP statements to reduce the size of your observations. When you create a temporary data set and include only the variables that you need, you can reduce the number of I/O operations that are required to process the data. See *SAS Language Reference: Dictionary* for more information on the DROP and KEEP statements.

Using LENGTH Statements

You can also use LENGTH statements to reduce the size of your observations. When you include only the necessary storage space for each variable, you can reduce the number of I/O operations that are required to process the data. Before you change the length of a numeric variable, however, see “Specifying Variable Lengths” on page 159. See *SAS Language Reference: Dictionary* for more information on the LENGTH statement.

Using the OBS= and FIRSTOBS= Data Set Options

You can also use the OBS= and FIRSTOBS= data set options to reduce the number of observations processed. When you create a temporary data set and include only the necessary observations, you can reduce the number of I/O operations that are required

to process the data. See *SAS Language Reference: Dictionary* for more information on the OBS= and FIRSTOBS= data set options.

Creating SAS Data Sets

If you process the same raw data repeatedly, it is usually more efficient to create a SAS data set. SAS can process SAS data sets more efficiently than it can process raw data files.

Another consideration involves whether you are using data sets created with previous releases of SAS. If you frequently process data sets created with previous releases, it is sometimes more efficient to convert that data set to a new one by creating it in the most recent version of SAS. See Chapter 34, “SAS 9.2 Compatibility with SAS Files from Earlier Releases,” on page 577 for more information.

Using Indexes

An index is an optional file that you can create for a SAS data file to provide direct access to specific observations. The index stores values in ascending value order for a specific variable or variables and includes information as to the location of those values within observations in the data file. In other words, an index allows you to locate an observation by the value of the indexed variable.

Without an index, SAS accesses observations sequentially in the order in which they are stored in the data file. With an index, SAS accesses the observation directly. Therefore, by creating and using an index, you can access an observation faster.

In general, SAS can use an index to improve performance in these situations:

- For WHERE processing, an index can provide faster and more efficient access to a subset of data.
- For BY processing, an index returns observations in the index order, which is in ascending value order, without using the SORT procedure.
- For the SET and MODIFY statements, the KEY= option allows you to specify an index in a DATA step to retrieve particular observations in a data file.

Note: An index exists to improve performance. However, an index conserves some resources at the expense of others. Therefore, you must consider costs associated with creating, using, and maintaining an index. See “Understanding SAS Indexes” on page 505 for more information about indexes and deciding whether to create one. Δ

Accessing Data through SAS Views

You can use the SQL procedure or a DATA step to create SAS views of your data. A SAS view is a stored set of instructions that subsets your data with fewer statements. Also, you can use a SAS view to group data from several data sets without creating a new one, saving both processing time and disk space. See Chapter 28, “SAS Views,” on page 529 and the *Base SAS Procedures Guide* for more details.

Using Engines Efficiently

If you do not specify an engine on a LIBNAME statement, SAS must perform extra processing steps in order to determine which engine to associate with the SAS library. SAS must look at all of the files in the directory until it has enough information to determine which engine to use. For example, the following statement is efficient because it explicitly tells SAS to use a specific engine for the libref FRUITS:

```
/* Engine specified. */

libname fruits v9 '/users/myid/mydir';
```

The following statement does not explicitly specify an engine. In the output, notice the NOTE about mixed engine types that is generated:

```
/* Engine not specified. */

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

Output 12.3 Output From the LIBNAME Statement

<pre>NOTE: Directory for library FRUITS contains files of mixed engine types. NOTE: Libref FRUITS was successfully assigned as follows: Engine: V9 Physical Name: /users/myid/mydir</pre>

Operating Environment Information: In the z/OS operating environment, you do not need to specify an engine for certain types of libraries. Δ

See Chapter 36, “SAS Engines,” on page 589 for more information about SAS engines.

Setting the BUFNO=, BUFSIZE=, CATCACHE=, and COMPRESS= System Options

The following SAS system options can help you reduce the number of disk accesses that are needed for SAS files, though they might increase memory usage.

BUFNO=

SAS uses the BUFNO= option to adjust the number of open page buffers when it processes a SAS data set. Increasing this option’s value can improve your application’s performance by allowing SAS to read more data with fewer passes; however, your memory usage increases. Experiment with different values for this option to determine the optimal value for your needs.

Note: You can also use the CBUFNO= system option to control the number of extra page buffers to allocate for each open SAS catalog. Δ

See system options in *SAS Language Reference: Dictionary* and the SAS documentation for your operating environment for more details on this option.

BUFSIZE=

When the Base SAS engine creates a data set, it uses the BUFSIZE= option to set the permanent page size for the data set. The page size is the amount of data that can be transferred for an I/O operation to one buffer. The default value for BUFSIZE= is determined by your operating environment. Note that the default is set to optimize the sequential access method. To improve performance for direct (random) access, you should change the value for BUFSIZE=.

Whether you use your operating environment's default value or specify a value, the engine always writes complete pages regardless of how full or empty those pages are.

If you know that the total amount of data is going to be small, you can set a small page size with the BUFSIZE= option, so that the total data set size remains small and you minimize the amount of wasted space on a page. In contrast, if you know that you are going to have many observations in a data set, you should optimize BUFSIZE= so that as little overhead as possible is needed. Note that each page requires some additional overhead.

Large data sets that are accessed sequentially benefit from larger page sizes because sequential access reduces the number of system calls that are required to read the data set. Note that because observations cannot span pages, typically there is unused space on a page.

“Calculating Data Set Size” on page 160 discusses how to estimate data set size.

See system options in *SAS Language Reference: Dictionary* and the SAS documentation for your operating environment for more details on this option.

CATCACHE=

SAS uses this option to determine the number of SAS catalogs to keep open at one time. Increasing its value can use more memory, although this might be warranted if your application uses catalogs that will be needed relatively soon by other applications. (The catalogs closed by the first application are cached and can be accessed more efficiently by subsequent applications.)

See system options in *SAS Language Reference: Dictionary* and the SAS documentation for your operating environment for more details on this option.

COMPRESS=

One further technique that can reduce I/O processing is to store your data as compressed data sets by using the COMPRESS= data set option. However, storing your data this way means that more CPU time is needed to decompress the observations as they are made available to SAS. But if your concern is I/O, and not CPU usage, compressing your data might improve the I/O performance of your application.

See *SAS Language Reference: Dictionary* for more details on this option.

Using the SASFILE Statement

The SASFILE global statement opens a SAS data set and allocates enough buffers to hold the entire data set in memory. Once it is read, data is held in memory, available to subsequent DATA and PROC steps, until either a second SASFILE statement closes the file and frees the buffers or the program ends, which automatically closes the file and frees the buffers.

Using the SASFILE statement can improve performance by

- reducing multiple open/close operations (including allocation and freeing of memory for buffers) to process a SAS data set to one open/close operation
- reducing I/O processing by holding the data in memory.

If your SAS program consists of steps that read a SAS data set multiple times and you have an adequate amount of memory so that the entire file can be held in real memory, the program should benefit from using the SASFILE statement. Also, SASFILE is especially useful as part of a program that starts a SAS server such as a SAS/SHARE server. See *SAS Language Reference: Dictionary* for more information on the SASFILE global statement.

Techniques for Optimizing Memory Usage

If memory is a critical resource, several techniques can reduce your dependence on increased memory. However, most of them also increase I/O processing or CPU usage.

However, by *increasing* memory available to SAS by increasing the value of the MEMSIZE= system option (or by using the MEMLEAVE= option, in some operating environments), you can decrease processing time because the amount of time that is spent on paging, or reading pages of data into memory, is reduced. The SORTSIZE= and SUMSIZE= system options enable you to limit the amount of memory that is available to sorting and summarization procedures.

You can also make tradeoffs between memory and other resources, as discussed in “Reducing CPU Time by Using More Memory or Reducing I/O” on page 158. To use the I/O subsystem most effectively, you must use more and larger buffers. However, these buffers share space with the other memory demands of your SAS session.

Operating Environment Information: The MEMSIZE= system option is not available in some operating environments. If MEMSIZE= is available in your operating environment, it might not increase memory. See the documentation for your operating environment for more information. Δ

Techniques for Optimizing CPU Performance

Reducing CPU Time by Using More Memory or Reducing I/O

Executing a single stream of code takes approximately the same amount of CPU time each time that code is executed. Optimizing CPU performance in these instances is usually a tradeoff. For example, you might be able to reduce CPU time by using more memory, because more information can be read and stored in one operation, but less memory is available to other processes.

Also, because the CPU performs all the processing that is needed to perform an I/O operation, an option or technique that reduces the number of I/O operations can also have a positive effect on CPU usage.

Storing a Compiled Program for Computation-Intensive DATA Steps

Another technique that can improve CPU performance is to store a DATA step that is executed repeatedly as a compiled program rather than as SAS statements. This is especially true for large DATA step jobs that are not I/O-intensive. For more information on storing compiled DATA steps, see Chapter 29, “Stored Compiled DATA Step Programs,” on page 539.

Reducing Search Time for SAS Executable Files

The PATH= system option specifies the list of directories (or libraries, in some operating environments) that contain SAS executable files. Your default configuration file specifies a certain order for these directories. You can rearrange the directory specifications in the PATH= option so that the most commonly accessed directories are listed first. Place the least commonly accessed directories last.

Operating Environment Information: The PATH= system option is not available in some operating environments. See the documentation for your operating environment for more information. Δ

Specifying Variable Lengths

When SAS processes the program data vector, it typically moves the data in one large operation rather than by individual variables. When data is properly aligned (in 8-byte boundaries), data movement can occur in as little as two clock cycles (a single load followed by a single store). SAS moves unaligned data by more complex means, at worst, a single byte at a time. This would be at least eight times slower for an 8-byte variable.

Many high performance RISC (Reduced Instruction Set Computer) processors pay a very large performance penalty for movement of unaligned data. When possible, leave numeric data at full width (eight bytes). Note that SAS must widen short numeric data for any arithmetic operation. On the other hand, short numeric data can save both memory and I/O. You must determine which method is most advantageous for your operating environment and situation.

Note: Alignment can be especially important when you process a data set by selecting only specific variables or when you use WHERE processing. Δ

Using Parallel Processing

SAS System 9 supports a new wave of SAS functionality related to parallel processing. Parallel processing means that processing is handled by multiple CPUs simultaneously. This technology takes advantage of SMP computers and provides performance gains for two types of SAS processes: threaded I/O and threaded application processing.

For information, see Chapter 13, "Support for Parallel Processing," on page 161.

Reducing CPU Time by Modifying Program Compilation Optimization

When SAS compiles a program, the code is optimized to remove redundant instructions, missing value checks, and repetitive computations for array subscripts. The code detects patterns of instruction and replaces them with more efficient sequences, and also performs optimizations that pertain to the SAS register. In most cases, performing the code-generation optimization is preferable. If you have a large DATA step program, performing code generation optimization can result in a significant increase in compilation time and overall execution time.

You can reduce or turn off the code generation optimization by using the CGOPTIMIZE= system option. Set the code generation optimization that you want SAS to perform using these CGOPTIMIZE= system option values:

- | | |
|---|---|
| 0 | performs no optimization during code compilation. |
| 1 | specifies to perform stage 1 optimization. Stage 1 optimization removes redundant instructions, missing value checks, and repetitive computations for array subscripts; detects patterns of instructions and replaces them with more efficient sequences. |
| 2 | specifies to perform stage 2 optimization. Stage 2 performs optimizations that pertain to the SAS register. Performing stage 2 optimization on large DATA step programs can result in a significant increase in compilation time. |

- 3 specifies to perform full optimization, which is a combination of stages 1 and 2. This is the default value.


For more information, see the CGOPTIMIZE= System Option in *SAS Language Reference: Dictionary*.

Calculating Data Set Size

If you have already applied optimization techniques but still experience lengthy processing times or excessive memory usage, the size of your data sets might be very large. In that case, further improvement might not be possible.

You can estimate the size of a data set by creating a dummy data set that contains the same variables as your data set. Run the CONTENTS procedure, which shows the size of each observation. Multiply the size by the number of observations in your data set to obtain the total number of bytes that must be processed. You can compare processing statistics with smaller data sets to determine whether the performance of the large data sets is in proportion to their size. If not, further optimization might still be possible.

Note: When you use this technique to calculate the size of a data set, you obtain only an estimate. Internal requirements, such as the storage of variable names, might cause the actual data set size to be slightly different. Δ



CHAPTER

13

Support for Parallel Processing

<i>Definition of Parallel Processing</i>	161
<i>Threaded I/O</i>	161
<i>Threaded Application Processing</i>	162

Definition of Parallel Processing

SAS 9 supports a new wave of SAS functionality related to parallel processing. *Parallel processing* refers to processing that is handled by multiple CPUs simultaneously. This technology takes advantage of hardware that has multiple CPUs, called *SMP computers*, and provides performance gains for two types of SAS processes:

- threaded I/O
- threaded application processing.

SMP computers have multiple CPUs and an operating environment that can spawn and manage multiple pieces of executable code called *threads*. A thread is a single, independent flow of control through a program or within a process. Threading takes advantage of multiple CPUs by dividing processing among the available CPUs. A thread-enabled operating environment provides support for threads; for example, each thread needs a context (like a register set and a program counter), a segment of code to execute, and a piece of memory to use in the process.

Even if your site does not use an SMP computer, SAS 9 can still provide increased performance. Some types of threading can be performed using a single CPU.

Threaded I/O

Some applications can process data faster than the data can be delivered to the application. When an application cannot keep the available CPUs busy, the application is said to be *I/O-bound*.

SAS supports threaded I/O for SAS applications by providing the SAS Scalable Performance Data (SPD) engine. The SPD engine boosts the performance of SAS applications that are I/O bound through parallel processing of partitioned data sets. Partitioned data sets can span disk drives but still be referenced as a single data set. In this way, the SPD engine can read many millions of observations into SAS applications very rapidly by spawning a thread for each data partition and evaluating WHERE expressions in multiple CPUs. SAS 9.1 support for multiple CPUs, for example on a Symmetric Multiprocessor (SMP) computer, and support for multiple disks per CPU make SPD engine's parallel I/O possible. See *SAS Scalable Performance Data Engine: Reference* for full details on this engine's capabilities.

The benefits of support for multiple disks in SAS 9.1 is not limited to use by the SPD engine. Multiple disks can be on an SMP computer, but they can also be a bank of disks on a single-CPU computer. Increasing the number of disks that can be accessed simultaneously increases the amount of data that can be delivered to an application. This is because reading or writing data to and from disk takes much more time than the associated CPU processing that directs the I/O. Even a single-CPU computer can support multiple disk drives, which boosts performance. When an SMP computer is used, the performance gain can be quite significant because each CPU can support multiple disk drives. However, multiple CPUs cannot speed up I/O from a single disk drive. The minimum configuration for I/O performance gain is at least one controller per CPU and at least two disk drives for each controller. For example, a site with four CPUs should have at least four controllers and eight disk drives.

Threaded Application Processing

Some applications receive data faster than they can perform the necessary processing on that data. These applications are sometimes referred to as *CPU-bound*. For CPU-bound applications, the solution is to increase processing power. Support for SMP computers provides access to threaded processing for CPU-bound applications. Even if your application is not currently CPU-bound, if you increase the amount of data that can be delivered to an application, you will naturally increase the need for faster processing of that data. Modifying your application to process data in threads solves this problem.

For SAS 9, certain procedures such as SORT and MEANS have been modified so that they can thread the processing through multiple CPUs, if they are available. In addition, threaded processing is being integrated into a variety of other SAS features in order to improve performance. For example, when you create an index, if sorting is required, SAS attempts to sort the data using the thread-enabled sort.


Some types of processing are not suited to threading, while other types of processing can benefit greatly. For example, sorting can be performed in multiple threads by dividing the sorting process into separately executable sorting processes. One or more of these threads can process data in each CPU. The sorted data from each thread is then joined together and written to the output data set. Sorting can be performed in threads, but the join process and the output process are nonthreadable processes. Even with applications that are highly threadable processes, simply providing additional disks and CPUs might not improve performance. That is, a particular algorithm can benefit by using four CPUs but cannot benefit an equal amount by adding four more CPUs.

For SAS procedures that are thread-enabled, new SAS system options are introduced with SAS 9:

CPUCOUNT=
specifies how many CPUs can be used.

THREAD|NOTTHREADS
controls whether to use threads.

For documentation on the SAS system options, see *SAS Language Reference: Dictionary*. In addition, the documentation for each thread-enabled procedure provides more information. See *Base SAS Procedures Guide*.



CHAPTER


14

Monitoring Performance Using Application Response Measurement (ARM)

What Is ARM? 163

What Is ARM?

Application Response Measurement (ARM) enables you to monitor the availability and performance of transactions within and across diverse applications. For more information, see *SAS Interface to Application Response Measurement (ARM): Reference*.



CHAPTER

15

The SAS Registry

<i>Introduction to the SAS Registry</i>	166
<i>What Is the SAS Registry?</i>	166
<i>Who Should Use the SAS Registry?</i>	166
<i>Where the SAS Registry Is Stored</i>	166
<i>Registry Files in the SASUSER and the SASHELP Libraries</i>	166
<i>How to Restore the Site Defaults</i>	167
<i>How Do I Display the SAS Registry?</i>	167
<i>Definitions for the SAS Registry</i>	167
<i>Managing the SAS Registry</i>	168
<i>Primary Concerns about Managing the SAS Registry</i>	168
<i>Backing Up the SASUSER Registry</i>	169
<i>Why Back Up the SASUSER Registry?</i>	169
<i>When SAS Resets to the Default Settings</i>	169
<i>Ways to Back Up the Registry</i>	169
<i>Using the Explorer to Back Up the SAS Registry</i>	170
<i>Using the Registry Editor to Back Up the SAS Registry</i>	170
<i>Recovering from Registry Failure</i>	170
<i>Using the SAS Registry to Control Color</i>	171
<i>Overview of Colors and the SAS Registry</i>	171
<i>Adding Colors Using the Registry Editor</i>	172
<i>Adding Colors Programmatically</i>	172
<i>Using the Registry Editor</i>	172
<i>When to Use the Registry Editor</i>	172
<i>Starting the Registry Editor</i>	173
<i>Finding Specific Data in the Registry</i>	173
<i>Changing a Value in the SAS Registry</i>	173
<i>Adding a New Value or Key to the SAS Registry</i>	174
<i>Deleting an Item from the SAS Registry</i>	174
<i>Renaming an Item in the SAS Registry</i>	175
<i>Displaying the SASUSER and SASHELP Registry Items Separately</i>	175
<i>Importing a Registry File</i>	175
<i>Exporting a Registry File</i>	176
<i>Configuring Your Registry</i>	176
<i>Configuring Universal Printing</i>	176
<i>Configuring SAS Explorer</i>	176
<i>Configuring Libraries and File Shortcuts with the SAS Registry</i>	177
<i>Fixing Library Reference (Libref) Problems with the SAS Registry</i>	178

Introduction to the SAS Registry

What Is the SAS Registry?

The SAS registry is the central storage area for configuration data for SAS. For example, the registry stores

- the libraries and file shortcuts that SAS assigns at startup
- the menu definitions for Explorer pop-up menus
- the printers that are defined for use
- configuration data for various SAS products.

This configuration data is stored in a hierarchical form, which works in a similar manner to the way directory-based file structures work under the operating environments in UNIX, Windows, VMS, and under the z/OS UNIX System Services (USS).

Note: Host printers are not referenced in the SAS registry. \triangle

Who Should Use the SAS Registry?

The SAS registry is designed for use by system administrators and experienced SAS users. This section provides an overview of registry tools, and describes how to import and export portions of the registry.

CAUTION:

If you make a mistake when you edit the registry, your system might become unstable or unusable. \triangle

Wherever possible, use the administrative tools, such as the New Library window, the PRTDEF procedure, Universal Print windows, and the Explorer Options window, to make configuration changes, rather than editing the registry directly. Using the administrative tools ensures that values are stored properly in the registry when you change the configuration.

CAUTION:

If you use the Registry Editor to change values, you will not be warned if any entry is incorrect. Incorrect entries can cause errors, and can even prevent you from starting a SAS session. \triangle

Where the SAS Registry Is Stored

Registry Files in the SASUSER and the SASHELP Libraries

Although the SAS registry is logically one data store, physically it consists of two different files located in both the SASUSER and SASHELP libraries. The physical filename for the registry is registry.sas7bitm. By default, these registry files are hidden in the SAS Explorer views of the SASHELP and SASUSER libraries.

- The SASHELP library registry file contains the site defaults. The system administrator usually configures the printers that a site uses, the global file shortcuts or libraries that will be assigned at startup, and any other configuration defaults for your site.

- The SASUSER library registry file contains the user defaults. When you change your configuration information through a specialized window such as the Print Setup window or the Explorer Options window, the settings are stored in the SASUSER library.

How to Restore the Site Defaults

If you want to restore the original site defaults to your SAS session, delete the registry.sas7bitm file from your SASUSER library and restart your SAS session.

How Do I Display the SAS Registry?

You can use one of the following three methods to view the SAS registry:

- Issue the REGEDIT command. This opens the SAS Registry Editor
- Select **Solutions ► Accessories ► Registry Editor**
- Submit the following line of code:

```
proc registry list;
run;
```

This method prints the registry to the SAS log, and it produces a large list that contains all registry entries, including subkeys. Because of the large size, it might take a few minutes to display the registry using this method.

For more information about how to view the SAS registry, see the REGISTRY Procedure in *Base SAS Procedures Guide*.

Definitions for the SAS Registry

The SAS registry uses *keys* and *subkeys* as the basis for its structure, instead of using directories and subdirectories like the file systems in DOS or UNIX. These terms and several others described here are frequently used when discussing the SAS Registry:

key An entry in the registry file that refers to a particular aspect of SAS. Each entry in the registry file consists of a key name, followed on the next line by one or more values. Key names are entered on a single line between square brackets ([and]).

The key can be a place holder without values or subkeys associated with it, or it can have many subkeys with associated values. Subkeys are delimited with a backslash (\). The length of a single key name or a sequence of key names cannot exceed 255 characters (including the square brackets and the backslash). Key names can contain any character except the backslash and are not case-sensitive.

The SAS Registry contains only one top-level key, called SAS_REGISTRY. All the keys under SAS_REGISTRY are subkeys.

subkey A key inside another key. Subkeys are delimited with a backslash (\). Subkey names are not case-sensitive. The following key contains one root key and two subkeys:

```
[ SAS_REGISTRY\HKEY_USER_ROOT\CORE ]
```

SAS_REGISTRY
is the root key.

HKEY_USER_ROOT
is a subkey of SAS_REGISTRY. In the SAS registry, there is one other subkey at this level it is HKEY_SYSTEM_ROOT.

CORE
is a subkey of HKEY_USER_ROOT, containing many default attributes for printers, windowing, and so on.

link a value whose contents reference a key. Links are designed for internal SAS use only. These values always begin with the word "link:".

value the names and content associated with a key or subkey. There are two components to a value, the value name and the value content, also known as a value datum.

Display 15.1 Section of the Registry Editor Showing Value Names and Value Data for the Subkey 'HTML'

Name	Data
coco	"D2,71,1E"
coco1	D2,73,1E

.SASXREG file a text file with the file extension .SASXREG that contains the text representation of the actual binary SAS Registry file.

Managing the SAS Registry

Primary Concerns about Managing the SAS Registry

CAUTION:

If you make a mistake when you edit the registry, your system might become unstable or unusable. Whenever possible, use the administrative tools, such as the New Library window, the PRTDEF procedure, Universal Print windows, and the Explorer Options window, to make configuration changes, rather than editing the registry. This is to ensure that values are stored properly in the registry when changing the configuration. \triangle

CAUTION:

If you use the Registry Editor to change values, you will not be warned if any entry is incorrect. Incorrect entries can cause errors, and can even prevent you from starting a SAS session. \triangle

Backing Up the SASUSER Registry

Why Back Up the SASUSER Registry?

The SASUSER* part of the registry contains personal settings. It is a good idea to back up the SASUSER part of the registry if you have made substantial customizations to your SAS session. Substantial customizations include the following:

- installing new printers
- modifying printer settings from the default printer settings that your system administrator provides for you
- changing localization settings
- altering translation tables with TRANTAB

When SAS Resets to the Default Settings

When SAS starts up, it automatically scans the registry file. SAS restores the registry to its original settings under two conditions:

- If SAS detects that the registry is corrupt, then SAS rebuilds the file.
- If you delete the registry file called `regstry.sas7bitm`, which is located in the SASUSER library, then SAS will restore the SASUSER registry to its default settings.

CAUTION:

Do not delete the registry file that is located in SASHELP; this will prevent SAS from starting. △

Ways to Back Up the Registry

There are two methods for backing up the registry and each achieves different results:

Method 1: Save a copy of the SASUSER registry file called `regstry.sas7bitm`.

The result is an exact copy of the registry at the moment you copied it. If you need to use that copy of the registry to restore a broken copy of the registry, then any changes to the registry after the copy date are lost. However, it is probably better to have this backup file than to revert to the original default registry.

Method 2: Use the Registry Editor or PROC REGISTRY to back up the parts of the SASUSER registry that have changed.

The result is a concatenated copy of the registry, which can be restored from the backup file. When you create the backup file using the `EXPORT=` statement in PROC REGISTRY, or by using the **Export Registry File** utility in the Registry Editor, SAS saves any portions of the registry that have been changed. When SAS restores this backup file to the registry, the backup file is concatenated with the current registry in the following way:

- Any completely new keys, subkeys, or values that were added to the SASUSER registry after the backup date are retained in the new registry.

* The SASHELP part of the registry contains settings that are common to all users at your site. SASHELP is write protected, and can be updated only by a system administrator.

- Any existing keys, subkeys, or values that were changed after SAS was initially installed, then changed again after the backup, are overwritten and revert to the backup file values after the restore.
- Any existing keys, subkeys, or values that retain the original default values, will have the default values after the restore.

Using the Explorer to Back Up the SAS Registry

To use the Explorer to back up the SAS Registry:

- 1 Start SAS Explorer with the **EXPLORER** command, or select **View ► Explorer**.
- 2 Select **Tools ► Options ► Explorer**. The Explorer Options window appears.
- 3 Select the **Members** tab.
- 4 Select **ITEMSTOR** in the **Type** list.
- 5 Click **Unhide**.

If **ITEMSTOR** does not have an icon associated with it in the **Type** list, you will be prompted to select an icon.

- 6 Open the Sasuser library in the Explorer window.
- 7 Right-click the **Registry.Itemstor** file.
- 8 Select **Copy** from the pop-up menu and copy the **Registry** file. SAS will name the file **Registry_copy**.

Note: You can also use a copy command from your operating environment to make a copy of your registry file for backup purposes. When viewed from outside SAS Explorer, the filename is **registry.sas7bitm**. Under z/OS, you cannot use the environment copy command to copy your registry file unless your SASUSER library is assigned to an HFS directory. △

Using the Registry Editor to Back Up the SAS Registry

Using the Registry Editor to back up the SAS registry is generally the preferred backup method, because it retains any new keys or values in case you must restore the registry from the backup.

To use the Registry Editor to back up the SAS Registry:

- 1 Open the Registry Editor with the **regedit** command.
- 2 Select the top-level key in the left pane of the registry window.
- 3 From the Registry Editor, select **File ► Export Registry File**. A Save As window appears.
- 4 Type a name for your registry backup file in the filename field. (SAS will apply the proper file extension name for your operating system.)
- 5 Click **Save**.

This saves the registry backup file in SASUSER. You can control the location of your registry backup file by specifying a different location in the Save As window.

Recovering from Registry Failure

This section gives instructions for restoring the registry with a backup file, and shows you how to repair a corrupt registry file.

To install the registry backup file that was created using SAS Explorer or an operating system copy command:

- 1 Change the name of your corrupt registry file to something else.

- 2 Rename your backup file to *registry.sas7bitm*, which is the name of your registry file.
- 3 Copy your renamed registry file to the SASUSER location where your previous registry file was located.
- 4 Restart your SAS session.

To restore a registry backup file created with the Registry Editor:

- 1 Open the Registry Editor with the **regedit** command.
- 2 Select **File ► Import Registry File**.
- 3 Select the registry file that you previously exported.
- 4 Click **Open**.
- 5 Restart SAS.

To restore a registry backup file created with PROC REGISTRY:

- 1 Open the Program editor and submit the following program to import the registry file that you created previously.

```
proc registry import=<registry file specification>;
run;
```

This imports the registry file to the SASUSER library.

- 2 If the file is not already properly named, then use Explorer to rename the registry file to *registry.sas7bitm*:
- 3 Restart SAS.

To attempt to repair a damaged registry:

- 1 Rename the damaged registry file to something other than “registry”; for example, *temp*.
- 2 Start your SAS session.
- 3 Define a library pointing to the location of the *temp* registry.

```
libname here '.'
```

- 4 Run the REGISTRY procedure and redefine the SASUSER registry:

```
proc registry setsasuser="here.temp";
run;
```

- 5 Start the Registry Editor with the **regedit** command. Select **Solutions ► Accessories ► Registry Editor ► View All**.
- 6 Edit any damaged fields under the HKEY_USER_ROOT key.
- 7 Close your SAS session and rename the modified registry back to the original name.
- 8 Open a new SAS session to see whether the changes fixed the problem.

Using the SAS Registry to Control Color

Overview of Colors and the SAS Registry

The SAS registry contains the RGB values for color names that are common to most Web browsers. These colors can be used for ODS and GRAPH output. The RGB value is a triplet (Red, Green, Blue), and each component has a range of 00 to FF (0 to 255).

The registry values for color are located in the COLORNAMES\HTML subkey.

Adding Colors Using the Registry Editor

You can create your own new color values by adding them to the registry in the `COLORNAMES\HTML` subkey:

- 1 Open the SAS Registry Editor using the `REGEDIT` command.
- 2 Select the `COLORNAMES\HTML` subkey.
- 3 Select **Edit** \blacktriangleright **New Binary Value**. A pop-up menu appears.
- 4 Enter the color name in the **Value Name** field and the RGB value in the **Value Data** field.
- 5 Click **OK**.

Adding Colors Programmatically

You can create your own new color values by adding them to the registry in the `COLORNAMES\HTML` subkey, using SAS code.

The easiest way is to first write the color values to a file in the layout that the `REGISTRY` procedure expects. Then you import the file by using the `REGISTRY` procedure. In this example, Spanish color names are added to the registry.

```
filename mycolors temp;
data _null_;
  file "mycolors";
  put "[colornames\html]";
  put ' "rojo"=hex:ff,00,00';
  put ' "verde"=hex:00,ff,00';
  put ' "azul"=hex:00,00,ff';
  put ' "blanco"=hex:ff,ff,ff';
  put ' "negro"=hex:00,00,00';
  put ' "anaranjado"=hex:ff,a5,00';
run;

proc registry import="mycolornames";
run;
```

After you add these colors to the registry, you can use these color names anywhere that you use the color names supplied by SAS. For example, you could use the color name in the `GOPTIONS` statement as shown in the following code:

```
goptions cback=anaranjado;
proc gtestit;
run;
```

Using the Registry Editor

When to Use the Registry Editor

The best way to view the contents of the registry is using the Registry Editor. The Registry Editor is a graphical alternative to `PROC REGISTRY`, an experienced SAS user might use the Registry Editor to do the following:

- View the contents of the registry, which shows keys and values stored in keys.
- Add, modify, and delete keys and values stored in the registry.
- Import registry files into the registry, starting at any key.

- Export the contents of the registry to a file, starting at any key.
- Uninstall a registry file.
- Compare a registry file to the SAS registry.

Many of the windows in the SAS windowing environment update the registry for you when you make changes to such items as your printer setting or your color preferences. Because these windows update the registry using the correct syntax and semantics, it is often best to use these alternatives when making adjustments to SAS.

Starting the Registry Editor

To run the Registry Editor, issue the **regedit** command on a SAS command line. You can also open the registry window by selecting **Solutions ► Accessories ► Registry Editor**.


Finding Specific Data in the Registry

In the Registry Editor window, double-click a folder icon that contains a registry key. This displays the contents of that key.

Another way to find things is to use the Find utility.

- 1 From the Registry Editor, select **Edit ► Find**.
- 2 Type all or part of the text string that you want to find, and click **Options** to specify whether you want to find a **key name**, a **value name**, or **data**.
- 3 Click **Find**.

Display 15.2 The Registry Editor Find Utility



Changing a Value in the SAS Registry

CAUTION:


Before modifying registry values, always back up the registry.sas7bitm file from SASUSER.

△

- 1 In the left pane of the Registry Editor window, click the key that you want to change. The values contained in the key appear in the right pane.

- 2 Double-click the value. The Registry Editor displays several types of windows, depending on the type of value you are changing.


Display 15.3 Example Window for Changing a Value in the SAS Registry



Adding a New Value or Key to the SAS Registry

- 1 In the SAS Registry Editor, right-click the key that you want to add the value to.
- 2 From the pop-up menu, select the **New** menu item with the type that you want to create.
- 3 Enter the values for the new key or value in the window that is displayed.

Display 15.4 Registry Editor with Pop-up Menu for Adding New Keys and Values



Deleting an Item from the SAS Registry

From the SAS Registry Editor:

- 1 Right-click the item that you want to delete.
- 2 Select **Delete** from the pop-up menu and confirm the deletion.

Renaming an Item in the SAS Registry

From the SAS Registry Editor:


- 1 Right-click the item you want to rename.
- 2 Select **Rename** from the context menu and enter the new name.

Displaying the SASUSER and SASHELP Registry Items Separately

After you open the Registry Editor, you can change your view from the default, which shows the registry contents without regard to the storage location. The other registry view displays both SASUSER and SASHELP items in separate trees in the Registry Editor's left pane.

- 1 Select **TOOLS ► Options ► Registry Editor** This opens the **Select Registry View** group box.
- 2 Select **View All** to display the SASUSER and SASHELP items separately in the Registry Editor's left pane.
 - The SASHELP portion of the registry will be listed under the HKEY_SYSTEM_ROOT folder in the left pane.
 - The SASUSER portion of the registry will be listed under the HKEY_USER_ROOT folder in the left pane.

Display 15.5 The Registry Editor in View Overlay Mode



Importing a Registry File

You usually import a registry file, or SASXREG file when you are restoring a backup registry file. A registry file can contain a complete registry or just part of a registry.

To import a registry file using the SAS Registry Editor:

- 1 Select **File ► Import Registry File**.
- 2 In the Open window, select the SASXREG file to import.

Note: In order to first create the backup registry file, you can use the REGISTRY Procedure or the **Export Registry File** menu choice in the Registry Editor. △

Exporting a Registry File

You usually export a registry file or SASXREG file, when you are preparing a backup registry file. You can export a complete registry or just part of a registry.

To export a registry file using the SAS Registry Editor:

- 1 In the left hand pane of the Registry Editor, select the key you want to export to a SASXREG file. To export the entire registry, select the top key.
- 2 Select **File** \blacktriangleright **Export Registry File**.
- 3 In the Save As window, give the export file a name.
- 4 Click Save.

Configuring Your Registry

Configuring Universal Printing

Universal Printers should be configured by using either the PRTDEF procedure or the Print Setup window. The REGISTRY procedure can be used to back up a printer definition and to restore a printer definition from a SASXREG file. Any other direct modification of the registry values should be done only under the guidance of SAS Technical Support.

Configuring SAS Explorer

While it is best to use the Explorer Options window to configure your Explorer settings, you can use the Registry Editor to view the current Explorer settings in the SAS registry. The Explorer Options Window is available from the **TOOLS** \blacktriangleright **Options** \blacktriangleright **Explorer** drop-down menu from within the Explorer. All the Explorer configuration data is stored in the registry under CORE\Explorer. The following table outlines the location of the most commonly used Explorer configuration data.

Registry Key	What portion of the Explorer it configures
CORE\EXPLORER\CONFIGURATION	the portions of the Explorer get initialized at startup.
CORE\EXPLORER\MENUS	the context menus that get displayed in the Explorer.
CORE\EXPLORER\KEYEVENTS	the valid key events for the 3270 interface. This key is used only on the mainframe platforms.
CORE\EXPLORER\ICONS	Which icons to display in the Explorer. If the icon value is -1, this causes the icon to be hidden in the Explorer.
CORE\EXPLORER\NEW	This subkey controls what types of objects are available from the File \blacktriangleright New menu in the Explorer.

Configuring Libraries and File Shortcuts with the SAS Registry

When you use the New Library window or the File Shortcut Assignment window to create a library reference (libref) or a file reference (fileref), these references are stored for future use when you click the **Enable at Startup** check box in either of these two windows.

Library references (librefs) and file references (filerrefs) are saved when you check “Enable at startup” and they are stored in the SAS registry, where it is possible to modify or delete them, as follows:

Deleting an “Enable at Startup” library reference

You can use the Registry Editor to delete an “Enable at Startup” library reference by deleting the corresponding key under CORE\OPTIONS\LIBNAMES*“your libref”*. However, it is best to delete your library reference by using the SAS Explorer, which removes this key from the registry when you delete the library reference.

Deleting an “Enable at Startup” file shortcut

You can use the Registry Editor to delete an “Enable at Startup” file shortcut by deleting the corresponding key under CORE\OPTIONS\FILEREFES*“your fileref”*. However, it is best to delete your library reference by using the SAS Explorer, which removes this key automatically when you delete the file shortcut.

Creating an “Enable at Startup” File Shortcut as a site default

A site administrator might want to create a file shortcut which is available to all users at a site. To do this, you first create a version of the file shortcut definition in the SASUSER registry. Then you modify it so that it can be used in the SASHELP registry.

Note: You need special permission to write to the SASHELP part of the SAS registry. △

- 1 Type the **DMFILEASSIGN** command. This opens the File Shortcut Assignment window.
- 2 Create the file shortcut that you want to use.
- 3 Check **Enable at Startup**.
- 4 Click **OK**.
- 5 Type the command **REGEDIT** after verifying that the file shortcut was created successfully.
- 6 Find and select the key CORE\OPTIONS\FILEREFES*“your fileref”*.
- 7 Select **File ► Export Registry File** and export the file.
- 8 Edit the exported file and replace all instances of HKEY_USER_ROOT with HKEY_SYSTEM_ROOT.
- 9 To apply your changes to the site’s SASHELP, use PROC REGISTRY. The following code will import the file:

```
proc registry import="yourfile.sasxreg" usesashelp;
run;
```

Creating an “Enable at Startup” Library as a site default

A site administrator might want to create a library which is available to all users at a site. To do this, the SASUSER version of the library definition needs to be migrated to SASHELP.

Note: You need special permission to write to the SASHELP part of the SAS registry. △

- 1 Type the `dmlibassign` command. This opens the New Library window.
- 2 Create the library reference that you want to use.
- 3 Select **Enable at Startup**.
- 4 Click **OK**.
- 5 Issue the `regedit` command after verifying that the library was created successfully.
- 6 Find and select the registry key `CORE\OPTIONS\LIBNAMES\<your libref>`.
- 7 Select **File ► Export Registry File**. The Save As window appears.
- 8 Select a location to store your registry file.
- 9 Enter a filename for your registry file in the **File name** field.
- 10 Click **Save** to export the file.
- 11 Right-click the file and select **Edit in NOTEPAD** to edit the file.
- 12 Edit the exported file and replace all instances of “HKEY_USER_ROOT” with “HKEY_SYSTEM_ROOT”.
- 13 To apply your changes to the site’s SASHELP use PROC REGISTRY. The following code will import the file:

```
proc registry import="yourfile.sasxreg" usesashelp;
run;
```

Fixing Library Reference (Libref) Problems with the SAS Registry

Library references (librefs) are stored in the SAS Registry. You might encounter a situation where a libref that previously worked, fails. In some situations, editing the registry is the fastest way to fix the problem. The following section describes what is involved in repairing a missing or failed libref.

If any permanent libref that is stored in the SAS Registry fails at startup, then the following note appears in the SAS Log:

NOTE: One or more library startup assignments were not restored.

The following errors are common causes of library assignment problems:

- Required field values for libref assignment in the SAS Registry are missing.
- Required field values for libref assignment in the SAS Registry are invalid. For example, library names are limited to eight characters, and engine values must match actual engine names.
- Encrypted password data for a libref has changed in the SAS Registry.

Note: You can also use the New Library window to add librefs. You can open this window by typing DMLIBASSIGN in the toolbar, or selecting **File ► New** from the Explorer window. △

CAUTION:

You can correct many libref assignment errors in the SAS Registry Editor. If you are unfamiliar with librefs or the SAS Registry Editor, then ask for technical support. Errors can be made easily in the SAS Registry Editor, and can prevent your libraries from being assigned at startup. △

To correct a libref assignment error using the SAS Registry Editor:

- 1 Select **Solutions ► Accessories ► Registry Editor** or issue the `regedit` command to open the Registry Editor.

- 2 Select one of the following paths, depending on your operating environment, and then make modifications to keys and key values as needed:


```
CORE\OPTIONS\LIBNAMES
```

or

```
CORE\OPTIONS\LIBNAMES\CONCATENATED
```

Note: These corrections are possible only for permanent librefs; that is, those that are created at start-up by using the New Library or File Shortcut Assignment window. Δ

For example, if you determine that a key for a permanent, concatenated library has been renamed to something other than a positive whole number, then you can rename that key again so that it is in compliance. Select the key, and then select **Rename** from the pop-up menu to begin the process.



CHAPTER

16

Printing with SAS

<i>Universal Printing</i>	183
<i>What Is Universal Printing?</i>	183
<i>Turning Universal Printing On and Off</i>	183
<i>Universal Printing Output Formats and Printers</i>	184
<i>Types of Output Formats</i>	184
<i>Universal Printing and ODS</i>	184
<i>Managing Printing Tasks with the Universal Printing User Interface</i>	185
<i>Overview of the Universal Printing Menu</i>	185
<i>Setting Up Printers</i>	186
<i>Print Setup Window</i>	186
<i>Change the Default Printer</i>	187
<i>Remove a Printer from the Selection List</i>	187
<i>Define a New Printer</i>	187
<i>Set Printer Properties for Your Default Printer</i>	191
<i>How to Specify a Printer for Your Session</i>	194
<i>Printing with Universal Printing</i>	195
<i>Print a Test Page</i>	195
<i>Print the Contents of an Active SAS Window</i>	195
<i>Working with Previewers</i>	197
<i>Define a New Previewer</i>	197
<i>Seeding the Print Previewer Command Box</i>	200
<i>Previewing Print Jobs</i>	200
<i>Set Page Properties</i>	200
<i>Configuring Universal Printing with Programming Statements</i>	204
<i>Introduction</i>	204
<i>System Options That Control Universal Printing</i>	204
<i>Defining Printers for Batch Mode</i>	205
<i>Defining New Printers and Previewers with the PRTDEF Procedure</i>	206
<i>Introduction</i>	206
<i>Creating a Data Set That Defines Multiple Printers</i>	206
<i>Installing a Printer Definition for Multiple Users</i>	206
<i>Adding, Modifying, and Deleting Printer Definitions</i>	207
<i>Creating a Ghostview Previewer Definition for Previewing PostScript Output</i>	208
<i>Exporting and Backing Up Printer Definitions</i>	208
<i>Sample Values for the Device Type, Destination, and Host Options Fields</i>	209
<i>Forms Printing</i>	210
<i>Overview of Forms Printing</i>	210
<i>Creating or Editing a Form</i>	210
<i>Using TrueType Fonts with Universal Printers and SAS/GRAPH Devices in SAS 9.2</i>	211
<i>Rendering TrueType Fonts</i>	211
<i>ODS Styles and TrueType Fonts</i>	212

Portability of TrueType Fonts	212
International Character Support	212
TrueType Fonts That Are Supplied by SAS	213
Other TrueType Fonts	214
Registering the TrueType Fonts with the SAS Registry	214
Registering TrueType Fonts for Microsoft Windows or UNIX	215
Registering TrueType Fonts for z/OS	215
Locating the TrueType Font Files	215
Using the TrueType Fonts	215
Specifying the TrueType Fonts with the Print Dialog Box	215
Specifying TrueType Fonts with SAS Program Statements	216
Specifying a Font with the SYSPRINTFONT Option	216
Examples of Specifying Fonts and Printing International Characters	217
Specifying a Font with SAS/GRAPH	217
Specifying a Font with PROC PRINT	217
Specifying a Font with PROC PRINT and a User-Defined ODS Template	218
Printing International Characters	219
Creating Scalable Vector Graphics Using Universal Printing	223
Overview of Scalable Vector Graphics in SAS	223
Scalable Vector Graphics in SAS	223
SVG Terminology	223
Why Create SVG Documents?	224
ODS Destinations That Produce SVG Documents	224
Fundamentals of Creating SVG Documents Using Universal Printers	224
Web Server Content Type for SVG Documents	225
Browser Support for Viewing SVG Documents	226
Browsers That Support SVG Documents	226
Notes on Using Adobe SVG Viewer	226
Notes on Using Mozilla Firefox	226
The SVG Universal Printers and the Output That They Create	227
Setting the Environment to Create Stand-alone SVG Documents	227
Overview of Setting the Environment to Create Stand-alone SVG Documents	227
SAS System Options that Affect Stand-alone SVG Documents	228
Setting the SVG Universal Printer	228
Scaling an SVG Document to the Viewport	228
Setting the ViewBox	229
Interaction between SAS SVG System Options and the SVG Tag Attributes	229
Creating a Static viewBox	230
Preserving the Aspect Ratio	231
Adding a Title to an SVG Document	231
Creating Stand-alone SVG Documents Using the ODS PRINTER Destination	231
Creating an SVG Document	231
Multi-Page SVG Documents in a Single File	233
Creating Separate Files for Multi-Page SVG Documents	234
Creating Overlaid Transparent SVG Documents	235
Embedding and Linking Stand-alone SVG Documents in HTML Documents	237
Embedding SVG Documents in HTML Documents	237
Linking to an SVG Document	237
An Example HTML Document That Links and Embeds a Stand-alone SVG Document	238
Images in SVG Documents	240
Printing an SVG Document from a Browser	241
Creating Portable Network Graphics Using Universal Printing	241
Overview of Portable Network Graphics (PNG) in SAS	241
Portable Network Graphics in SAS	241

<i>The PNG Universal Printers</i>	241
<i>Producing PNG Images for ODS Destinations</i>	241
<i>Web Browsers and Viewers That Support PNG Files</i>	242
<i>Creating PNG Files Using ODS Printer</i>	242

Universal Printing

What Is Universal Printing?

Universal Printing is a printing system that provides both interactive and batch printing capabilities to SAS applications and procedures on all the operating environments that are supported by SAS.

When Universal Printing is ON, SAS routes all printing through Universal Printing services. All Universal Printing features are controlled by system options, thereby enabling you to control many print features even in batch mode. For more information on these system options, see “System Options That Control Universal Printing” on page 204.

Note: Before the introduction of Universal Printing, SAS supported a utility for print jobs known as Forms. Forms printing is still available if you select **File ► Print Setup** from the menu, then check the **Use Forms** check box. This turns off Universal Printing menus and functionality. For more information, see “Examples of Specifying Fonts and Printing International Characters” on page 217. △

Turning Universal Printing On and Off

Universal Printing is available on all operating environments that SAS supports. By default, Universal Printing is turned ON on all operating environments except Windows.

The UNIVERSALPRINT system option must be set for Universal Printing to be ON and used by SAS. This option can be set only in the SAS configuration file or at start up. You cannot turn Universal Printing menus on or off after start up.

Note: When you use the PRINTERPATH= system option to specify a printer, the print job is controlled by Universal Printing. △

Operating Environment Information: In the Windows operating environment, you must set the UNIVERSALPRINT system option in order to use Universal Printing. You must also set the UPRINTMENUSWITCH system option to use the Universal Printing windows. To use Universal Printing, include the following system options when starting SAS:

```
-uprint -uprintmenuswitch
```

△

Universal Printing Output Formats and Printers

Types of Output Formats

In addition to sending print jobs to a printer, you can also direct output to external files that are widely recognized by different types of printers and print viewing software programs. You can use Universal Printing to produce the following commonly recognized file types.

Table 16.1 Available Print Output Formats

Type	Full Name	Description
GIF	Graphics Interchange Format	A simple raster image format designed for the online transmission and interchange of graphic data. The format is widely used to display images on the World Wide Web.
PCL	Printer Control Language	Developed by Hewlett-Packard as a language that applications use to control a wide range of printer features across a number of printing devices. Universal Printing currently supports PCL5, PCL5e, and PCL5c levels of the language.
PDF	Portable Document Format	A file format developed by Adobe Systems for viewing and printing a formatted document. To view a file in PDF format, you need Adobe Reader, a free application distributed by Adobe Systems. <i>Note:</i> Adobe Acrobat is not required to produce PDF files with Universal Printing. \triangle
PNG	Portable Network Graphics	A raster image format that was designed as a replacement for the older simple GIF format and the more complex TIFF format. As with GIF, one of the major uses of PNG is to display images on the World Wide Web. PNG has four major advantages over GIF on the Web: gamma correction, two-dimensional interlacing, variable transparency (alpha channel), and more than 256 colors.
PS	PostScript	A page description language also developed by Adobe Systems. PostScript is primarily a language for printing documents on laser printers, but it can be adapted to produce images on other types of devices.
SVG	Scalable Vector Graphics	A language for describing two-dimensional graphics and graphical applications in XML.

You can access the list of available printers from either the Print dialog box or the SAS Registry path CORE\PRINTING\PRINTERS.

Universal Printing and ODS

The ODS PRINTER destination can use Universal Printing whether the UNIVERSALPRINT system option is on or off. The PRINTER destinations used by the ODS PRINTER statement are described in *SAS Output Delivery System: User's Guide*.

The Output Delivery System (ODS) uses Universal Printing for the following ODS statements.

Table 16.2 ODS Destinations That Use the Universal Printing Interface

ODS PRINTER Destination	Description
ODS PRINTER PRINTER= option	Uses the selected printer.
ODS PDF statement	Uses Universal Printing's PDF printer.
ODS PS statement	Uses Universal Printing's PostScript Level 1 printer.
ODS PCL statement	Uses Universal Printing's PCL5 printer.

Operating Environment Information: In the Windows operating environment, the ODS PRINTER destination uses the Windows system printers unless SAS is started with the UNIVERSALPRINT system option, or when you specify a printer with the PRINTERPATH= system option. If Universal Printing is enabled in Windows, SAS overrides the use of the Windows system printer and causes ODS to use Universal Printing. △

For more information on ODS, see *SAS Output Delivery System: User's Guide*.

Managing Printing Tasks with the Universal Printing User Interface

Overview of the Universal Printing Menu

Most SAS Universal Printing windows are accessible from the **File** menu.

The following display shows the File menu containing the Universal Printing choices of Page Setup, Print Setup, Print Preview, and Print.

Display 16.1 File Menu Displaying Universal Printing Options




Table 16.3 Menu Choices or Commands to Open Universal Printing Windows

Menu Choice	Equivalent Command
Page Setup	DMPAGESETUP
Print Setup	DMPRINTSETUP
Print Preview	DMPRTPREVIEW
Print	DMPRINT

Operating Environment Information: In the Windows operating environment, SAS uses the Windows print windows as the default. To access the Universal Printing user interface, the UPRINTMENUSWITCH system option must also be set along with the UNIVERSALPRINT system option to configure all menus in SAS to use the Universal Printing windows. To do this, include

```
-uprint -uprintmenuswitch
```

in the string that you use to invoke SAS in Windows. Δ

You can open all Universal Printing windows by entering commands at the command line or into the command box in the menu bar. The following table lists the commands that you use to do the most common tasks.

Table 16.4 Commands to Open Universal Printing Windows

Action	Command
Print the current window	DMPRINT
Change the default printer	DMPRINT, or DMPRINTSETUP
Create a new printer or previewer definition	DMPRTCREATE PRINTER, or DMPRTCREATE PREVIEWER
Modify, add, remove, or test printer definitions	DMPRINTSETUP
Show default printer properties sheet	DMPRINTPROPS
Show page properties sheet	DMPAGESETUP
Print preview the current window	DMPRTPREVIEW

Setting Up Printers

Print Setup Window

The **File ▶ Print Setup** menu selection opens the Print Setup window, where you can perform the following tasks:

- Change the default printer.
- Remove a printer from the selection list.
- Print a test page.
- Open the Printer Properties window.
- Launch the New Printer wizard.

Alternatively, you can issue the DMPRINTSETUP command.

Change the Default Printer

To change the default printer device for this SAS session and future SAS sessions, follow these steps:

- 1 Select **File ► Print Setup**. The Print Setup window appears.
- 2 Select the new default device from the list of printers in the **Printer** field.
- 3 Click **OK**.

Alternatively, you can issue the DMPRINTSETUP command.

Remove a Printer from the Selection List

To remove a printer from the selection list:

- 1 Select **File ► Print Setup**. The Print Setup window appears.
- 2 Select the printer you want to delete from the list of printers in the **Printer** field.
- 3 Click **Remove**.

Note: Only your system administrator can remove printers that the administrator has defined for your site. If you select a printer that was defined by your system administrator, the **Remove** button will be unavailable. △

Alternatively, you can issue the DMPRINTSETUP command.

Define a New Printer

While Universal Printing provides you with predefined printers, you can also add your own printers with the Define a New Printer wizard, which guides you step-by-step through the process of installing a printer.

To start the New Printer wizard and define a new printer, follow these steps:

- 1 Select **File ► Print Setup** and click **New**. The following window appears.

Display 16.2 Printer Definition Window to Enter Name and Description

Define a New Printer: Step 1 of 5

Enter a name and description for the printer. This will be associated with attributes specified in this dialog.

For example:

Name: myPrinter
Description: HP Laserjet in room R3283

Name:

Description:

Help < Previous Next > Cancel

Alternatively, you can either issue the `DMPRTCREATE PRINTER` command or issue the `DMPRINTSETUP` command and click **New**.

- 2 Enter the name and a description for the new printer (127 character maximum, no backslashes).
- 3 Click **Next** to proceed to Step 2 of the wizard. Select a printer model. If your exact printer model is not available, select a general model that is compatible with your printer. For example, for the HP LaserJet printer series, select PCL5.


Note: General models might provide fewer options than specific models. △

Display 16.3 Printer Definition Window to Select Printer Model



- 4 Click **Next** to proceed to Step 3 of the wizard. The following window appears:

Display 16.4 Printer Definition Window to Select Output Device



- 5 Select the **Device type** for your print output. The device type selections are host-dependent.

If you select **Catalog**, **Disk**, **Ftp**, **Socket**, or **Pipe** as the device type, then you must specify a destination.

If you select a device type of **Printer**, then a destination might not be required, because some operating environments use the Host options box to route output.


Note: Examples for your operating system of Device Type, Destination, and Host options are also provided in “Sample Values for the Device Type, Destination, and Host Options Fields” on page 209. △

- 6 Enter the **Destination** for your file. The destination is the target location for your device type. For example, if your device type is **disk**, then your destination is an operating environment-specific filename. With some system device types, the destination might be blank and you can specify the target location using the **Host options** box.
- 7 Select or enter any host-specific options for the device that you chose. This field might be optional for your operating environment. For a list of host options, refer to the FILENAME statement information for your operating environment.

Note: The **Destination** and **Host Options** lists can also be populated using the REGISTRY procedure. Click the **Help** button in step 3 to refer to the “Populating Destination and Host Option Lists” topic, which contains more details. △

- 8 Click **Next** to proceed to Step 4 of the wizard, in which you select from a list of installed print previewers. If no previewers are defined, proceed to the next step of the wizard.

Display 16.5 Printer Definition Window to Select Previewer




If the previewer selection box appears, select the previewer for this printer. If you do not need a previewer, choose **None** or leave the field blank.

Note: You can add a previewer to any printer through the DMPRTCREATE PREVIEWER command. For more information, see “Define a New Previewer” on page 197. △

Note: It is not required that printers and print previewers share a common language. △

- 9 Click **Next** to proceed to Step 5 of the wizard. The following window appears:

Display 16.6 Printer Definition Window to Complete Process

10 Click **Previous** to change any information. Click **Finish** when you have completed your printer definition.

You have now finished setting your default printer.

After you have returned to the Print Setup window, you can test your default printer by clicking **Print Test Page**.

Set Printer Properties for Your Default Printer

Printer properties that you can change include the following:

- the printer name and description
- the printer destination device and its properties
- the default font for the printer
- advanced features such as translation tables, printer resolution, and the print previewer associated with the printer


To change printer properties for your default printer, follow these steps:

1 Select **File ► Print Setup** and choose **Properties**. The Printer Properties window appears.


Alternatively, you can issue the DMPRTPROPS command.

2 From the Printer Properties window, select the tab that contains the information that you need to modify.

- In the **Name** tab, you can modify the printer name and the printer description.

Display 16.7 Printer Properties Window Displaying Name Tab


- The **Destination** tab enables you to designate the device type, destination, and host options for the printer. See “Sample Values for the Device Type, Destination, and Host Options Fields” on page 209 for examples.

Display 16.8 Printer Properties Window Displaying Destination Tab

- The **Font** tab controls the available font options. The selections available in the drop-down boxes are printer specific. The font size is in points.


Note: This window enables you to set attributes for the default fonts. Typically, procedure output is controlled by the fonts specified by the ODS style or by program statements that specify font attributes. Δ

Display 16.9 Printer Properties Window Displaying Font Tab



- The **Advanced** tab lists the Resolution, Protocol, Translate table, Buffer size, Previewer, and Preview command options for the printer. The information in the drop-down fields is printer specific.

Display 16.10 Printer Properties Window Displaying Advanced Tab



Resolution

specifies the resolution for the printed output in dots per inch (dpi).

Protocol

provides the mechanism for converting the output to a format that can be processed by a protocol converter that connects the EBCDIC host mainframe

to an ASCII device. Protocol is required in the z/OS operating environment, and if you must use one, select one of the protocol converters that are listed.

Translate table

manages the transfer of data between an EBCDIC host and an ASCII device. Normally, the driver selects the correct table for your locale; the translate table needs to be specified only when you require nonstandard translation.

Buffer size

controls the size of the output buffer or record length. If the buffer size is left blank, a default size will be used.

Previewer

specifies the Previewer definition to use when Print Preview is requested. The Previewer box contains the previewer application that you have defined. See “Define a New Previewer” on page 197.

Preview command

is the command that will be used to open an external printer language viewer. For example, if you want Ghostview as your previewer, type **ghostview %s**. When a Preview Command is entered into a Printer definition, the printer definition becomes a previewer definition. The Preview Command must a valid command. When the command is executed as part of the preview process the %s will be replaced with the name of a temporary file that contains the input for the preview command.

Note: The **Previewer** and **Preview Command** fields are mutually exclusive. When you enter a command path into the **Preview Command** field, the **Previewer** box will appear dimmed. Δ

How to Specify a Printer for Your Session

The **PRINTERPATH=** system option enables you to specify a Universal Printer to use for the current SAS session. This printer specification is not retained across SAS sessions. The **PRINTERPATH=** system option is primarily used in batch mode, when there is no windowing environment in which to set the default printer. This option accepts a character string as its value—for example:

```
options printerpath=myprinter;
options printerpath="Print PostScript to disk";
```

Note: If the printer name contains blanks, you must enclose it in quotation marks. Δ

You can get a list of valid character strings from two places:

- the list of printers in the **Printer** field of the Print Setup window.
- the list of defined printers in the registry under **SAS_REGISTRY\CORE\PRINTING\PRINTERS**.

You can also override the printer destination by specifying a fileref with the **PRINTERPATH=** system option:

```
options printerpath= (myprinter printout);
filename printout ...;
```

Printing with Universal Printing

Print a Test Page

To print a test page, follow these steps:

- 1 Select **File ► Print Setup** and choose **Print Test Page** to open the Print Setup window.
- 2 Select the printer for which you would like a test page from **Printer** listview.
- 3 Click **Print Test Page**.

Alternatively, you can issue the DMPRINTSETUP command.


Print the Contents of an Active SAS Window

To print the contents of a window in SAS, follow these steps:

- 1 Click inside the window to make it active.
- 2 Select **File ► Print**. A print window appears. Your print window might differ from the window that follows.

Alternatively, you can issue the DMPRINT command.

Display 16.11 Print Window



- 3 If the **Use Forms** check box is visible, clear it in order to use Universal Printing.
- 4 From the **Printer** group box, select the name of the printer definition.
- 5 Enter the number of copies that you want.
- 6 If you want to save your print job to a file,

- a Select **Print to File**.
- b Select **OK**. The **File Selection** window appears.
- c Select an existing file or enter a new filename.

Note: If you print to an existing file, the contents of the file are either overwritten or appended, depending on whether you choose **replace** or **append** from the print window that is open. Most viewers for GIF, PDF, SVG, and PNG will not view appended files. When **append** is selected, a merged PDF is not produced. \triangle

7 Set additional printing options.

The fields in the Page Options area provide choices according to the content in the SAS window that you are trying to print. By default, SAS prints the entire contents of the selected window.

Table 16.5 Page Options

Item to Print	Do This
Selected lines of text in a window <i>Note:</i> not available on z/OS	Select the text that you want to print, and then open the Print window. In the Page Options box, check the Print Selected Text box.
The page that is currently displayed in the window	Select Current page .
A range of pages or other individual pages	Select Range and enter the page numbers in the Pages field. Separate individual page numbers and page ranges with either a comma (,) or a blank. You can enter page ranges in any of these formats: <ul style="list-style-type: none"> <input type="checkbox"/> n–m prints all pages from n to m, inclusive. <input type="checkbox"/> –n prints all pages from page 1 to page n. <input type="checkbox"/> n– prints all pages from page n to the last page.
In color	Check the Print Color box.
Line numbers	Check the Print Line Numbers box.
Page numbers	Check the Print Page Numbers box.
A graph	Use the DMPRINT command, or select File \blacktriangleright Print. Verify that the Use SAS/GRAPH Drivers check box is deselected in order to use Universal Printing.

8 Click **OK** to print.

Working with Previewers

Define a New Previewer


Previewers enable you to preview a print job. SAS does not set a default previewer application. To use the Print Preview feature in SAS, you or your system administrator must first define a previewer for your system.

Operating Environment Information: Print Previewers are not supported on z/OS. △

Previewers can be defined using the New Previewer wizard. To use the New Previewer wizard to define a new print previewer, follow these steps:

- 1 Issue the DMPRTCREATE PREVIEWER command. The following window appears:

Display 16.12 Previewer Definition Window to Enter Name and Description



Define a New Previewer: Step 1 of 4

Enter a name and description for the previewer. This will be associated with attributes specified in this dialog.

For example:

Name: myPreviewer
Description: Adobe viewer for PCL Printers

Name:

Description:

Help < Previous Next > Cancel

- 2 Enter the name and a description for the new previewer (127 character maximum, no backslashes).
- 3 Click **next** to proceed to Step 2 of the wizard.

Display 16.13 Previewer Definition Window to Enter Previewer Language

- 4 Select the printer model that you want to associate with your previewer definition. The PostScript, PCL or PDF language generated for the model must be a language that your external viewer package supports. For best results, select the generic models such as PostScript Level 1 (Color) or PCL 5.
- 5 Click **next** to proceed to Step 3 of the wizard.

Display 16.14 Previewer Definition Window to Enter Command to Open Previewer Application

- 6 Enter the command or commands used to open the previewer application, followed by %s where you would normally put the input filename. For example, if the command for starting your previewer is “ghostview,” then you would type **ghostview %s** in the text field.

You can pre-populate or *seed* a list of commands used to invoke a print preview application. For more information, see “Seeding the Print Previewer Command Box” on page 200.

Note: The %s can be used as many times as needed in the commands for starting the viewer. However, the start command needs to be the fully qualified command if it is not in the machine’s command path. △

- 7 Click **Next** to proceed to Step 4 of the wizard.


Display 16.15 Previewer Definition Window to Complete Process



- 8 Click **Previous** to correct any information. Click **Finish** when you have finished defining your default previewer.

The newly defined previewer will display a previewer icon in the Print Setup window.

Display 16.16 Print Setup Window Displaying New Previewer



This previewer application can be tested with the **Print Test Page** button on the Print Setup window.

Seeding the Print Previewer Command Box

Print Preview is supported by print previewer applications such as Ghostview, gv, and Adobe Reader. The Preview command box that appears in the Previewer Definition wizard (Display 16.14 on page 198) and on the **Advanced** tab of the Printer Properties window (Display 16.10 on page 193) can be prepopulated or seeded with a list of commands. These commands are used to invoke print previewer applications that are available at your site. Users and administrators can manually update the registry, or define and import a registry file that contains a list of previewer commands. This is an example of a registry file.

```
[CORE\PRINTING\PREVIEW COMMANDS]
"1"="/usr/local/gv %s"
"2"="/usr/local/ghostview %s"
```

Previewing Print Jobs

You can use the print preview feature if a print viewer is installed for the designated printer. Print Preview is always available from the File menu in SAS. You can also issue the DMPRTPREVIEW command.

Set Page Properties

For your current SAS session, you can customize how your printed output appears in the Page Setup window. Depending on which printer you have currently set, some of the Page Setup options that are described in the following steps might be unavailable.

To customize your printed output, follow these steps:

- 1 Select **File ► Page Setup**. The Page Setup window appears.


Alternatively, you can issue the DMPAGESETUP command.

- 2 Select a tab to open windows that control various aspects of your printed output. Descriptions of the tabbed windows follow.

The Page Setup window consists of four tabs: **General**, **Orientation**, **Margins**, and **Paper**.

- The **General** tab enables you to change the options for Binding, Collate, Duplex, and Color Printing.

Display 16.17 Page Setup Window Displaying General Tab



Binding

specifies the binding edge (Long Edge or Short Edge) to use with duplexed output. This sets the Binding option.

Collate

specifies whether the printed output should be collated. This sets the Collate option.

Duplex


specifies whether the printed output should be single-sided or double-sided. This sets the Duplex option.

Color Printing

specifies whether output should be printed in color. This sets the Colorprinting option.


- The **Orientation** tab enables you to change the output's orientation on the page. The default is Portrait. This tab sets the ORIENTATION option.

Display 16.18 Page Setup Window Displaying Orientation Tab




- The **Margin** tab enables you to change the top, bottom, left and right margins for your pages. The value range depends upon the type of printer that you are using. The values that are specified on this tab set the TOPMARGIN, BOTTOMMARGIN, LEFTMARGIN, and RIGHTMARGIN options.

Display 16.19 Page Setup Window Displaying Margins Tab



- The **Paper** tab specifies the Size, Type, Source, and Destination of the paper used for the printed output.

Display 16.20 Page Setup Window Displaying Paper Tab



Size

specifies the size of paper to use by setting the PAPERSIZE option. Paper sizes include Letter, Legal, A4, and so on.

Type

specifies the type of paper to use. Examples of choices include Standard, Glossy, and Transparency. This sets the PAPERTYPE option.

Source

designates which input paper tray is to be used. This sets the PAPERSOURCE option.

Destination

specifies the bin or output paper tray that is to be used for the resulting output. This sets the PAPERDEST option.

Note: Page settings are stored in the SAS registry. Although your page settings should remain in effect from one SAS session to another, changing default printers could lose, change, or disable some of the settings. If you change printers during a SAS session, check the Page Setup window to ensure that all of your settings are valid for your new default printer. Δ

Configuring Universal Printing with Programming Statements

Introduction

Universal Printing windows give you control over most printing functions through a graphical user interface. (You can also write a program that uses various SAS language elements to configure your printing environment and to control print jobs.)

System Options That Control Universal Printing

Universal Printing can be configured in batch mode and interactive mode by setting system option values within an OPTIONS statement. The following system options control Universal Printing.

Table 16.6 System Options That Control Universal Printing

System Option	Description
BINDING=	Specifies the binding edge for the printer.
BOTTOMMARGIN=	Specifies the size of the margin at the bottom of the page for printing.
COLLATE	Specifies the collation of multiple copies for output for the printer.
COLORPRINTING	Specifies color printing, if it is supported.
COPIES =	Specifies the number of copies to make when printing.
DUPLEX	Specifies double-sided printing, if it is supported.
LEFTMARGIN=	Specifies the size of the margin on the left side of the page.

System Option	Description
ORIENTATION=	Specifies the paper orientation to use (either portrait or landscape).
PAPERDEST=	Specifies the bin or output paper tray to receive printed output.
PAPERSIZE=	Specifies the paper size to use when printing.
PAPERSOURCE=	Specifies the input paper tray to use for printing.
PAPERTYPE=	Specifies the type of paper to use for printing.
PRINTERPATH=	Specifies a printer for Universal Printing print jobs (see Note on page 205).
RIGHTMARGIN=	Specifies the size of the margin on the right side of the page.
SYSPRINTFONT=	Specifies the default font to use when printing.
TOPMARGIN=	Specifies the size of the margin at the top of the page.

Note: The PRINTERPATH= system option specifies which printer will be used.

- If the PRINTERPATH= system option is blank, then the default printer will be used.
- If the PRINTERPATH= system option is not blank, then Universal Printing will be used.

In the Windows environment, the default printer is the current Windows system printer or the printer specified by the SYSPRINT system option; therefore, Universal Printing is not used. In all other operating environments, Universal Printing is used and the default printer is PostScript Level 1. To change this printer, use either the SYSPRINT system option or the DMPRINTSETUP command to open the Print Setup window. Δ

Defining Printers for Batch Mode

Printer definitions can be created for an individual, or for all SAS users at a site, by using the PRTDEF procedure. The PRTDEF procedure can be used to do many of the same printer management activities that you can also do with the Universal Printing windows. The PRTDEF procedure is especially useful if you use SAS in batch mode, where the Universal Printing windows are unavailable.

Only system administrators or others who have write permission to the SASHELP library can use the PRTDEF procedure to create printer definitions for all SAS users at a site.

To define one or more printers with the PRTDEF procedure, you first create a SAS data set that contains variables corresponding to printer attributes. The PRTDEF procedure reads the data set and converts the variable attributes into one or more printer definitions in the SAS registry.

See *Base SAS Procedures Guide* for more information about the PRTDEF procedure.

Defining New Printers and Previewers with the PRTDEF Procedure

Introduction

These examples show you how to use the PRTDEF procedure to define new printers and to manage your installed printers and previewers.

After a program statement containing the PRTDEF procedure runs successfully, the printers or previewers that have been defined will be displayed in the Print Setup window. A complete set of all available printers and previewers will be displayed in the Printer name list. Printer definitions can also be seen in the Registry Editor window under CORE\PRINTING\PRINTERS.

Creating a Data Set That Defines Multiple Printers

When you create a data set to use with the PRTDEF procedure to define a printer, you must specify the name, model, device and, destination variables.

See the PRTDEF procedure in *Base SAS Procedures Guide* for the names of the optional variables that you can also use.

The following code creates a data set to use with the PRTDEF procedure:

```
data printers;
    input name $& model $& device $& dest $&;
    datalines;
Myprinter          PostScript Level 1   PRINTER   printer1
Laserjet           PCL 5 Printer        PIPE      lp -dprinter5
Color LaserJet     PostScript Level 2   PIPE      lp -dprinter2
;
run;
```

After you create the data set containing the variables, you run a SAS program that contains the PRTDEF procedure. The PRTDEF procedure defines the printers that are named in the data set by creating the appropriate entries in the SAS registry.

```
proc prtdef data=printers usesashelp;
run;
```

Note: The USESASHELP option specifies that the printer definitions are to be placed in the SASHELP library, where they are available to all users. If the USESASHELP option is not specified, then the printer definitions are placed in the current SASUSER library, where they are available to the local user only, and the printers that are defined will be available only in the local SASUSER directory. However, to use the USESASHELP option, you must have permission to write to the SASHELP catalog. Δ

Installing a Printer Definition for Multiple Users

This example creates a Tektronix Phaser 780 printer definition with a Ghostview print previewer in the SASHELP registry. The bottom margin is set to two centimeters, the font size to 14 point, and the paper size to A4.

```
data tek780;
    name = "Tek780";
    desc = "Test Lab Phaser 780P";
    model = "Tek Phaser 780 Plus";
    device = "PRINTER";
```

```

dest = "testlab3";
preview = "Ghostview";
units = "cm";
bottom = 2;
fontsize = 14;
papersiz = "ISO A4";
run;

proc prtdef data=tek780 usesashelp;
run;

```

Note: You must specify a preview command either in the Preview Definition Wizard (Display 16.10 on page 193) or on the **Advanced** tab of the Printer Properties window (Display 16.14 on page 198) or by using the PRDEF procedure. Δ

Adding, Modifying, and Deleting Printer Definitions

This example uses the PRINTERS data set to add, modify, and delete printer definitions. See the PRTDEF procedure in *Base SAS Procedures Guide* for more variables that you can use to define a printer. The following list describes the variables used in the example:

- The MODEL variable specifies the printer prototype to use when defining this printer.
- The DEVICE variable specifies the type of I/O device to use when sending output to the printer.
- The DEST variable specifies the output destination for the printer.
- The OPCODE variable specifies what action (Add, Delete, or Modify) to perform on the printer definition.
- The first Add operation creates a new printer definition for Color PostScript in the registry and the second Add operation creates a new printer definition for ColorPS in the registry.
- The Mod operation modifies the existing printer definition for LaserJet 5 in the registry.
- The Del operation deletes the printer definitions for printers named "Gray PostScript" and "test" from the registry.

The following example creates a printer definition in the SASHELP library. Because the definition is in SASHELP, the definition becomes available to all users. Special system administration privileges are required to write to the SASHELP library. An individual user can create a personal printer definition by specifying the SASUSER library instead.

```

data printers;
length name $ 80
       model $ 80
       device $ 8
       dest $ 80
       opcode $ 3;

input opcode $ & name $ & model $ & device $ & dest $ &;
datalines;
add Color PostScript   PostScript Level 2 (Color)      DISK   sasprt.ps
mod LaserJet 5         PCL 5                                DISK   sasprt.pcl
del Gray PostScript    PostScript Level 2 (Gray Scale)  DISK   sasprt.ps
del test               PostScript Level 2 (Color)      DISK   sasprt.ps

```

```

add ColorPS          PostScript Level 2 (Color)      DISK    sasprt.ps
;

proc prtdef data=printers list usesashelp;
run;

```

Note: If the end user modifies and saves new attributes for an administrator-defined printer in the SASHELP library, the printer will become a user-defined printer in the SASUSER library. Values that are specified by the user will override the values that were set by the administrator. If the user-defined printer definition is deleted, the administrator-defined printer will reappear. Δ

Creating a Ghostview Previewer Definition for Previewing PostScript Output

This example creates the GSVIEW data set. The variables in the GSVIEW data set have values that the PRTDEF procedure uses to produce the print previewer definition in the SAS registry.

- The NAME variable specifies the printer name that will be associated with the rest of the attributes in the printer definition data record.
- The DESC variable specifies the description of the printer.
- The MODEL variable specifies the printer prototype to use when defining this printer.
- The VIEWER variable specifies the host system commands for print preview.

Note: The **ghostview %s** command needs to be the fully qualified command if it is not in the machine's command path. Δ

Note: You must specify a preview command either in the Preview Definition Wizard (Display 16.10 on page 193) or on the **Advanced** tab of the Printer Properties window (Display 16.14 on page 198). Δ

- The DEVICE variable should always be DUMMY.
- DEST should be blank to specify that output is not returned.

```

data gsview;
  name = "Ghostview";
  desc = "Print Preview with Ghostview";
  model= "PostScript Level 2 (Color)";
  viewer = 'ghostview %s';
  device = "dummy";
  dest = " ";
run;
proc prtdef data=gsview list replace;
run;

```

Exporting and Backing Up Printer Definitions

The PRTEXP procedure enables you to back up your printer definitions as a SAS data set that can be restored with the PRTDEF procedure.

The PRTEXP procedure has the following syntax.

```

PROC PRTEXP [USESASHELP] [OUT=dataset]
[SELECT | EXCLUDE] printer_1 printer_2 ... printer_n;

```

The following example shows how to back up four printer definitions (named PDF, postscript, PCL5, and PCL5c) using the PRTDEF procedure.

```
proc prtexp out=printers;
select PDF postscript PCL5 PCL5c;
run;
```

For more information, see the PRTEXP procedure in *Base SAS Procedures Guide*.

Sample Values for the Device Type, Destination, and Host Options Fields

The following list provides examples of the printer values for device type, destination and host options. Because these values can be dependent on each other, and the values can vary by operating environment, several examples are shown. You might want to refer to this list when you are installing or modifying a printer or when you change the destination of your output.

- Device Type: Printer
 - z/OS
 - Device type: Printer
 - Destination: (leave blank)
 - Host options: sysout=*class-value* dest=*printer-name*
 - UNIX and Windows
 - Device type: Printer
 - Destination: *printer name*
 - Host options: (leave blank)
 - OpenVMS
 - Device type: Printer
 - Destination: *printer name*
 - Host options: passall=yes queue=*printer-name*
- Device Type: Pipe

Note: A sample command to send output to an lp-defined printer queue on a UNIX host is `lp -ddest Δ`

 - UNIX
 - Device Type: Pipe
 - Destination: *command*
 - Host options: (leave blank)
- Device Type: FTP

Note: An example of a nodename is `pepper.unx Δ`

 - z/OS
 - Device type: FTP
 - Destination: ftp.out
 - Host options: host=*'nodename'* recfm=vb prompt
 - Device type: Printer
 - Destination: *printer name*
 - Host options: (leave blank)
 - Windows
 - Device type: FTP

- Destination: ftp.out
 - Host options: host='nodename' prompt
- UNIX and OpenVMS
 - Device type: FTP
 - Destination: host='nodename' prompt
 - Host options: (leave blank)
- Device Type: Socket

Note: An example of an lp destination queue is *lp286nc0.prt:9100* Δ

 - UNIX and OpenVMS
 - Device type: Socket
 - Destination: *destination-queue*
 - Host options: (leave blank)

Forms Printing

Overview of Forms Printing

Before Universal Printing was introduced, SAS provided a utility for print jobs called a form. A *form* was a standard template that let you control such things as line size and margin information for pages in a print job. While Universal Printing is easier to use and has more features than the simple controls offered in forms printing, SAS still supports forms.

Printing with forms is still available through the Print window. You can switch to forms print mode by selecting **File** \blacktriangleright **Print** and selecting **Use Forms**.

Note: Forms printing is not available in batch mode. Δ

Creating or Editing a Form

If your organization has legacy reports that need to be printed using forms, you might have to use the FORM window to create or edit a form. SAS still supports the ability to create or edit forms, though Universal Printing provides more features, and is the recommended method of printing.

You can create or edit a form by entering the FSFORM command:

FSFORM<catalog-name.>form-name

If you do not specify a catalog-name, SAS uses the SASUSER.PROFILE catalog. If the form name that you specify does not exist, SAS creates a new form.

If you are creating a new form, SAS displays the Printer Selection frame. If you are editing an existing form, SAS displays the Text Body and Margin Information frame.

To move between the FORMS frames, you can do the following:

- Use the NEXTSCR command to scroll to the next frame and the PREVSCR command to scroll to the previous frame.
- Enter an equal sign (=) and the number of the frame that you want to go to. For example, =1 displays the Text Body and Margin Information frame, and =2 displays the Carriage Control Information frame.

- Select the name of the frame from the **Tools** menu.
- Select **Next Screen** or **Previous Screen** from the **Tools** menu.

You can move between fields on a frame with the TAB key.

After you have finished defining or editing your form, issue the END command to save your changes and exit the FORM window.

Note: Turning on Forms by checking the **Use Forms** check box in the print window turns Universal Printing off for printing non-graphic windows. Δ

Operating Environment Information: For more information on printing with Forms, see the documentation for your operating environment. Δ

Using TrueType Fonts with Universal Printers and SAS/GRAPH Devices in SAS 9.2

Rendering TrueType Fonts

In SAS 9.2, TrueType fonts are rendered using the FreeType library or by exercising the font-rendering capabilities of the host.

Rendering TrueType fonts with the FreeType library is the preferred method because fonts can be rendered for output in all of the operating environments that SAS supports. All Universal Printers and many SAS/GRAPH devices use the FreeType library to render TrueType fonts.

The following SAS output methods and devices use the FreeType library:

- SAS/GRAPH devices
 - UGIF, UPCL5, UPCL5C, UPCL5E, UPDF, UPNG, UPNGT, UPSL, UPSLC
 - GIF160, GIF260, GIF373, GIF570
 - Note:* Text is rendered at very small sizes. Δ
 - GIF, GIFANIM, GIF733, HTML, WEBFRAME
 - TIFFBII, TIFFBMM, TIFFB, TIFFP, TIFFG3, TIFFB300, TIFFP300
 - SASBMP
 - JPEG
 - PCL5, PCL5C, PCL5E
 - PNG, PNGT, PNG300
 - PSL, PSCOLOR, PSLEPSF, PSLEPSFC
 - SVG, SVGT, SVGVIEW, and SVGZ
 - SASEMF, SASWMF
 - SASPRTC, SASPRTG, SASPRTM printer graphics device drivers
- ODS printing and Universal Printing
 - GIF
 - PCL5, PCL5C, PCL5E
 - PDF, PDFA
 - PNG, PNGT

- PostScript
 - SVG, SVGT, SVGZ, SVGView, and SGVnotip
- ODS RTF
- ODS HTML

The following SAS/GRAPH devices use host font rendering only:

- ActiveX, ACTXIMG, JAVA, JAVAIMG

Note: These are client devices. \triangle

- BMP, BMP20
- DIB
- EMF, WMF
- ZGIF, ZGIF733, ZGIFANIM, ZHTML, ZWEBFRAM, ZJPEG, ZPNG, ZSASBMP
- ZTIFFB, ZTIFFBII, ZTIFFBMM, ZTIFFG3, ZTIFFG4, ZTIFFP

The following SAS/GRAPH devices also use host font rendering:

- GIF160, GIF260, GIF373, GIF570

Note: Text is rendered at very small sizes. \triangle

- GIF, GIFANIM, GIF733, HTML, WEBFRAME
- TIFFB, TIFFP, TIFFG3, TIFFBII, TIFFBMM, TIFFB300, TIFFP300
- JPEG

Note: You can set **OPTIONS FONTRENDERING= FREETYPE_POINTS | HOST_PIXELS** to change the rendering method for devices that support both FreeType library and host rendering. \triangle

Note: With devices that use host rendering in a UNIX operating environment, the TrueType fonts must be installed for the display in order for the fonts to work. For more information, see the *Configuration Guide for SAS 9.2 Foundation for UNIX Environments*. \triangle

ODS Styles and TrueType Fonts

By default, many SAS/GRAPH device drivers and all Universal Printers will generate output by using ODS styles, and these ODS styles use TrueType fonts. If no style is specified, the default style is the style for the listing output unless the ODS LISTING destination is closed. If you want the appearance of graphs to be compatible with graphs that were generated before SAS 9.2, you can set the GSTYLE system option to NOGSTYLE.

Portability of TrueType Fonts

TrueType fonts are portable across operating environments and are always available in Microsoft Windows environments. A few TrueType fonts are included with some versions of UNIX X Windows.

International Character Support

TrueType fonts support a wide range of international characters. For more information about SAS code that uses TrueType fonts, see “Examples of Specifying Fonts and Printing International Characters” on page 217.

TrueType Fonts That Are Supplied by SAS

When you install SAS, in addition to SAS Monospace and SAS Monospace Bold, a number of TrueType fonts are available based on choices that were made during the installation. TrueType fonts that are supplied by SAS can be categorized into four groups: Windows Glyph List (WGL) Pan-European character set fonts that are compatible with Microsoft; graphic symbol; multilingual; and monolingual Asian, as shown in the following tables.

Table 16.7 Windows Glyph List (WGL) Pan-European Character Set Fonts That Are Compatible with Microsoft

Font Name	Font Description	Compatibility with Microsoft Font
Albany AMT	sans serif	Arial
Thorndale AMT	serif	Times New Roman
Cumberland AMT	serif fixed ¹	Courier New

1 Fixed refers to uniform spacing.

Table 16.8 Graphic Symbol TrueType Fonts

Font Name	Font Description	Compatibility with Microsoft Font	Compatibility with Adobe Type1 Font
Symbol MT	192 symbols	Symbol	Symbol
Monotype Sorts	205 Wingding characters	Wingdings	ITC Zapf Dingbats
Arial Symbol	42 symbols ¹	Not applicable	Not applicable
Arial Symbol Bold	42 symbols ¹	Not applicable	Not applicable
Arial Symbol Bold Italic	42 symbols ¹	Not applicable	Not applicable
Arial Symbol Italic	42 symbols ¹	Not applicable	Not applicable
Times New Roman Symbol	42 symbols ¹	Not applicable	Not applicable
Times New Roman Symbol Bold	42 symbols ¹	Not applicable	Not applicable
Times New Roman Symbol Bold Italic	42 symbols ¹	Not applicable	Not applicable
Times New Roman Symbol Italic	42 symbols ¹	Not applicable	Not applicable

1 These fonts have special glyphs for the Latin characters 0, <, =, C, D, L, M, N, P, R, S, U, V, W, X, Z, and a-z. All other characters are undefined and might be rendered as a rectangle. For example, with ODS destinations such as HTML, the rectangle is replaced with the matching Latin1 character when it is displayed in Internet Explorer.

Table 16.9 Multilingual TrueType Fonts

Language Supported	Font Name	Font Description
Japanese	Monotype Sans WT J ¹	sans serif
	Thorndale Duospace WT J ¹	serif
Korean	Monotype Sans WT K ¹	sans serif
	Thorndale Duospace WT K ¹	serif
Simplified Chinese	Monotype Sans WT SC ¹	sans serif
	Thorndale Duospace WT SC ¹	serif
Traditional Chinese	Monotype Sans WT TC ¹	sans serif
	Thorndale Duospace WT TC ¹	serif

1 Both Thorndale Duospace WT and Monostype Sans WT support HKSCS (Hong Kong Supplement Character Set) and substantially support the GB18030 standard.

Table 16.10 Monolingual Asian TrueType Fonts

Language Supported	Font Name	Character Set
Japanese	MS Gothic, MS UI Gothic, MS PGothic	Shift JIS
	MS Mincho, MS PMincho	Shift JIS
Korean	Gulim, GulimChe, Dotum, DotumChe	KSC5601
	Batang, BatangChe, Gungsuh, GungsuhChe	KSC5601
Simplified Chinese	SimSun, NSimSun	GB2312
	Sim Hei	GB2312
Traditional Chinese ¹	HeiT	Big5
	MingLiU, MingLiU_HKSCS, PMingLiU	Big5

1 HeiT, MingLiU, MingLiU_HKSCS, and PMingLiU support the HKSCS2004 (Hong Kong Supplemental Character Set) characters.

The fonts that are supplied by SAS and any currently installed system fonts are automatically registered in the SAS registry when you install SAS. Fonts that are installed after you install SAS must be registered manually in the SAS registry. To register other TrueType Fonts, see “Other TrueType Fonts” on page 214.

Other TrueType Fonts

Registering the TrueType Fonts with the SAS Registry

To use TrueType fonts that are not registered when SAS is installed, use the FONTREG procedure to register these fonts in the SAS Registry. For more information, see the FONTREG procedure in the *Base SAS Procedures Guide*.

Note: The fonts that are supplied by SAS and the fonts that are installed by Microsoft are automatically registered in the SAS registry when you install SAS. Fonts that are installed after you install SAS must be registered manually in the SAS registry. △

Registering TrueType Fonts for Microsoft Windows or UNIX

Submit the following SAS program. The FONTPATH statement specifies the directory that contains the TrueType fonts and **pathname** is the directory path of the TrueType fonts.

```
proc fontreg;
    fontpath 'pathname';
run;
```

Note: In Microsoft Windows environments, TrueType fonts are usually located in either the **C:\WINNT\Fonts** or **C:\Windows\Fonts** directory. For all other operating environments, contact your system administrator for the location of the TrueType font files. Δ

Registering TrueType Fonts for z/OS

Use the FONTFILE statement instead of the FONTPATH statement to specify the fixed block sequential file that contains a TrueType font.

```
proc fontreg;
    fontfile 'filename';
run;
```

Note: When you add TrueType fonts to a z/OS system, the TrueType font file must be allocated as a sequential data set with a fixed block record format and a record length of 1. Δ

Locating the TrueType Font Files

After you update the SAS Registry, you can use the SAS Registry Editor to view a list of any registered fonts by typing REGEDIT on the SAS command line or by submitting a REGISTRY procedure statement that specifies the location of the files.

```
proc registry list startat='core\printing\freetype\fonts';
run;
```

When SAS is invoked, the FONTSLOC= system option points to the location of the fonts that are supplied by SAS. Use the following OPTIONS procedure statement to display the location of the SAS fonts.

```
proc options option=fontsloc;
run;
```

Using the TrueType Fonts

Specifying the TrueType Fonts with the Print Dialog Box

After you have updated the SAS Registry, the TrueType fonts are available for use within SAS. To access the TrueType fonts interactively when Universal Printing is enabled, follow these steps:

- 1 Select **File** ► **Print**
- 2 Select an appropriate printer, such as PDF or PCL5.
- 3 Click the **Properties** button.
- 4 Click the **Font** tab.

This window contains drop-down boxes for Font, Style, Weight, Size (in points), and Character Set.

- 5 Click the arrow on the right side of the Font box and scroll through the list of available fonts.

TrueType fonts are indicated by a three-character tag enclosed in angle brackets. For example, the TrueType font Albany AMT will be listed as <ttf> Albany AMT. The <ttf> tag makes the distinction between the different types of fonts with the same name, such as <ttf> Symbol and a Symbol font that resides on a physical printer. Fonts that do not have a <ttf> tag reside in the printer's memory. If you designate PDF as the printer, the font is in the Adobe Reader rather than on a physical printer.

- 6 Select the font that you want to use.
- 7 Click **OK** to return to the Print dialog box.
- 8 Click **OK** again to create your output.

Specifying TrueType Fonts with SAS Program Statements

To choose a printer that supports TrueType fonts when you're using SAS in a batch or line-mode environment, submit the following OPTIONS statement.

```
options printerpath=printername;
```

printername can be a printer name such as PDF, PNG, SVG, PCL5, PCL5e, or GIF.

Then, submit your program code with the font that you want to use. For example, if you want to use the TrueType font Albany AMT on a TITLE statement, include the following line of code in your SAS program.

```
Title1 f="Albany AMT" "Text in Albany AMT";
```

When you use the SAS/GRAPH graphics device drivers SASPRTC, SASPRTG, and SASPRTM, you can specify attributes, such as style or weight, by using the forward slash (/) as a delimiter.

```
Title1 f="Albany AMT/Italic/Bold" "Text in Bold Italic Albany AMT";
```

For ODS templates, the attributes are specified after the text size parameter. See "Specifying a Font with PROC PRINT and a User-Defined ODS Template" on page 218 for a complete example.

Note: You should use the <ttf> tag only when it is necessary, for example, to distinguish between a TrueType font and another type of font with the same name. Δ

Specifying a Font with the SYSPRINTFONT Option

The SYSPRINTFONT= system option sets the default font that you want to use for printing to windows such as the Program Editor, the Log, and Output windows. For example, you could use the SYSPRINTFONT= system option to print your output in the Albany AMT font by submitting the following OPTIONS statement.

```
options sysprintfont=("Albany AMT");
```

You can also use the SYSPRINTFONT= system option to specify the weight and size of a font. For example, the following code specifies an Arial font that uses bold face, is italicized, and has a size of 14 points.

```
options sysprintfont=("Arial" bold italic 14);
```

You can override the default font by explicit font specifications or ODS styles.

For more information, see the SYSPRINTFONT= system option in the *SAS Language Reference: Dictionary*.


Examples of Specifying Fonts and Printing International Characters

Specifying a Font with SAS/GRAPH

The following example creates an output file, **sasprt.pdf**, with a title in the Albany AMT font that uses bold face and is italicized.

```
goptions reset=all device=sasprtc;
title1 color=black f="Albany AMT/Italic/Bold" "Sample Title in Bold Italic Albany AMT";
options printerpath=pdf;
proc gplot data=sashelp.class;
  plot height*weight;
run;
quit;
```

Display 16.21 Gplot Output



Specifying a Font with PROC PRINT

The following example produces an output file, **print1.pdf**, with the titles in the Albany AMT, Thorndale AMT, and Cumberland AMT fonts.

```

filename new 'print1.pdf';
options printerpath=(PDF new) obs=5;
ods printer;
proc print data=sashelp.class;
  title1 f='Albany AMT' h=2 'TrueType Albany AMT';
  title2 f='Thorndale AMT' h=3 'Thorndale AMT';
  title3 f='Cumberland AMT' 'Cumberland AMT ';
run;
ods printer close;

```

Display 16.22 PDF Output Using PROC PRINT

TrueType Albany AMT
 Thorndale AMT
 Cumberland AMT

Obs	Name	Sex	Age	Height	Weight
1	Alfred	M	14	69.0	112.5
2	Alice	F	13	56.5	84.0
3	Barbara	F	13	65.3	98.0
4	Carol	F	14	62.8	102.5
5	Henry	M	14	63.5	102.5

Specifying a Font with PROC PRINT and a User-Defined ODS Template

The following example creates a template of font styles and then produces a PDF file.

```

filename out 'print2.pdf';

options printerpath=(pdf out);
proc template;
define style New_style / store = SASUSER.TEMPLAT;
parent = styles.printer;
replace fonts /
'docFont'          = ("Thorndale Duospace WT J", 12pt)
'headingFont'      = ("Albany AMT", 10pt, bold)
'headingEmphasisFont' = ("Albany AMT", 10pt, bold italic)
'TitleFont'        = ("Albany AMT", 12pt, italic bold)
'TitleFont2'       = ("Albany AMT", 11pt, italic bold)
'FixedFont'        = ("Cumberland AMT", 11pt)
'BatchFixedFont'   = ("Cumberland AMT", 6pt)
'FixedHeadingFont' = ("Cumberland AMT", 9pt, bold)
'FixedStrongFont'  = ("Cumberland AMT", 9pt, bold)
'FixedEmphasisFont' = ("Cumberland AMT", 9pt, italic)

```

```
'EmphasisFont'      = ("Albany AMT", 10pt, italic)
'StrongFont'         = ("Albany AMT", 10pt, bold);
end;
run;

ods printer style=New_style;
proc print data=sashelp.class;
  title1 'Proc Print';
  title2 'Templated ODS output';
run;
ods printer close;
```

Display 16.23 PDF Output Using PROC PRINT and a User-Defined ODS Template

Proc Print Templated ODS output

Obs	Name	Sex	Age	Height	Weight
1	Alfred	M	14	69.0	112.5
2	Alice	F	13	56.5	84.0
3	Barbara	F	13	65.3	98.0
4	Carol	F	14	62.8	102.5
5	Henry	M	14	63.5	102.5

Printing International Characters

Example 1

The following example produces an output file, **titles.png**. It prints ten Unicode characters.

```
filename new 'titles.png';
options printerpath=(png new);
ods printer;

proc gslide;
  title1 "Printing Unicode code points";
  title2 "double exclamation mark" f="Monotype Sans WT SC/Unicode" h=2 '203C'x;
  title3 "French Franc symbol " f="Monotype Sans WT SC/Unicode" h=3 '20A3'x;
```

```

title4 "Lira symbol " f="Monotype Sans WT SC/Unicode" h=3 '20A4'x;
title4 "Rupee symbol " f="Monotype Sans WT SC/Unicode" h=3 '20A8'x;
title5 "Euro symbol " f="Monotype Sans WT SC/Unicode" h=3 '20Ac'x;
title6 "Fraction, one third " f="Monotype Sans WT SC/Unicode" h=3 '2153'x;
title7 "Fraction, one fifth " f="Monotype Sans WT SC/Unicode" h=3 '2155'x;
title8 "Fraction one eighth " f="Monotype Sans WT SC/Unicode" h=3 '215B'x;
title9 "Black Florette " f="Monotype Sans WT SC/Unicode" h=3 '273F'x;
title10 "Black Star " f="Monotype Sans WT SC/Unicode" h=3 '2605'x;
run;
quit;
ods printer close;

```

Display 16.24 Output of Ten Unicode Characters

Printing Unicode code points

double exclamation mark !!

French Franc symbol ₣

Rupee symbol ₨

Euro symbol €

Fraction, one third $\frac{1}{3}$

Fraction, one fifth $\frac{1}{5}$

Fraction one eighth $\frac{1}{8}$

Black Florette ☀

Black Star ★

Example 2

The following example produces an output file, **utf8.gif**. It must be run with a UTF-8 server and requires a TrueType font that contains the characters that are used. The table of character names and the associated codes can be found on the Unicode Web site at <http://www.unicode.org/charts>.

```

proc template;
define style utf8_style / store = SASUSER.TEMPLAT;
parent = styles.printer;
replace fonts /
'docFont'          = ("Monotype Sans WT J", 12pt)
'headingFont'     = ("Monotype Sans WT J", 10pt, bold)
'headingEmphasisFont' = ("Monotype Sans WT J", 10pt, bold italic)
'TitleFont'       = ("Monotype Sans WT J", 12pt, italic bold)
'TitleFont2'      = ("Monotype Sans WT J", 11pt, italic bold)
'FixedFont'       = ("Thorndale Duospace WT J", 11pt)

```

```

'BatchFixedFont'      = ("Thorndale Duospace WT J", 6pt)
'FixedHeadingFont'   = ("Thorndale Duospace WT J", 9pt, bold)
'FixedStrongFont'    = ("Thorndale Duospace WT J", 9pt, bold)
'FixedEmphasisFont'  = ("Thorndale Duospace WT J", 9pt, italic)
'EmphasisFont'       = ("Monotype Sans WT J", 10pt, italic)
'StrongFont'          = ("Monotype Sans WT J", 10pt, bold);
end;
run;

%macro utf8chr(ucs2);
  kcvf(&ucs2, 'ucs2b', 'utf8');
%mend utf8chr;
%macro namechar(name, char);
  name="&name"; code=upcase("&char"); char=%utf8chr("&char"x); output;
%mend namechar;

data uft8char;
length name $40;
%namechar(Registered Sign, 00AE);
%namechar(Cent Sign, 00A2);
%namechar(Pound Sign, 00A3);
%namechar(Currency Sign, 00A4);
%namechar(Yen Sign, 00A5);
%namechar(Rupee Sign, 20A8);
%namechar(Euro Sign, 20Ac);
%namechar(Dong Sign, 20Ab);
%namechar(Euro-currency Sign, 20A0);
%namechar(Colon Sign, 20A1);
%namechar(Cruzeiro Sign, 20A2);
%namechar(French Franc Sign, 20A3);
%namechar(Lira Sign, 20A4);
run;

options printerpath=(gif out);
filename out 'utf8.gif';

ods printer style=utf8_style;
  proc print;
  run;
ods printer close;

```

Display 16.25 Output of International Characters

Obs	name	code	char
1	Registered Sign	00AE	®
2	Cent Sign	00A2	¢
3	Pound Sign	00A3	£
4	Currency Sign	00A4	¤
5	Yen Sign	00A5	¥
6	Rupee Sign	20A8	₹
7	Euro Sign	20AC	€
8	Dong Sign	20AB	₫
9	Euro-currency Sign	20A0	₠
10	Colon Sign	20A1	₡
11	Cruzeiro Sign	20A2	₧
12	French Franc Sign	20A3	₣
13	Lira Sign	20A4	₤

Creating Scalable Vector Graphics Using Universal Printing

Overview of Scalable Vector Graphics in SAS

Scalable Vector Graphics in SAS

Scalable Vector Graphics (SVG) is an XML language for describing two-dimensional vector graphics. SAS creates SVG documents based on the W3C recommendation for SVG documents. SAS SVG documents are created using the UNICODE standard encoding.

Note: Animation is not supported in SAS 9.2. \triangle

SAS can create SVG documents by using either Universal Printers or SAS/GRAPH device drivers. SVG documents that are created by Universal Printers are stand-alone SVG documents. SVG documents that are created by SAS/GRAPH device drivers can be either stand-alone SVG documents or SVG documents that are embedded in HTML documents.

A stand-alone SVG document can be embedded in an HTML document by reference as an external link, embedded inline, or as a CSS2 or XSL property reference. For information about embedding SVG documents in Web pages, see the topic on using SVG documents in Web pages in the SVG 1.1 specification on the W3 SVG Web site <http://www.w3.org/TR/SVG>.

SVG documents created by SAS/GRAPH can contain tooltips and have drill-down capabilities. SVG documents created by Universal Printers are stand-alone SVG documents.

The information provided here is limited to creating SVG documents using Universal Printers. For complete information about SVG documents, see the W3C recommendations for SVG documents. For information about creating SVG documents using SAS/GRAPH, see *SAS/GRAPH: Reference*.

For detailed information about the SVG standard, see the W3 documentation at <http://www.w3.org/TR/SVG>.

SVG Terminology

SVG canvas

the space upon which the SVG document is rendered.

viewBox

specifies the coordinate system and the area of the SVG document that is visible in the viewport.

viewport

a finite rectangular space within the SVG canvas where an SVG document is rendered. In SAS, the viewport is determined by the value of the PAPER SIZE= system option for a scalable viewport and by the SVGWIDTH= and SVGHEIGHT= system options for a static viewport.

viewport coordinate system or viewport space

the starting X and Y coordinates and the width and height values of the viewport.

user coordinate system or user space
 the starting X and Y coordinates and the width and height values of the area of the document to display in the viewport.

user units
 is equal to one unit of measurement that is defined in your environment's coordinate system. In many cases, the coordinate system uses pixels. Check with your system administrator to determine the unit of measure that is used in your environment.

Why Create SVG Documents?

SVG documents display clearly at any size in any viewer or browser that supports SVG. SVG documents are ideal for producing documents to display on a computer monitor, PDA, or cell phone; or documents to be printed. Because it is a vector graphic, a single SVG document can be transformed to any screen resolution without compromising the clarity of the document. In comparison, a multiple raster graphic image might require using different screen resolutions in order to display the image at various screen resolutions and sizes.

An SVG document might also be smaller in file size than the same image created by a raster graphic Universal Printer, such as GIF or PNG.

ODS Destinations That Produce SVG Documents

You create SVG documents using the ODS Listing, HTML, or Printer destinations. SVG documents created by the ODS Listing and HTML destinations are useful for creating graphs using SAS/GRAPH. Using the ODS Listing destination results in a stand-alone graph. Using the ODS HTML destination results in an SVG document that is embedded in an HTML document using `<embed>` tags. SAS/GRAPH output can also be created using the ODS Printer destination with SVG printers. However, support for tooltips and drilldown paths are not available with the ODS Printer destination. For information about creating SVG documents for SAS/GRAPH devices, see *SAS/GRAPH: Reference*.

SVG documents that are created by the ODS Printer destination are created by an SVG Universal Printer, and they are stand-alone documents. These documents can contain output from both SAS/GRAPH and non-graphic procedures such as tables.

Fundamentals of Creating SVG Documents Using Universal Printers

You can create SVG documents by using either programming statements or SVG Universal Printers in the Print dialog box.

To create SVG documents programmatically, specify an SVG Universal Printer as the value of the `PRINTERPATH=` system option and specify the ODS `PRINTER` statement, as shown below.

```
options printerpath=svg;
ods listing close;
ods printer;

...more SAS code...

ods printer close;
ods listing;
```

Alternatively, you can specify the SVG printer in the ODS `PRINTER` statement and eliminate the `OPTIONS` statement, as shown below.

```
ods listing close;
ods printer printer=svg;

...more SAS code...


ods printer close;
ods listing;
```

SAS has several system options that enable you to modify various aspects of your SVG document, such as a specific SVG Universal Printer, the height and width of the SVG document, the size of the viewBox, or whether a multi-page SVG document contains navigational controls. By using the NEWFILE option in the ODS PRINTER statement, you can create an SVG document for the output from each procedure or DATA step.

For more information, see the PRINTERPATH= system option in *SAS Language Reference: Dictionary* and the ODS Printer statement in the *SAS Output Delivery System: User's Guide*.

To create SVG documents by using the Print dialog box, you select an SVG printer from the Name drop-down list box.

Display 16.26 Print Dialog Box For Printing an SVG Document



Web Server Content Type for SVG Documents

If the mime content type setting for your Web server does not have the correct setting for SVG documents, your Web browser might render SVG documents as text files, or SVG documents might be unreadable.

To ensure that SVG documents are rendered correctly, you can configure your Web server to use this mime content type:

```
image/svg+xml
```

Browser Support for Viewing SVG Documents

Browsers That Support SVG Documents

In order to view SVG documents, you need a viewer or browser that supports Scalable Vector Graphics. Some browsers, such as Mozilla Firefox, have built-in support for SVG documents. Other browsers, such as Microsoft Internet Explorer, require an SVG plug-in to view SVG documents. One such plug-in is available from Adobe Systems, Inc.

The following table lists some browsers and viewers that support SVG documents.

Table 16.11 SVG Browser Support

Browser or Viewer	Company
Adobe SVG Viewer 3 *	Adobe Systems, Inc.
Batik SVG Toolkit	Apache Software Foundation
eSVG Viewer and IDE	eSVG Viewer for PC, PDA, Mobile
GPAC Project	GPAC
Mozilla Firefox	Mozilla Foundation
Opera	Opera Software
TinyLine	TinyLine

* Adobe SVG Viewer 3 works in Internet Explorer 7 but is no longer supported by Adobe Systems Inc.

Notes on Using Adobe SVG Viewer

- You can right-click on an SVG document to open a pop-up menu for the document. From here you can zoom in on or out of the SVG document, view its source, or copy it.
- You can use keyboard shortcuts to pan and zoom in on or out of the document. To pan, hold the Alt key and click and drag the left mouse button. If the scroll lock is enabled and the focus is on the Adobe SVG Viewer, you can use the arrow keys to pan the image.
- To zoom in at the location of the mouse pointer, hold the Ctrl key and click the left mouse button. To zoom out, hold the Ctrl and Shift keys while you click the left mouse button.
- There is a known problem that causes an SVG graph to be stretched out using an incorrect aspect ratio when you print an SVG document that is displayed in the Adobe SVG Viewer 3. To work around this problem, use the **x** at the top of the window to close the Print Preview window. Do not use the **Close** button. Preview the SVG document again. The graph now appears with the correct aspect ratio and can be printed by selecting **Print** from the **Print Preview** window.

Notes on Using Mozilla Firefox

- Compressed SVG documents using the SVGZ Universal Printer are not supported.
- Zooming and panning features are not currently implemented.
- If you select **View ► Page Style ► No Style**, all graphs appear as a black rectangle.

- Firefox does not support font embedding. To avoid font mapping problems in your SVG document, you can set the NOFONTEMBEDDING system option. If the FONTEMBEDDING option is set when an SVG document is created and the SVG document is subsequently viewed in Firefox, Firefox uses the default font setting that is defined on the **Contents** tab in the Firefox Options dialog box.

The SVG Universal Printers and the Output That They Create

The following table lists the SAS SVG Universal Printers.

SVG	produces SVG 1.1 documents.
SVGt	produces SVG 1.1 documents that are transparent (no background).
SVGnotip	produces SVG 1.1 documents without tooltips.
SVGZ	produces compressed SVG 1.1 documents.
SVGView	produces SVG1.1 documents with navigational controls when the SVG file contains multiple pages.

SVG prototype printers are available in the SAS Registry under CORE\PRINTING\PROTOTYPES. You can define your own SVG printer using the PRTDEF procedure. For more information, see *The PRTDEF Procedure in Base SAS Procedures Guide*.

Setting the Environment to Create Stand-alone SVG Documents

Overview of Setting the Environment to Create Stand-alone SVG Documents

As shown in “Fundamentals of Creating SVG Documents Using Universal Printers” on page 224, an SVG Universal Printer must be specified either as the printer value using the PRINTERPATH= system option or the ODS Printer statement. You can set any of the SVG system options when SAS is invoked in a SAS program by using the OPTIONS statement, or by using the SAS System Options window.

SAS SVG documents can be created easily by using default values for SVG system options (except for the PRINTERPATH= system option) that establish the SVG environment. This section explains the SVG system options and how they effect stand-alone SVG documents.

SAS System Options that Affect Stand-alone SVG Documents

You can use the following SVG system to set the environment for creating SVG documents:

Table 16.12 SAS System Options That Affect SVG Documents

Task	System Option
Set the paper size to use for Universal Printing.	PAPERSIZE=
Specify the name of an SVG printer to create a stand-alone SVG document.	PRINTERPATH=
Set the height of the SVG document. If the SVG document has embedded SVG documents, the height value affects only the outermost SVG document.	SVGHEIGHT=
Set the width of the SVG document. If the SVG document has embedded SVG documents, the width value affects only the outermost SVG document.	SVGWIDTH=
Set the x-axis coordinate for the lower, left corner of an embedded SVG document.	SVGX=
Set the y-axis coordinate for lower, left corner of an embedded SVG document.	SVGY=
Specify the X and Y coordinates, and the width and height that are used to set the viewBox for the outermost SVG document; specify the coordinates of the area of the document that is displayed in the viewport.	SVGVIEWBOX=
Specify whether to force uniform scaling of an SVG document.	SVGPRESERVEASPECTRATIO=
Set the title that appears in the title bar of the SVG document.	SVGTITLE=
Specify whether to display navigational controls in a multi-page SVG document.	SVGCONTROLBUTTONS

Setting the SVG Universal Printer

You set the SVG Universal Printer by setting the PRINTERPATH= system option to one of the SVG Universal Printers. You can set the PRINTERPATH= system option at any time. The following OPTIONS statement sets the Universal Printer to create compressed SVG documents:

```
options printerpath=svgz;
```

If the UPRINT system option is specified when SAS starts, the value of the PAPERSIZE= system option is set to the default value of 600 x 800 pixels.

For more information, see the following topics:

- “The SVG Universal Printers and the Output That They Create” on page 227
- PRINTERPATH= System Option in *SAS Language Reference: Dictionary*

Scaling an SVG Document to the Viewport

To scale an SVG document to the viewport, you can use the default value of null for the SVGHEIGHT= and SVGWIDTH= system options. A null value equates to the value of 100%, which scales the SVG document to the size of the viewport. In addition, the value of the SVGVIEWBOX= system option must be null.

For more information, see the following system option in *SAS Language Reference: Dictionary*:

- SVGHEIGHT= System Option
- SVGWIDTH= System Option
- SVGVIEWBOX= System Option

Setting the ViewBox

The `viewBox` attribute on the `<svg>` element is a set of four numbers: the starting X coordinate, the starting Y coordinate, the height of the SVG document, and the width of the SVG document. SAS sets the `viewBox` attribute value from the value of the `SVGVIEWBOX=` system option. If that option has no value, SAS uses the value of the `PAPERSIZE=` system option to set the height and the width arguments of the `viewBox` attribute. The starting coordinate values are set to 0.

When the `SVGVIEWBOX=`, `SVGHEIGHT=`, and `SVGWIDTH=` system options have a null value (the default value for each of these system options), the SVG document scales to the size of the viewport. If you specify a value for the `SVGVIEWBOX=` system option, the SVG document is scaled to the dimensions specified in the `SVGVIEWBOX=` option.

If you specify the `SVGHEIGHT=` option and the `SVGWIDTH=` option using percentage units, the SVG document scales to the size of the browser window whenever the browser window changes size. If these options are specified using units other than percentage, such as in, cm, or px, the SVG document is a static size and does not scale to the browser window when the window changes size.

For more information, see the following topics:

- ODS PRINTER Statement in *SAS Output Delivery System: User's Guide*
- `SVGVIEWBOX=` System Option in *SAS Language Reference: Dictionary*
- `PAPERSIZE=` System Option in *SAS Language Reference: Dictionary*
- "Creating a Static `viewBox`" on page 230

Interaction between SAS SVG System Options and the SVG Tag Attributes

SAS uses the values of the `SVGHEIGHT=`, `SVGWIDTH=`, `SVGVIEWBOX=`, `SVGPRESERVEASPECTRATIO=`, `SVGX=`, and `SVGY=` system options as values for their respective attributes on the outermost `<svg>` element: `height`, `width`, `viewBox`, and `preserveAspectRatio`. For example, if you specify `SVGWIDTH="400"` and `SVGHEIGHT="300"`, SAS creates the `<svg>` element with the attributes `width="400"` and `height="300"`. The values of the `SVGX=` and `SVGY=` system options are used only on embedded `<svg>` tags for the `x` and `y` attributes.

All of these system options have a null default value. When the `SVGVIEWBOX=` system option is null, SAS determines the `viewBox` size based on the value of the `PAPERSIZE=` system option. Therefore, if you do not specify a value for any of these system options, the only `<svg>` attribute that SAS sets is the `viewBox` attribute using the SAS SVG system options. Other `<svg>` attributes, such as `version` and `xmlns` are set by SAS and not by using system options.

SAS creates the following `<svg>` tag when all of the SAS SVG system options are set to their default values:

```
<svg xmlns="http://www.w3.org/2000/svg"
      xmlns:xlink="http://www.w3.org/1999/xlink"
      xml:space="preserve" onload='Init(evt)'\
      version="1.1" baseProfile="full" viewBox="0 0 800 600">
```

The `SVGPRESERVEASPECTRATIO=` system option is used to set the `preserveAspectRatio` attribute in the `<svg>` tag and has an effect only when the

viewBox attribute has also been specified in an SVG document. By default, SVG documents produced by SAS contain a viewBox attribute that is set to scale to the size of the paper size of the SVG printer.


Negative values can be specified for the values of SVG options. However, if a negative value is specified for the SVGHEIGHT= option or the SVGWIDTH= option, or the height or width arguments in the SVGVIEWBOX= option, the SVG document is not rendered by the browser. It can be useful to specify negative values for the x and y arguments of the SVGVIEWBOX= option to place the origin of the SVG document. A negative argument in the SVGVIEWBOX= option shifts the output to the right. A negative value of the SVGVIEWBOX= option shifts the placement of the document downward.

Creating a Static viewBox

A static viewBox is a viewBox that cannot be changed. When the viewport changes, such as when you resize your browser window, the viewBox remains the same size. To create a static viewBox, you specify the same width and height values for the PAPERSIZE=, SVGWIDTH=, and SVGHEIGHT= system options. The PAPERSIZE= system option sets the viewBox. The SVGWIDTH= and SVGHEIGHT= system options set the size of the SVG document. If the SVGHEIGHT= and SVGWIDTH= options are specified using percentage unit, the SVG document scales to the size of the browser window when the browser window changes size. Display 16.27 on page 230 shows a static viewBox created by using the following system options:

```
options printerpath=svg papersize=("8cm" "5cm") svgwidth="8cm" svgheight="5cm"
        svgtitle="SVG Cowboy Hat";
```

Display 16.27 A Static Viewbox



To reset the SVGWIDTH=, SVGHEIGHT=, and SVGPRESERVEASPECTRATIO= system options to null, specify two single quotation marks or two double quotation marks with no space between them:

```
options printerpath=svg svgwidth="" svgheight="" svgpreserveaspectratio="";
```

Preserving the Aspect Ratio

When you change the size of the viewBox, you can use the SVGPRESERVEASPECTRATIO= system option to specify whether you want to preserve the aspect ratio of the SVG document and how to place the SVG document in the viewport. You can set this option by using one of the following assignments:

```
SVGPRESERVEASPECTRATIO=align | meetOrSlice | NONE | ""
```

```
SVGPRESERVEASPECTRATIO=""align meetOrSlice
```

The first argument, *align*, specifies whether to force uniform scaling by specifying the alignment method to use. For example, you can use the **xMidyMid** value to align the midpoint X value of the viewBox to the midpoint X value of the viewport, which centers the document horizontally.

The second argument, *meetOrSlice*, specifies how to scale the SVG document to the viewBox. The value for this argument can be either **meet** or **slice**. If you specify **meet**, the SVG document is scaled up as much as possible while meeting other criteria. The viewport displays some unused space. If you specify **slice**, the SVG document is scaled down as much as possible while meeting other criteria. In the latter case, some of the SVG document appears to be cut off. The SVG document is still complete, but you cannot see all of it in the viewport. You can use your browser controls to move the SVG document around in the viewport.

For more information, see the SVGPRESERVEASPECTRATIO= System Option in *SAS Language Reference: Dictionary*.

Adding a Title to an SVG Document

You use the SVGTITLE= system option to add a title to the title bar of a window when the browser displays only the SVG document. If the SVG document is embedded in an HTML page, the svgtitle attribute on the <svg> tag has no effect. The static viewport example in the previous topic shows a title in the browser title bar.

For more information, see the SVGTITLE= System Option in *SAS Language Reference: Dictionary*.

Creating Stand-alone SVG Documents Using the ODS PRINTER Destination

Creating an SVG Document

To create an SVG document, the least you need to do is set the PRINTERPATH= system option to an SVG Universal printer and specify the ODS PRINTER statement in your SAS program, or specify the PRINTER= option in the ODS PRINTER statement:

```
options printerpath=svg;
ods listing close;
ods printer;
data hat;
  do x=-5 to 5 by .25;
    do y=-5 to 5 by .25;
      z=sin(sqrt(x*x+y*y));
      output;
    end;
  end;
end;
```

```

proc g3d data=hat;
  plot y*x=z/ctop=red;
  title 'Cowboy Hat with G3D';
run;
quit;
ods printer close;
ods listing;

```

In this example, no specific SVG system option values were set to size the SVG document; therefore, the viewBox is the default size specified by the PAPERSIZE= system option. The SVG document scales to the viewport because no value was specified for the SVGWIDTH= and SVGHEIGHT= system options. The following is the <svg> tag that SAS creates:

```


<svg xmlns="http://www.w3.org/2000/svg"
  xmlns:xlink="http://www.w3.org/1999/xlink" xml:space="preserve"
  onload='Init(evt)' version="1.1" baseProfile="full" viewBox="0 0 850 1100">

```

SAS creates a single SVG document named sasprt.svg and stores it in a specific location, depending on your operating environment. Under Windows, the file is stored in the current directory. Under UNIX and OpenVMS, the file is stored in your home directory. Under z/OS, the file is stored as a z/OS UNIX System Services Hierarchical File System (HFS) file, or as a z/OS data set. If the SVG file is written to a z/OS data set, it is written to PDSE library *userid.SASPR.T.SVG*. You can use the FILE= option in the ODS PRINTER statement to specify a different filename.

The following figure is the SVG file that was created for the Cowboy Hat program using the Adobe Acrobat SVG plug-in for Microsoft Internet Explorer:

Display 16.28 Cowboy Hat as an SVG File



When you use the SVG, SVGnotip, SVGt, SVGView, and SVGZ Universal Printers, SAS creates a single SVG document. Depending on the size of the SVG document, the

browser might display the complete SVG document. Check the documentation for your browser to determine whether your browser has controls for viewing SVG documents. In the Adobe SVG Viewer plug-in for Internet Explorer, you can press the Alt key and the left mouse button to pan and move to different pages in a continuous, multi-page SVG document.

Multi-Page SVG Documents in a Single File

When a DATA step or procedure creates a new page in the output, a new SVG page is created in an SVG document. SAS creates either one file with multiple pages or multiple SVG files with one file for each SVG document page. The SVGCONTROLBUTTONS system option and the NEWFILE= option in the ODS PRINTER statement control whether a multi-page SVG document is one continuous file (with controls to navigate the pages in the file) or multiple SVG files.

SAS creates a single file, multi-page SVG document with navigational controls when the NEWFILE= option of the ODS PRINTER statement is a value other than PAGE and one of the following set of system options is specified:


- the PRINTERPATH= system option is set to SVG or SVGZ, and the SVGCONTROLBUTTONS system option is set.
- the PRINTERPATH= system option is set to SVGView.

The SVGView Universal Printer sets the SVGCONTROLBUTTONS system option.

If the SVGCONTROLBUTTONS system option is not specified or the Universal Printer is not SVGView, the SVG document is created in a continuous-page layout. To navigate the document, you would use your browser controls.


The navigation controls enable you to go to the next page, the previous page, the first page, or the last page; to display an index of all pages; or to hide or show the controls.

Display 16.29 First Page of a Multi-page SVG File With Navigation Controls



To display an index of all pages in the SVG file, select the **Index** button. To go to a specific page from the index, select the thumbnail image of the page.

Display 16.30 Index of a Multi-page SVG File With Navigation Controls



You can hide the control buttons by selecting the **SVG Controls** button. The tooltip displays when the cursor is over the control. To show the navigation controls again, click in the top area of the output when you see the tooltip **Click to toggle SVG control button bar**. This is useful when you want to print a page in the document without the SVG controls.

Display 16.31 A Multi-page SVG File That Hides the Navigational Controls

Obs	LIBNAME	MEMNAME	MEMLABEL	TYPEMEM	NAME	TYPE	LENGTH	VARNUM
1	WORK	SCHOOL			Class	2	8	4
2	WORK	SCHOOL			Grade	1	8	3

Obs	LABEL	FORMAT	FORMATL	FORMATD	INFORMAT	INFORML	INFORMD	JUST
1			0	0		0	0	0
2			0	0		0	0	1

Obs	NPOS	NOBS	ENGINE	DELOBS	IDXUSAGE	MEMTYPE	IDXCOUNT	PROTECT	FLAGS
1	32	7	V9	0	NONE	DATA	0	---	---
2	8	7	V9	0	NONE	DATA	0	---	---

Obs	COMPRESS	REUSE	SORTED	SORTEDBY	CHARSET	COLLATE	NODUPKEY
1	NO	NO	NO
2	NO	NO	NO

Obs	NODUPREC	ENCRYPT	POINTOBS	GENMAX	GENNUM	GENNEXT
1	NO	NO	YES	0	.	.
2	NO	NO	YES	0	.	.

For information about the NEWFILE= option, see ODS PRINTER Statement in the *SAS Output Delivery System: User's Guide*.

Creating Separate Files for Multi-Page SVG Documents

You can create a separate file for each page in an SVG document by specifying the NEWFILE=PAGE option in the ODS PRINTER statement. A new page is created when

a procedure explicitly starts a new page and not when the page size is exceeded. The first file is named `sasprt.svg`. Subsequent filenames have a number appended, starting with the number 1: `sasprt1.svg`, `sasprt2.svg`, and so on.

The following code creates two files: `sasprt.svg`, which contains the CONTENTS procedure output; and `svgprt1.svg`, which contains the PRINT procedure output:

```
options printerpath=svg;
ods listing close;
ods printer newfile=page;

data school;
  input Name $ Y Grade Class $ SSN;
datalines;
Phillip 1 85 Math 234107589
Jose 1 90 English 190873452
Cameron 2 70 Math 257902348
Tia 2 71 English 234567823
Johanna 3 98 French 876345290
Frauke 3 75 History 231987222
Carlotta 4 89 PE 87654321
;
run;

proc contents data=school out=schout(drop=crdate modate);
  title 'School Dataset';
run;

proc print data=schout;
run;

ods printer close;
ods listing;
```

For information about the `NEWFILE=` option, see the ODS PRINTER Statement in *SAS Output Delivery System: User's Guide*.

Creating Overlaid Transparent SVG Documents

You use the SVTt Universal Printer to create a transparent SVG document in which the pages are transparent and can be overlaid. The following is a SAS program that overlays a bar chart on a map of the United States:

```
%let name=map069;
filename odsout '.';

data bardata;
  input x y;
  datalines;
1 3
2 5
3 8
4 12
;
run;
```

```

goptions reset=all;
ods listing close;
options printerpath=svgt nodate nonumber;
ods printer file='map069.svg';

goptions border;

goptions gunit=pct htitle=5 htext=3 ftext="arial" cback=lightblue;
pattern1 v=s c=cornsilk;
title "SAS/Graph gmap and overlaid gchart with printerpath=svgt";
proc gmap data=maps.us map=maps.us ;
    id state;
    choro state / levels=1 nolegend coutline=graycc
        des="" name("&name");
run;

quit;

goptions dev=svgt hsize=3 vsize=1.5 horigin=2 vorigin=2 autosize=on cback=white;

/* you must use the default ods style, for transparency to work */


title;
goptions gunit=pct htext=12 ftext="arial/bo";
pattern1 v=s c=red;
axis1 label=none value=none major=none minor=none style=0;
axis2 color=blue label=none offset=(7,7) value=(color=blue);
proc gchart data=bardata;
    vbar x / discrete
        type=sum sumvar=y
        descending
        raxis=axis1
        maxis=axis2
        width=9
        space=5
        coutline=blue
        noframe
        des="" name("&name.b");
run;

quit;

ods printer close;

```

This code creates the following image:



Embedding and Linking Stand-alone SVG Documents in HTML Documents

Embedding SVG Documents in HTML Documents

When you embed SVG documents in an HTML document, the height and width attributes of the `<EMBED>` tag become the dimensions of the viewport. If you use the default values for the SVG system options when you create your SVG document, the SVG document scales to the size of the viewport. This is because there is no default value of the `SVGHEIGHT=` and `SVGWIDTH=` system options, which effectively equates to specifying a value of 100%. A value of 100% for these system options scales the SVG document to 100% of the viewport.

If you do not specify height and width attributes on the embed tag, the viewport dimensions are determined by the browser, and the embedded document might not render as you expected it to render.

For information about embedding SVG documents in HTML documents using the ODS HTML destination, see *SAS/GRAPH: Reference*.

Linking to an SVG Document

If you link to an SVG document in an HTML document and are using the default values for the SVG system options, the SVG document opens in the browser window and scales to the size of the viewable area in the window.


An Example HTML Document That Links and Embeds a Stand-alone SVG Document

The following simple HTML document links and embeds the cowboy hat SVG document:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN">
<html>
<head>
  <title>Linking and Embedding an SVG Document in an HTML Document</title>
</head>
<body>
<p>Linking to an SVG document:</p>
<a href="sasprt.svg">Cowboy Hat</a>
<p>Embed the SVG document:</p>
<embed src="sasprt.svg" type="image/svg+xml" height="400" width="300">
</body>
</html>
```


Microsoft Internet Explorer with the Adobe SVG Viewer 3 displays the HTML documents with the link and the embedded SVG document:

Display 16.32 An HTML Document Displaying a Link to a Stand-alone SVG Document and an Embedded SVG Document



The viewport has a height of 400 pixels and a width of 300 pixels. Because the default SVG system option values were used, the SVG document scales to 100% of the viewport.

If you click on the Cowboy Hat link, the browser displays the following SVG document:

Display 16.33 A Stand-alone SVG Document Displayed as a Result of Clicking an HTML Link

The viewport is the displayable area in the browser window and the SVG document scales to 100% of the viewport.

Images in SVG Documents

When you create an SVG document that contains an image, the image is created as a PNG file. The filename has the form *printer-destination-filename_counter.png*. The counter is an integer that is incremented each time a new image is created. For example, if the printer destination filename is *sasprt.svg*, the first image is named *sasprt_1.png*. The second image is named *sasprt_2.png*, and so on.

When a .png file for an SVG document is saved, SAS writes a note in the SAS Log indicating the path to the image.

Printing an SVG Document from a Browser

SVG document printing is controlled by the browser. The browser prints only what is displayed in the browser window.

Creating Portable Network Graphics Using Universal Printing

Overview of Portable Network Graphics (PNG) in SAS

Portable Network Graphics in SAS

Portable Network Graphics (PNG) is a raster image format that was designed to replace GIF and TIFF image formats that are viewed on the World Wide Web. PNG images that are created with a SAS Universal Printer or a SAS/GRAPH device driver use the PNG Reference Library, Version 1.2.8, also known as libpng. The PNG device is the default graphics device for ODS HTML destinations.

The PNG Universal Printers

SAS provides three PNG Universal Printers.

PNG	produces PNG images at 96 dpi using a paper size of 600x800 pixels
PNGt	produces PNG images at 96 dpi using a paper size of 600x800 pixels with a transparent background
PNG300	produces PNG images at 300 dpi using a paper size of 1875x2500 pixels

Producing PNG Images for ODS Destinations

Use the ODS Listing, Printer, or HTML destinations to create PNG images. For PNG images created by SAS/GRAPH devices, use the ODS Listing and ODS HTML destinations. For PNG images created by Universal Printers, use the ODS Printer destination. For information about creating PNG images by SAS/GRAPH devices, see *SAS/GRAPH: Reference*.

To create PNG images using Universal Printers, specify a PNG Universal Printer as the value of the PRINTERPATH system option, and the ODS PRINTER statement as shown in the following example:

```
options printerpath=png;
ods listing close;
ods printer;

...more SAS code...

ods printer close;
ods listing;
```

SAS creates a PNG file with the name `sasprt.png`.

Web Browsers and Viewers That Support PNG Files

Although many browsers support PNG images, the following browsers, using the specified version or later, support most PNG image capabilities:

- Microsoft Internet Explorer 7.01b
- Mozilla Firefox 1.5.0.4
- Netscape Navigator 6
- Irfan View for Windows
- Microsoft Photo Editor
- Windows Picture and Fax Viewer


For more information about browsers and viewers that support PNG images, see the PNG Web pages at www.libpng.org.

Creating PNG Files Using ODS Printer


To create a PNG image in SAS using one of the PNG Universal Printers, specify the PNG printer in the `PRINTERPATH=` system option and the ODS `PRINTER` statement as shown in the following example:

```
options printerpath=png;
ods listing close;
ods printer;
data hat;
  do x=-5 to 5 by .25;
    do y=-5 to 5 by .25;
      z=sin(sqrt(x*x+y*y));
      output;
    end;
  end;
proc g3d data=hat;
  plot y*x=z/ctop=red;
  title 'Cowboy Hat with G3D';
run;
quit;
ods printer close;
ods listing;
```

The following output is the PNG graphic displayed in Internet Explorer:




You can create PNG images for any procedure out. You are limited only by what fits on the paper size that is used to define the size of the image.



Windowing Environment Concepts

Chapter 17.....**Introduction to the SAS Windowing Environment** 247

Chapter 18.....**Managing Your Data in the SAS Windowing Environment** 263



CHAPTER

17

Introduction to the SAS Windowing Environment

<i>What Is the SAS Windowing Environment?</i>	247
<i>Main Windows in the SAS Windowing Environment</i>	248
<i>Overview of SAS Windows</i>	248
<i>SAS Explorer Window</i>	250
<i>Uses of the SAS Explorer Window</i>	250
<i>Open the SAS Explorer Window</i>	250
<i>Display SAS Explorer with and without a Tree View</i>	250
<i>Program Editor Window</i>	251
<i>Uses of the Program Editor Window</i>	251
<i>Open the Program Editor Window</i>	251
<i>View a Program in the Program Editor Window</i>	251
<i>Log Window</i>	252
<i>Uses of the Log Window</i>	252
<i>Open the Log Window</i>	252
<i>View Log Output</i>	252
<i>Output Window</i>	253
<i>Uses of the Output Window</i>	253
<i>Open the Output Window</i>	253
<i>View Listing Output</i>	253
<i>Results Window</i>	254
<i>Uses of the Results Window</i>	254
<i>Open the Results Window</i>	255
<i>View Output in the Results Window</i>	255
<i>Navigating in the SAS Windowing Environment</i>	255
<i>Overview of SAS Navigation</i>	255
<i>Drop-Down Menus in SAS</i>	256
<i>Pop-Up Menus in SAS</i>	257
<i>Toolbars in SAS</i>	257
<i>The Command Line</i>	258
<i>Getting Help in SAS</i>	259
<i>Type Help in the Command Line</i>	259
<i>Open the Help Menu from the Toolbar</i>	259
<i>Click Help in Individual SAS Windows</i>	260
<i>List of SAS Windows and Window Commands</i>	260

What Is the SAS Windowing Environment?

SAS provides a graphical user interface that makes SAS easier to use. Collectively, all the windows in SAS are called the SAS windowing environment.

The SAS windowing environment contains the windows that you use to create SAS programs, but you will also find other windows that enable you to manipulate data or change your SAS settings without writing a single line of code.

You might find the SAS windowing environment a convenient alternative to writing a SAS program when you want to work with a SAS data set, or control some aspect of your SAS session.

Main Windows in the SAS Windowing Environment

Overview of SAS Windows

SAS windows have several features that operate in a similar manner across all operating environments: drop-down menus, pop-up menus, toolbars, and online Help. You can customize many features of the SAS windowing environment, including toolbars, icons, menus, and so on.

The five main windows in the SAS windowing environment are the Explorer, Results, Program Editor, Log, and Output windows.


Note: The arrangement of your SAS windows depends on your operating environment. For example, in the Microsoft Windows operating environment, the Enhanced Editor window appears instead of the Program Editor. \triangle

When you first invoke SAS, the Program Editor, Log, and Output windows display. When you execute a SAS program output displays in the output windows and lists in the Results window. If you use a PUT statement in your program, then by default output is written to the SAS Log.

Note: The Microsoft Windows operating environment was used to create the examples in this section. Menus and toolbars in other operating environments have a similar appearance and behavior. \triangle


Operating Environment Information: If you are using Microsoft Windows, the active window determines which items are available on the main menu bar. \triangle

The following display shows one example of the arrangement of SAS windows.

Display 17.1 One View of the SAS Windowing Environment (Microsoft Windows)

The Enhanced Editor and the Log windows are displayed, and the Output window is positioned behind these two windows.

The following display shows the Log, Results, and Explorer windows.

Display 17.2 Another View of the SAS Windowing Environment

SAS Explorer Window

Uses of the SAS Explorer Window

The Explorer window enables you to manage your files in the windowing environment. You can use the SAS Explorer to do the following tasks:

- View lists of your SAS files.
- Create new SAS files.
- View, add, or delete libraries.
- Create shortcuts to external files.
- Open any SAS file and view its contents.
- Move, copy, and delete files.
- Open related windows, such as the New Library window.

Open the SAS Explorer Window

You can open SAS Explorer in the following ways:


Command: Type **EXPLORER** in the command line and press Enter.

Menu: Select **View ► Explorer**.


Display SAS Explorer with and without a Tree View

You can display the Explorer window with or without a tree view of its contents. Displaying the Explorer with a tree view enables you to view the hierarchy of the files. To display the tree view, select **Show Tree** from the **View** menu. To turn tree view off, deselect **Show Tree** in the menu. The following display shows SAS Explorer with the tree view turned on. Click the expansion icons to view other files in the hierarchy.

Display 17.3 SAS Explorer with Tree View Turned On



The following display shows SAS Explorer with the tree view turned off.

Display 17.4 SAS Explorer with Tree View Turned Off

Note: You can resize the Explorer window by dragging an edge or a corner of the window. You can resize the left and right panes of the Explorer window by clicking the split bar between the two panes and dragging it to the right or left. △

Program Editor Window

Uses of the Program Editor Window

The Program Editor window enables you to enter, edit, submit, and save SAS programs.

Open the Program Editor Window

You can open the Program Editor window in the following ways:

Command: Type **PROGRAM** or **PGM** in the command line and press Enter.


Menu: Select **View ► Program Editor**.

Note: To open your SAS programs in the SAS windowing environment, you can drag and drop them onto the Program Editor window. △

View a Program in the Program Editor Window

The following display shows a SAS program that is written in the Program Editor window.

Display 17.5 Example of the Program Editor Window



```

Program Editor - (Untitled)
options linesize=80 pagesize=64 nodate pageno=1;
ods listing;
ods html;

data Emplist;
  input Name $ 1-8 EmpNumber $ 10-13 Division $ 15-25;
  datalines;
Jacobson 4358 Agriculture
Lincoln 3422 Chemistry
Demara 4599 Physics
Vojtyla 5593 Ethics
;
run;

proc sort data=Emplist;
  by Name;
run;

proc print data=Emplist;
  title 'Employee Listing';
run;

ods listing close;
ods html close;

```

Log Window

Uses of the Log Window

The Log window enables you to view messages about your SAS session and your SAS programs. If the program you submit has unexpected results, then the Log helps you to identify the error. You can also use a PUT statement to write program output to the Log.

Note: To keep the lines of your log from wrapping when your window is maximized, use the LINESIZE= system option. Δ

Open the Log Window

You can open the Log window in the following ways:

Command: Type **Log** in the command line and press Enter.

Menu: Select **View ► Log**.

View Log Output

The following is an example of Log output.

Display 17.6 Example of Output in the Log Window

```

Log - (Untitled)
120 options linesize=80 pagesize=64 nodate pageno=1;
121 ods listing;
122 ods html;
NOTE: Writing HTML Body file: sashtml.htm
123
124 data EmplList;
125     input Name $ 1-8 EmpNumber $ 10-13 Division $ 15-25;
126     datalines;

NOTE: The data set WORK.EMPLIST has 4 observations and 3 variables.
NOTE: DATA statement used (Total process time):
      real time           0.07 seconds
      cpu time             0.01 seconds

131 ;
132 run;
133
134 proc sort data=EmplList;
135     by Name;
136 run;

NOTE: There were 4 observations read from the data set WORK.EMPLIST.
NOTE: The data set WORK.EMPLIST has 4 observations and 3 variables.
NOTE: PROCEDURE SORT used (Total process time):
      real time           0.09 seconds
      cpu time             0.01 seconds

137
138 proc print data=EmplList;
139     title 'Employee Listing';
140 run;

NOTE: There were 4 observations read from the data set WORK.EMPLIST.
NOTE: PROCEDURE PRINT used (Total process time):

```

Output Window

Uses of the Output Window

The Output window enables you to view listing output from your SAS programs. By default, the Output window is positioned behind the other windows. When you create output, the Output window automatically moves to the front of your display.

Note: To keep the lines of your output from wrapping when your window is maximized, use the `LINESIZE=` system option. △

Open the Output Window


You can open the Output window in the following ways:

- Command:
- Type **OUTPUT** or **OUT** in the command line and press Enter.
 - Type **LISTING** or **LST** in the command line and press Enter.

Menu: Select **View ► Output**.

View Listing Output

The following display shows the listing output from executing a SAS program.


Display 17.7 Example of Listing Output in the Output Window


Obs	Name	Emp Number	Division
1	Denara	4599	Physics
2	Jacobson	4358	Agriculture
3	Lincoln	3422	Chemistry
4	Wojtyla	5593	Ethics

Results Window

Uses of the Results Window

The Results window enables you to view HTML output from a SAS program. To do this, you must specify HTML output in the **Results** tab of the Preferences window. To access the Preferences window, select **Tools** \blacktriangleright **Options** \blacktriangleright **Preferences** from the menu. In the **Results** tab, select Listing and HTML output:

Display 17.8 The Preferences Window

The Results window uses a tree structure to list various types of output that might be available when you run SAS. You can view, save, or print individual items of output.

By default, the Results window is positioned behind the Explorer window and it is empty until you submit a SAS program. Then the Results window moves to the front of your display.

Open the Results Window

You can open the Results window in the following ways:


Command: Type **ODSRESULTS** in the command line and press Enter.

Menu: Select **View ► Results**.

View Output in the Results Window

The left pane of the following display shows the Results window, and the right pane shows the Results Viewer where HTML output is displayed. The Results window lists the files that were created when the SAS program executed. The icons indicate that both listing output and HTML output are available for viewing.

Display 17.9 Results Window and Results Viewer



Navigating in the SAS Windowing Environment


Overview of SAS Navigation

SAS windows have several features that work in a similar manner across all operating environments: drop-down menus, pop-up menus, toolbars, and online Help. You can customize many of these features by selecting **Tools ► Customize** from the menu. For specific information about these features, see the documentation for your operating environment.

Drop-Down Menus in SAS


A drop-down menu contains a list of options that appears when you select a menu item in the menu bar. The following example shows the drop-down menu for online Help.

Display 17.10 The Online Help Drop-down Menu



Drop-down menu choices change as you change the windows you are using. For example, if you select **Explorer** from the **View** drop-down menu, and then select **View** again, the drop-down menu lists the **View** options that are available when the Explorer window is active. The following display shows the **View** drop-down menu when the Explorer window is active.


Display 17.11 View Options when the Explorer Window is Active



If you select **Program Editor** from the **View** drop-down menu, and then select **View** again, the drop-down menu lists the **View** options that are available when the Program

Editor window is active. The following display shows the **view** drop-down menu when the Program Editor window is active.


Display 17.12 View Options when the Program Editor Window is Active



Pop-Up Menus in SAS

In SAS windows, a pop-up menu displays when you right-click an item. For example, when you right-click the Libraries icon in the Explorer window, the following pop-up menu displays.

Display 17.13 Example of a Pop-up Menu



A pop-up menu remains visible until you make a selection from the menu or until you click an area outside of the pop-up menu area.

Toolbars in SAS

A toolbar displays a block of window buttons or icons. When you click items in the toolbar, a function or an action is started. For example, clicking a picture of a printer in

a toolbar starts a print process. The toolbar displays icons for many of the actions that you perform most often in a particular window.

Operating Environment Information: SAS in the z/OS operating environment does not have a toolbar. See *SAS Companion for z/OS* for more information. Δ

The toolbar that you see depends on which window is active. For example, when the Program Editor window is active, the following toolbar displays.

Display 17.14 Example of the SAS Toolbar when the Program Editor Window is Active




When you position your cursor at one of the items in the toolbar, a text window appears that identifies the purpose of the icon.

The Command Line

The command line is located to the left of the toolbar. In the command line, you can enter commands, such as those that open SAS windows and that retrieve help information. The following is an example of entering the **footnotes** command to display the Footnotes window.

Display 17.15 The Footnotes Window




The following is an example of the command line with **Help footnotes** entered.

Display 17.16 Example of the Command Line



When you press Enter, SAS searches online Help and the following window displays.

Display 17.17 Results for the Footnotes Help Entry

Click an item to access Help for that topic.

Getting Help in SAS

Type Help in the Command Line

When you type Help in the command line, help for the active window displays. When you type **Help <item>** (for example, **Help footnotes**), you can access help for the item that you entered.

Open the Help Menu from the Toolbar

When you open the Help menu, you can select from the following choices:

- | | |
|-----------------------------------|---|
| Using This Window | opens a Help system window that describes the current active window. |
| SAS Help and Documentation | opens the SAS Help and Documentation system. Help is available for Base SAS and other SAS products that are installed on your system. You can find information by clicking an item in the table of contents or by searching for the item and then clicking the results. |
| Getting Started with SAS Software | opens the <i>Getting Started with SAS</i> tutorial. This is a good way to learn the basics of how to use SAS. |
| Learning SAS Programming | enables you to use SAS online training if you have an online training license. The software provides 50–60 hours of instruction for beginning as well as experienced SAS programmers. |
| SAS on the Web | provides links to the SAS Web site (if you have Web access), where you can do the following: <ul style="list-style-type: none"> □ Contact Technical Support. |

- Find information about Training Services.
- Read Frequently Asked Questions (FAQs).
- Send feedback.
- Access the Customer Support Center.
- Browse the SAS home page.

About SAS®9 provides version and release information about SAS.

Click Help in Individual SAS Windows

When you open a SAS window, you can press the HELP key (usually F1) from your keyboard to display information about that window.

List of SAS Windows and Window Commands

The basic SAS windows consist of the Explorer, Results, Program Editor, Enhanced Editor (Windows operating environment), Log, and Output windows, but there are more than thirty other windows to help you with such tasks as printing and fine-tuning your SAS session.

The following table lists all portable SAS windows, window descriptions, and the commands that open the windows.


Table 17.1 List of SAS Windows, Descriptions, and Window Commands

Window Name	Description	Window Commands
Documents	Displays your ODS Documents in a hierarchical tree structure.	ODSDOCUMENTS
Explorer	Provides a central access point to data such as catalogs, libraries, data sets and host files.	ACCESS, BUILD, CATALOG, DIR, EXPLORER, FILENAME, LIBNAME, V6CAT, V6DIR, V6FILENAME, V6LIBNAME
Explorer Options	Enables you to add or delete file types, change or add pop-up menu items, select folders that appear in the Explorer, and display member details.	DMEXPOPTS
File Shortcut Assignment	Assigns a file shortcut to a file using a graphical user interface.	DMFILEASSIGN
Find	Enables you to search for an expression in a SAS library.	EXPFIND
FOOTNOTES	Enables you to enter, browse, and modify footnotes for output.	FOOTNOTES
FSBrowse	Enables you to select a data set for browsing.	FSBROWSE
FSEdit	Enables you to select a data set to be processed by the FSEdit procedure.	FSEdit

Window Name	Description	Window Commands
FSForm	Enables you to customize a form for sending output to the printer.	FSFORM <i>formname</i>
FSLetter	Enables you to edit or create catalog entries.	FSLETTER
FSList	Enables you to browse external files in a SAS session.	FSLIST
FSView	Enables you to browse, edit, or create a SAS data set, displaying the data set as a table with rows and columns.	FSVIEW
Help	Displays help information about SAS.	HELP
Keys	Enables you to browse, alter, and save function key settings.	KEYS
Log	Displays messages and SAS statements for the current SAS session.	LOG
Metafind	Enables you to search for metadata objects in repositories by using Uniform Resource Identifiers (URIs).	METAFIND
Metadata Server Connections	Enables you to import, export, add, remove, reorder, and test metadata server connections.	METACON
New Library	Enables you to create a new SAS library and assign a libref.	DMLIBASSIGN
NOTEPAD	Enables you to create and store notepads of text.	NOTEPAD, NOTE, FILEPAD <i>filename</i>
Options (SAS system options)	Enables you to view and change some SAS system options.	OPTIONS
Output	Displays procedure output in listing format.	OUTPUT, OUT, LISTING, LIST, LST
Page Setup	Enables you to specify page setup options that apply to Universal Printing jobs.	DMPAGESETUP
Password	Enables you to edit, assign, or clear passwords for a particular data set.	SETPASSWORD (followed by a two-level data set name)
Print	Enables you to print the content of an active SAS window through Universal Printing.	DMPRINT

Window Name	Description	Window Commands
Program Editor	Enables you to enter, edit, and submit SAS statements and save source files.	PROGRAM, PGM
Properties	Shows details that are associated with the current data set.	VAR <i>libref.SAS-data-set</i> , V6VAR <i>libref.SAS-data-set</i>
Registry Editor	Enables you to edit the SAS registry and to customize aspects of the SAS windowing environment.	REGEDIT
Results	Lists the procedure output that is produced by SAS.	ODSRESULTS
SAS/AF	Displays windowing applications that are created by SAS/AF software.	AF, AFA
SAS/ASSIST	Displays the primary menu of SAS/ASSIST software, which simplifies the use of SAS.	ASSIST
SASCOLOR	Enables you to change default colors for the different window elements in your SAS windows.	SASCOLOR
SQL QUERY	Enables you to build, run, and save queries without being familiar with Structured Query Language (SQL).	QUERY
System Options	Enables you to change SAS system option settings.	OPTIONS
Templates	Enables you to browse and edit template source code.	ODSTEMPLATES
TITLES	Enables you to enter, browse, and modify titles for output.	TITLES
VIEWTABLE	Enables you to browse, edit, or create tables (data sets).	VIEWTABLE, VT

Note: Some additional SAS windows that are specific to your operating environment might also be available. Refer to the SAS documentation for your operating environment for more information. \triangle



CHAPTER

18

Managing Your Data in the SAS Windowing Environment

<i>Introduction to Managing Your Data in the SAS Windowing Environment</i>	263
<i>Managing Data with SAS Explorer</i>	264
<i>Introduction to Managing Data with SAS Explorer</i>	264
<i>Viewing Libraries and Data Sets</i>	264
<i>Assign File Shortcuts</i>	265
<i>Rename a SAS Data Set</i>	266
<i>Copy a SAS Data Set</i>	267
<i>Sorting Data Sets in a Library</i>	267
<i>View the Properties of a SAS Data Set</i>	267
<i>Working with VIEWTABLE</i>	268
<i>Overview of VIEWTABLE</i>	268
<i>View the Contents of a SAS Data Set</i>	268
<i>Move Columns in a Table</i>	269
<i>Temporarily Change Column Headings</i>	270
<i>Sort by Values of a Column</i>	271
<i>Edit Cell Values</i>	272
<i>Subsetting Data by Using the WHERE Expression</i>	273
<i>Subset Rows of a Table</i>	273
<i>Clear the WHERE Expression</i>	276
<i>Exporting a Subset of Data</i>	276
<i>Overview of Exporting Data</i>	276
<i>Export Data</i>	276
<i>Importing Data into a Table</i>	279
<i>Overview of Importing Data</i>	279
<i>Import a Standard File</i>	279
<i>Import a Non-standard File</i>	282

Introduction to Managing Your Data in the SAS Windowing Environment

The SAS windowing environment contains windows that enable you to perform common data manipulation and make changes without writing code.

If you are not familiar with SAS or with writing code in the SAS language, then you might find the windowing environment helpful. With the windowing environment, you can open a data set, point to rows and columns in your data, and then click menu items to reorganize and perform analyses on the information.

For more information about the SAS windowing environment, select **SAS Help and Documentation** from the **Help** menu after you invoke a SAS session.

Managing Data with SAS Explorer

Introduction to Managing Data with SAS Explorer

You can use SAS Explorer to view and manage data sets. Data sets are stored in libraries, which are storage locations for SAS files and catalogs. By default, SAS defines several libraries for you:

Sashelp	is a library created by SAS that stores the text for Help windows, default function-key definitions, window definitions, and menus.
Maps	is a library created by SAS that presents graphical representations of geographical or other areas.
Sasuser	is a permanent SAS library that is created at the beginning of your first SAS session. This library contains a Profile catalog that stores the customized features or settings that you specify for SAS. (You can store other SAS files in this library.)
Work	is a library that is created by SAS at the beginning of each SAS session or SAS job. Unless you have specified a User library, any newly created SAS file that has a one-level name will be placed in the Work library by default, and will be deleted at the end of the current session or job.


Viewing Libraries and Data Sets

Libraries and data sets are represented in SAS by large icons, small icons, or as a list. You can change this representation by selecting an option from the **View** menu:

- To view large icons, select **Large Icons** from the **View** menu.
- To view small icons, select **Small Icons** from the **View** menu.
- To view data sets in a list, select **List** from the **View** menu.


The following display shows large icon representation for the Sashelp library:

Display 18.1 Sashelp Library Represented by Large Icons



If you select the Sashelp library and then select **View ► Details**, the contents of the Sashelp library displays, along with the size and type of the data sets:

Display 18.2 Detailed View of the Sashelp Library



If you double-click a table in this list, the data set opens. The VIEWTABLE window, which is a SAS table viewer and editor, appears and is populated with the data from the table.

Assign File Shortcuts


A file shortcut is also known as a file reference or fileref. Filerefs save you programming time by enabling you to assign a nickname to a commonly used file. You

can use the `FILENAME` statement to create a fileref, or you can use the File Shortcut Assignment window from SAS Explorer.

To assign a fileref to a file, follow these steps:

- 1 With the SAS Explorer window active, select **File shortcuts** from the left pane.
- 2 Right-click in the right pane to select **New File Shortcut** from the menu.
- 3 In the File Shortcut Assignment window that appears, type the name of the fileref you want to use in the **Name** field.
- 4 Type the full pathname for the file in the **File** field.

The following display shows the File Shortcut Assignment window:



By default, filerefs you create are temporary and can be used in the current SAS session only. Selecting **Enable at Startup** from the File Shortcut Assignment window, however, assigns the fileref to the file whenever you start a new SAS session.

Rename a SAS Data Set

You can rename any data set in a SAS library as long as it is not write protected. To rename a data set, follow these steps:

- 1 Open SAS Explorer and select a library.
The contents of the library appears in the right pane.
- 2 Right-click the data set that you want to rename.
- 3 Select **Rename** from the menu, and type the new name of the data set.
- 4 Click **OK**.

Copy a SAS Data Set

You can copy a SAS data set to another library or catalog, or you can duplicate the data set in the same directory as the original data set. To copy or duplicate a data set, follow these steps:

- 1 Open SAS Explorer and select a library.
The contents of the library appears in the right pane.
- 2 Right-click the data set you want to copy or duplicate.
- 3 From the menu that displays, choose **Copy** to copy a data set to another library or catalog, or choose **Duplicate** to copy the data set to the same library or catalog.
- 4 If you choose **Copy**, do the following:
 - a Left-click the library in the left pane of SAS Explorer to select the library or catalog to which the data set will be copied.
 - b In the right pane, right-click the mouse and select **Paste** from the menu that displays.
A copy of the data set now resides in the new directory.
- 5 If you choose **Duplicate**, then the Duplicate window appears. In the Duplicate window, SAS appends **_copy** to the data set name—for example, **data-set-name_copy**.
Do one of the following:
 - Keep the name and click **OK**.
 - Create another name for your duplicated data set and click **OK**.

Sorting Data Sets in a Library


Data sets in SAS Explorer are sorted automatically by name. You can change the sort order of the data sets by size or type by clicking the **Size** or **Type** column. To return data sets to their original order, select the **Refresh** option from the **View** menu.

View the Properties of a SAS Data Set

You can view the properties of a data set by using the Properties window. To view properties, follow these steps:

- 1 Open SAS Explorer and select a library.
The contents of the library appear in the right pane.
- 2 Right-click the data set that you want to view.
- 3 Select **Properties** from the menu.

The following window appears:



- 4 In the **Description** field of the **General** tab, you can enter a description of the data set. To save the description click **OK**.
- 5 Select other tabs to display additional information about the data set.

Working with VIEWTABLE


Overview of VIEWTABLE

To manipulate data interactively, you can use SAS software's table editor, the VIEWTABLE window. In the VIEWTABLE window, you can create a new table, and browse or edit an existing table.

View the Contents of a SAS Data Set

To view the contents of a data set with VIEWTABLE, follow these steps:

- 1 Open SAS Explorer and select a library.
The contents of the library displays.
- 2 Double-click a data set within the library.
The VIEWTABLE window appears and is populated with data from the data set:




3 Scroll the VIEWTABLE window to view all of the data in your data set.

Move Columns in a Table

Within the VIEWTABLE window, you can rearrange columns in your table. To move columns in your table,


- 1 Left-click a column heading for the column you want to move.
- 2 Drag and drop the heading onto another column heading. In this example, if you left-click the heading Product Type, and then drag and drop Product Type onto Country, the Product Type column moves to the right of the Country column.



Temporarily Change Column Headings

Within the VIEWTABLE window, you can temporarily change column headings. To temporarily change column headings, follow these steps:

- 1 Right-click the heading for the column that you want to change, and then select **Column Attributes** from the menu. In this example, the column heading **Country** is listed in the **Name** field of the Column Attributes window.



- 2 In the **Label** field of the Column Attributes window, type the new name of the column heading and then click **Apply**.

The column heading in VIEWTABLE changes to the new name. In this example the label was changed to **Country of Origin**.

	Country of Origin	Product Type	State/Province	County	Actual Sales	Predicted Sales	Product
22	U.S.A.	FURNITURE	California		\$646.00	\$950.64	BED
23	U.S.A.	FURNITURE	California		\$2,392.24	\$2,356.88	BED
24	U.S.A.	FURNITURE	California		\$913.92	\$1,538.16	BED
25	U.S.A.	OFFICE	California		\$1,090.72	\$375.36	CHAIR
26	U.S.A.	OFFICE	California		\$2,151.52	\$1,025.44	CHAIR
27	U.S.A.	OFFICE	California		\$1,221.28	\$265.20	CHAIR
28	U.S.A.	OFFICE	California		\$1,370.88	\$2,283.44	CHAIR
29	U.S.A.	OFFICE	California		\$2,547.28	\$1,925.76	CHAIR
30	U.S.A.	OFFICE	California		\$1,105.68	\$1,889.04	CHAIR
31	U.S.A.	OFFICE	California		\$1,977.44	\$606.56	CHAIR
32	U.S.A.	OFFICE	California		\$2,069.92	\$949.28	CHAIR
33	U.S.A.	OFFICE	California		\$329.12	\$2,258.96	CHAIR
34	U.S.A.	OFFICE	California		\$1,784.32	\$2,275.28	CHAIR
35	U.S.A.	OFFICE	California		\$1,640.16	\$1,908.08	CHAIR
36	U.S.A.	OFFICE	California		\$350.88	\$146.88	CHAIR
37	U.S.A.	OFFICE	California		\$829.60	\$2,431.68	DESK
38	U.S.A.	OFFICE	California		\$2,584.00	\$2,132.48	DESK
39	U.S.A.	OFFICE	California		\$2,222.24	\$2,883.20	DESK
40	U.S.A.	OFFICE	California		\$446.08	\$2,766.24	DESK
41	U.S.A.	OFFICE	California		\$1,490.56	\$2,256.24	DESK
42	U.S.A.	OFFICE	California		\$299.20	\$1,225.36	DESK

- 3 Click **Close** to close the Column Attributes window.

Sort by Values of a Column

You can sort your table in ascending or descending order, based on the values in a column. You can sort data permanently or create a sorted copy of your table.


To sort your table, follow these steps:

- 1 Right-click the heading of the column on which you want to sort, and select **Sort** from the menu.
- 2 Select **Ascending** or **Descending** from the menu.
- 3 When the following warning message appears, click **Yes** to create a sorted copy of the table.




Note: If you selected **Edit Mode** after opening the table and clicking a data cell, this window does not appear. SAS updates the original table. △

- 4 In the Sort window, type the name of the new sorted table. In this example, the name of the sorted table is **FurnSalesSorted**.



5 Click **OK**.

The rows in the new table are sorted in descending order by values of Actual Sales.



The screenshot shows the SAS Explorer window with a table titled "VIEWTABLE: Furniture sales data". The table has the following columns: Country of Origin, Product Type, State/Province, County, Actual Sales, Predicted Sales, and Product. The data is sorted by Actual Sales in descending order.


	Country of Origin	Product Type	State/Province	County	Actual Sales	Predicted Sales	Product
1	U.S.A.	FURNITURE	California		\$3,515.60	\$795.60	SOFA
2	U.S.A.	FURNITURE	California		\$3,440.80	\$680.00	BED
3	U.S.A.	FURNITURE	California		\$3,425.50	\$632.40	BED
4	U.S.A.	OFFICE	California		\$3,415.30	\$812.60	DESK
5	U.S.A.	OFFICE	California		\$3,413.60	\$2,002.60	DESK
6	U.S.A.	FURNITURE	California		\$3,410.20	\$793.90	SOFA
7	U.S.A.	OFFICE	California		\$3,394.90	\$821.10	DESK
8	U.S.A.	FURNITURE	California		\$3,394.90	\$780.30	SOFA
9	U.S.A.	OFFICE	California		\$3,384.70	\$924.80	DESK
10	U.S.A.	FURNITURE	California		\$3,379.60	\$2,864.90	SOFA
11	U.S.A.	OFFICE	California		\$3,376.20	\$2,284.80	CHAIR
12	U.S.A.	FURNITURE	California		\$3,372.80	\$756.50	SOFA
13	U.S.A.	FURNITURE	California		\$3,372.80	\$834.70	BED
14	U.S.A.	OFFICE	California		\$3,369.40	\$2,823.70	DESK
15	U.S.A.	FURNITURE	California		\$3,367.70	\$1,754.40	BED
16	U.S.A.	FURNITURE	California		\$3,359.20	\$1,530.00	BED
17	U.S.A.	OFFICE	California		\$3,359.20	\$3,105.90	CHAIR
18	U.S.A.	OFFICE	California		\$3,355.80	\$1,110.10	DESK
19	U.S.A.	OFFICE	California		\$3,337.10	\$3,139.90	CHAIR
20	U.S.A.	OFFICE	California		\$3,325.20	\$912.90	CHAIR
21	U.S.A.	OFFICE	California		\$3,323.50	\$2,312.00	DESK

Edit Cell Values

By default, VIEWTABLE opens existing tables in browse mode, which protects the table data. To edit the table, you need to switch to edit mode. To switch to edit mode, follow these steps:

- 1 With the table open, select **Edit ► Edit Mode** from the **Edit** menu.
- 2 Left-click a cell in the table, and the value in the cell is highlighted.

In this example, the third cell in the fourth row is selected.



The screenshot shows the same SAS Explorer window with the "VIEWTABLE: Furniture sales data" table. The cell containing "U.S.A." in the fourth row (row 4) is highlighted, indicating it is selected for editing.

	Country of Origin	Product Type	State/Province	County	Actual Sales	Predicted Sales	Product
1	U.S.A.	FURNITURE	California		\$3,515.60	\$795.60	SOFA
2	U.S.A.	FURNITURE	California		\$3,440.80	\$680.00	BED
3	U.S.A.	FURNITURE	California		\$3,425.50	\$632.40	BED
4	U.S.A.	OFFICE	California		\$3,415.30	\$812.60	DESK
5	U.S.A.	OFFICE	California		\$3,413.60	\$2,002.60	DESK
6	U.S.A.	FURNITURE	California		\$3,410.20	\$793.90	SOFA
7	U.S.A.	OFFICE	California		\$3,394.90	\$821.10	DESK
8	U.S.A.	FURNITURE	California		\$3,394.90	\$780.30	SOFA
9	U.S.A.	OFFICE	California		\$3,384.70	\$924.80	DESK
10	U.S.A.	FURNITURE	California		\$3,379.60	\$2,864.90	SOFA
11	U.S.A.	OFFICE	California		\$3,376.20	\$2,284.80	CHAIR
12	U.S.A.	FURNITURE	California		\$3,372.80	\$756.50	SOFA
13	U.S.A.	FURNITURE	California		\$3,372.80	\$834.70	BED
14	U.S.A.	OFFICE	California		\$3,369.40	\$2,823.70	DESK
15	U.S.A.	FURNITURE	California		\$3,367.70	\$1,754.40	BED
16	U.S.A.	FURNITURE	California		\$3,359.20	\$1,530.00	BED
17	U.S.A.	OFFICE	California		\$3,359.20	\$3,105.90	CHAIR
18	U.S.A.	OFFICE	California		\$3,355.80	\$1,110.10	DESK
19	U.S.A.	OFFICE	California		\$3,337.10	\$3,139.90	CHAIR
20	U.S.A.	OFFICE	California		\$3,325.20	\$912.90	CHAIR
21	U.S.A.	OFFICE	California		\$3,323.50	\$2,312.00	DESK

- 3 Type a new value in the cell and press **Enter**.

In this example, the cell has been updated with the value North Carolina.

	Country of Origin	Product Type	State/Province	County	Actual Sales	Predicted Sales	Product
1	U.S.A.	FURNITURE	California		\$3,515.60	\$795.60	SOFA
2	U.S.A.	FURNITURE	California		\$3,440.80	\$680.00	BED
3	U.S.A.	FURNITURE	California		\$3,425.50	\$632.40	BED
4	U.S.A.	OFFICE	North Carolina		\$3,415.30	\$812.60	DESK
5	U.S.A.	OFFICE	California		\$3,413.60	\$2,002.60	DESK
6	U.S.A.	FURNITURE	California		\$3,410.20	\$793.90	SOFA
7	U.S.A.	OFFICE	California		\$3,394.90	\$821.10	DESK
8	U.S.A.	FURNITURE	California		\$3,394.90	\$780.30	SOFA
9	U.S.A.	OFFICE	California		\$3,384.70	\$924.80	DESK
10	U.S.A.	FURNITURE	California		\$3,379.60	\$2,864.90	SOFA
11	U.S.A.	OFFICE	California		\$3,376.20	\$2,284.80	CHAIR
12	U.S.A.	FURNITURE	California		\$3,372.80	\$756.50	SOFA
13	U.S.A.	FURNITURE	California		\$3,372.80	\$834.70	BED
14	U.S.A.	OFFICE	California		\$3,369.40	\$2,823.70	DESK
15	U.S.A.	FURNITURE	California		\$3,367.70	\$1,754.40	BED
16	U.S.A.	FURNITURE	California		\$3,359.20	\$1,530.00	BED
17	U.S.A.	OFFICE	California		\$3,359.20	\$3,105.90	CHAIR
18	U.S.A.	OFFICE	California		\$3,355.80	\$1,110.10	DESK
19	U.S.A.	OFFICE	California		\$3,337.10	\$3,139.90	CHAIR
20	U.S.A.	OFFICE	California		\$3,325.20	\$912.90	CHAIR
21	U.S.A.	OFFICE	California		\$3,323.50	\$2,312.00	DESK

- 4 Select **File** ► **Close** from the **File** menu.
- 5 When prompted to save pending changes to the table, click **Yes** to save your changes or **No** to disregard changes.

Note: If you make changes in one row and then edit cells in another row, the changes in the first row are automatically saved. When you select **File** ► **Close**, you are prompted to save the pending changes to the second row. △

Subsetting Data by Using the WHERE Expression

Subset Rows of a Table


In the VIEWTABLE window, you can subset the display to show only those rows that meet one or more conditions. To subset rows of a table, follow these steps:

- 1 In the Explorer window, open a library and double-click the table that you want to subset.

In this example, the Furniture sales data table is selected.


	Country of Origin	State/Province	County	Actual Sales	Predicted Sales	Product Type	Product
1	U.S.A.	California		\$3,515.60	\$795.60	FURNITURE	SOFA
2	U.S.A.	California		\$3,440.80	\$680.00	FURNITURE	BED
3	U.S.A.	California		\$3,425.50	\$632.40	FURNITURE	BED
4	U.S.A.	California		\$3,415.30	\$812.60	OFFICE	DESK
5	U.S.A.	California		\$3,413.60	\$2,002.60	OFFICE	DESK
6	U.S.A.	California		\$3,410.20	\$793.90	FURNITURE	SOFA
7	U.S.A.	California		\$3,394.90	\$821.10	OFFICE	DESK
8	U.S.A.	California		\$3,394.90	\$780.30	FURNITURE	SOFA
9	U.S.A.	California		\$3,384.70	\$924.80	OFFICE	DESK
10	U.S.A.	California		\$3,379.60	\$2,864.50	FURNITURE	SOFA
11	U.S.A.	California		\$3,376.20	\$2,284.80	OFFICE	CHAIR
12	U.S.A.	California		\$3,372.80	\$766.50	FURNITURE	SOFA
13	U.S.A.	California		\$3,372.80	\$834.70	FURNITURE	BED
14	U.S.A.	California		\$3,369.40	\$2,823.70	OFFICE	DESK
15	U.S.A.	California		\$3,367.70	\$1,754.40	FURNITURE	BED
16	U.S.A.	California		\$3,359.20	\$1,530.00	FURNITURE	BED
17	U.S.A.	California		\$3,359.20	\$3,105.90	OFFICE	CHAIR
18	U.S.A.	California		\$3,355.80	\$1,110.10	OFFICE	DESK
19	U.S.A.	California		\$3,337.10	\$3,139.90	OFFICE	CHAIR
20	U.S.A.	California		\$3,325.20	\$912.80	OFFICE	CHAIR
21	U.S.A.	California		\$3,323.50	\$2,312.00	OFFICE	DESK

- Right-click any table cell that is not a heading and select **Where** from the menu. The WHERE EXPRESSION window opens.



- In the **Available Columns** list, select a column, and then select an operator from the **Operators** menu.

In this example, **COUNTRY** is selected from the Available Columns list, and **EQ** (equal to) is selected from the Operators menu. Note that the WHERE expression is being built in the **where** box at the bottom of the window.




- In the **Available Columns** list, select another value to complete the WHERE expression.


In this example, **<LOOKUP distinct values>** is selected.

- 5 In the Lookup Distinct Values window that appears, select a value.

In this example, **Canada** is selected.



Note that the complete WHERE expression appears in the **where** box at the bottom of the window.



- 6 Click **OK** to close the WHERE EXPRESSION window.

In this example, VIEWTABLE displays only rows where the value of Country is Canada.

	Country	State/Province	County	Actual Sales	Predicted Sales	Product Type	Product
18433	Canada	British Columbia		\$708.80	\$572.80	FURNITURE	SOFA
18434	Canada	British Columbia		\$1,008.80	\$600.80	FURNITURE	SOFA
18435	Canada	British Columbia		\$391.20	\$1,004.00	FURNITURE	SOFA
18436	Canada	British Columbia		\$350.40	\$1,074.40	FURNITURE	SOFA
18437	Canada	British Columbia		\$1,245.60	\$902.40	FURNITURE	SOFA
18438	Canada	British Columbia		\$332.00	\$842.40	FURNITURE	SOFA
18439	Canada	British Columbia		\$195.20	\$480.80	FURNITURE	SOFA
18440	Canada	British Columbia		\$1,024.00	\$1,201.60	FURNITURE	SOFA
18441	Canada	British Columbia		\$816.00	\$324.00	FURNITURE	SOFA
18442	Canada	British Columbia		\$464.80	\$442.40	FURNITURE	SOFA
18443	Canada	British Columbia		\$1,111.20	\$524.80	FURNITURE	SOFA
18444	Canada	British Columbia		\$510.40	\$591.20	FURNITURE	SOFA
18445	Canada	British Columbia		\$372.00	\$589.60	FURNITURE	BED
18446	Canada	British Columbia		\$624.80	\$404.80	FURNITURE	BED
18447	Canada	British Columbia		\$350.40	\$409.60	FURNITURE	BED
18448	Canada	British Columbia		\$978.40	\$776.00	FURNITURE	BED
18449	Canada	British Columbia		\$90.40	\$772.00	FURNITURE	BED
18450	Canada	British Columbia		\$395.20	\$301.60	FURNITURE	BED
18451	Canada	British Columbia		\$1,084.00	\$802.40	FURNITURE	BED
18452	Canada	British Columbia		\$443.20	\$339.20	FURNITURE	BED
18453	Canada	British Columbia		\$13.60	\$404.80	FURNITURE	BED

Clear the WHERE Expression

You can clear the WHERE expression that you used to subset your data, and redisplay all of the data in the table. To do this, follow these steps:

- 1 Right-click anywhere in the table except in a column heading.
- 2 Select **WHERE Clear** from the pop-up menu.

The VIEWTABLE window removes any existing subsets of data that were created with the WHERE expression, and displays all of the rows of the table.

Exporting a Subset of Data

Overview of Exporting Data

The Export Wizard reads data from a SAS data set and writes it to an external file. You can export SAS data to a variety of formats. The formats that are available depend on your operating environment and the SAS products that you have installed.


Export Data

To export data, follow these steps:

- 1 Select **File ► Export Data** from the File menu.


The Export Wizard - Select library and member window appears.
- 2 Select the SAS data set from which you want to export data.

In this example, **WORK** is selected as the library, and **FURNSALESORTED** is the member name.




- 3 Click **Next** and the Export Wizard - Select export type window appears.
- 4 Select the type of data source to which you want to export files. Note that **Standard data source** is selected by default.

In this example, **Microsoft Excel 97/2000/2002/2003 workbook** is selected.



- 5 Click **Next** to display the Connect to MS Excel window.
- 6 In the **Workbook** field, enter the workbook name that will contain the exported file and then click **OK**.

In this example, **Myworkbook** is entered as the name of the workbook.



- 7 When the Export Wizard - Select table window appears, enter a name for the table you are exporting.

In this example, **Mytable** is the table name.



- 8 Click **Next**.
- 9 If you want SAS to create a file of PROC EXPORT statements for later use, then enter a filename that will contain the SAS statements.



10 Click **Finish** to complete this task.

Importing Data into a Table

Overview of Importing Data

Whether your data is stored in a standard file format or in your own special file format, you can use the Import Wizard to import data into a SAS table. The types of files that you can import depend on your operating environment.

Import a Standard File


To import a standard file, follow these steps:

- 1 Select **File ► Import data** from the toolbar.


The Import Wizard - Select import type window appears.

- 2 Select the type of file that you are importing by selecting a data source from the **Select a data source** menu. Note that **Standard data source** is selected by default.

In this example, **Microsoft Excel 97/2000/2002/2003 Workbook** is selected.



- 3 Click **Next** to continue.
- 4 In the Connect to MS Excel window, enter the pathname of the file that you want to export, and then click **OK**.




- 5 In the Import Wizard - Select table window, enter the name of the table that you want to import.




- 6 Click **Next** to continue.
- 7 In the Import Wizard - Select library and member window that appears, enter a location in which to store the imported file.

In this example, **WORK** is selected as the library, and **FURNSASLESORTED** is selected as the member name.




- 8 Click **Next** to continue.
- 9 If you want SAS to create a file of PROC IMPORT statements for later use, then enter a filename that will contain the SAS statements.



- 10 Click **Finish** to complete this task.


Import a Non-standard File

If your data is not in standard format, you can use the External File Interface (EFI) facility to import data. This tool enables you to define your file format and offers you a range of format options. To use EFI, select User-defined file format in the Import Wizard and follow the directions for describing your data file.



DATA Step Concepts

- Chapter 19*.....**DATA Step Processing** 285
- Chapter 20*.....**Reading Raw Data** 311
- Chapter 21*.....**BY-Group Processing in the DATA Step** 329
- Chapter 22*.....**Reading, Combining, and Modifying SAS Data Sets** 341
- Chapter 23*.....**Using DATA Step Component Objects** 391
- Chapter 24*.....**Array Processing** 429



CHAPTER

19

DATA Step Processing

<i>Why Use a DATA Step?</i>	285
<i>Overview of DATA Step Processing</i>	286
<i>Flow of Action</i>	286
<i>The Compilation Phase</i>	288
<i>The Execution Phase</i>	288
<i>Processing a DATA Step: A Walkthrough</i>	289
<i>Sample DATA Step</i>	289
<i>Creating the Input Buffer and the Program Data Vector</i>	289
<i>Reading a Record</i>	290
<i>Writing an Observation to the SAS Data Set</i>	291
<i>Reading the Next Record</i>	292
<i>When the DATA Step Finishes Executing</i>	293
<i>About DATA Step Execution</i>	294
<i>The Default Sequence of Execution in the DATA Step</i>	294
<i>Changing the Default Sequence of Execution</i>	295
<i>Using Statements to Change the Default Sequence of Execution</i>	295
<i>Using Functions to Change the Default Sequence of Execution</i>	295
<i>Altering the Flow for a Given Observation</i>	295
<i>Step Boundary—How To Know When Statements Take Effect</i>	297
<i>What Causes a DATA Step to Stop Executing</i>	298
<i>About Creating a SAS Data Set with a DATA Step</i>	299
<i>Creating a SAS Data File or a SAS View</i>	299
<i>Sources of Input Data</i>	299
<i>Reading Raw Data: Examples</i>	300
<i>Example 1: Reading External File Data</i>	300
<i>Example 2: Reading Instream Data Lines</i>	300
<i>Example 3: Reading Instream Data Lines with Missing Values</i>	301
<i>Example 4: Using Multiple Input Files in Instream Data</i>	301
<i>Reading Data from SAS Data Sets</i>	302
<i>Generating Data from Programming Statements</i>	302
<i>Writing a Report with a DATA Step</i>	303
<i>Example 1: Creating a Report without Creating a Data Set</i>	303
<i>Example 2: Creating a Customized Report</i>	304
<i>Example 3: Creating an HTML Report Using ODS and the DATA Step</i>	308
<i>The DATA Step and ODS</i>	310

Why Use a DATA Step?

Using the DATA step is the primary method for creating a SAS data set with Base SAS software. A *DATA step* is a group of SAS language statements that begin with a

DATA statement and contains other programming statements that manipulate existing SAS data sets or create SAS data sets from raw data files.

You can use the DATA step for

- creating SAS data sets (SAS data files or SAS views)
- creating SAS data sets from input files that contain raw data (external files)
- creating new SAS data sets from existing ones by subsetting, merging, modifying, and updating existing SAS data sets
- analyzing, manipulating, or presenting your data
- computing the values for new variables
- report writing, or writing files to disk or tape
- retrieving information
- file management


Note: A DATA step creates a *SAS data set*. This data set can be a SAS data file or a SAS view. A SAS data file stores data values while a SAS view stores instructions for retrieving and processing data. When you can use a SAS view as a SAS data file, as is true in most cases, this documentation uses the broader term SAS data set. \triangle

Overview of DATA Step Processing

Flow of Action

When you submit a DATA step for execution, it is first compiled and then executed. The following figure shows the flow of action for a typical SAS DATA step.

Figure 19.1 Flow of Action in the DATA Step



The Compilation Phase

When you submit a DATA step for execution, SAS checks the syntax of the SAS statements and compiles them, that is, automatically translates the statements into machine code. In this phase, SAS identifies the type and length of each new variable, and determines whether a variable type conversion is necessary for each subsequent reference to a variable. During the compile phase, SAS creates the following three items:

input buffer is a logical area in memory into which SAS reads each record of raw data when SAS executes an INPUT statement. Note that this buffer is created only when the DATA step reads raw data. (When the DATA step reads a SAS data set, SAS reads the data directly into the program data vector.)

program data vector (PDV) is a logical area in memory where SAS builds a data set, one observation at a time. When a program executes, SAS reads data values from the input buffer or creates them by executing SAS language statements. The data values are assigned to the appropriate variables in the program data vector. From here, SAS writes the values to a SAS data set as a single observation.

Along with data set variables and computed variables, the PDV contains two automatic variables, `_N_` and `_ERROR_`. The `_N_` variable counts the number of times the DATA step begins to iterate. The `_ERROR_` variable signals the occurrence of an error caused by the data during execution. The value of `_ERROR_` is either 0 (indicating no errors exist), or 1 (indicating that one or more errors have occurred). SAS does not write these variables to the output data set.

descriptor information is information that SAS creates and maintains about each SAS data set, including data set attributes and variable attributes. It contains, for example, the name of the data set and its member type, the date and time that the data set was created, and the number, names and data types (character or numeric) of the variables.

The Execution Phase

By default, a simple DATA step iterates once for each observation that is being created. The flow of action in the Execution Phase of a simple DATA step is described as follows:

- 1 The DATA step begins with a DATA statement. Each time the DATA statement executes, a new iteration of the DATA step begins, and the `_N_` automatic variable is incremented by 1.
- 2 SAS sets the newly created program variables to missing in the program data vector (PDV).
- 3 SAS reads a data record from a raw data file into the input buffer, or it reads an observation from a SAS data set directly into the program data vector. You can use an INPUT, MERGE, SET, MODIFY, or UPDATE statement to read a record.
- 4 SAS executes any subsequent programming statements for the current record.
- 5 At the end of the statements, an output, return, and reset occur automatically. SAS writes an observation to the SAS data set, the system automatically returns to the top of the DATA step, and the values of variables created by INPUT and assignment statements are reset to missing in the program data vector. Note that

variables that you read with a SET, MERGE, MODIFY, or UPDATE statement are not reset to missing here.

- 6 SAS counts another iteration, reads the next record or observation, and executes the subsequent programming statements for the current observation.
- 7 The DATA step terminates when SAS encounters the end-of-file in a SAS data set or a raw data file.

Note: The figure shows the default processing of the DATA step. You can place data-reading statements (such as INPUT or SET), or data-writing statements (such as OUTPUT), in any order in your program. Δ

Processing a DATA Step: A Walkthrough

Sample DATA Step

The following statements provide an example of a DATA step that reads raw data, calculates totals, and creates a data set:

```
data total_points (drop=TeamName); ❶
  input TeamName $ ParticipantName $ Event1 Event2 Event3; ❷
  TeamTotal + (Event1 + Event2 + Event3); ❸
  datalines;
Knights Sue      6  8  8
Kings Jane       9  7  8
Knights John     7  7  7
Knights Lisa     8  9  9
Knights Fran     7  6  6
Knights Walter  9  8 10
;
```

- ❶ The DROP= data set option prevents the variable TeamName from being written to the output SAS data set called TOTAL_POINTS.
- ❷ The INPUT statement describes the data by giving a name to each variable, identifying its data type (character or numeric), and identifying its relative location in the data record.
- ❸ The SUM statement accumulates the scores for three events in the variable TeamTotal.

Creating the Input Buffer and the Program Data Vector

When DATA step statements are compiled, SAS determines whether to create an input buffer. If the input file contains raw data (as in the example above), SAS creates an input buffer to hold the data before moving the data to the program data vector (PDV). (If the input file is a SAS data set, however, SAS does not create an input buffer. SAS writes the input data directly to the PDV.)

The PDV contains all the variables in the input data set, the variables created in DATA step statements, and the two variables, _N_ and _ERROR_, that are automatically generated for every DATA step. The _N_ variable represents the number

of times the DATA step has iterated. The `_ERROR_` variable acts like a binary switch whose value is 0 if no errors exist in the DATA step, or 1 if one or more errors exist. The following figure shows the Input Buffer and the program data vector after DATA step compilation.

Figure 19.2 Input Buffer and Program Data Vector

Input Buffer																			
1										2									
1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0

Program Data Vector							
TeamName	ParticipantName	Event1	Event2	Event3	TeamTotal	_N_	_ERROR_
		.	.	.	0	1	0
Drop						Drop	Drop

Variables that are created by the INPUT and the Sum statements (TeamName, ParticipantName, Event1, Event2, Event3, and TeamTotal) are set to missing initially. Note that in this representation, numeric variables are initialized with a period and character variables are initialized with blanks. The automatic variable `_N_` is set to 1; the automatic variable `_ERROR_` is set to 0.

The variable TeamName is marked Drop in the PDV because of the `DROP=` data set option in the DATA statement. Dropped variables are not written to the SAS data set. The `_N_` and `_ERROR_` variables are dropped because automatic variables created by the DATA step are not written to a SAS data set. See Chapter 4, “SAS Variables,” on page 25 for details about automatic variables.

Reading a Record

SAS reads the first data line into the input buffer. The input *pointer*, which SAS uses to keep its place as it reads data from the input buffer, is positioned at the beginning of the buffer, ready to read the data record. The following figure shows the position of the input pointer in the input buffer before SAS reads the data.

Figure 19.3 Position of the Pointer in the Input Buffer Before SAS Reads Data


Input Buffer																			
1										2									
1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0
K	n	i	g	h	t	s		S	u	e					6			8	

↑

The INPUT statement then reads *data values* from the record in the input buffer and writes them to the PDV where they become *variable values*. The following figure shows

both the position of the pointer in the input buffer, and the values in the PDV after SAS reads the first record.

Figure 19.4 Values from the First Record are Read into the Program Data Vector



Program Data Vector

TeamName	ParticipantName	Event1	Event2	Event3	TeamTotal	_N_	_ERROR_
Knights	Sue	6	8	8	0	1	0
Drop						Drop	Drop

After the INPUT statement reads a value for each variable, SAS executes the Sum statement. SAS computes a value for the variable TeamTotal and writes it to the PDV. The following figure shows the PDV with all of its values before SAS writes the observation to the data set.

Figure 19.5 Program Data Vector with Computed Value of the Sum Statement

Program Data Vector

TeamName	ParticipantName	Event1	Event2	Event3	TeamTotal	_N_	_ERROR_
Knights	Sue	6	8	8	22	1	0
Drop						Drop	Drop

Writing an Observation to the SAS Data Set

When SAS executes the last statement in the DATA step, all values in the PDV, except those marked to be dropped, are written as a single observation to the data set TOTAL_POINTS. The following figure shows the first observation in the TOTAL_POINTS data set.

Figure 19.6 The First Observation in Data Set TOTAL_POINTS

Output SAS Data Set TOTAL_POINTS: 1st observation

ParticipantName	Event1	Event2	Event3	TeamTotal
Sue	6	8	8	22

SAS then returns to the DATA statement to begin the next iteration. SAS resets the values in the PDV in the following way:

- The values of variables created by the INPUT statement are set to missing.
- The value created by the Sum statement is automatically retained.
- The value of the automatic variable `_N_` is incremented by 1, and the value of `_ERROR_` is reset to 0.

The following figure shows the current values in the PDV.

Figure 19.7 Current Values in the Program Data Vector

Program Data Vector

TeamName	ParticipantName	Event1	Event2	Event3	TeamTotal	_N_	_ERROR_
		.	.	.	22	2	0
Drop						Drop	Drop

Reading the Next Record

SAS reads the next record into the input buffer. The INPUT statement reads the data values from the input buffer and writes them to the PDV. The Sum statement adds the values of Event1, Event2, and Event3 to TeamTotal. The value of 2 for variable `_N_` indicates that SAS is beginning the second iteration of the DATA step. The following figure shows the input buffer, the PDV for the second record, and the SAS data set with the first two observations.

Figure 19.8 Input Buffer, Program Data Vector, and First Two Observations

Input Buffer

1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5
C	a	r	d	i	n	a	l	s		J	a	n	e		9		7		8					

↑

Program Data Vector

TeamName	ParticipantName	Event1	Event2	Event3	TeamTotal	_N_	_ERROR_
Cardinals	Jane	9	7	8	46	2	0
Drop						Drop	Drop

Output SAS Data Set TOTAL_POINTS: 1st and 2nd observations

ParticipantName	Event1	Event2	Event3	TeamTotal
Sue	6	8	8	22
Jane	9	7	8	46

As SAS continues to read records, the value in TeamTotal grows larger as more participant scores are added to the variable. _N_ is incremented at the beginning of each iteration of the DATA step. This process continues until SAS reaches the end of the input file.

When the DATA Step Finishes Executing

The DATA step stops executing after it processes the last input record. You can use PROC PRINT to print the output in the TOTAL_POINTS data set:

Output 19.1 Output from the Walkthrough DATA Step

Total Team Scores						1
Obs	Participant Name	Event1	Event2	Event3	Team Total	
1	Sue	6	8	8	22	
2	Jane	9	7	8	46	
3	John	7	7	7	67	
4	Lisa	8	9	9	93	
5	Fran	7	6	6	112	
6	Walter	9	8	10	139	

About DATA Step Execution

The Default Sequence of Execution in the DATA Step

The following table outlines the default sequence of execution for statements in a DATA step. The DATA statement begins the step and identifies usually one or more SAS data sets that the step will create. (You can use the keyword `_NULL_` as the data set name if you do not want to create an output data set.) Optional programming statements process your data. SAS then performs the default actions at the end of processing an observation.

Table 19.1 Default Execution for Statements in a DATA Step

Structure of a DATA Step	Action
DATA statement	begins the step counts iterations
Data-reading statements: *	
INPUT	describes the arrangement of values in the input data record from a raw data source
SET	reads an observation from one or more SAS data sets
MERGE	joins observations from two or more SAS data sets into a single observation
MODIFY	replaces, deletes, or appends observations in an existing SAS data set in place
UPDATE	updates a master file by applying transactions
Optional SAS programming statements, for example:	further processes the data for the current observation
FirstQuarter=Jan+Feb+Mar; if RetailPrice < 500;	computes the value for FirstQuarter for the current observation subsets by value of variable RetailPrice for the current observation
Default actions at the end of processing an observation	
At end of DATA step:	writes an observation to a SAS data set
Automatic write, automatic return	returns to the DATA statement
At top of DATA step:	
Automatic reset	resets values to missing in program data vector

* The table shows the default processing of the DATA step. You can alter the sequence of statements in the DATA step. You can code optional programming statements, such as creating or reinitializing a constant, before you code a data-reading statement.

Note: You can also use functions to read and process data. For information about how statements and functions process data differently, see “Using Functions to

Manipulate Files” in *SAS Language Reference: Dictionary*. For specific information about SAS functions, see the SAS I/O Files and External Files categories in “Functions and CALL Routines by Category” in *SAS Language Reference: Dictionary*. △

Changing the Default Sequence of Execution

Using Statements to Change the Default Sequence of Execution

You can change the default sequence of execution to control how your program executes. SAS language statements offer you a lot of flexibility to do this in a DATA step. The following list shows the most common ways to control the flow of execution in a DATA step program.

Table 19.2 Common Methods That Alter the Sequence of Execution

Task	Possible Methods
Read a record	merge, modify, join data sets read multiple records to create a single observation randomly select records for processing read from multiple external files read selected fields from a record by using statement or data set options
Process data	use conditional logic retain variable values
Write an observation	write to a SAS data set or to an external file control when output is written to a data set write to multiple files

For more information, see the individual statements in *SAS Language Reference: Dictionary*.

Using Functions to Change the Default Sequence of Execution

You can also use functions to read and process data. For information about how statements and functions process data differently, see “Using Functions to Manipulate Files” in *SAS Language Reference: Dictionary*. For specific information about SAS functions, see the SAS I/O Files and External Files categories in “Functions and CALL Routines by Category” in *SAS Language Reference: Dictionary*.

Altering the Flow for a Given Observation

You can use statements, statement options, and data set options to alter the way SAS processes specific observations. The following table lists SAS language elements and their effects on processing.

Table 19.3 Language Elements That Alter Programming Flow

SAS Language Element	Function
subsetting IF statement	stops the current iteration when a condition is false, does not write the current observation to the data set, and returns control to the top of the DATA step.
IF-THEN/ELSE statement	executes a SAS statement for observations that meet the current condition and continues with the next statement.
DO loops	cause parts of the DATA step to be executed multiple times.
LINK and RETURN statements	alter the flow of control, execute statements following the label specified, and return control of the program to the next statement following the LINK statement.
HEADER= option in the FILE statement	alters the flow of control whenever a PUT statement causes a new page of output to begin; statements following the label specified in the HEADER= option are executed until a RETURN statement is encountered, at which time control returns to the point from which the HEADER= option was activated.
GO TO statement	alters the flow of execution by branching to the label that is specified in the GO TO statement. SAS executes subsequent statements then returns control to the beginning of the DATA step.
EOF= option in an INFILE statement	alters the flow of execution when the end of the input file is reached; statements following the label that is specified in the EOF= option are executed at that time.
N automatic variable in an IF-THEN construct	causes parts of the DATA step to execute only for particular iterations.
SELECT statement	conditionally executes one of a group of SAS statements.
OUTPUT statement in an IF-THEN construct	outputs an observation before the end of the DATA step, based on a condition; prevents automatic output at the bottom of the DATA step.
DELETE statement in an IF-THEN construct	deletes an observation based on a condition and causes a return to the top of the DATA step.

SAS Language Element	Function
ABORT statement in an IF-THEN construct	stops execution of the DATA step and instruct SAS to resume execution with the next DATA or PROC step. It can also stop executing a SAS program altogether, depending on the options specified in the ABORT statement and on the method of operation.
WHERE statement or WHERE= data set option	causes SAS to read certain observations based on one or more specified criteria.

Step Boundary—How To Know When Statements Take Effect

Understanding step boundaries is an important concept in SAS programming because step boundaries determine when SAS statements take effect. SAS executes program statements only when SAS crosses a default or a step boundary. Consider the following DATA steps:

```
data _null_; ❶
  set allscores(drop=score5-score7);
  title 'Student Test Scores'; ❷

data employees; ❸
  set employee_list;
run;
```

- ❶ The DATA statement begins a DATA step and is a step boundary.
- ❷ The TITLE statement is in effect for both DATA steps because it appears before the boundary of the first DATA step. (Because the TITLE statement is a global statement,
- ❸ The DATA statement is the default boundary for the first DATA step.

The TITLE statement in this example is in effect for the first DATA step as well as for the second because the TITLE statement appears before the boundary of the first DATA step. This example uses the default step boundary **data employees;**.

The following example shows an OPTIONS statement inserted after a RUN statement.

```
data scores; ❶
  set allscores(drop=score5-score7);
run; ❷

options firstobs=5 obs=55; ❸

data test;
  set alltests;
run;
```

The OPTIONS statement specifies that the first observation that is read from the input data set should be the 5th, and the last observation that is read should be the 55th. Inserting a RUN statement immediately before the OPTIONS statement causes the first DATA step to reach its boundary (**run;**) before SAS encounters the OPTIONS statement. The OPTIONS statement settings, therefore, are put into effect for the second DATA step only.

- ❶ The DATA statement is a step boundary.
- ❷ The RUN statement is the boundary for the first DATA step.
- ❸ The OPTIONS statement affects the second DATA step only.

Following the statements in a DATA step with a RUN statement is the simplest way to make the step begin to execute, but a RUN statement is not always necessary. SAS recognizes several step boundaries for a SAS step:

- another DATA statement
- a PROC statement
- a RUN statement.

Note: For SAS programs executed in interactive mode, a RUN statement is required to signal the step boundary for the last step you submit. Δ

- the semicolon (with a DATALINES or CARDS statement) or four semicolons (with a DATALINES4 or CARDS4 statement) after data lines
- an ENDSAS statement
- in noninteractive or batch mode, the end of a program file containing SAS programming statements
- a QUIT statement (for some procedures).

When you submit a DATA step during interactive processing, it does not begin running until SAS encounters a step boundary. This fact enables you to submit statements as you write them while preventing a step from executing until you have entered all the statements.

What Causes a DATA Step to Stop Executing

DATA steps stop executing under different circumstances, depending on the type and number of sources of input.

Table 19.4 Causes that Stop DATA Step Execution

Data Read	Data Source	SAS Statements	DATA Step Stops
no data			after only one iteration
any data			when it executes STOP or ABORT when the data is exhausted
raw data	instream data lines	INPUT statement	after the last data line is read
	one external file	INPUT and INFILE statements	when end-of-file is reached
	multiple external files	INPUT and INFILE statements	when end-of-file is first reached on any of the files
observations sequentially	one SAS data set	SET and MODIFY statements	after the last observation is read

Data Read	Data Source	SAS Statements	DATA Step Stops
	multiple SAS data sets	one SET, MERGE, MODIFY, or UPDATE statement	when all input data sets are exhausted
	multiple SAS data sets	multiple SET, MERGE, MODIFY, or UPDATE statements	when end-of-file is reached by any of the data-reading statements

A DATA step that reads observations from a SAS data set with a SET statement that uses the POINT= option has no way to detect the end of the input SAS data set. (This method is called direct or random access.) Such a DATA step usually requires a STOP statement.

A DATA step also stops when it executes a STOP or an ABORT statement. Some system options and data set options, such as OBS=, can cause a DATA step to stop earlier than it would otherwise.

About Creating a SAS Data Set with a DATA Step

Creating a SAS Data File or a SAS View

You can create either a SAS data file, a data set that holds actual data, or a SAS view, a data set that references data that is stored elsewhere. By default, you create a SAS data file. To create a SAS view instead, use the VIEW= option in the DATA statement. With a SAS view you can, for example, process monthly sales figures without having to edit your DATA step. Whenever you need to create output, the output from a SAS view reflects the current input data values.

The following DATA statement creates a SAS view called MONTHLY_SALES.

```
data monthly_sales / view=monthly_sales;
```

The following DATA statement creates a data file called TEST_RESULTS.

```
data test_results;
```

Sources of Input Data

You select data-reading statements based on the source of your input data. There are at least six sources of input data:

- raw data in an external file
- raw data in the jobstream (instream data)
- data in SAS data sets
- data that is created by programming statements
- data that you can remotely access through an FTP protocol, TCP/IP socket, a SAS catalog entry, or through a URL
- data that is stored in a Database Management System (DBMS) or other vendor's data files.

Usually DATA steps read input data records from only one of the first three sources of input. However, DATA steps can use a combination of some or all of the sources.

Reading Raw Data: Examples

Example 1: Reading External File Data

The components of a DATA step that produce a SAS data set from raw data stored in an external file are outlined here.

```
data weight; ❶
  infile 'your-input-file'; ❷
  input IDnumber $ Week1 Week16; ❸
  WeightLoss=Week1-Week16; ❹
run; ❺

proc print data=weight; ❻
run; ❼
```

- ❶ Begin the DATA step and create a SAS data set called WEIGHT.
- ❷ Specify the external file that contains your data.
- ❸ Read a record and assign values to three variables.
- ❹ Calculate a value for variable WeightLoss.
- ❺ Execute the DATA step.
- ❻ Print data set WEIGHT using the PRINT procedure.
- ❼ Execute the PRINT procedure.

Example 2: Reading Instream Data Lines

This example reads raw data from instream data lines.

```
data weight2; ❶
  input IDnumber $ Week1 Week16; ❷
  WeightLoss2=Week1-Week16; ❸
  datalines; ❹
2477 195 163
2431 220 198
2456 173 155
2412 135 116
; ❺

proc print data=weight2; ❻
run; ❼
```

- ❶ Begin the DATA step and create SAS data set WEIGHT2.
- ❷ Read a data line and assign values to three variables.
- ❸ Calculate a value for variable WeightLoss2.
- ❹ Begin the data lines.
- ❺ Signal end of data lines with a semicolon and execute the DATA step.

- ⑥ Print data set WEIGHT2 using the PRINT procedure.
- ⑦ Execute the PRINT procedure.

Example 3: Reading Instream Data Lines with Missing Values

You can also take advantage of options in the INFILE statement when you read instream data lines. This example shows the use of the MISSEVER option, which assigns missing values to variables for records that contain no data for those variables.

```
data weight2;
  infile datalines missover; ①
  input IDnumber $ Week1 Week16;
  WeightLoss2=Week1-Week16;
  datalines; ②
2477 195 163
2431
2456 173 155
2412 135 116
; ③

proc print data=weight2; ④
run; ⑤
```

- ① Use the MISSEVER option to assign missing values to variables that do not contain values in records that do not satisfy the current INPUT statement.
- ② Begin data lines.
- ③ Signal end of data lines and execute the DATA step.
- ④ Print data set WEIGHT2 using the PRINT procedure.
- ⑤ Execute the PRINT procedure.

Example 4: Using Multiple Input Files in Instream Data

This example shows how to use multiple input files as instream data to your program. This example reads the records in each file and creates the ALL_ERRORS SAS data set. The program then sorts the observations by Station, and creates a sorted data set called SORTED_ERRORS. The print procedure prints the results.

```
options pageno=1 nodate linesize=60 pagesize=80;

data all_errors;
  length filelocation $ 60;
  input filelocation; /* reads instream data */
  infile daily filevar=filelocation
          filename=daily end=done;
  do while (not done);
    input Station $ Shift $ Employee $ NumberOfFlaws;
    output;
  end;
  put 'Finished reading ' daily=;
  datalines;
. . .myfile_A. . .
. . .myfile_B. . .
```

```

. . .myfile_C. . .
;

proc sort data=all_errors out=sorted_errors;
  by Station;
run;

proc print data = sorted_errors;
  title 'Flaws Report sorted by Station';
run;

```

Output 19.2 Multiple Input Files in Instream Data

Flaws Report sorted by Station					1
Obs	Station	Shift	Employee	Number OfFlaws	
1	Amherst	2	Lynne	0	
2	Goshen	2	Seth	4	
3	Hadley	2	Jon	3	
4	Holyoke	1	Walter	0	
5	Holyoke	1	Barb	3	
6	Orange	2	Carol	5	
7	Otis	1	Kay	0	
8	Pelham	2	Mike	4	
9	Stanford	1	Sam	1	
10	Suffield	2	Lisa	1	

Reading Data from SAS Data Sets

This example reads data from one SAS data set, generates a value for a new variable, and creates a new data set.

```

data average_loss; ❶
  set weight;      ❷
  Percent=round((AverageLoss * 100) / Week1); ❸
run;               ❹

```

- ❶ Begin the DATA step and create a SAS data set called AVERAGE_LOSS.
- ❷ Read an observation from SAS data set WEIGHT.
- ❸ Calculate a value for variable Percent.
- ❹ Execute the DATA step.

Generating Data from Programming Statements

You can create data for a SAS data set by generating observations with programming statements rather than by reading data. A DATA step that reads no input goes through only one iteration.

```

data investment; ❶
  begin='01JAN1990'd;
  end='31DEC2009'd;

```

```

do year=year(begin) to year(end); ❷
    Capital+2000 + .07*(Capital+2000);
    output; ❸
end;
put 'The number of DATA step iterations is '_n_'; ❹
run; ❺

proc print data=investment; ❻
    format Capital dollar12.2; ❼
run; ❽

```

- ❶ Begin the DATA step and create a SAS data set called INVESTMENT.
- ❷ Calculate a value based on a \$2,000 capital investment and 7% interest each year from 1990 to 2009. Calculate variable values for one observation per iteration of the DO loop.
- ❸ Write each observation to data set INVESTMENT.
- ❹ Write a note to the SAS log proving that the DATA step iterates only once.
- ❺ Execute the DATA step.
- ❻ To see your output, print the INVESTMENT data set with the PRINT procedure.
- ❼ Use the FORMAT statement to write numeric values with dollar signs, commas, and decimal points.
- ❽ Execute the PRINT procedure.

Writing a Report with a DATA Step

Example 1: Creating a Report without Creating a Data Set

You can use a DATA step to generate a report without creating a data set by using `_NULL_` in the DATA statement. This approach saves system resources because SAS does not create a data set. The report can contain both TITLE statements and FOOTNOTE statements. If you use a FOOTNOTE statement, be sure to include FOOTNOTE as an option in the FILE statement in the DATA step.

```

title1 'Budget Report'; ❶
title2 'Mid-Year Totals by Department';
footnote 'compiled by Manager,
Documentation Development Department'; ❷

data _null_; ❸
    set budget; ❹
    file print footnote; ❺
    MidYearTotal=Jan+Feb+Mar+Apr+May+Jun; ❻
    if _n_=1 then ❼
        do;
            put @5 'Department' @30 'Mid-Year Total';
        end;
    put @7 Department @35 MidYearTotal; ❽
run; ❾

```

- ❶ Define titles.
- ❷ Define the footnote.
- ❸ Begin the DATA step. `_NULL_` specifies that no data set will be created.
- ❹ Read one observation per iteration from data set BUDGET.
- ❺ Name the output file for the PUT statements and use the PRINT fileref. By default, the PRINT fileref specifies that the file will contain carriage-control characters and titles. The FOOTNOTE option specifies that each page of output will contain a footnote.
- ❻ Calculate a value for the variable MidYearTotal on each iteration.
- ❼ Write variable name headings for the report on the first iteration only.
- ❽ Write the current values of variables Department and MidYearTotal for each iteration.
- ❾ Execute the DATA step.

The example above uses the FILE statement with the PRINT fileref to produce listing output. If you want to print to a file, specify a fileref or a complete filename. Use the PRINT option if you want the file to contain carriage-control characters and titles. The following example shows how to use the FILE statement in this way.

```
file 'external-file' footnote print;
```

You can also use the `data _null_;` statement to write to an external file. For more information about writing to external files, see the FILE statement in *SAS Language Reference: Dictionary*, and the SAS documentation for your operating environment.

Example 2: Creating a Customized Report

You can create very detailed, fully customized reports by using a DATA step with PUT statements. The following example shows a customized report that contains three distinct sections: a header, a table, and a footer. It contains existing SAS variable values, constant text, and values that are calculated as the report is written.

Output 19.3 Sample of a Customized Report

Expense Report										
Around The World Retailers										
EMPLOYEE BUSINESS, TRAVEL, AND TRAINING EXPENSE REPORT										
Employee Name: ALEJANDRO MARTINEZ			Destination: CARY, NC			Departure Date: 11JUL2007				
Department: SALES & MARKETING			Purpose of Trip/Activity: MARKETING TRAINING			Return Date: 16JUL2007				
Trip ID#: 93-0002519						Activity from: 11JUL2007				
						to: 16JUL2007				
EXPENSE DETAIL	SUN 07/11	MON 07/12	TUE 07/13	WED 07/14	THU 07/15	FRI 07/16	SAT 07/17	TOTALS	PAID BY COMPANY	PAID BY EMPLOYEE
Lodging, Hotel	92.96	92.96	92.96	92.96	92.96			464.80	464.80	
Telephone	4.57	4.73						9.30		9.30
Personal Auto 36 miles @.28/mile	5.04					5.04		10.08		10.08
Car Rental, Taxi, Parking, Tolls		35.32	35.32	35.32	35.32	35.32		176.60	176.60	
Airlines, Bus, Train (Attach Stub)	485.00					485.00		970.00	970.00	
Dues										
Registration Fees	75.00							75.00		75.00
Other (explain below)						5.00		5.00		5.00
Tips (excluding meal tips)	3.00					3.00		6.00		6.00
Meals										
Breakfast						7.79		7.79		7.79
Lunch										
Dinner	36.00	28.63	36.00	36.00	30.00			166.63		166.63
Business Entertainment										
TOTAL EXPENSES	641.57	176.64	179.28	179.28	173.28	541.15		1891.20	1611.40	279.80
Travel Advance to Employee										\$0.00
Reimbursement due Employee (or ATWR)										\$279.80
Other: (i.e. miscellaneous expenses and/or names of employees sharing receipt.)										
16JUL1999 CAR RENTAL INCLUDE \$5.00 FOR GAS										
APPROVED FOR PAYMENT BY: Authorizing Manager: _____ Emp. # _____										
Employee Signature: _____ Emp. # 1118										
Charge to Division: ATW Region: TX Dept: MKT Acct: 6003 Date: 27JUL2007										

The code shown below generates the report example (you must create your own input data). It is beyond the scope of this discussion to fully explain the code that generated the report example. For a complete explanation of this example, see the *SAS Guide to Report Writing: Examples*.


```

        airlin1-airlin10 dues1-dues10
        regfeel-regfeel10 other1-other10
        tips1-tips10 meals1-meals10
        bkfst1-bkfst10 lunch1-lunch10
        dinner1-dinner10 busent1-busent10
        total1-total10;
array misc{8} $ misc1-misc8;
array mday{7} mday1-mday7;
dptday=weekday(dptdate);
mday{dptday}=dptdate;
if dptday>1 then
    do dayofwk=1 to (dptday-1);
        mday{dayofwk}=dptdate-(dptday-dayofwk);
    end;
if dptday<7 then
    do dayofwk=(dptday+1) to 7;
        mday{dayofwk}=dptdate+(dayofwk-dptday);
    end;
if rptdate=. then rptdate="&sysdate9"d;

tripnum=substr(tripid,4,2)||'-'||substr(scan(tripid,1),6);

put // @1 'Around The World Retailers' //

        @1 'EMPLOYEE BUSINESS, TRAVEL, AND TRAINING EXPENSE REPORT' ///

        @1 'Employee Name: ' @16 empname
        @44 'Destination: ' @57 dest
        @106 'Departure Date:' @122 dptdate /

        @4 'Department: ' @16 dept
        @44 'Purpose of Trip/Activity: ' @70 purpose
        @109 'Return Date:' @122 rtrndate /

        @6 'Trip ID#: ' @16 tripnum
        @107 'Activity from:' @122 actdate1 /

        @118 'to:' @122 actdate2 //
        @1 '+-----+-----+-----+-----+-----+'
            '-----+-----+-----+-----+-----+' /

        @1 '|
            ' TUE | WED | THU | FRI | SAT | MON |
            ' PAID BY PAID BY' /

        @1 '| EXPENSE DETAIL
            ' | ' mday1 mmddyy5. ' | ' mday2 mmddyy5.
            ' | ' mday3 mmddyy5. ' | ' mday4 mmddyy5.
            ' | ' mday5 mmddyy5. ' | ' mday6 mmddyy5.
            ' | ' mday7 mmddyy5.
        @100 '| TOTALS | COMPANY EMPLOYEE' ;
do i=1 to 15;

        if i=1 or i=10 or i=15 then

```

```

put @1 '|-----|-----|-----|-----|-----|-----|';
      '-----|-----|-----|-----|-----|-----|';
if i=3 then
  put @1 '| ' i category. @16 automile 4.0 @21 'miles @'
      @28 permile 3.2 @31 '/mile' @37 '| ' @;
  else put @1 '| ' i category. @37 '| ' @;
col=38;
do j=1 to 10;
  if j<9 then put @col expenses{i,j} blanks8. '| ' @;
  else if j=9 then put @col expenses{i,j} blanks8. @;
  else put @col expenses{i,j} blanks8.;
  col+9;
  if j=8 then col+2;
end;
end;
Put @1 '+-----+-----+-----+-----+-----+-----+'
      '-----+-----+-----+-----+-----+-----+' //

@1 'Travel Advance to Employee .....
  '.....
@121 empadv dollar8.2 //

@1 'Reimbursement due Employee (or ATWR) .....
  '.....
@121 reimburs dollar8.2 //

@1 'Other: (i.e. miscellaneous expenses and/or names of '
  'employees sharing receipt.)' //;
do j=1 to 8;
  put @1 misc{j} ;
end;
put / @1 'APPROVED FOR PAYMENT BY: Authorizing Manager:'
@48 ' _____'
@100 'Emp. # _____' ///

@27 'Employee Signature:'
@48 ' _____'
@100 'Emp. # ' empid ///

@6 'Charge to Division:' @26 div $cuscore.
@39 'Region:' @48 region $cuscore.
@59 'Dept:' @66 deptchg $cuscore.
@79 'Acct:' @86 acct nuscore.
@100 'Date:' @107 rptdate /
_page_;

run;

```

Example 3: Creating an HTML Report Using ODS and the DATA Step

```

options nodate pageno=1 linesize=64 pagesize=60;
ods html body='your_file.html';

```

```

title 'Leading Grain Producers';
title2 'for 1996';

proc format;
  value $cntry 'BRZ'='Brazil'
              'CHN'='China'
              'IND'='India'
              'INS'='Indonesia'
              'USA'='United States';

run;

data _null_;
  length Country $ 3 Type $ 5;
  input Year country $ type $ Kilotons;
  format country $cntry.;
  label type='Grain';

file print ods=(variables=(country
                          type
                          kilotons));

put _ods_;

  datalines;
1996 BRZ  Wheat    3302
1996 BRZ  Rice     10035
1996 BRZ  Corn     31975
1996 CHN  Wheat   109000
1996 CHN  Rice    190100
1996 CHN  Corn    119350
1996 IND  Wheat    62620
1996 IND  Rice    120012
1996 IND  Corn     8660
1996 INS  Wheat    .
1996 INS  Rice     51165
1996 INS  Corn     8925
1996 USA  Wheat    62099
1996 USA  Rice     7771
1996 USA  Corn    236064
;

run;

ods html close;

```

Display 19.1 HTML File Produced by ODS


**Leading Grain Producers
for 1996**

Country	Grain	Kilotons
Brazil	Wheat	3302
Brazil	Rice	10035
Brazil	Corn	31975
China	Wheat	109000
China	Rice	190100
China	Corn	119350
India	Wheat	62620
India	Rice	120012
India	Corn	8660
Indonesia	Wheat	.
Indonesia	Rice	51165
Indonesia	Corn	8925
United States	Wheat	62099
United States	Rice	7771
United States	Corn	236064

The DATA Step and ODS

The Output Delivery System (ODS) is a method of delivering output in a variety of formats and making these formats easy to access. ODS provides templates that define the structure of the output from DATA steps and from PROC steps. The DATA step enables you to use the ODS option in a FILE statement and in a PUT statement.

ODS combines raw data with one or more templates to produce several types of output called output objects. Output objects are sent to “destinations” such as the output destination, the listing destination, the printer destination, or Hypertext Markup Language (HTML). For more information, see “Output Delivery System” on page 128. For complete information about ODS, see the *SAS Output Delivery System: User’s Guide*.



CHAPTER 20

Reading Raw Data

<i>Definition of Reading Raw Data</i>	311
<i>Ways to Read Raw Data</i>	312
<i>Types of Data</i>	312
<i>Definitions</i>	312
<i>Numeric Data</i>	313
<i>Character Data</i>	314
<i>Sources of Raw Data</i>	315
<i>Instream Data</i>	315
<i>Instream Data Containing Semicolons</i>	316
<i>External Files</i>	316
<i>Reading Raw Data with the INPUT Statement</i>	316
<i>Choosing an Input Style</i>	316
<i>List Input</i>	317
<i>Modified List Input</i>	317
<i>Column Input</i>	318
<i>Formatted Input</i>	319
<i>Named Input</i>	319
<i>Additional Data-Reading Features</i>	320
<i>How SAS Handles Invalid Data</i>	321
<i>Reading Missing Values in Raw Data</i>	322
<i>Representing Missing Values in Input Data</i>	322
<i>Special Missing Values in Numeric Input Data</i>	322
<i>Reading Binary Data</i>	323
<i>Definitions</i>	323
<i>Using Binary Informats</i>	324
<i>Reading Column-Binary Data</i>	325
<i>Definition</i>	325
<i>How to Read Column-Binary Data</i>	325
<i>Description of Column-Binary Data Storage</i>	326

Definition of Reading Raw Data

raw data

is unprocessed data that has not been read into a SAS data set. You can use a DATA step to read raw data into a SAS data set from two sources:

- instream data
- an external file.

Note: Raw data does not include Database Management System (DBMS) files. You must license SAS/ACCESS software to access data stored in DBMS files. See Chapter

32, “About SAS/ACCESS Software,” on page 561 for more information about SAS/ACCESS features. \triangle

Ways to Read Raw Data

You can read raw data by using one of the following items:

- SAS statements
- SAS functions
- External File Interface (EFI)
- Import Wizard

When you read raw data with a DATA step, you can use a combination of the INPUT, DATALINES, and INFILE statements. SAS automatically reads your data when you use these statements. For more information on these statements, see “Reading Raw Data with the INPUT Statement” on page 316.

You can also use SAS functions to manipulate external files and to read records of raw data. These functions provide more flexibility in handling raw data. For a description of available functions, see the External File and SAS File I/O categories in “Functions and CALL Routines by Category” in *SAS Language Reference: Dictionary*. For further information about how statements and functions manipulate files differently, see “Functions and CALL Routines” in *SAS Language Reference: Dictionary*.

If your operating environment supports a graphical user interface, you can use the EFI or the Import Wizard to read raw data. The EFI is a point-and-click graphical interface that you can use to read and write data that is not in SAS software’s internal format. By using EFI, you can read data from an external file and write it to a SAS data set, and you can read data from a SAS data set and write it to an external file. See *SAS/ACCESS Interface to PC Files: Reference* for more information about EFI.

Note: If the data file that you are passing to EFI is password protected, you will be prompted multiple times for your login ID and password. \triangle

The Import Wizard guides you through the steps to read data from an external data source and write it to a SAS data set. As a wizard, it is a series of windows that present simple choices to guide you through a process. See *SAS/ACCESS Interface to PC Files: Reference* for more information on the wizard.

Operating Environment Information: Using external files with your SAS jobs requires that you specify filenames with syntax that is appropriate to your operating environment. See the SAS documentation for your operating environment for more information. \triangle

Types of Data

Definitions

data values
are character or numeric values.

numeric value

contains only numbers, and sometimes a decimal point, a minus sign, or both. When they are read into a SAS data set, numeric values are stored in the floating-point format native to the operating environment. Nonstandard numeric values can contain other characters as numbers; you can use formatted input to enable SAS to read them.

character value

is a sequence of characters.

standard data

are character or numeric values that can be read with list, column, formatted, or named input. Examples of standard data include:

- **ARKANSAS**
- **1166.42**

nonstandard data

is data that can be read only with the aid of informats. Examples of nonstandard data include numeric values that contain commas, dollar signs, or blanks; date and time values; and hexadecimal and binary values.

Numeric Data

Numeric data can be represented in several ways. SAS can read standard numeric values without any special instructions. To read nonstandard values, SAS requires special instructions in the form of informats. Table 20.1 on page 313 shows standard, nonstandard, and invalid numeric data values and the special tools, if any, that are required to read them. For complete descriptions of all SAS informats, see *SAS Language Reference: Dictionary*.

Table 20.1 Reading Different Types of Numeric Data

Example of Numeric Data	Description	Solution Required to Read
Standard Numeric Data		
	Standard Numeric Data	
	23 input right aligned	None needed
23	input not aligned	None needed
23	input left aligned	None needed
00023	input with leading zeros	None needed
23.0	input with decimal point	None needed
2.3E1	in E-notation, 2.30 (ss1)	None needed
230E-1	in E-notation, 230x10 (ss-1)	None needed
-23	minus sign for negative numbers	None needed
Nonstandard Numeric Data		
2 3	embedded blank	COMMA. or BZ. informat
- 23	embedded blank	COMMA. or BZ. informat
2,341	comma	COMMA. informat
(23)	parentheses	COMMA. informat

Example of Numeric Data	Description	Solution Required to Read
C4A2	hexadecimal value	HEX. informat
1MAR90	date value	DATE. informat
	Invalid Numeric Data	
23 -	minus sign follows number	Put minus sign before number or solve programmatically. ¹
..	double instead of single periods	Code missing values as a single period or use the ?? modifier in the INPUT statement to code any invalid input value as a missing value.
J23	not a number	Read as a character value, or edit the raw data to change it to a valid number.

¹ It might be possible to use the S370FZDT*w.d* informat, but positive values require the trailing plus sign (+).

Remember the following rules for reading numeric data:

- Parentheses or a minus sign preceding the number (without an intervening blank) indicates a negative value.
- Leading zeros and the placement of a value in the input field do not affect the value assigned to the variable. Leading zeros and leading and trailing blanks are not stored with the value. Unlike some languages, SAS does not read trailing blanks as zeros by default. To cause trailing blanks to be read as zeros, use the BZ. informat described in *SAS Language Reference: Dictionary*.
- Numeric data can have leading and trailing blanks but cannot have embedded blanks (unless they are read with a COMMA. or BZ. informat).
- To read decimal values from input lines that do not contain explicit decimal points, indicate where the decimal point belongs by using a decimal parameter with column input or an informat with formatted input. See the full description of the INPUT statement in *SAS Language Reference: Dictionary* for more information. An explicit decimal point in the input data overrides any decimal specification in the INPUT statement.

Character Data

A value that is read with an INPUT statement is assumed to be a character value if one of the following is true:

- A dollar sign (\$) follows the variable name in the INPUT statement.
- A character informat is used.
- The variable has been previously defined as character: for example, in a LENGTH statement, in the RETAIN statement, by an assignment statement, or in an expression.

Input data that you want to store in a character variable can include any character. Use the guidelines in the following table when your raw data includes leading blanks and semicolons.

Table 20.2 Reading Instream Data and External Files Containing Leading Blanks and Semicolons

Characters in the Data	What to Use	Reason
leading or trailing blanks that you want to preserve	formatted input and the \$CHARw. informat	List input trims leading and trailing blanks from a character value before the value is assigned to a variable.
semicolons in instream data	DATALINES4 or CARDS4 statements and four semicolons (;;;;) to mark the end of the data	With the normal DATALINES and CARDS statements, a semicolon in the data prematurely signals the end of the data.
delimiters, blank characters, or quoted strings	DSD option, with DLM= or DLMSTR= option on the INFILE statement	These options enable SAS to read a character value that contains a delimiter within a quoted string; these options can also treat two consecutive delimiters as a missing value and remove quotation marks from character values.

Remember the following when reading character data:

- In a DATA step, when you place a dollar sign (\$) after a variable name in the INPUT statement, character data that is read from data lines remains in its original case. If you want SAS to read data from data lines as uppercase, use the CAPS system option or the \$UPCASE informat.
- If the value is shorter than the length of the variable, SAS adds blanks to the end of the value to give the value the specified length. This process is known as padding the value with blanks.

Sources of Raw Data

Instream Data

The following example uses the INPUT statement to read in instream data:

```
data weight;
  input PatientID $ Week1 Week8 Week16;
  loss=Week1-Week16;
  datalines;
2477 195 177 163
2431 220 213 198
2456 173 166 155
2412 135 125 116
;
```

Note: A semicolon appearing alone on the line immediately following the last data line is the convention that is used in this example. However, a PROC statement, DATA statement, or global statement ending in a semicolon on the line immediately following the last data line also submits the previous DATA step. \triangle

Instream Data Containing Semicolons

The following example reads in instream data containing semicolons:

```
data weight;
  input PatientID $ Week1 Week8 Week16;
  loss=Week1-Week16;
  datalines4;
24;77 195 177 163
24;31 220 213 198
24;56 173 166 155
24;12 135 125 116
;;;
```

External Files

The following example shows how to read in raw data from an external file using the INFILE and INPUT statements:

```
data weight;
  infile file-specification or path-name;
  input PatientID $ Week1 Week8 Week16;
  loss=Week1-Week16;
run;
```

Note: See the SAS documentation for your operating environment for information on how to specify a file with the INFILE statement. \triangle

Reading Raw Data with the INPUT Statement

Choosing an Input Style

The INPUT statement reads raw data from instream data lines or external files into a SAS data set. You can use the following different input styles, depending on the layout of data values in the records:

- list input
- column input
- formatted input
- named input.

You can also combine styles of input in a single INPUT statement. For details about the styles of input, see the INPUT statement in *SAS Language Reference: Dictionary*.

List Input

List input uses a scanning method for locating data values. Data values are not required to be aligned in columns but must be separated by at least one blank (or other defined delimiter). List input requires only that you specify the variable names and a dollar sign (\$), if defining a character variable. You do not have to specify the location of the data fields.

An example of list input follows:

```
data scores;
  length name $ 12;
  input name $ score1 score2;
  datalines;
Riley 1132 1187
Henderson 1015 1102
;
```

List input has several restrictions on the type of data that it can read:

- Input values must be separated by at least one blank (the default delimiter) or by the delimiter specified with the DLM= or DLMSTR= option in the INFILE statement. If you want SAS to read consecutive delimiters as if there is a missing value between them, specify the DSD option in the INFILE statement.
- Blanks cannot represent missing values. A real value, such as a period, must be used instead.
- To read and store a character input value longer than 8 bytes, define a variable's length by using a LENGTH, INFORMAT, or ATTRIB statement before the INPUT statement, or by using modified list input, which consists of an informat and the colon modifier in the INPUT statement. See "Modified List Input" on page 317 for more information.
- Character values cannot contain embedded blanks when the file is delimited by blanks.
- Fields must be read in order.
- Data must be in standard numeric or character format.

Note: Nonstandard numeric values, such as packed decimal data, must use the formatted style of input. See "Formatted Input" on page 319 for more information. Δ

Modified List Input

A more flexible version of list input, called modified list input, includes format modifiers. The following format modifiers enable you to use list input to read nonstandard data by using SAS informats:

- The & (ampersand) format modifier enables you to read character values that contains one or more embedded blanks with list input and to specify a character informat. SAS reads until it encounters two consecutive blanks, the defined length of the variable, or the end of the input line, whichever comes first.
- The : (colon) format modifier enables you to use list input but also to specify an informat after a variable name, whether character or numeric. SAS reads until it encounters a blank column, the defined length of the variable (character only), or the end of the data line, whichever comes first.
- The ~ (tilde) format modifier enables you to read and retain single quotation marks, double quotation marks, and delimiters within character values.

The following is an example of the `:` and `~` format modifiers. You must use the DSD option in the INFILE statement. Otherwise, the INPUT statement ignores the `~` format modifier.

```
data scores;
  infile datalines dsd;
  input Name : $9. Score1-Score3 Team ~ $25. Div $;
  datalines;
Smith,12,22,46,"Green Hornets, Atlanta",AAA
Mitchel,23,19,25,"High Volts, Portland",AAA
Jones,09,17,54,"Vulcans, Las Vegas",AA
;

proc print data=scores noobs;
run;
```

Output 20.1 Output from Example with Format Modifiers

Name	Score1	Score2	Score3	Team	Div
Smith	12	22	46	"Green Hornets, Atlanta"	AAA
Mitchel	23	19	25	"High Volts, Portland"	AAA
Jones	9	17	54	"Vulcans, Las Vegas"	AA

Column Input

Column input enables you to read standard data values that are aligned in columns in the data records. Specify the variable name, followed by a dollar sign (\$) if it is a character variable, and specify the columns in which the data values are located in each record:

```
data scores;
  infile datalines trunccover;
  input name $ 1-12 score2 17-20 score1 27-30;
  datalines;
Riley          1132      987
Henderson      1015      1102
;
```

Note: Use the TRUNCOVER option in the INFILE statement to ensure that SAS handles data values of varying lengths appropriately. Δ

To use column input, data values must be:

- in the same field on all the input lines
- in standard numeric or character form.

Note: You cannot use an informat with column input. Δ

Features of column input include the following:

- Character values can contain embedded blanks.
- Character values can be from 1 to 32,767 characters long.
- Placeholders, such as a single period (.), are not required for missing data.
- Input values can be read in any order, regardless of their position in the record.

- Values or parts of values can be reread.
- Both leading and trailing blanks within the field are ignored.
- Values do not need to be separated by blanks or other delimiters.

Formatted Input

Formatted input combines the flexibility of using informats with many of the features of column input. By using formatted input, you can read nonstandard data for which SAS requires additional instructions. Formatted input is typically used with pointer controls that enable you to control the position of the input pointer in the input buffer when you read data.

The INPUT statement in the following DATA step uses formatted input and pointer controls. Note that \$12. and COMMA5. are informats and +4 and +6 are column pointer controls.

```
data scores;
    input name $12. +4 score1 comma5. +6 score2 comma5.;
    datalines;
Riley          1,132      1,187
Henderson      1,015      1,102
;
```

Note: You can also use informats to read data that is not aligned in columns. See “Modified List Input” on page 317 for more information. Δ

Important points about formatted input are:

- Character values can contain embedded blanks.
- Character values can be from 1 to 32,767 characters long.
- Placeholders, such as a single period (.), are not required for missing data.
- With the use of pointer controls to position the pointer, input values can be read in any order, regardless of their positions in the record.
- Values or parts of values can be reread.
- Formatted input enables you to read data stored in nonstandard form, such as packed decimal or numbers with commas.

Named Input

You can use named input to read records in which data values are preceded by the name of the variable and an equal sign (=). The following INPUT statement reads the data lines containing equal signs.

```
data games;
    input name=$ score1= score2=;
    datalines;
name=riley score1=1132 score2=1187
;
```

```
proc print data=games;
run;
```

Note: When an equal sign follows a variable in an INPUT statement, SAS expects that data remaining on the input line contains only named input values. You cannot switch to another form of input in the same INPUT statement after using named input.

Also, note that any variable that exists in the input data but is not defined in the INPUT statement generates a note in the SAS log indicating a missing field. Δ

Additional Data-Reading Features

In addition to different styles of input, there are many tools to meet the needs of different data-reading situations. You can use options in the INFILE statement in combination with the INPUT statement to give you additional control over the reading of data records. Table 20.3 on page 320 lists common data-reading tasks and the appropriate features available in the INPUT and INFILE statements.

Table 20.3 Additional Data-Reading Features

Input Data Feature	Goal	Use
multiple records	create a single observation	#n or / line pointer control in the INPUT statement with a DO loop.
a single record	create multiple observations	trailing @@ in the INPUT statement. trailing @ with multiple INPUT and OUTPUT statements.
variable-length data fields and records	read delimited data	list input with or without a format modifier in the INPUT statement and the TRUNCOVER, DLM=, DLMSTR=, or DSD options in the INFILE statement.
	read non-delimited data	\$VARYINGw. informat in the INPUT statement and the LENGTH= and TRUNCOVER options in the INFILE statement.
a file with varying record layouts		IF-THEN statements with multiple INPUT statements, using trailing @ or @@ as necessary.
hierarchical files		IF-THEN statements with multiple INPUT statements, using trailing @ as necessary.
more than one input file or to control the program flow at EOF		EOF= or END= option in an INFILE statement. multiple INFILE and INPUT statements.

Input Data Feature	Goal	Use
		FILEVAR=option in an INFILE statement.
		FILENAME statement with concatenation, wildcard, or piping.
only part of each record		LINESIZE=option in an INFILE statement.
some but not all records in the file		FIRSTOBS=and OBS= options in an INFILE statement; FIRSTOBS= and OBS= system options; #n line pointer control.
instream data lines	control the reading with special options	INFILE statement with DATALINES and appropriate options.
starting at a particular column		@ column pointer controls.
leading blanks	maintain them	\$CHARw. informat in an INPUT statement.
a delimiter other than blanks (with list input or modified list input with the colon modifier)		DLM= or DLMSTR= option, DSD option, or both in an INFILE statement.
the standard tab character		DLM= or DLMSTR= option in an INFILE statement; or the EXPANDTABS option in an INFILE statement.
missing values (with list input or modified list input with the colon modifier)	create observations without compromising data integrity; protect data integrity by overriding the default behavior	TRUNCOVER option in an INFILE statement; DLM= or DLMSTR= options, DSD option, or both might also be needed.

For further information on data-reading features, see the INPUT and INFILE statements in *SAS Language Reference: Dictionary*.

How SAS Handles Invalid Data

An input value is invalid if it has any of the following characteristics:

- It requires an informat that is not specified.
- It does not conform to the informat specified.
- It does not match the input style used; for example, if it is read as standard numeric data (no dollar sign or informat) but does not conform to the rules for standard SAS numbers.
- It is out of range (too large or too small).

Operating Environment Information: The range for numeric values is operating environment-specific. See the SAS documentation for your operating environment for more information. △

If SAS reads a data value that is incompatible with the type specified for that variable, SAS tries to convert the value to the specified type, as described in “How SAS Handles Invalid Data” on page 321. If conversion is not possible, an error occurs, and SAS performs the following actions:

- sets the value of the variable being read to missing or to the value specified with the INVALIDDATA= system option
- prints an invalid data note in the SAS log
- sets the automatic variable `_ERROR_` to 1 for the current observation.
- prints the input line and column number containing the invalid value in the SAS log. If a line contains unprintable characters, it is printed in hexadecimal form. A scale is printed above the input line to help determine column numbers

Reading Missing Values in Raw Data

Representing Missing Values in Input Data

Many collections of data contain some missing values. SAS can recognize these values as missing when it reads them. You use the following characters to represent missing values when reading raw data:

numeric missing values

are represented by a single decimal point (.). All input styles except list input also allow a blank to represent a missing numeric value.

character missing values

are represented by a blank, with one exception: list input requires that you use a period (.) to represent a missing value.

special numeric missing values

are represented by two characters: a decimal point (.) followed by either a letter or an underscore (_).

For more information about missing values, see Chapter 5, “Missing Values,” on page 49.

Special Missing Values in Numeric Input Data

SAS enables you to differentiate among classes of missing values in numeric data. For numeric variables, you can designate up to 27 special missing values by using the letters A through Z, in either upper- or lowercase, and the underscore character (_).

The following example shows how to code missing values by using a MISSING statement in a DATA step:

```
data test_results;
  missing a b c;
  input name $8. Answer1 Answer2 Answer3;
  datalines;
Smith    2 5 9
Jones    4 b 8
Carter   a 4 7
Reed     3 5 c
;
```

```
proc print;
run;
```

Note that you must use a period when you specify a special missing numeric value in an expression or assignment statement, as in the following:

```
x=.d;
```

However, you do not need to specify each special missing numeric data value with a period in your input data. For example, the following DATA step, which uses periods in the input data for special missing values, produces the same result as the input data without periods:

```
data test_results;
  missing a b c;
  input name $8. Answer1 Answer2 Answer3;
  datalines;
Smith    2 5 9
Jones    4 .b 8
Carter   .a 4 7
Reed     3 5 .c
;

proc print;
run;
```

Output for both examples is shown here:

Output 20.2 Output of Data with Special Missing Numeric Values

The SAS System				
Obs	name	Answer1	Answer2	Answer3
1	Smith	2	5	9
2	Jones	4	B	8
3	Carter	A	4	7
4	Reed	3	5	C

Note: SAS displays and prints special missing values that use letters in uppercase. △

Reading Binary Data

Definitions

binary data

is numeric data that is stored in binary form. Binary numbers have a base of two and are represented with the digits 0 and 1.

packed decimal data

are binary decimal numbers that are encoded by using each byte to represent two decimal digits. Packed decimal representation stores decimal data with exact

precision; the fractional part of the number must be determined by using an informat or format because there is no separate mantissa and exponent.

zoned decimal data

are binary decimal numbers that are encoded so that each digit requires one byte of storage. The last byte contains the number's sign as well as the last digit. Zoned decimal data produces a printable representation.

Using Binary Informats

SAS can read binary data with the special instructions supplied by SAS informats. You can use formatted input and specify the informat in the INPUT statement. The informat you choose is determined by the following factors:

- the type of number being read: binary, packed decimal, zoned decimal, or a variation of one of these
- the type of system on which the data was created
- the type of system that you use to read the data.

Different computer platforms store numeric binary data in different forms. The ordering of bytes differs by platforms that are referred to as either “big endian” or “little endian.” For more information, see “Byte Ordering for Integer Binary Data on Big Endian and Little Endian Platforms” in *SAS Language Reference: Dictionary*.

SAS provides a number of informats for reading binary data and corresponding formats for writing binary data. Some of these informats read data in native mode, that is, by using the byte-ordering system that is standard for the system on which SAS is running. Other informats force the data to be read by the IBM 370 standard, regardless of the native mode of the system on which SAS is running. The informats that read in native or IBM 370 mode are listed in the following table.

Table 20.4 Informats for Native or IBM 370 Mode

Description	Native Mode Informats	IBM 370 Mode Informats
ASCII Character	<i>\$w.</i>	<i>\$ASCIIw.</i>
ASCII Numeric	<i>w.d</i>	<i>\$ASCIIw.</i>
EBCDIC Character	<i>\$w.</i>	<i>\$EBCDICw.</i>
EBCDIC Numeric (Standard)	<i>w.d</i>	<i>S370FFw.d</i>
Integer Binary	<i>IBw.d</i>	<i>S370FIBw.d</i>
Positive Integer Binary	<i>PIBw.d</i>	<i>S370FPIBw.d</i>
Real Binary	<i>RBw.d</i>	<i>S370FRBw.d</i>
Unsigned Integer Binary	<i>PIBw.d</i>	<i>S370FIBUw.d,</i> <i>S370FPIBw.d</i>
Packed Decimal	<i>PDw.d</i>	<i>S370FPDw.d</i>
Unsigned Packed Decimal	<i>PKw.d</i>	<i>S370FPDUw.d</i> or <i>PKw.d</i>
Zoned Decimal	<i>ZDw.d</i>	<i>S370FZDw.d</i>
Zoned Decimal Leading Sign	<i>S370FZDLw.d</i>	<i>S370FZDLw.d</i>
Zoned Decimal Separate Leading Sign	<i>S370FZDSw.d</i>	<i>S370FZDSw.d</i>

Description	Native Mode Informats	IBM 370 Mode Informats
Zoned Decimal Separate Trailing Sign	S370FZDTw.d	S370FZDTw.d
Unsigned Zoned Decimal	ZDw.d	S370FZDUw.d

If you write a SAS program that reads binary data and that will be run on only one type of system, you can use the native mode informats and formats. However, if you want to write SAS programs that can be run on multiple systems that use different byte-storage systems, use the IBM 370 informats. The IBM 370 informats enable you to write SAS programs that can read data in this format and that can be run in any SAS environment, regardless of the standard for storing numeric data.* The IBM 370 informats can also be used to read data originally written with the corresponding native mode formats on an IBM mainframe.

Note: Any time a text file originates from anywhere other than the local encoding environment, it might be necessary to specify the ENCODING= option on either EBCDIC or ASCII systems.

When you read an EBCDIC text file on an ASCII platform, it is recommended that you specify the ENCODING= option in the FILENAME or INFILE statement. However, if you use the DSD and the DLM= or DLMSTR= options on the FILENAME or INFILE statement, the ENCODING= option is a requirement because these options require certain characters in the session encoding (such as quotes, commas, and blanks).

The use of encoding-specific informats should be reserved for use with true binary files that contain both character and non-character fields. △

For complete descriptions of all SAS formats and informats, including how numeric binary data is written, see *SAS Language Reference: Dictionary*.

Reading Column-Binary Data

Definition

column-binary data storage

is an older form of data storage that is no longer widely used and is not needed by most SAS users. Column-binary data storage compresses data so that more than 80 items of data can be stored on a single “virtual” punched card. The advantage is that this method enables you to store more data in the same amount of space. Because card-image data sets remain in existence, SAS provides informats for reading column-binary data. See “Description of Column-Binary Data Storage” on page 326 for a more detailed explanation of column-binary data storage.

How to Read Column-Binary Data

To read column-binary data with SAS, you need to know:

- how to select the appropriate SAS column-binary informat

* For example, using the IBM 370 informats, you could download data that contain binary integers from a mainframe to a PC and then use the S370FIB informats to read the data.

- how to set the RECFM= and LRECL= options in the INFILE statement
- how to use pointer controls.

The following table lists and describes SAS column-binary informats.

Table 20.5 SAS Informats for Reading Column-Binary Data

Informat Name	Description
\$CBw.	reads standard character data from column-binary files
CBw.	reads standard numeric data from column-binary files
PUNCH.d	reads whether a row is punched
ROWw.d	reads a column-binary field down a card column

To read column-binary data, you must set two options in the INFILE statement:

- Set RECFM= to F for fixed.
- Set the LRECL= to 160, because each card column of column-binary data is expanded to two bytes before the fields are read.

For example, to read column-binary data from a file, use an INFILE statement in the following form before the INPUT statement that reads the data:

```
infile file-specification or path-name
       recfm=f lrecl=160;
```

Note: The expansion of each column of column-binary data into two bytes does *not* affect the position of the column pointer. You use the absolute column pointer control @, as usual, because the informats automatically compute the true location on the doubled record. If a value is in column 23, use the pointer control @23 to move the pointer there. \triangle


Description of Column-Binary Data Storage


The arrangement and numbering of rows in a column on physical punched cards originated with the Hollerith system of encoding characters and numbers. It was based on using a pair of values to represent either a character or a numeric digit. In the Hollerith system, each column on a card had a maximum of two punches, one punch in the zone portion, and one in the digit portion. These punches corresponded to a pair of values, and each pair of values corresponded to a specific alphabetic character or sign and numeric digit.

In the zone portion of the punched card, which is the first three rows, the zone component of the pair can have the values 12, 11, 0 (or 10), or not punched. In the digit portion of the card, which is the fourth through the twelfth rows, the digit component of the pair can have the values 1 through 9, or not punched.

The following figure shows the multi-punch combinations corresponding to letters of the alphabet.

Figure 20.2 Column-Binary Representation on a “Virtual” Punched Card





CHAPTER

21

BY-Group Processing in the DATA Step

<i>Definitions for BY-Group Processing</i>	329
<i>Syntax for BY-Group Processing</i>	330
<i>Understanding BY Groups</i>	331
<i>BY Groups with a Single BY Variable</i>	331
<i>BY Groups with Multiple BY Variables</i>	331
<i>Invoking BY-Group Processing</i>	332
<i>Determining Whether the Data Requires Preprocessing for BY-Group Processing</i>	332
<i>Preprocessing Input Data for BY-Group Processing</i>	333
<i>Sorting Observations for BY-Group Processing</i>	333
<i>Indexing for BY-Group Processing</i>	333
<i>How the DATA Step Identifies BY Groups</i>	334
<i>Processing Observations in a BY Group</i>	334
<i>How SAS Determines FIRST.VARIABLE and LAST.VARIABLE</i>	334
<i>Grouping Observations by State, City, Zip Code, and Street</i>	334
<i>Grouping Observations by City, State, Zip Code, and Street</i>	335
<i>Grouping Observations: Another Example</i>	335
<i>Processing BY-Groups in the DATA Step</i>	336
<i>Overview</i>	336
<i>Processing BY-Groups Conditionally</i>	337
<i>Data Not in Alphabetic or Numeric Order</i>	338
<i>Data Grouped by Formatted Values</i>	338

Definitions for BY-Group Processing

BY-group processing

is a method of processing observations from one or more SAS data sets that are grouped or ordered by values of one or more common variables. The most common use of BY-group processing in the DATA step is to combine two or more SAS data sets by using the BY statement with a SET, MERGE, MODIFY, or UPDATE statement.

BY variable

names a variable or variables by which the data set is sorted or indexed. All data sets must be ordered or indexed on the values of the BY variable if you use the SET, MERGE, or UPDATE statements. If you use MODIFY, data does not need to be ordered. However, your program might run more efficiently with ordered data. All data sets that are being combined must include one or more BY variables. The position of the BY variable in the observations does not matter.

BY value

is the value or formatted value of the BY variable.

BY group

includes all observations with the same BY value. If you use more than one variable in a BY statement, a BY group is a group of observations with the same combination of values for these variables. Each BY group has a unique combination of values for the variables.

FIRST.variable and *LAST.variable*

are variables that SAS creates for each BY variable. SAS sets *FIRST.variable* when it is processing the first observation in a BY group, and sets *LAST.variable* when it is processing the last observation in a BY group. These assignments enable you to take different actions, based on whether processing is starting for a new BY group or ending for a BY group. For more information, see “How the DATA Step Identifies BY Groups” on page 334.

For more information about BY-Group processing, see Chapter 22, “Reading, Combining, and Modifying SAS Data Sets,” on page 341. See also *Combining and Modifying SAS Data Sets: Examples*.

Syntax for BY-Group Processing

Use one of the following forms for BY-group processing:

BY *variable(s)*;

BY <DESCENDING> *variable(s)* <NOTSORTED> <GROUPFORMAT>;

where

variable

names each variable by which the data set is sorted or indexed.

Note: All data sets must be ordered or indexed on the values of the BY variable if you process them using the SET, MERGE, or UPDATE statements. If you use the MODIFY statement, your data does not need to be ordered. However, your program might run more efficiently with ordered data. All data sets that are being combined must include the BY variable or variables. The position of the BY variable in the observations does not matter. Δ

GROUPFORMAT

uses the formatted values, instead of the internal values, of the BY variables to determine where BY-groups begin and end, and therefore how *FIRST.variable* and *LAST.variable* are assigned. Although the GROUPFORMAT option can appear anywhere in the BY statement, the option applies to *all* variables in the BY statement.

DESCENDING

indicates that the data sets are sorted in descending order (largest to smallest) by the variable that is specified. If you have more than one variable in the BY group, DESCENDING applies only to the variable that immediately follows it.

NOTSORTED

specifies that observations with the same BY value are grouped together but are not necessarily stored in alphabetical or numeric order.

For complete information about the BY statement, see *SAS Language Reference: Dictionary*.

Understanding BY Groups

BY Groups with a Single BY Variable

The following figure represents the results of processing your data with the single BY variable ZipCode. The input SAS data set contains street names, cities, states, and ZIP codes that are arranged in an order that you can use with the following BY statement:

```
by ZipCode;
```

The figure shows five BY groups each containing the BY variable ZipCode. The data set is shown with the BY variable ZipCode printed on the left for easy reading, but the position of the BY variable in the observations does not matter.

Figure 21.1 BY Groups for the Single BY Variable ZipCode

BY variable				
ZipCode	State	City	Street	
33133	FL	Miami	Rice St	} BY group
33133	FL	Miami	Thomas Ave	
33133	FL	Miami	Surrey Dr	
33133	FL	Miami	Trade Ave	
33146	FL	Miami	Nervia St	} BY group
33146	FL	Miami	Corsica St	
33801	FL	Lakeland	French Ave	} BY group
33809	FL	Lakeland	Egret Dr	} BY Group
85730	AZ	Tucson	Domenic Ln	} BY group
85730	AZ	Tucson	Gleeson Pl	

The first BY group contains all observations with the smallest BY value, which is 33133; the second BY group contains all observations with the next smallest BY value, which is 33146, and so on.

BY Groups with Multiple BY Variables

The following figure represents the results of processing your data with two BY variables, State and City. This example uses the same data set as in “BY Groups with a Single BY Variable” on page 331, and is arranged in an order that you can use with the following BY statement:

```
by State City;
```

The figure shows three BY groups. The data set is shown with the BY variables State and City printed on the left for easy reading, but the position of the BY variables in the observations does not matter.

Figure 21.2 BY Groups for the BY Variables State and City

BY variables		Street	ZipCode	
State	City			
AZ	Tucson	Domenic Ln	85730	} BY group
AZ	Tucson	Gleeson Pl	85730	
FL	Lakeland	French Ave	33801	} BY group
FL	Lakeland	Egret Dr	33809	
FL	Miami	Nervia St	33146	} BY group
FL	Miami	Rice St	33133	
FL	Miami	Corsica St	33146	
FL	Miami	Thomas Ave	33133	
FL	Miami	Surrey Dr	33133	
FL	Miami	Trade Ave	33133	

The observations are arranged so that the observations for Arizona occur first. The observations within each value of State are arranged in order of the value of City. Each BY group has a unique combination of values for the variables State and City. For example, the BY value of the first BY group is **AZ Tucson**, and the BY value of the second BY group is **FL Lakeland**.

Invoking BY-Group Processing

You can invoke BY-group processing in both DATA steps and PROC steps by using a BY statement. For example, the following DATA step program uses the SET statement to combine observations from three SAS data sets by interleaving the files. The BY statement shows how the data is ordered.

```
data all_sales;
  set region1 region2 region3;
  by State City Zip;
  ... more SAS statements ...
run;
```

This section describes BY-group processing for the DATA step. For information about BY-group processing with procedures, see *Base SAS Procedures Guide*.

Determining Whether the Data Requires Preprocessing for BY-Group Processing

Before you process one or more SAS data sets using grouped or ordered data with the SET, MERGE, or UPDATE statements, you must check the data to determine whether they require preprocessing. They require no preprocessing if the observations in all of the data sets occur in one of the following patterns:

- ascending or descending numeric order

- ascending or descending character order
- not alphabetically or numerically ordered, but grouped in some way, such as by calendar month or by a formatted value

If the observations are not in the order that you want, you must either sort the data set or create an index for it before using BY-group processing.

If you use the MODIFY statement in BY-group processing, you do not need to presort the input data. Presorting, however, can make processing more efficient and less costly.

You can use PROC SQL views in BY-group processing. For complete information, see *SAS 9.2 SQL Procedure User's Guide*.

SAS/ACCESS Users: If you use SAS views or librefs, refer to the SAS/ACCESS documentation for your operating environment for information about using BY groups in your SAS programs.

Preprocessing Input Data for BY-Group Processing

Sorting Observations for BY-Group Processing

You can use the SORT procedure to change the physical order of the observations in the data set. You can either replace the original data set, or create a new, sorted data set by using the OUT= option of the SORT procedure. In this example, PROC SORT rearranges the observations in the data set INFORMATION based on ascending values of the variables State and ZipCode, and replaces the original data set.

```
proc sort data=information;
  by State ZipCode;
run;
```

As a general rule, when you use PROC SORT, specify the variables in the BY statement in the same order that you plan to specify them in the BY statement in the DATA step. For a detailed description of the default sorting orders for numeric and character variables, see the SORT procedure in *Base SAS Procedures Guide*.

Note: The BY statement honors the linguistic collation of sorted data when you use the SORT procedure with the SORTSEQ=LINGUISTIC option. Δ

Indexing for BY-Group Processing

You can also ensure that observations are processed in ascending numeric or character order by creating an index based on one or more variables in the SAS data set. If you specify a BY statement in a DATA step, SAS looks for an appropriate index. If it finds the index, SAS automatically retrieves the observations from the data set in indexed order.

Note: Because creating and maintaining indexes require additional resources, you should determine whether using them significantly improves performance. Depending on the nature of the data in your SAS data set, using PROC SORT to order data values can be more advantageous than indexing. For an overview of indexes, see “Understanding SAS Indexes” on page 505. Δ

How the DATA Step Identifies BY Groups

Processing Observations in a BY Group

In the DATA step, SAS identifies the beginning and end of each BY group by creating two temporary variables for each BY variable: *FIRST.variable* and *LAST.variable*. These temporary variables are available for DATA step programming but are not added to the output data set. Their values indicate whether an observation is

- the first one in a BY group
- the last one in a BY group
- neither the first nor the last one in a BY group
- both first and last, as is the case when there is only one observation in a BY group

You can take actions conditionally, based on whether you are processing the first or the last observation in a BY group.

How SAS Determines FIRST.VARIABLE and LAST.VARIABLE

When an observation is the first in a BY group, SAS sets the value of *FIRST.variable* to 1 for the variable whose value changed, as well as for all of the variables that follow in the BY statement. For all other observations in the BY group, the value of *FIRST.variable* is 0. Likewise, if the observation is the last in a BY group, SAS sets the value of *LAST.variable* to 1 for the variable whose value changes on the next observation, as well as for all of the variables that follow in the BY statement. For all other observations in the BY group, the value of *LAST.variable* is 0. For the last observation in a data set, the value of all *LAST.variable* variables are set to 1.

Grouping Observations by State, City, Zip Code, and Street

This example shows how SAS uses the *FIRST.variable* and *LAST.variable* to flag the beginning and end of four BY groups: State, City, ZipCode, and Street. Six temporary variables are created within the program data vector. These variables can be used during the DATA step, but they do not become variables in the new data set.

In the figure that follows, observations in the SAS data set are arranged in an order that can be used with this BY statement:

```
by State City ZipCode;
```

SAS creates the following temporary variables: *FIRST.State*, *LAST.State*, *FIRST.City*, *LAST.City*, *FIRST.ZipCode*, and *LAST.ZipCode*.

Observations in Four BY Groups				Corresponding FIRST. and LAST. Values					
State	City	ZipCode	Street	FIRST. State	LAST. State	FIRST. City	LAST. City	FIRST. ZipCode	LAST. ZipCode
AZ	Tucson	85730	Glen Pl	1	1	1	1	1	1
FL	Miami	33133	Rice St	1	0	1	0	1	0

Observations in Four BY Groups				Corresponding FIRST. and LAST. Values					
FL	Miami	33133	Tom Ave	0	0	0	0	0	0
FL	Miami	33133	Surrey Dr	0	0	0	0	0	1
FL	Miami	33146	Nervia St	0	0	0	0	1	0
FL	Miami	33146	Corsica St	0	1	0	1	0	1
OH	Miami	45056	Myrtle St	1	1	1	1	1	1

Grouping Observations by City, State, Zip Code, and Street

This example shows how SAS uses the *FIRST.variable* and *LAST.variable* to flag the beginning and end of four BY groups: City, State, ZipCode, and Street. Six temporary variables are created within the program data vector. These variables can be used during the DATA step, but they do not become variables in the new data set.

In the figure that follows, observations in the SAS data set are arranged in an order that can be used with this BY statement:

```
by City State ZipCode;
```

SAS creates the following temporary variables: FIRST.City, LAST.City, FIRST.State, LAST.State, FIRST.ZipCode, and LAST.ZipCode.

Observations in Four BY Groups				Corresponding FIRST. and LAST. Values					
City	State	ZipCode	Street	FIRST. City	LAST. City	FIRST. State	LAST. State	FIRST. ZipCode	LAST. ZipCode
Miami	FL	33133	Rice St	1	0	1	0	1	0
Miami	FL	33133	Tom Ave	0	0	0	0	0	0
Miami	FL	33133	Surrey Dr	0	0	0	0	0	1
Miami	FL	33146	Nervia St	0	0	0	0	1	0
Miami	FL	33146	Corsica St	0	0	0	1	0	1
Miami	OH	45056	Myrtle St	0	1	1	1	1	1
Tucson	AZ	85730	Glen Pl	1	1	1	1	1	1

Grouping Observations: Another Example

The value of *FIRST.variable* can be affected by a change in a previous value, even if the current value of the variable remains the same.

In this example, the value of *FIRST.variable* and *LAST.variable* are dependent on sort order, and not just by the value of the BY variable. For observation 3, the value of FIRST.Y is set to 1 because BLUEBERRY is a new value for **y**. This change in **y** causes FIRST.Z to be set to 1 as well, even though the value of Z did not change.

```

options pageno=1 nodate linesize=80 pagesize=60;

data testfile;
  input x $ y $ 9-17 z $ 19-26;
  datalines;
apple  banana    coconut
apple  banana    coconut
apricot blueberry citron
;

data _null_;
  set testfile;
  by x y z;
  if _N_=1 then put 'Grouped by X Y Z';
  put _N_= x= first.x= last.x= first.y= last.y= first.z= last.z= ;
run;

data _null_;
  set testfile;
  by y x z;
  if _N_=1 then put 'Grouped by Y X Z';
  put _N_= x= first.x= last.x= first.y= last.y= first.z= last.z= ;
run;

```

Output 21.1 Partial SAS Log Showing the Results of Processing with BY Variables

```

Grouped by X Y Z
_N_=1 x=Apple FIRST.x=1 LAST.x=0 FIRST.y=1 LAST.y=0 FIRST.z=1 LAST.z=0
_N_=2 x=Apple FIRST.x=0 LAST.x=0 FIRST.y=0 LAST.y=1 FIRST.z=0 LAST.z=1
_N_=3 x=Apple FIRST.x=0 LAST.x=1 FIRST.y=1 LAST.y=1 FIRST.z=1 LAST.z=1
_N_=4 x=Apricot FIRST.x=1 LAST.x=1 FIRST.y=1 LAST.y=1 FIRST.z=1 LAST.z=1

Grouped by Y X Z
_N_=1 x=Apple FIRST.x=1 LAST.x=0 FIRST.y=1 LAST.y=0 FIRST.z=1 LAST.z=0
_N_=2 x=Apple FIRST.x=0 LAST.x=1 FIRST.y=0 LAST.y=1 FIRST.z=0 LAST.z=1
_N_=3 x=Apple FIRST.x=1 LAST.x=1 FIRST.y=1 LAST.y=0 FIRST.z=1 LAST.z=1
_N_=4 x=Apricot FIRST.x=1 LAST.x=1 FIRST.y=0 LAST.y=1 FIRST.z=1 LAST.z=1

```

Processing BY-Groups in the DATA Step

Overview

The most common use of BY-group processing in the DATA step is to combine two or more SAS data sets using a BY statement with a SET, MERGE, MODIFY, or UPDATE statement. (If you use a SET, MERGE, or UPDATE statement with the BY statement, your observations must be grouped or ordered.) When processing these statements, SAS reads one observation at a time into the program data vector. With BY-group processing, SAS selects the observations from the data sets according to the values of the BY variable or variables. After processing all the observations from one BY group, SAS expects the next observation to be from the next BY group.

The BY statement modifies the action of the SET, MERGE, MODIFY, or UPDATE statement by controlling when the values in the program data vector are set to missing. During BY-group processing, SAS retains the values of variables until it has copied the last observation it finds for that BY group in any of the data sets. Without the BY statement, the SET statement sets variables to missing when it reads the last observation from any data set, and the MERGE statement does not set variables to missing after the DATA step starts reading observations into the program data vector.

Processing BY-Groups Conditionally

You can process observations conditionally by using the subsetting IF or IF-THEN statements, or the SELECT statement, with the temporary variables *FIRST.variable* and *LAST.variable* (set up during BY-group processing). For example, you can use them to perform calculations for each BY group and to write an observation when the first or the last observation of a BY group has been read into the program data vector.

The following example computes annual payroll by department. It uses IF-THEN statements and the values of *FIRST.variable* and *LAST.variable* automatic variables to reset the value of PAYROLL to 0 at the beginning of each BY group and to write an observation after the last observation in a BY group is processed.

```
options pageno=1 nodate linesize=80 pagesize=60;

data salaries;
  input Department $ Name $ WageCategory $ WageRate;
  datalines;
BAD Carol Salaried 20000
BAD Elizabeth Salaried 5000
BAD Linda Salaried 7000
BAD Thomas Salaried 9000
BAD Lynne Hourly 230
DDG Jason Hourly 200
DDG Paul Salaried 4000
PPD Kevin Salaried 5500
PPD Amber Hourly 150
PPD Tina Salaried 13000
STD Helen Hourly 200
STD Jim Salaried 8000
;

proc print data=salaries;
run;

proc sort data=salaries out=temp;
  by Department;
run;

data budget (keep=Department Payroll);
  set temp;
  by Department;
  if WageCategory='Salaried' then YearlyWage=WageRate*12;
  else if WageCategory='Hourly' then YearlyWage=WageRate*2000;

  /* SAS sets FIRST.variable to 1 if this is a new          */
```

```

        /* department in the BY group.                */
        if first.Department then Payroll=0;
        Payroll+YearlyWage;

        /* SAS sets LAST.variable to 1 if this is the last */
        /* department in the current BY group.            */
        if last.Department;
run;

proc print data=budget;
    format Payroll dollar10.;
    title 'Annual Payroll by Department';
run;

```

Output 21.2 Output from Conditional BY-Group Processing

Annual Payroll by Department			1
Obs	Department	Payroll	
1	BAD	\$952,000	
2	DDG	\$448,000	
3	PPD	\$522,000	
4	STD	\$496,000	

Data Not in Alphabetic or Numeric Order

In BY-group processing, you can use data that is arranged in an order other than alphabetic or numeric, such as by calendar month or by category. To do this, use the `NOTSORTED` option in a `BY` statement when you use a `SET` statement. The `NOTSORTED` option in the `BY` statement tells SAS that the data is not in alphabetic or numeric order, but that it is arranged in groups by the values of the `BY` variable. You cannot use the `NOTSORTED` option with the `MERGE` statement, the `UPDATE` statement, or when the `SET` statement lists more than one data set.

This example assumes that the data is grouped by the character variable `MONTH`. The subsetting `IF` statement conditionally writes an observation, based on the value of `LAST.month`. This `DATA` step writes an observation only after processing the last observation in each `BY` group.

```

data total_sale(drop=sales);
    set region.sales
    by month notsorted;
    total+sales;
    if last.month;
run;

```

Data Grouped by Formatted Values

Use the `GROUPFORMAT` option in the `BY` statement to ensure that

- formatted values are used to group observations when a FORMAT statement and a BY statement are used together in a DATA step
- the *FIRST.variable* and *LAST.variable* are assigned by the formatted values of the variable

The GROUPFORMAT option is valid only in the DATA step that creates the SAS data set. It is particularly useful with user-defined formats. The following example illustrates the use of the GROUPFORMAT option.


```
proc format;
  value range
    low -55 = 'Under 55'
    55-60  = '55 to 60'
    60-65  = '60 to 65'
    65-70  = '65 to 70'
    other  = 'Over 70';
run;

proc sort data=class out=sorted_class;
  by height;
run;

data _null_;
  format height range.;
  set sorted_class;
  by height groupformat;
  if first.height then
    put 'Shortest in ' height 'measures ' height:best12.;
run;
```

SAS writes the following output to the log:

```
Shortest in Under 55 measures 51.3
Shortest in 55 to 60 measures 56.3
Shortest in 60 to 65 measures 62.5
Shortest in 65 to 70 measures 65.3
Shortest in Over 70 measures 72
```

CHAPTER 22

Reading, Combining, and Modifying SAS Data Sets

<i>Definitions for Reading, Combining, and Modifying SAS Data Sets</i>	343
<i>Overview of Tools</i>	343
<i>Reading SAS Data Sets</i>	344
<i>Reading a Single SAS Data Set</i>	344
<i>Reading from Multiple SAS Data Sets</i>	344
<i>Controlling the Reading and Writing of Variables and Observations</i>	344
<i>Combining SAS Data Sets: Basic Concepts</i>	345
<i>What You Need to Know before Combining Information Stored In Multiple SAS Data Sets</i>	345
<i>The Four Ways That Data Can Be Related</i>	345
<i>One-to-One</i>	346
<i>One-to-Many and Many-to-One</i>	346
<i>Many-to-Many</i>	347
<i>Access Methods: Sequential versus Direct</i>	348
<i>Overview</i>	348
<i>Sequential Access</i>	348
<i>Direct Access</i>	348
<i>Overview of Methods for Combining SAS Data Sets</i>	349
<i>Concatenating</i>	349
<i>Interleaving</i>	349
<i>One-to-One Reading and One-to-One Merging</i>	350
<i>Match-Merging</i>	351
<i>Updating</i>	351
<i>Overview of Tools for Combining SAS Data Sets</i>	352
<i>Using Statements and Procedures</i>	352
<i>Using Error Checking</i>	353
<i>How to Prepare Your Data Sets</i>	354
<i>Knowing the Structure and Contents of the Data Sets</i>	354
<i>Looking at Sources of Common Problems</i>	354
<i>Ensuring Correct Order</i>	355
<i>Testing Your Program</i>	355
<i>Combining SAS Data Sets: Methods</i>	356
<i>Concatenating</i>	356
<i>Definition</i>	356
<i>Syntax</i>	356
<i>DATA Step Processing during Concatenation</i>	356
<i>Example 1: Concatenation Using the DATA Step</i>	357
<i>Example 2: Concatenation Using SQL</i>	358
<i>Appending Files</i>	358
<i>Efficiency</i>	359
<i>Interleaving</i>	359
<i>Definition</i>	359

Syntax	359
Sort Requirements	359
DATA Step Processing during Interleaving	360
Example 1: Interleaving in the Simplest Case	360
Example 2: Interleaving with Duplicate Values of the BY variable	361
Example 3: Interleaving with Different BY Values in Each Data Set	362
Comments and Comparisons	363
One-to-One Reading	363
Definition	363
Syntax	363
DATA Step Processing during a One-to-One Reading	364
Example 1: One-to-One Reading: Processing an Equal Number of Observations	364
Comments and Comparisons	365
One-to-One Merging	366
Definition	366
Syntax	366
DATA Step Processing during One-to-One Merging	366
Example 1: One-to-One Merging with an Equal Number of Observations	367
Example 2: One-to-One Merging with an Unequal Number of Observations	367
Example 3: One-to-One Merging with Duplicate Values of Common Variables	368
Example 4: One-to-One Merging with Different Values of Common Variables	369
Comments and Comparisons	370
Match-Merging	370
Definition	370
Syntax	370
DATA Step Processing during Match-Merging	371
Example 1: Combining Observations Based on a Criterion	371
Example 2: Match-Merge with Duplicate Values of the BY Variable	372
Example 3: Match-Merge with Nonmatched Observations	373
Updating with the UPDATE and the MODIFY Statements	374
Definitions	374
Syntax of the UPDATE Statement	374
Syntax of the MODIFY Statement	375
DATA Step Processing with the UPDATE Statement	375
Updating with Nonmatched Observations, Missing Values, and New Variables	376
Sort Requirements for the UPDATE Statement	376
Using an Index with the MODIFY Statement	376
Choosing between UPDATE or MODIFY with BY	376
Primary Uses of the MODIFY Statement	377
Example 1: Using UPDATE for Basic Updating	377
Example 2: Using UPDATE with Duplicate Values of the BY Variable	378
Example 3: Using UPDATE for Processing Nonmatched Observations, Missing Values, and New Variables	379
Example 4: Updating a MASTER Data Set by Adding an Observation	381
Error Checking When Using Indexes to Randomly Access or Update Data	382
The Importance of Error Checking	382
Error-Checking Tools	382
Example 1: Routing Execution When an Unexpected Condition Occurs	383
Overview	383
Input Data Sets	383
Original Program	384
Resulting Log	384
Resulting Data Set	384
Revised Program	385

<i>Resulting Log</i>	385
<i>Correctly Updated MASTER Data Set</i>	386
<i>Example 2: Using Error Checking on All Statements That Use KEY=</i>	386
<i>Overview</i>	386
<i>Input Data Sets</i>	386
<i>Original Program with Logic Error</i>	386
<i>Resulting Log</i>	388
<i>Resulting Data Set</i>	388
<i>Revised Program</i>	388
<i>Resulting Log</i>	390
<i>Correctly Created COMBINE Data Set</i>	390

Definitions for Reading, Combining, and Modifying SAS Data Sets

In the context of DATA step processing, the terms reading, combining and modifying have these meanings:

Reading a SAS data set

refers to opening a SAS data set and bringing an observation into the program data vector for processing.

Combining SAS data sets

refers to reading data from two or more SAS data sets and processing them by

- concatenating
- interleaving
- one-to-one reading
- one-to-one merging
- match-merging
- updating a master data set with a transaction data set

The methods for combining SAS data sets are defined in “Combining SAS Data Sets: Methods” on page 356.

Modifying SAS data sets

refers to using the MODIFY statement to update information in a SAS data set in place. The MODIFY statement can save disk space because it modifies data in place, without creating a copy of the data set. You can modify a SAS data set with programming statements or with information that is stored in another data set.

Overview of Tools

The primary tools that are used for reading, combining, and modifying SAS data sets are four statements: SET, MERGE, MODIFY, and UPDATE. This section describes these tools and shows examples. For complete information about these statements see the *SAS Language Reference: Dictionary*.

Reading SAS Data Sets

Reading a Single SAS Data Set

To read data from an existing SAS data set, use a SET statement. In this example, the DATA step creates data set PERM.TOUR155_PEAKECOST by reading data from data set PERM.TOUR155_BASIC_COST and by calculating values for the three new variables Total_Cost, Peak_Cost, and Average_Night_Cost.

```
data perm.tour155_peakcost;
  set perm.tour155_basic_cost;
  Total_Cost=AirCost+LandCost;
  Peak_Cost=(AirCost*1.15);
  Average_Night_Cost=LandCost/Nights;
run;
```

Reading from Multiple SAS Data Sets

You can read from multiple SAS data sets and combine and modify data in different ways. You can, for example, combine two or more input data sets to create one output data set, merge data from two or more input data sets that share a common variable, and update a master file based on transaction records.

For details about reading from multiple SAS data sets, see “Combining SAS Data Sets: Methods” on page 356.

Controlling the Reading and Writing of Variables and Observations

If you do not instruct it to do otherwise, SAS writes all variables and all observations from input data sets to output data sets. You can, however, control which variables and observations you want to read and write by using SAS statements, data set options, and functions. The statements and data set options that you can use are listed in the following table.

Table 22.1 Statements and Options That Control Reading and Writing

Task	Statements	Data Set Options	System Options
Control variables	DROP	DROP=	
	KEEP	KEEP=	
	RENAME	RENAME=	
Control observations	WHERE	WHERE=	FIRSTOBS=
	subsetting IF	FIRSTOBS=	OBS=
	DELETE	OBS=	
	REMOVE		

Task	Statements	Data Set Options	System Options
	OUTPUT		

Use statements or data set options (such as KEEP= and DROP=) to control the variables and observations you want to write to the output data set. The WHERE statement is an exception: it controls which observations are read into the program data vector based on the value of a variable. You can use data set options (including WHERE=) on input or output data sets, depending on their function and what you want to control. You can also use SAS system options to control your data.

Combining SAS Data Sets: Basic Concepts

What You Need to Know before Combining Information Stored In Multiple SAS Data Sets

Many applications require input data to be in a specific format before the data can be processed to produce meaningful results. The data typically comes from multiple sources and might be in different formats. Therefore, you often, if not always, have to take intermediate steps to logically relate and process data before you can analyze it or create reports from it.

Application requirements vary, but there are common factors for all applications that access, combine, and process data. Once you have determined what you want the output to look like, you must

- determine how the input data is related
- ensure that the data is properly sorted or indexed, if necessary
- select the appropriate access method to process the input data
- select the appropriate SAS tools to complete the task

The Four Ways That Data Can Be Related

Relationships among multiple sources of input data exist when each of the sources contains common data, either at the physical or logical level. For example, employee data and department data could be related through an employee ID variable that shares common values. Another data set could contain numeric sequence numbers whose partial values logically relate it to a separate data set by observation number.

You must be able to identify the existing relationships in your data. This knowledge is crucial for understanding how to process input data in order to produce desired results. All related data fall into one of these four categories, characterized by how observations relate among the data sets:

- one-to-one
- one-to-many
- many-to-one
- many-to-many

To obtain the results you want, you should understand how each of these methods combines observations, how each method treats duplicate values of common variables,

and how each method treats missing values or nonmatched values of common variables. Some of the methods also require that you preprocess your data sets by sorting them or by creating indexes. See the description of each method in “Combining SAS Data Sets: Methods” on page 356.

One-to-One

In a one-to-one relationship, typically a single observation in one data set is related to a single observation from another based on the values of one or more selected variables. A one-to-one relationship implies that each value of the selected variable occurs no more than once in each data set. When you work with multiple selected variables, this relationship implies that each combination of values occurs no more than once in each data set.

In the following example, observations in data sets SALARY and TAXES are related by common values for EmployeeNumber.

Figure 22.1 One-to-One Relationship

SALARY			TAXES	
EmployeeNumber	Salary		EmployeeNumber	TaxBracket
1234	55000	↔	1111	0.18
3333	72000	↔	1234	0.28
4876	32000	↔	3333	0.32
5489	17000	↔	4222	0.18
			4876	0.24

One-to-Many and Many-to-One

A one-to-many or many-to-one relationship between input data sets implies that one data set has at most one observation with a specific value of the selected variable, but the other input data set can have more than one occurrence of each value. When you work with multiple selected variables, this relationship implies that each combination of values occurs no more than once in one data set, but can occur more than once in the other data set. The order in which the input data sets are processed determines whether the relationship is one-to-many or many-to-one.

In the following example, observations in data sets ONE and TWO are related by common values for variable A. Values of A are unique in data set ONE but not in data set TWO.

Figure 22.2 One-to-Many Relationship

ONE			TWO		
A	B	C	A	E	F
1	5	6	1	2	0
3	3	4	1	3	99
			1	4	88
			1	5	77
			2	1	66
			2	2	55
			3	4	44

In the following example, observations in data sets ONE, TWO, and THREE are related by common values for variable ID. Values of ID are unique in data sets ONE and THREE but not in TWO. For values 2 and 3 of ID, a one-to-many relationship exists between observations in data sets ONE and TWO, and a many-to-one relationship exists between observations in data sets TWO and THREE.

Figure 22.3 One-to-Many and Many-to-One Relationships

ONE		TWO		THREE	
ID	Name	ID	Sales	ID	Quota
1	Joe Smith	1	28000	1	15000
2	Sally Smith	2	30000	2	7000
3	Cindy Long	2	40000	3	15000
4	Sue Brown	3	15000	4	5000
5	Mike Jones	3	20000	5	8000
		3	25000		
		4	35000		
		5	40000		

Many-to-Many

The many-to-many category implies that multiple observations from each input data set can be related based on values of one or more common variables.

In the following example, observations in data sets BREAKDOWN and MAINTENANCE are related by common values for variable Vehicle. Values of Vehicle are not unique in either data set. A many-to-many relationship exists between observations in these data sets for values AAA and CCC of Vehicle.

Figure 22.4 Many-to-Many Relationship

BREAKDOWN		MAINTENANCE	
Vehicle	BreakDownDate	Vehicle	MaintenanceDate
AAA	02MAR99	AAA	03JAN99
AAA	20MAY99	AAA	05APR99
AAA	19JUN99	AAA	10AUG99
AAA	29NOV99	CCC	28JAN99
BBB	04JUL99	CCC	16MAY99
CCC	31MAY99	CCC	07OCT99
CCC	24DEC99	DDD	24FEB99
		DDD	22JUN99
		DDD	19SEP99

Access Methods: Sequential versus Direct

Overview

Once you have established data relationships, the next step is to determine the best mode of data access to relate the data. You can access observations sequentially in the order in which they appear in the physical file. Or you can access them directly, that is, you can go straight to an observation in a SAS data set without having to process each observation that precedes it.

Sequential Access

The simplest and perhaps most common way to process data with a DATA step is to read observations in a data set sequentially. You can read observations sequentially using the SET, MERGE, UPDATE, or MODIFY statements. You can also use the SAS File I/O functions, such as OPEN, FETCH, and FETCHOBS.

Direct Access

Direct access allows a program to access specific observations based on one of two methods:

- by an observation number
- by the value of one or more variables through a simple or composite index.

To access observations directly by their observation number, use the POINT= option with the SET or MODIFY statement. The POINT= option names a variable whose current value determines which observation a SET or MODIFY statement reads.

To access observations directly based on the values of one or more specified variables, you must first create an index for the variables and then read the data set using the KEY= statement option with the SET or MODIFY statement. An *index* is a separate structure that contains the data values of the key variable or variables, paired with a location identifier for the observations containing the value.

Note: You can also use the SAS File I/O functions such as CUROBS, NOTE, POINT, and FETCHOBS to access observations by observation number. Δ

Overview of Methods for Combining SAS Data Sets

You can use these methods to combine SAS data sets:

- concatenating
- interleaving
- one-to-one reading
- one-to-one merging
- match merging
- updating

Concatenating

The following figure shows the results of concatenating two SAS data sets. Concatenating the data sets appends the observations from one data set to another data set. The DATA step reads DATA1 sequentially until all observations have been processed, and then reads DATA2. Data set COMBINED contains the results of the concatenation. Note that the data sets are processed in the order in which they are listed in the SET statement.

Figure 22.5 Concatenating Two Data Sets

DATA1		DATA2		COMBINED
Year		Year		Year
1991		1991		1991
1992		1992		1992
1993	+	1993	=	1993
1994		1994		1994
1995		1995		1995

```
data combined;
  set data1 data2;
run;
```

1991
1992
1993
1994
1995

Interleaving

The following figure shows the results of interleaving two SAS data sets. Interleaving intersperses observations from two or more data sets, based on one or more common variables. Data set COMBINED shows the result.

Figure 22.6 Interleaving Two Data Sets

DATA1		DATA2		COMBINED
Year		Year		Year
1991		1992		1991
1992		1993		1992
1993	+	1994	=	1992
1994		1995		1993
1995		1996		1994

```
data combined;
  set data1 data2;
  by Year;
run;
```

1991		1992		1994
1992		1993		1995
1993		1994		1995
1994		1995		1996

One-to-One Reading and One-to-One Merging

The following figure shows the results of one-to-one reading and one-to-one merging. One-to-one reading combines observations from two or more SAS data sets by creating observations that contain all of the variables from each contributing data set. Observations are combined based on their relative position in each data set, that is, the first observation in one data set with the first in the other, and so on. The DATA step stops after it has read the last observation from the smallest data set. One-to-one merging is similar to a one-to-one reading, with two exceptions: you use the MERGE statement instead of multiple SET statements, and the DATA step reads all observations from all data sets. Data set COMBINED shows the result.

Figure 22.7 One-to-One Reading and One-to-One Merging

DATA1		DATA2		COMBINED
VarX		VarY		VarX VarY
X1		Y1		X1 Y1
X2		Y2		X2 Y2
X3	+	Y3	=	X3 Y3
X4		Y4		X4 Y4
X5		Y5		X5 Y5

```
data combined;
  set data1;
  set data2;
run;
```

```
data combined;
  merge data1 data2;
run;
```

Match-Merging

The following figure shows the results of match-merging. Match-merging combines observations from two or more SAS data sets into a single observation in a new data set based on the values of one or more common variables. Data set COMBINED shows the results.

Figure 22.8 Match-Merging Two Data Sets

DATA1			DATA2			COMBINED		
Year	VarX		Year	VarY		Year	VarX	VarY
1991	X1		1991	Y1		1991	X1	Y1
1992	X2		1991	Y2		1991	X1	Y2
1993	X3	+	1993	Y3	=	1992	X2	.
1994	X4		1994	Y4		1993	X3	Y3
1995	X5		1995	Y5		1994	X4	Y4
						1995	X5	Y5

```
data combined;
  merge data1 data2;
  by Year;
run;
```

Updating

The following figure shows the results of updating a master data set. Updating uses information from observations in a transaction data set to delete, add, or alter information in observations in a master data set. You can update a master data set by using the UPDATE statement or the MODIFY statement. If you use the UPDATE statement, your input data sets must be sorted by the values of the variables listed in the BY statement. (In this example, MASTER and TRANSACTION are both sorted by Year.) If you use the MODIFY statement, your input data does not need to be sorted.

UPDATE replaces an existing file with a new file, allowing you to add, delete, or rename columns. MODIFY performs an update in place by rewriting only those records that have changed, or by appending new records to the end of the file.

Note that by default, UPDATE and MODIFY do not replace nonmissing values in a master data set with missing values from a transaction data set.

Table 22.2 Statements or Procedures for Combining SAS Data Sets

Statement or Procedure	Action Performed	Access Method		Can Use with BY statement	Comments
		Sequential	Direct		
BY	Controls the operation of a SET, MERGE, UPDATE, or MODIFY statement in the DATA step and sets up special grouping variables.	NA	NA	NA	BY-group processing is a means of processing observations that have the same values of one or more variables.
MERGE	Reads observations from two or more SAS data sets and joins them into a single observation.	X		X	When using MERGE with BY, the data must be sorted or indexed on the BY variable.
MODIFY	Processes observations in a SAS data set in place. (Contrast with UPDATE.)	X	X	X	Sorted or indexed data are not required for use with BY, but are recommended for performance.
SET	Reads an observation from one or more SAS data sets.	X	X	X	Use KEY= or POINT= statement options for directly accessing data.
UPDATE	Applies transactions to observations in a master SAS data set. UPDATE does not update observations in place; it produces an updated copy of the current data set.	X		X	Both the master and transaction data sets must be sorted by or indexed on the BY variable.
PROC APPEND	Adds the observations from one SAS data set to the end of another SAS data set.	X			
PROC SQL ¹	Reads an observation from one or more SAS data sets; reads observations from up to 32 SAS data sets and joins them into single observations; manipulates observations in a SAS data set in place; easily produces a Cartesian product.	X	X	X	All three access methods are available in PROC SQL, but the access method is chosen by the internal optimizer.

¹ PROC SQL is the SAS implementation of Structured Query Language. In addition to expected SQL capabilities, PROC SQL includes additional capabilities specific to SAS, such as the use of formats and SAS macro language.

Using Error Checking

You can use the `_IORC_` automatic variable and the `SYSRC` autocall macro to perform error checking in a DATA step. Use these tools with the `MODIFY` statement or with the

SET statement and the KEY= option. For more information about these tools, see “Error Checking When Using Indexes to Randomly Access or Update Data” on page 382.

How to Prepare Your Data Sets

Before combining SAS data sets, follow these guidelines to produce the results you want:

- Know the structure and the contents of the data sets.
- Look at sources of common problems.
- Ensure that observations are in the correct order, or that they can be retrieved in the correct order (for example, by using an index).
- Test your program.

Knowing the Structure and Contents of the Data Sets

To help determine how your data is related, look at the structure of the data sets. To see the data set structure, execute the DATASETS procedure, the CONTENTS procedure, or access the SAS Explorer window in your windowing environment to display the descriptor information. Descriptor information includes the number of observations in each data set, the name and attributes of each variable, and which variables are included in indexes. To print a sample of the observations, use the PRINT procedure or the REPORT procedure.

You can also use functions such as VTYPE, VLENGTH, and VLENGTHX to show specific descriptor information. For complete information about these functions, see *SAS Language Reference: Dictionary*.

Looking at Sources of Common Problems

If your program does not execute correctly, review your input data for the following errors:

- variables that have the same name but that represent different data*
SAS includes only one variable of a given name in the new data set. If you are merging two data sets that have variables with the same names but different data, the values from the last data set that was read are written over the values from other data sets.

To correct the error, you can rename variables before you combine the data sets by using the RENAME= data set option in the SET, UPDATE, or MERGE statement, or you can use the DATASETS procedure.

- common variables with the same data but different attributes*
The way SAS handles these differences depends on which attributes are different:

- type attribute*

If the type attribute is different, SAS stops processing the DATA step and issues an error message stating that the variables are incompatible.

To correct this error, you must use a DATA step to re-create the variables. The SAS statements you use depend on the nature of the variable.

- length attribute*

If the length attribute is different, SAS takes the length from the first data set that contains the variable. In the following example, all data sets that are listed in the MERGE statement contain the variable Mileage. In

QUARTER1, the length of the variable Mileage is four bytes; in QUARTER2, it is eight bytes and in QUARTER3 and QUARTER4, it is six bytes. In the output data set YEARLY, the length of the variable Mileage is four bytes, which is the length derived from QUARTER1.

```
data yearly;
  merge quarter1 quarter2 quarter3 quarter4;
  by Account;
run;
```

To override the default and set the length yourself, specify the appropriate length in a LENGTH statement that *precedes* the SET, MERGE, or UPDATE statement.

Note: If the length of a variable changes as a result of combining data sets, SAS prints a warning message to the log and issues a nonzero return code (for example, on z/OS, SYSRC=4). If you expect truncation of data—for example, when removing insignificant blanks from the end of character values, the warning is expected and you do not want SAS to issue a nonzero return code. In this case, you can turn this warning off by setting the VARLENCHK system option to NOWARN. For more information, see the VARLENCHK system option in *SAS Language Reference: Dictionary*. Δ

□ *label, format, and informat attributes*

If any of these attributes are different, SAS takes the attribute from the first data set that contains the variable with that attribute. However, any label, format, or informat that you explicitly specify overrides a default. If all data sets contain explicitly specified attributes, the one specified in the first data set overrides the others. To ensure that the new output data set has the attributes you prefer, use an ATTRIB statement.

You can also use the SAS File I/O functions such as VLABEL, VLABELX, and other Variable Information functions to access this information. For complete information about these functions, see *SAS Language Reference: Dictionary*.

Ensuring Correct Order

If you use BY-group processing with the UPDATE, SET, and MERGE statements to combine data sets, ensure that the observations in the data sets are sorted in the order of the variables that are listed in the BY statement, or that the data sets have an appropriate index. If you use BY-group processing in a MODIFY statement, your data does not need to be sorted, but sorting the data improves efficiency. The BY variable or variables must be common to both data sets, and they must have the same attributes. For more information, see Chapter 21, “BY-Group Processing in the DATA Step,” on page 329.

Testing Your Program

As a final step in preparing your data sets, you should test your program. Create small temporary SAS data sets that contain a sample of observations that test all of your program’s logic. If your logic is faulty and you get unexpected output, you can use the DATA step debugger to debug your program. For complete information about the DATA Step Debugger, see *SAS Language Reference: Dictionary*.

Combining SAS Data Sets: Methods

Concatenating

Definition

Concatenating data sets is the combining of two or more data sets, one after the other, into a single data set. The number of observations in the new data set is the sum of the number of observations in the original data sets. The order of observations is sequential. All observations from the first data set are followed by all observations from the second data set, and so on.

In the simplest case, all input data sets contain the same variables. If the input data sets contain different variables, observations from one data set have missing values for variables defined only in other data sets. In either case, the variables in the new data set are the same as the variables in the old data sets.

Syntax

Use this form of the SET statement to concatenate data sets:

```
SET data-set(s);
```

where

data-set

specifies any valid SAS data set name.

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

DATA Step Processing during Concatenation

Compilation phase

SAS reads the descriptor information of each data set that is named in the SET statement and then creates a program data vector that contains all the variables from all data sets as well as variables created by the DATA step.

Execution — Step 1

SAS reads the first observation from the first data set into the program data vector. It processes the first observation and executes other statements in the DATA step. It then writes the contents of the program data vector to the new data set.

The SET statement does not reset the values in the program data vector to missing, except for variables whose value is calculated or assigned during the DATA step. Variables that are created by the DATA step are set to missing at the beginning of each iteration of the DATA step. Variables that are read from a data set are not.

Execution — Step 2

SAS continues to read one observation at a time from the first data set until it finds an end-of-file indicator. The values of the variables in the program data vector are then set to missing, and SAS begins reading observations from the second data set, and so on, until it reads all observations from all data sets.

Example 1: Concatenation Using the DATA Step

In this example, each data set contains the variables Common and Number, and the observations are arranged in the order of the values of Common. Generally, you concatenate SAS data sets that have the same variables. In this case, each data set also contains a unique variable to show the effects of combining data sets more clearly. The following shows the ANIMAL and the PLANT input data sets in the library that is referenced by the libref EXAMPLE:

ANIMAL				PLANT			
OBS	Common	Animal	Number	OBS	Common	Plant	Number
1	a	Ant	5	1	g	Grape	69
2	b	Bird		2	h	Hazelnut	55
3	c	Cat	17	3	i	Indigo	
4	d	Dog	9	4	j	Jicama	14
5	e	Eagle		5	k	Kale	5
6	f	Frog	76	6	l	Lentil	77

The following program uses a SET statement to concatenate the data sets and then prints the results:

```
libname example 'SAS-library';

data example.concatenation;
  set example.animal example.plant;
run;

proc print data=example.concatenation;
  var Common Animal Plant Number;
  title 'Data Set CONCATENATION';
run;
```

Output 22.1 Concatenated Data Sets (DATA Step)

Data Set CONCATENATION						1
Obs	Common	Animal	Plant	Number		
1	a	Ant		5		
2	b	Bird		.		
3	c	Cat		17		
4	d	Dog		9		
5	e	Eagle		.		
6	f	Frog		76		
7	g		Grape	69		
8	h		Hazelnut	55		
9	i		Indigo	.		
10	j		Jicama	14		
11	k		Kale	5		
12	l		Lentil	77		

The resulting data set CONCATENATION has 12 observations, which is the sum of the observations from the combined data sets. The program data vector contains all variables from all data sets. The values of variables found in one data set but not in another are set to missing.

Example 2: Concatenation Using SQL

You can also use the SQL language to concatenate tables. In this example, SQL reads each row in both tables and creates a new table named COMBINED. The following shows the YEAR1 and YEAR2 input tables:

YEAR1	YEAR2
Date1	Date2
1996	
1997	1997
1998	1998
1999	1999
	2000
	2001

The following SQL code creates and prints the table COMBINED.

```
proc sql;
  title 'SQL Table COMBINED';
  create table combined as
    select * from year1
  union all
    select * from year2;
  select * from combined;
quit;
```

Output 22.2 Concatenated Tables (SQL)

SQL Table COMBINED		1
	Year	

	1996	
	1997	
	1998	
	1999	
	1997	
	1998	
	1999	
	2000	
	2001	

Appending Files

Instead of concatenating data sets or tables, you can append them and produce the same results as concatenation. SAS concatenates data sets (DATA step) and tables (SQL) by reading each row of data to create a new file. To avoid reading all the records, you can append the second file to the first file by using the APPEND procedure:

```
proc append base=year1 data=year2;
run;
```

The YEAR1 file will contain all rows from both tables.

Note: You cannot use PROC APPEND to add observations to a SAS data set in a sequential library. Δ

Efficiency

If no additional processing is necessary, using PROC APPEND or the APPEND statement in PROC DATASETS is more efficient than using a DATA step to concatenate data sets.

Interleaving

Definition

Interleaving uses a SET statement and a BY statement to combine multiple data sets into one new data set. The number of observations in the new data set is the sum of the number of observations from the original data sets. However, the observations in the new data set are arranged by the values of the BY variable or variables and, within each BY group, by the order of the data sets in which they occur. You can interleave data sets either by using a BY variable or by using an index.

Syntax

Use this form of the SET statement to interleave data sets when you use a BY variable:

```
SET data-set(s);
```

```
BY variable(s);
```

where

data-set

specifies a one-level name, a two-level name, or one of the special SAS data set names.

variable

specifies each variable by which the data set is sorted. These variables are referred to as BY variables for the current DATA or PROC step.

Use this form of the SET statement to interleave data sets when you use an index:

```
SET data-set-1 . . . data-set-n KEY= index;
```

where

data-set

specifies a one-level name, a two-level name, or one of the special SAS data set names.

index

provides nonsequential access to observations in a SAS data set, which are based on the value of an index variable or key.

For a complete description of the SET statement, including SET with the KEY= option, see *SAS Language Reference: Dictionary*.

Sort Requirements

Before you can interleave data sets, the observations must be sorted or grouped by the same variable or variables that you use in the BY statement, or you must have an appropriate index for the data sets.

DATA Step Processing during Interleaving

Compilation phase

- SAS reads the descriptor information of each data set that is named in the SET statement and then creates a program data vector that contains all the variables from all data sets as well as variables created by the DATA step.
- SAS creates the *FIRST.variable* and *LAST.variable* for each variable listed in the BY statement.

Execution — Step 1

SAS compares the first observation from each data set that is named in the SET statement to determine which BY group should appear first in the new data set. It reads all observations from the first BY group from the selected data set. If this BY group appears in more than one data set, it reads from the data sets in the order in which they appear in the SET statement. The values of the variables in the program data vector are set to missing each time SAS starts to read a new data set and when the BY group changes.

Execution — Step 2

SAS compares the next observations from each data set to determine the next BY group and then starts reading observations from the selected data set in the SET statement that contains observations for this BY group. SAS continues until it has read all observations from all data sets.

Example 1: Interleaving in the Simplest Case

In this example, each data set contains the BY variable Common, and the observations are arranged in order of the values of the BY variable. The following shows the ANIMAL and the PLANT input data sets in the library that is referenced by the libref EXAMPLE:

ANIMAL			PLANT		
OBS	Common	Animal	OBS	Common	Plant
1	a	Ant	1	a	Apple
2	b	Bird	2	b	Banana
3	c	Cat	3	c	Coconut
4	d	Dog	4	d	Dewberry
5	e	Eagle	5	e	Eggplant
6	f	Frog	6	f	Fig

The following program uses SET and BY statements to interleave the data sets, and prints the results:

```
data example.interleaving;
  set example.animal example.plant;
  by Common;
run;

proc print data=example.interleaving;
  title 'Data Set INTERLEAVING';
run;
```

Output 22.3 Interleaved Data Sets

Data Set INTERLEAVING				1
Obs	common	animal	plant	
1	a	Ant		
2	a		Apple	
3	b	Bird		
4	b		Banana	
5	c	Cat		
6	c		Coconut	
7	d	Dog		
8	d		Dewberry	
9	e	Eagle		
10	e		Eggplant	
11	f	Frog		
12	f		Fig	

The resulting data set INTERLEAVING has 12 observations, which is the sum of the observations from the combined data sets. The new data set contains all variables from both data sets. The value of variables found in one data set but not in the other are set to missing, and the observations are arranged by the values of the BY variable.

Example 2: Interleaving with Duplicate Values of the BY variable

If the data sets contain duplicate values of the BY variables, the observations are written to the new data set in the order in which they occur in the original data sets. This example contains duplicate values of the BY variable Common. The following shows the ANIMAL1 and PLANT1 input data sets:

ANIMAL1			PLANT1		
OBS	Common	Animall	OBS	Common	Plant1
1	a	Ant	1	a	Apple
2	a	Ape	2	b	Banana
3	b	Bird	3	c	Coconut
4	c	Cat	4	c	Celery
5	d	Dog	5	d	Dewberry
6	e	Eagle	6	e	Eggplant

The following program uses SET and BY statements to interleave the data sets, and prints the results:

```
data example.interleaving2;
  set example.animall example.plant1;
  by Common;
run;

proc print data=example.interleaving2;
  title 'Data Set INTERLEAVING2: Duplicate BY Values';
run;
```

Output 22.4 Interleaved Data Sets with Duplicate Values of the BY Variable

Data Set INTERLEAVING2: Duplicate BY Values				1
Obs	Common	Animal1	Plant1	
1	a	Ant		
2	a	Ape		
3	a		Apple	
4	b	Bird		
5	b		Banana	
6	c	Cat		
7	c		Coconut	
8	c		Celery	
9	d	Dog		
10	d		Dewberry	
11	e	Eagle		
12	e		Eggplant	

The number of observations in the new data set is the sum of the observations in all the data sets. The observations are written to the new data set in the order in which they occur in the original data sets.

Example 3: Interleaving with Different BY Values in Each Data Set

The data sets ANIMAL2 and PLANT2 both contain BY values that are present in one data set but not in the other. The following shows the ANIMAL2 and the PLANT2 input data sets:

ANIMAL2			PLANT2		
OBS	Common	Animal2	OBS	Common	Plant2
1	a	Ant	1	a	Apple
2	c	Cat	2	b	Banana
3	d	Dog	3	c	Coconut
4	e	Eagle	4	e	Eggplant
			5	f	Fig

This program uses SET and BY statements to interleave these data sets, and prints the results:

```
data example.interleaving3;
  set example.animal2 example.plant2;
  by Common;
run;

proc print data=example.interleaving3;
  title 'Data Set INTERLEAVING3: Different BY Values';
run;
```

Output 22.5 Interleaving Data Sets with Different BY Values

Data Set INTERLEAVING3: Different BY Values				1
Obs	Common	Animal2	Plant2	
1	a	Ant		
2	a		Apple	
3	b		Banana	
4	c	Cat		
5	c		Coconut	
6	d	Dog		
7	e	Eagle		
8	e		Eggplant	
9	f		Fig	

The resulting data set has nine observations arranged by the values of the BY variable.

Comments and Comparisons

- In other languages, the term *merge* is often used to mean *interleave*. SAS reserves the term *merge* for the operation in which observations from two or more data sets are combined into one observation. The observations in interleaved data sets are not combined; they are copied from the original data sets in the order of the values of the BY variable.
- If one table has multiple rows with the same BY value, the DATA step preserves the order of those rows in the result.
- To use the DATA step, the input tables must be appropriately sorted or indexed. SQL does not require the input tables to be in order.

One-to-One Reading**Definition**

One-to-one reading combines observations from two or more data sets into one observation by using two or more SET statements to read observations independently from each data set. This process is also called *one-to-one matching*. The new data set contains all the variables from all the input data sets. The number of observations in the new data set is the number of observations in the smallest original data set. If the data sets contain common variables, the values that are read in from the last data set replace the values that were read in from earlier data sets.

Syntax

Use this form of the SET statement for one-to-one reading:

```
SET data-set-1;
```

```
SET data-set-2;
```

where

data-set-1

specifies a one-level name, a two-level name, or one of the special SAS data set names. *data-set-1* is the first file that the DATA step reads.

data-set-2

specifies a one-level name, a two-level name, or one of the special SAS data set names. *data-set-2* is the second file that the DATA step reads.

CAUTION:

Use care when you combine data sets with multiple SET statements. Using multiple SET statements to combine observations can produce undesirable results. Test your program on representative samples of the data sets before using this method to combine them. \triangle

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

DATA Step Processing during a One-to-One Reading

Compilation phase

SAS reads the descriptor information of each data set named in the SET statement and then creates a program data vector that contains all the variables from all data sets as well as variables created by the DATA step.

Execution — Step 1

When SAS executes the first SET statement, SAS reads the first observation from the first data set into the program data vector. The second SET statement reads the first observation from the second data set into the program data vector. If both data sets contain the same variables, the values from the second data set replace the values from the first data set, even if the value is missing. After reading the first observation from the last data set and executing any other statements in the DATA step, SAS writes the contents of the program data vector to the new data set. The SET statement does not reset the values in the program data vector to missing, except for those variables that were created or assigned values during the DATA step.

Execution — Step 2

SAS continues reading from one data set and then the other until it detects an end-of-file indicator in one of the data sets. SAS stops processing with the last observation of the shortest data set and does not read the remaining observations from the longer data set.

Example 1: One-to-One Reading: Processing an Equal Number of Observations

The SAS data sets ANIMAL and PLANT both contain the variable Common, and are arranged by the values of that variable. The following shows the ANIMAL and the PLANT input data sets:

ANIMAL			PLANT		
OBS	Common	Animal	OBS	Common	Plant
1	a	Ant	1	a	Apple
2	b	Bird	2	b	Banana

3	c	Cat	3	c	Coconut
4	d	Dog	4	d	Dewberry
5	e	Eagle	5	e	Eggplant
6	f	Frog	6	g	Fig

The following program uses two SET statements to combine observations from ANIMAL and PLANT, and prints the results:

```
data twosets;
  set animal;
  set plant;
run;

proc print data=twosets;
  title 'Data Set TWOSSETS - Equal Number of Observations';
run;
```

Output 22.6 Data Set Created from Two Data Sets That Have Equal Observations

Data Set TWOSSETS - Equal Number of Observations				1
Obs	Common	Animal	Plant	
1	a	Ant	Apple	
2	b	Bird	Banana	
3	c	Cat	Coconut	
4	d	Dog	Dewberry	
5	e	Eagle	Eggplant	
6	g	Frog	Fig	

Each observation in the new data set contains all the variables from all the data sets. Note, however, that the Common variable value in observation 6 contains a “g.” The value of Common in observation 6 of the ANIMAL data set was overwritten by the value in PLANT, which was the data set that SAS read last.

Comments and Comparisons

- The results that are obtained by reading observations using two or more SET statements are similar to those that are obtained by using the MERGE statement with no BY statement. However, with one-to-one reading, SAS stops processing before all observations are read from all data sets if the number of observations in the data sets is not equal.
- Using multiple SET statements with other DATA step statements makes the following applications possible:
 - merging one observation with many
 - conditionally merging observations
 - reading from the same data set twice

One-to-One Merging

Definition

One-to-one merging combines observations from two or more SAS data sets into a single observation in a new data set. To perform a one-to-one merge, use the MERGE statement without a BY statement. SAS combines the first observation from all data sets in the MERGE statement into the first observation in the new data set, the second observation from all data sets into the second observation in the new data set, and so on. In a one-to-one merge, the number of observations in the new data set equals the number of observations in the largest data set that was named in the MERGE statement.

If you use the MERGENOBY= SAS system option, you can control whether SAS issues a message when MERGE processing occurs without an associated BY statement.

Syntax

Use this form of the MERGE statement to merge SAS data sets:

```
MERGE data-set(s);
```

where

data-set

names at least two existing SAS data sets.

CAUTION:

Avoid using duplicate values or different values of common variables. One-to-one merging with data sets that contain duplicate values of common variables can produce undesirable results. If a variable exists in more than one data set, the value from the last data set that is read is the one that is written to the new data set. The variables are combined exactly as they are read from each data set. Using a one-to-one merge to combine data sets with different values of common variables can also produce undesirable results. If a variable exists in more than one data set, the value from the last data set read is the one that is written to the new data set even if the value is missing. Once SAS has processed all observations in a data set, all subsequent observations in the new data set have missing values for the variables that are unique to that data set. \triangle

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

DATA Step Processing during One-to-One Merging

Compilation phase

SAS reads the descriptor information of each data set that is named in the MERGE statement and then creates a program data vector that contains all the variables from all data sets as well as variables created by the DATA step.

Execution — Step 1

SAS reads the first observation from each data set into the program data vector, reading the data sets in the order in which they appear in the MERGE statement. If two data sets contain the same variables, the values from the second data set

replace the values from the first data set. After reading the first observation from the last data set and executing any other statements in the DATA step, SAS writes the contents of the program data vector to the new data set. Only those variables that are created or assigned values during the DATA step are set to missing.

Execution — Step 2

SAS continues until it has read all observations from all data sets.

Example 1: One-to-One Merging with an Equal Number of Observations

The SAS data sets ANIMAL and PLANT both contain the variable Common, and the observations are arranged by the values of Common. The following shows the ANIMAL and the PLANT input data sets:

ANIMAL			PLANT		
OBS	Common	Animal	OBS	Common	Plant
1	a	Ant	1	a	Apple
2	b	Bird	2	b	Banana
3	c	Cat	3	c	Coconut
4	d	Dog	4	d	Dewberry
5	e	Eagle	5	e	Eggplant
6	f	Frog	6	g	Fig

The following program merges these data sets and prints the results:

```
data combined;
  merge animal plant;
run;

proc print data=combined;
  title 'Data Set COMBINED';
run;
```

Output 22.7 Merged Data Sets That Have an Equal Number of Observations

Data Set COMBINED					1
Obs	Common	Animal	Plant		
1	a	Ant	Apple		
2	b	Bird	Banana		
3	c	Cat	Coconut		
4	d	Dog	Dewberry		
5	e	Eagle	Eggplant		
6	g	Frog	Fig		

Each observation in the new data set contains all variables from all data sets. If two data sets contain the same variables, the values from the second data set replace the values from the first data set, as shown in observation 6.

Example 2: One-to-One Merging with an Unequal Number of Observations

The SAS data sets ANIMAL1 and PLANT1 both contain the variable Common, and the observations are arranged by the values of Common. The PLANT1 data set has

fewer observations than the ANIMAL1 data set. The following shows the ANIMAL1 and the PLANT1 input data sets:

ANIMAL1			PLANT1		
OBS	Common	Animal	OBS	Common	Plant
1	a	Ant	1	a	Apple
2	b	Bird	2	b	Banana
3	c	Cat	3	c	Coconut
4	d	Dog			
5	e	Eagle			
6	f	Frog			

The following program merges these unequal data sets and prints the results:

```
data combined1;
    merge animall plant1;
run;

proc print data=combined1;
    title 'Data Set COMBINED1';
run;
```

Output 22.8 Merged Data Sets That Have an Unequal Number of Observations

Data Set COMBINED1					1
Obs	Common	Animal	Plant		
1	a	Ant	Apple		
2	b	Bird	Banana		
3	c	Cat	Coconut		
4	d	Dog			
5	e	Eagle			
6	f	Frog			

Note that observations 4 through 6 contain missing values for the variable Plant.

Example 3: One-to-One Merging with Duplicate Values of Common Variables

The following example shows the undesirable results that you can obtain by using one-to-one merging with data sets that contain duplicate values of common variables. The value from the last data set that is read is the one that is written to the new data set. The variables are combined exactly as they are read from each data set. In the following example, the data sets ANIMAL1 and PLANT1 contain the variable Common, and each data set contains observations with duplicate values of Common. The following shows the ANIMAL1 and the PLANT1 input data sets:

ANIMAL1			PLANT1		
OBS	Common	Animal	OBS	Common	Plant
1	a	Ant	1	a	Apple
2	a	Ape	2	b	Banana
3	b	Bird	3	c	Coconut

4	c	Cat	4	c	Celery
5	d	Dog	5	d	Dewberry
6	e	Eagle	6	e	Eggplant

The following program produces the data set MERGE1 data set and prints the results:

```

/* This program illustrates undesirable results. */
data merge1;
  merge animal1 plant1;
run;

proc print data=merge1;
  title 'Data Set MERGE1';
run;

```

Output 22.9 Undesirable Results with Duplicate Values of Common Variables

Data Set MERGE1				1
Obs	Common	Animal1	Plant1	
1	a	Ant	Apple	
2	b	Ape	Banana	
3	c	Bird	Coconut	
4	c	Cat	Celery	
5	d	Dog	Dewberry	
6	e	Eagle	Eggplant	

The number of observations in the new data set is six. Note that observations 2 and 3 contain undesirable values. SAS reads the second observation from data set ANIMAL1. It then reads the second observation from data set PLANT1 and replaces the values for the variables Common and Plant1. The third observation is created in the same way.

Example 4: One-to-One Merging with Different Values of Common Variables

The following example shows the undesirable results obtained from using the one-to-one merge to combine data sets with different values of common variables. If a variable exists in more than one data set, the value from the last data set that is read is the one that is written to the new data set even if the value is missing. Once SAS processes all observations in a data set, all subsequent observations in the new data set have missing values for the variables that are unique to that data set. In this example, the data sets ANIMAL2 and PLANT2 have different values of the Common variable. The following shows the ANIMAL2 and the PLANT2 input data sets:

ANIMAL2			PLANT2		
OBS	Common	Animal	OBS	Common	Plant
1	a	Ant	1	a	Apple
2	c	Cat	2	b	Banana
3	d	Dog	3	c	Coconut
4	e	Eagle	4	e	Eggplant
			5	f	Fig

The following program produces the data set MERGE2 and prints the results:

```

/* This program illustrates undesirable results. */
data merge2;
  merge animal2 plant2;
run;

proc print data=merge2;
  title 'Data Set MERGE2';
run;

```

Output 22.10 Undesirable Results with Different Values of Common Variables

Data Set MERGE2				1
Obs	Common	Animal2	Plant2	
1	a	Ant	Apple	
2	b	Cat	Banana	
3	c	Dog	Coconut	
4	e	Eagle	Eggplant	
5	f		Fig	

Comments and Comparisons

The results from a one-to-one merge are similar to the results obtained from using two or more SET statements to combine observations. However, with the one-to-one merge, SAS continues processing all observations in all data sets that were named in the MERGE statement.

Match-Merging

Definition

Match-merging combines observations from two or more SAS data sets into a single observation in a new data set according to the values of a common variable. The number of observations in the new data set is the sum of the largest number of observations in each BY group in all data sets. To perform a match-merge, use the MERGE statement with a BY statement. Before you can perform a match-merge, all data sets must be sorted by the variables that you specify in the BY statement or they must have an index.

Syntax

Use this form of the MERGE statement to match-merge data sets:

```
MERGE data-set(s);
```

```
BY variable(s);
```

where

data-set

names at least two existing SAS data sets from which observations are read.

variable

names each variable by which the data set is sorted or indexed. These variables are referred to as BY variables.

For a complete description of the MERGE and the BY statements, see *SAS Language Reference: Dictionary*.

DATA Step Processing during Match-Merging

Compilation phase

SAS reads the descriptor information of each data set that is named in the MERGE statement and then creates a program data vector that contains all the variables from all data sets as well as variables created by the DATA step. SAS creates the *FIRST.variable* and *LAST.variable* for each variable that is listed in the BY statement.

Execution – Step 1

SAS looks at the first BY group in each data set that is named in the MERGE statement to determine which BY group should appear first in the new data set. The DATA step reads into the program data vector the first observation in that BY group from each data set, reading the data sets in the order in which they appear in the MERGE statement. If a data set does not have observations in that BY group, the program data vector contains missing values for the variables unique to that data set.

Execution – Step 2

After processing the first observation from the last data set and executing other statements, SAS writes the contents of the program data vector to the new data set. SAS retains the values of all variables in the program data vector except those variables that were created by the DATA step; SAS sets those values to missing. SAS continues to merge observations until it writes all observations from the first BY group to the new data set. When SAS has read all observations in a BY group from all data sets, it sets all variables in the program data vector (except those created by SAS) to missing. SAS looks at the next BY group in each data set to determine which BY group should appear next in the new data set.

Execution – Step 3

SAS repeats these steps until it reads all observations from all BY groups in all data sets.

Example 1: Combining Observations Based on a Criterion

The SAS data sets ANIMAL and PLANT each contain the BY variable Common, and the observations are arranged in order of the values of the BY variable. The following shows the ANIMAL and the PLANT input data sets:

ANIMAL			PLANT		
OBS	Common	Animal	OBS	Common	Plant
1	a	Ant	1	a	Apple
2	b	Bird	2	b	Banana
3	c	Cat	3	c	Coconut
4	d	Dog	4	d	Dewberry
5	e	Eagle	5	e	Eggplant
6	f	Frog	6	f	Fig

The following program merges the data sets according to the values of the BY variable Common, and prints the results:

```
data combined;
  merge animal plant;
  by Common;
run;

proc print data=combined;
  title 'Data Set COMBINED';
run;
```

Output 22.11 Data Sets Combined by Match-Merging

Data Set COMBINED				1
Obs	Common	Animal	Plant	
1	a	Ant	Apple	
2	b	Bird	Banana	
3	c	Cat	Coconut	
4	d	Dog	Dewberry	
5	e	Eagle	Eggplant	
6	f	Frog	Fig	

Each observation in the new data set contains all the variables from all the data sets.

Example 2: Match-Merge with Duplicate Values of the BY Variable

When SAS reads the last observation from a BY group in one data set, SAS retains its values in the program data vector for all variables that are unique to that data set until all observations for that BY group have been read from all data sets. In the following example, the data sets ANIMAL1 and PLANT1 contain duplicate values of the BY variable Common. The following shows the ANIMAL1 and the PLANT1 input data sets:

ANIMAL1			PLANT1		
OBS	Common	Animall	OBS	Common	Plant1
1	a	Ant	1	a	Apple
2	a	Ape	2	b	Banana
3	b	Bird	3	c	Coconut
4	c	Cat	4	c	Celery
5	d	Dog	5	d	Dewberry
6	e	Eagle	6	e	Eggplant

The following program produces the merged data set MATCH1, and prints the results:

```
data match1;
  merge animall plant1;
  by Common;
run;

proc print data=match1;
  title 'Data Set MATCH1';
run;
```

Output 22.12 Match-Merged Data Set with Duplicate BY Values

Data Set MATCH1				1
Obs	Common	Animal1	Plant1	
1	a	Ant	Apple	
2	a	Ape	Apple	
3	b	Bird	Banana	
4	c	Cat	Coconut	
5	c	Cat	Celery	
6	d	Dog	Dewberry	
7	e	Eagle	Eggplant	

In observation 2 of the output, the value of the variable Plant1 is retained until all observations in the BY group are written to the new data set. Match-merging also produced duplicate values in ANIMAL1 for observations 4 and 5.

Example 3: Match-Merge with Nonmatched Observations

When SAS performs a match-merge with nonmatched observations in the input data sets, SAS retains the values of all variables in the program data vector even if the value is missing. The data sets ANIMAL2 and PLANT2 do not contain all values of the BY variable Common. The following shows the ANIMAL2 and the PLANT2 input data sets:

ANIMAL2			PLANT2		
OBS	Common	Animal2	OBS	Common	Plant2
1	a	Ant	1	a	Apple
2	c	Cat	2	b	Banana
3	d	Dog	3	c	Coconut
4	e	Eagle	4	e	Eggplant
			5	f	Fig

The following program produces the merged data set MATCH2, and prints the results:

```
data match2;
  merge animal2 plant2;
  by Common;
run;

proc print data=match2;
  title 'Data Set MATCH2';
run;
```

Output 22.13 Match-Merged Data Set with Nonmatched Observations

Data Set MATCH2				1
Obs	Common	Animal2	Plant2	
1	a	Ant	Apple	
2	b		Banana	
3	c	Cat	Coconut	
4	d	Dog		
5	e	Eagle	Eggplant	
6	f		Fig	

As the output shows, all values of the variable Common are represented in the new data set, including missing values for the variables that are in one data set but not in the other.

Updating with the UPDATE and the MODIFY Statements

Definitions

Updating a data set refers to the process of applying changes to a master data set. To update data sets, you work with two input data sets. The data set containing the original information is the *master data set*, and the data set containing the new information is the *transaction data set*.

You can update data sets by using the UPDATE statement or the MODIFY statement:

UPDATE	uses observations from the transaction data set to change the values of corresponding observations from the master data set. You must use a BY statement with the UPDATE statement because all observations in the transaction data set are keyed to observations in the master data set according to the values of the BY variable.
MODIFY	can replace, delete, and append observations in an existing data set. Using the MODIFY statement can save disk space because it modifies data in place, without creating a copy of the data set.

The number of observations in the new data set is the sum of the number of observations in the master data set and the number of unmatched observations in the transaction data set.

For complete information about the UPDATE and the MODIFY statements, see *SAS Language Reference: Dictionary*.

Syntax of the UPDATE Statement

Use this form of the UPDATE statement to update a master data set:

```
UPDATE master-data-set transaction-data-set;
```

```
BY variable-list;
```

where

```
master-data-set
```

names the SAS data set that is used as the master file.

transaction-data-set

names the SAS data set that contains the changes to be applied to the master data set.

variable-list

specifies the variables by which observations are matched.

If the transaction data set contains duplicate values of the BY variable, SAS applies both transactions to the observation. The last values that are copied into the program data vector are written to the new data set. If your data is in this form, use the MODIFY statement instead of the UPDATE statement to process your data.

CAUTION:

Values of the BY variable must be unique for each observation in the master data set. If the master data set contains two observations with the same value of the BY variable, the first observation is updated and the second observation is ignored. SAS writes a warning message to the log when the DATA step executes. △

For complete information about the UPDATE statement, see *SAS Language Reference: Dictionary*.

Syntax of the MODIFY Statement

This form of the MODIFY statement is used in the examples that follow:

```
MODIFY master-data-set;
```

```
BY variable-list;
```

where

master-data-set

specifies the SAS data set that you want to modify.

variable-list

names each variable by which the data set is ordered.

Note: The MODIFY statement does not support changing the descriptor portion of a SAS data set, such as adding a variable. △

For complete information about the MODIFY statement, see *SAS Language Reference: Dictionary*.

DATA Step Processing with the UPDATE Statement*Compilation phase*

- SAS reads the descriptor information of each data set that is named in the UPDATE statement and creates a program data vector that contains all the variables from all data sets as well as variables created by the DATA step.
- SAS creates the *FIRST.variable* and *LAST.variable* for each variable that is listed in the BY statement.

Execution – Step 1

SAS looks at the first observation in each data set that is named in the UPDATE statement to determine which BY group should appear first. If the transaction BY value precedes the master BY value, SAS reads from the transaction data set only and sets the variables from the master data set to missing. If the master BY value precedes the transaction BY value, SAS reads from the master data set only and sets the unique variables from the transaction data set to missing. If the BY

values in the master and transaction data sets are equal, it applies the first transaction by copying the nonmissing values into the program data vector.

Execution – Step 2

After completing the first transaction, SAS looks at the next observation in the transaction data set. If SAS finds one with the same BY value, it applies that transaction too. The first observation then contains the new values from both transactions. If no other transactions exist for that observation, SAS writes the observation to the new data set and sets the values in the program data vector to missing. SAS repeats these steps until it has read all observations from all BY groups in both data sets.

Updating with Nonmatched Observations, Missing Values, and New Variables

In the UPDATE statement, if an observation in the master data set does not have a corresponding observation in the transaction data set, SAS writes the observation to the new data set without modifying it. Any observation from the transaction data set that does not correspond to an observation in the master data set is written to the program data vector and becomes the basis for an observation in the new data set. The data in the program data vector can be modified by other transactions before it is written to the new data set. If a master data set observation does not need updating, the corresponding observation can be omitted from the transaction data set.

SAS does not replace existing values in the master data set with missing values if those values are coded as periods (for numeric variables) or blanks (for character variables) in the transaction data set. To replace existing values with missing values, you must either create a transaction data set in which missing values are coded with the special missing value characters, or use the UPDATEMODE=NOMISSINGCHECK statement option.

With UPDATE, the transaction data set can contain new variables to be added to all observations in the master data set.

To view a sample program, see “Example 3: Using UPDATE for Processing Nonmatched Observations, Missing Values, and New Variables” on page 379.

Sort Requirements for the UPDATE Statement

If you do not use an index, both the master data set and the transaction data set must be sorted by the same variable or variables that you specify in the BY statement that accompanies the UPDATE statement. The values of the BY variable should be unique for each observation in the master data set. If you use more than one BY variable, the combination of values of all BY variables should be unique for each observation in the master data set. The BY variable or variables should be ones that you never need to update.

Note: The MODIFY statement does not require sorted files. However, sorting the data improves efficiency. Δ

Using an Index with the MODIFY Statement

The MODIFY statement maintains the index. You do not have to rebuild the index like you do for the UPDATE statement.

Choosing between UPDATE or MODIFY with BY

Using the UPDATE statement is comparable to using MODIFY with BY to apply transactions to a data set. While MODIFY is a more powerful tool with several other

applications, UPDATE is still the tool of choice in some cases. The following table helps you choose whether to use UPDATE or MODIFY with BY.

Table 22.3 MODIFY with BY versus UPDATE

Issue	MODIFY with BY	UPDATE
Disk space	saves disk space because it updates data in place	requires more disk space because it produces an updated copy of the data set
Sort and index	sorted input data sets are not required, although for good performance, it is strongly recommended that both data sets be sorted and that the master data set be indexed	requires only that both data sets be sorted
When to use	use only when you expect to process a SMALL portion of the data set	use if you expect to need to process most of the data set
Where to specify the modified data set	specify the updated data set in both the DATA and the MODIFY statements	specify the updated data set in the DATA and the UPDATE statements
Duplicate BY-values	allows duplicate BY-values in both the master and the transaction data sets	allows duplicate BY-values in the transaction data set only (If duplicates exist in the master data set, SAS issues a warning.)
Scope of changes	cannot change the data set descriptor information, so changes such as adding or deleting variables, variable labels, and so on, are not valid	can make changes that require a change in the descriptor portion of a data set, such as adding new variables, and so on
Error checking	has error-checking capabilities using the <code>_IORC_</code> automatic variable and the <code>SYSRC</code> autocall macro	needs no error checking because transactions without a corresponding master record are not applied but are added to the data set
Data set integrity	data might be only partially updated due to an abnormal task termination	no data loss occurs because UPDATE works on a copy of the data

For more information about tools for combining SAS data sets, see Table 22.2 on page 353.

Primary Uses of the MODIFY Statement

The MODIFY statement has three primary uses:

- modifying observations in a single SAS data set.
- modifying observations in a single SAS data set directly, either by observation number or by values in an index.
- modifying observations in a master data set, based on values in a transaction data set. MODIFY with BY is similar to using the UPDATE statement.

Several of the examples that follow demonstrate these uses.

Example 1: Using UPDATE for Basic Updating

In this example, the data set MASTER contains original values of the variables Animal and Plant. The data set NEWPLANT is a transaction data set with new values

of the variable Plant. The following shows the MASTER and the NEWPLANT input data sets:

MASTER				NEWPLANT		
OBS	Common	Animal	Plant	OBS	Common	Plant
1	a	Ant	Apple	1	a	Apricot
2	b	Bird	Banana	2	b	Barley
3	c	Cat	Coconut	3	c	Cactus
4	d	Dog	Dewberry	4	d	Date
5	e	Eagle	Eggplant	5	e	Escarole
6	f	Frog	Fig	6	f	Fennel

The following program updates MASTER with the transactions in the data set NEWPLANT, writes the results to UPDATE_FILE, and prints the results:

```
data update_file;
  update master newplant;
  by common;
run;

proc print data=update_file;
  title 'Data Set Update_File';
run;
```

Output 22.14 Master Data Set Updated by Transaction Data Set

Data Set Update_File				1
Obs	Common	Animal	Plant	
1	a	Ant	Apricot	
2	b	Bird	Barley	
3	c	Cat	Cactus	
4	d	Dog	Date	
5	e	Eagle	Escarole	
6	f	Frog	Fennel	

Each observation in the new data set contains a new value for the variable Plant.

Example 2: Using UPDATE with Duplicate Values of the BY Variable

If the master data set contains two observations with the same value of the BY variable, the first observation is updated and the second observation is ignored. SAS writes a warning message to the log. If the transaction data set contains duplicate values of the BY variable, SAS applies both transactions to the observation. The last values copied into the program data vector are written to the new data set. The following shows the MASTER1 and the DUPPLANT input data sets.

MASTER1				DUPPLANT		
OBS	Common	Animall	Plantl	OBS	Common	Plantl
1	a	Ant	Apple	1	a	Apricot

2	b	Bird	Banana	2	b	Barley
3	b	Bird	Banana	3	c	Cactus
4	c	Cat	Coconut	4	d	Date
5	d	Dog	Dewberry	5	d	Dill
6	e	Eagle	Eggplant	6	e	Escarole
7	f	Frog	Fig	7	f	Fennel

The following program applies the transactions in DUPPLANT to MASTER1 and prints the results:

```
data update1;
  update master1 dupplant;
  by Common;
run;

proc print data=update1;
  title 'Data Set Update1';
run;
```

Output 22.15 Updating Data Sets with Duplicate BY Values

Data Set Update1				1
Obs	Common	Animal1	Plant1	
1	a	Ant	Apricot	
2	b	Bird	Barley	
3	b	Bird	Banana	
4	c	Cat	Cactus	
5	d	Dog	Dill	
6	e	Eagle	Escarole	
7	f	Frog	Fennel	

When this DATA step executes, SAS generates a warning message stating that there is more than one observation for a BY group. However, the DATA step continues to process, and the data set UPDATE1 is created.

The resulting data set has seven observations. Observations 2 and 3 have duplicate values of the BY variable Common. However, the value of the variable PLANT1 was not updated in the second occurrence of the duplicate BY value.

Example 3: Using UPDATE for Processing Nonmatched Observations, Missing Values, and New Variables

In this example, the data set MASTER2 is a master data set. It contains a missing value for the variable Plant2 in the first observation, and not all of the values of the BY variable Common are included. The transaction data set NONPLANT contains a new variable Mineral, a new value of the BY variable Common, and missing values for several observations. The following shows the MASTER2 and the NONPLANT input data sets:

MASTER2				NONPLANT			
OBS	Common	Animal2	Plant2	OBS	Common	Plant2	Mineral
1	a	Ant		1	a	Apricot	Amethyst
2	c	Cat	Coconut	2	b	Barley	Beryl

3	d	Dog	Dewberry	3	c	Cactus	
4	e	Eagle	Eggplant	4	e		
5	f	Frog	Fig	5	f	Fennel	
				6	g	Grape	Garnet

The following program updates the data set MASTER2 and prints the results:

```
data update2_file;
  update master2 nonplant;
  by Common;
run;

proc print data=update2_file;
  title 'Data Set Update2_File';
run;
```

Output 22.16 Results of Updating with New Variables, Nonmatched Observations, and Missing Values

Data Set Update2_File					1
Obs	Common	Animal2	Plant2	Mineral	
1	a	Ant	Apricot	Amethyst	
2	b		Barley	Beryl	
3	c	Cat	Cactus		
4	d	Dog	Dewberry		
5	e	Eagle	Eggplant		
6	f	Frog	Fennel		
7	g		Grape	Garnet	

As shown, all observations now include values for the variable Mineral. The value of Mineral is set to missing for some observations. Observations 2 and 6 in the transaction data set did not have corresponding observations in MASTER2, and they have become new observations. Observation 3 from the master data set was written to the new data set without change, and the value for Plant2 in observation 4 was not changed to missing. Three observations in the new data set have updated values for the variable Plant2.

The following program uses the UPDATEMODE statement option in the UPDATE statement, and prints the results:

```
data update2_file;
  update master2 nonplant updatemode=nomissingcheck;
  by Common;
run;

proc print data=update2_file;
  title 'Data Set Update2_File - UPDATEMODE Option';
run;
```

Output 22.17 Results of Updating with the UPDATEMODE Option

Data Set Update2_File - UPDATEMODE Option					1
Obs	Common	Animal2	Plant2	Mineral	
1	a	Ant	Apricot	Amethyst	
2	b		Barley	Beryl	
3	c	Cat	Cactus		
4	d	Dog	Dewberry		
5	e	Eagle			
6	f	Frog	Fennel		
7	g		Grape	Garnet	

The value of Plant2 in observation 5 is set to missing because the UPDATEMODE=NOMISSINGCHECK option is in effect.

For detailed examples for updating data sets, see *Combining and Modifying SAS Data Sets: Examples*.

Example 4: Updating a MASTER Data Set by Adding an Observation

If the transaction data set contains an observation that does not match an observation in the master data set, you must alter the program. The Year value in observation 5 of TRANSACTION has no match in MASTER. The following shows the MASTER and the TRANSACTION input data sets:

MASTER				TRANSACTION			
OBS	Year	VarX	VarY	OBS	Year	VarX	VarY
1	1985	x1	y1	1	1991	x2	
2	1986	x1	y1	2	1992	x2	y2
3	1987	x1	y1	3	1993	x2	
4	1988	x1	y1	4	1993		y2
5	1989	x1	y1	5	1995	x2	y2
6	1990	x1	y1				
7	1991	x1	y1				
8	1992	x1	y1				
9	1993	x1	y1				
10	1994	x1	y1				

You must use an explicit OUTPUT statement to write a new observation to a master data set. (The default action for a DATA step using a MODIFY statement is REPLACE, not OUTPUT.) Once you specify an explicit OUTPUT statement, you must also specify a REPLACE statement. The following DATA step updates data set MASTER, based on values in TRANSACTION, and adds a new observation. This program also uses the _IORC_ automatic variable for error checking. (For more information about error checking, see “Error Checking When Using Indexes to Randomly Access or Update Data” on page 382.)

Output 22.18 Modified MASTER Data Set

Updated Master Data Set -- MODIFY				1
One Observation Added				
Obs	Year	VarX	VarY	
1	1985	x1	y1	
2	1986	x1	y1	
3	1987	x1	y1	
4	1988	x1	y1	
5	1989	x1	y1	
6	1990	x1	y1	
7	1991	x2	y1	
8	1992	x2	y2	
9	1993	x2	y2	
10	1994	x1	y1	
11	1995	x2	y2	

SAS added a new observation, observation 11, to the MASTER data set and updated observations 7, 8, and 9.

Error Checking When Using Indexes to Randomly Access or Update Data

The Importance of Error Checking

When reading observations with the SET statement and KEY= option or with the MODIFY statement, error checking is imperative for several reasons. The most important reason is that these tools use nonsequential access methods, and so there is no guarantee that an observation will be located that satisfies the request. Error checking enables you to direct execution to specific code paths, depending on the outcome of the I/O operation. Your program will continue execution for expected conditions and terminate execution when unexpected results occur.

Error-Checking Tools

Two tools have been created to make error checking easier when you use the MODIFY statement or the SET statement with the KEY= option to process SAS data sets:

- _IORC_ automatic variable
- SYSRC autocall macro.

IORC is created automatically when you use the MODIFY statement or the SET statement with KEY=. The value of _IORC_ is a numeric return code that indicates the status of the I/O operation from the most recently executed MODIFY or SET statement with KEY=. Checking the value of this variable enables you to detect abnormal I/O conditions and to direct execution down specific code paths instead of having the application terminate abnormally. For example, if the KEY= variable value does match between two observations, you might want to combine them and output an observation. If they don't match, however, you might want only to write a note to the log.

Because the values of the _IORC_ automatic variable are internal and subject to change, the SYSRC macro was created to enable you to test for specific I/O conditions

while protecting your code from future changes in `_IORC_` values. When you use `SYSRC`, you can check the value of `_IORC_` by specifying one of the mnemonics listed in the following table.

Table 22.4 Most Common Mnemonic Values of `_IORC_` for DATA Step Processing

Mnemonic Value	Meaning of Return Code	When Return Code Occurs
<code>_DSENMR</code>	The TRANSACTION data set observation does not exist in the MASTER data set.	MODIFY with BY is used and no match occurs.
<code>_DSEMTR</code>	Multiple TRANSACTION data set observations with the same BY variable value do not exist in the MASTER data set.	MODIFY with BY is used and consecutive observations with the same BY values do not find a match in the first data set. In this situation, the first observation that fails to find a match returns <code>_DSENMR</code> . The subsequent observations return <code>_DSEMTR</code> .
<code>_DSENMOM</code>	No matching observation was found in the MASTER data set.	SET or MODIFY with KEY= finds no match.
<code>_SENOCHN</code>	The output operation was unsuccessful.	the KEY= option in a MODIFY statement contains duplicate values.
<code>_SOK</code>	The I/O operation was successful.	a match is found.

Example 1: Routing Execution When an Unexpected Condition Occurs

Overview

This example shows how to prevent an unexpected condition from terminating the DATA step. The goal is to update a master data set with new information from a transaction data set. This application assumes that there are no duplicate values for the common variable in either data set.

Note: This program works as expected only if the master and transaction data sets contain no consecutive observations with the same value for the common variable. For an explanation of the behavior of MODIFY with KEY= when duplicates exist, see the MODIFY statement in *SAS Language Reference: Dictionary*. Δ

Input Data Sets

The TRANSACTION data set contains three observations: two updates to information in MASTER and a new observation about PartNumber value 6 that needs to be added. MASTER is indexed on PartNumber. There are no duplicate values of PartNumber in MASTER or TRANSACTION. The following shows the MASTER and the TRANSACTION input data sets:

MASTER			TRANSACTION		
OBS	PartNumber	Quantity	OBS	PartNumber	AddQuantity

1	1	10	1	4	14
2	2	20	2	6	16
3	3	30	3	2	12
4	4	40			
5	5	50			

Original Program

The objective is to update the MASTER data set with information from the TRANSACTION data set. The program reads TRANSACTION sequentially. MASTER is read directly, not sequentially, using the MODIFY statement and the KEY= option. Only observations with matching values for PartNumber, which is the KEY= variable, are read from MASTER.

```
data master; ❶
  set transaction; ❷
  modify master key=PartNumber; ❸
  Quantity = Quantity + AddQuantity; ❹
run;
```

- ❶ Open the MASTER data set for update.
- ❷ Read an observation from the TRANSACTION data set.
- ❸ Match observations from the MASTER data set based on the values of PartNumber.
- ❹ Update the information on Quantity by adding the new values from the TRANSACTION data set.

Resulting Log

This program has correctly updated one observation but it stopped when it could not find a match for PartNumber value 6. The following lines are written to the SAS log:

```
ERROR: No matching observation was found in MASTER data set.
PartNumber=6 AddQuantity=16 Quantity=70 _ERROR_=1
_IORC_=1230015 _N_=2
NOTE: The SAS System stopped processing this step because
of errors.
NOTE: The data set WORK.MASTER has been updated. There were
1 observations rewritten, 0 observations added and 0
observations deleted.
```

Resulting Data Set

The MASTER file was incorrectly updated. The updated master has five observations. One observation was updated correctly, a new one was not added, and a second update was not made. The following shows the incorrectly updated MASTER data set:

MASTER		
OBS	PartNumber	Quantity
1	1	10
2	2	20
3	3	30
4	4	54
5	5	50

Revised Program

The objective is to apply two updates and one addition to MASTER, preventing the DATA step from stopping when it does not find a match in MASTER for the PartNumber value 6 in TRANSACTION. By adding error checking, this DATA step is allowed to complete normally and produce a correctly revised version of MASTER. This program uses the `_IORC_` automatic variable and the `SYSRC` autocall macro in a `SELECT` group to check the value of the `_IORC_` variable and execute the appropriate code based on whether a match is found.

```

data master; ❶
  set transaction; ❷
  modify master key=PartNumber; ❸
select(_iorc_); ❹
  when(%sysrc(_sok)) do;
    Quantity = Quantity + AddQuantity;
    replace;
  end;
  when(%sysrc(_dsenom)) do;
    Quantity = AddQuantity;
    _error_ = 0;
    output;
  end;
  otherwise do;
    put 'ERROR: Unexpected value for _IORC_ = ' _iorc_;
    put 'Program terminating. DATA step iteration # ' _n_;
    put _all_;
    stop;
  end;
end;
run;

```

- ❶ Open the MASTER data set for update.
- ❷ Read an observation from the TRANSACTION data set.
- ❸ Match observations from the MASTER data set based on the value of PartNumber.
- ❹ Take the correct course of action based on whether a matching value for PartNumber is found in MASTER. Update Quantity by adding the new values from TRANSACTION. The `SELECT` group directs execution to the correct code. When a match occurs (`_SOK`), update Quantity and replace the original observation in MASTER. When there is no match (`_DSENUM`), set Quantity equal to the AddQuantity amount from TRANSACTION, and append a new observation. `_ERROR_` is reset to 0 to prevent an error condition that would write the contents of the program data vector to the SAS log. When an unexpected condition occurs, write messages and the contents of the program data vector to the log, and stop the DATA step.

Resulting Log

The DATA step executed without error and observations were appropriately updated and added. The following lines are written to the SAS log:

```

NOTE: The data set WORK.MASTER has been updated. There were
      2 observations rewritten, 1 observations added and 0
      observations deleted.

```

Correctly Updated MASTER Data Set

MASTER contains updated quantities for PartNumber values 2 and 4 and a new observation for PartNumber value 6. The following shows the correctly updated MASTER data set:

MASTER		
OBS	PartNumber	Quantity
1	1	10
2	2	32
3	3	30
4	4	54
5	5	50
6	6	16

Example 2: Using Error Checking on All Statements That Use KEY=

Overview

This example shows how important it is to use error checking on all statements that use the KEY= option when reading data.

Input Data Sets

The MASTER and DESCRIPTION data sets are both indexed on PartNumber. The ORDER data set contains values for all parts in a single order. Only ORDER contains the PartNumber value 8. The following shows the MASTER, ORDER, and DESCRIPTION input data sets:

MASTER			ORDER	
OBS	PartNumber	Quantity	OBS	PartNumber
1	1	10	1	2
2	2	20	2	4
3	3	30	3	1
4	4	40	4	3
5	5	50	5	8
			6	5
			7	6

DESCRIPTION		
OBS	PartNumber	PartDescription
1	4	Nuts
2	3	Bolts
3	2	Screws
4	6	Washers

Original Program with Logic Error

The objective is to create a data set that contains the description and number in stock for each part in a single order, except for the parts that are not found in either of

the two input data sets, MASTER and DESCRIPTION. A transaction data set contains the part numbers of all parts in a single order. One data set is read to retrieve the description of the part and another is read to retrieve the quantity that is in stock.

The program reads the ORDER data set sequentially and then uses SET with the KEY= option to read the MASTER and DESCRIPTION data sets directly, based on the key value of PartNumber. When a match occurs, an observation is written that contains all the necessary information for each value of PartNumber in ORDER. This first attempt at a solution uses error checking for only one of the two SET statements that use KEY= to read a data set.

```

data combine; ❶
  length PartDescription $ 15;
  set order; ❷
  set description key=PartNumber; ❷
  set master key=PartNumber; ❷
  select(_iorc_); ❸
    when(%sysrc(_sok)) do;
      output;
    end;
    when(%sysrc(_dsenom)) do;
      PartDescription = 'No description';
      _error_ = 0;
      output;
    end;
    otherwise do;
      put 'ERROR: Unexpected value for _IORC_ = ' _iorc_;
      put 'Program terminating.';
      put _all_;
      stop;
    end;
  end;
run;

```

- ❶ Create the COMBINE data set.
- ❷ Read an observation from the ORDER data set. Read an observation from the DESCRIPTION and the MASTER data sets based on a matching value for PartNumber, the key variable. Note that no error checking occurs after an observation is read from DESCRIPTION.
- ❸ Take the correct course of action, based on whether a matching value for PartNumber is found in MASTER or DESCRIPTION. (This logic is based on the erroneous assumption that this SELECT group performs error checking for both of the preceding SET statements that contain the KEY= option. It actually performs error checking for only the most recent one.) The SELECT group directs execution to the correct code. When a match occurs (_SOK), the value of PartNumber in the observation that is being read from MASTER matches the current PartNumber value from ORDER. So, output an observation. When there is no match (_DSENUM), no observations in MASTER contain the current value of PartNumber, so set the value of PartDescription appropriately and output an observation. _ERROR_ is reset to 0 to prevent an error condition that would write the contents of the program data vector to the SAS log. When an unexpected condition occurs, write messages and the contents of the program data vector to the log, and stop the DATA step.

Resulting Log

This program creates an output data set but executes with one error. The following lines are written to the SAS log:

```
PartNumber=1 PartDescription=Nuts Quantity=10 _ERROR_=1
_IORC_=0 _N_=3
PartNumber=5 PartDescription=No description Quantity=50
_ERROR_=1 _IORC_=0 _N_=6
NOTE: The data set WORK.COMBINE has 7 observations and 3 variables.
```

Resulting Data Set

The following shows the incorrectly created COMBINE data set. Observation 5 should not be in this data set. PartNumber value 8 does not exist in either MASTER or DESCRIPTION, so no Quantity should be listed for it. Also, observations 3 and 7 contain descriptions from observations 2 and 6, respectively.

```
COMBINE
```

OBS	PartNumber	PartDescription	Quantity
1	2	Screws	20
2	4	Nuts	40
3	1	Nuts	10
4	3	Bolts	30
5	8	No description	30
6	5	No description	50
7	6	No description	50

Revised Program

To create an accurate output data set, this example performs error checking on both SET statements that use the KEY= option:

```
data combine(drop=Foundes); ❶
  length PartDescription $ 15;
  set order; ❷
  Foundes = 0; ❸
  set description key=PartNumber; ❹
  select(_iorc_); ❺
    when(%sysrc(_sok)) do;
      Foundes = 1;
    end;
    when(%sysrc(_dsenom)) do;
      PartDescription = 'No description';
      _error_ = 0;
    end;
    otherwise do;
      put 'ERROR: Unexpected value for _IORC_=' _iorc_;
      put 'Program terminating. Data set accessed is DESCRIPTION';
      put _all_;
      _error_ = 0;
      stop;
    end;
  end;
set master key=PartNumber; ❻
select(_iorc_); ❼
```

```

        when(%sysrc(_sok)) do;
            output;
        end;
        when(%sysrc(_dsenom)) do;
            if not Foundes then do;
                _error_ = 0;
                put 'WARNING: PartNumber ' PartNumber ' is not in'
                    ' DESCRIPTION or MASTER.';
            end;
            else do;
                Quantity = 0;
                _error_ = 0;
                output;
            end;
        end;
        otherwise do;
            put 'ERROR: Unexpected value for _IORC_ = ' _iorc_;
            put 'Program terminating. Data set accessed is MASTER';
            put _all_;
            _error_ = 0;
            stop;
        end;
    end;      /* ends the SELECT group */
run;

```

- ❶ Create the COMBINE data set.
- ❷ Read an observation from the ORDER data set.
- ❸ Create the variable Foundes so that its value can be used later to indicate when a PartNumber value has a match in the DESCRIPTION data set.
- ❹ Read an observation from the DESCRIPTION data set, using PartNumber as the key variable.
- ❺ Take the correct course of action based on whether a matching value for PartNumber is found in DESCRIPTION. The SELECT group directs execution to the correct code based on the value of _IORC_. When a match occurs (_SOK), the value of PartNumber in the observation that is being read from DESCRIPTION matches the current value from ORDER. Foundes is set to 1 to indicate that DESCRIPTION contributed to the current observation. When there is no match (_DSENO), no observations in DESCRIPTION contain the current value of PartNumber, so the description is set appropriately. _ERROR_ is reset to 0 to prevent an error condition that would write the contents of the program data vector to the SAS log. Any other _IORC_ value indicates that an unexpected condition has been met, so messages are written to the log and the DATA step is stopped.
- ❻ Read an observation from the MASTER data set, using PartNumber as a key variable.
- ❼ Take the correct course of action based on whether a matching value for PartNumber is found in MASTER. When a match is found (_SOK) between the current PartNumber value from ORDER and from MASTER, write an observation. When a match isn't found (_DSENO) in MASTER, test the value of Foundes. If Foundes is not true, then a value was not found in DESCRIPTION either, so write a message to the log but do not write an observation. If Foundes is true, however, the value is in DESCRIPTION but not MASTER. So write an observation but set Quantity to 0. Again, if an unexpected condition occurs, write a message and stop the DATA step.

Resulting Log

The DATA step executed without error. Six observations were correctly created and the following message was written to the log:


```
WARNING: PartNumber 8 is not in DESCRIPTION or MASTER.  
NOTE: The data set WORK.COMBINE has 6 observations  
and 3 variables.
```

Correctly Created COMBINE Data Set

The following shows the correctly updated COMBINE data set. Note that COMBINE does not contain an observation with the PartNumber value 8. This value does not occur in either MASTER or DESCRIPTION.

```
COMBINE
```

OBS	PartNumber	PartDescription	Quantity
1	2	Screws	20
2	4	Nuts	40
3	1	No description	10
4	3	Bolts	30
5	5	No description	50
6	6	Washers	0



CHAPTER 23

Using DATA Step Component Objects

<i>Introduction to DATA Step Component Objects</i>	392
<i>Using the Hash Object</i>	392
<i>Why Use the Hash Object?</i>	392
<i>Declaring and Instantiating a Hash Object</i>	393
<i>Initializing Hash Object Data Using a Constructor</i>	393
<i>Defining Keys and Data</i>	395
<i>Non-Unique Key and Data Pairs</i>	397
<i>Storing and Retrieving Data</i>	397
<i>Example 1: Using the ADD and FIND Methods to Store and Retrieve Data</i>	398
<i>Example 2: Loading a Data Set and Using the FIND Method to Retrieve Data</i>	399
<i>Maintaining Key Summaries</i>	399
<i>Replacing and Removing Data in the Hash Object</i>	402
<i>Saving Hash Object Data in a Data Set</i>	404
<i>Comparing Hash Objects</i>	405
<i>Using Hash Object Attributes</i>	405
<i>Using the Hash Iterator Object</i>	406
<i>About the Hash Iterator Object</i>	406
<i>Declaring and Instantiating a Hash Iterator Object</i>	406
<i>Example: Retrieving Hash Object Data by Using the Hash Iterator</i>	407
<i>Using the Java Object</i>	409
<i>About the Java Object</i>	409
<i>CLASSPATH and Java Options</i>	409
<i>Restrictions and Requirements for Using the Java Object</i>	410
<i>Declaring and Instantiating a Java Object</i>	411
<i>Accessing Object Fields</i>	411
<i>Accessing Object Methods</i>	412
<i>Type Issues</i>	412
<i>Java Objects and Arrays</i>	414
<i>Passing Java Object Arguments</i>	415
<i>Exceptions</i>	417
<i>Java Standard Output</i>	417
<i>Java Object Examples</i>	418
<i>Example 1: Calling a Simple Java Method</i>	418
<i>Example 2: Creating a User Interface</i>	418
<i>Example 3: Creating a Custom Class Loader</i>	422

Introduction to DATA Step Component Objects

SAS provides these five predefined component objects for use in a DATA step:

hash and hash iterator objects	enable you to quickly and efficiently store, search, and retrieve data based on lookup keys. The hash object keys and data are DATA step variables. Key and data values can be directly assigned constant values or values from a SAS data set. For information about the hash and hash iterator object language elements, see “Hash and Hash Iterator Object Language Elements” in <i>SAS Language Reference: Dictionary</i> .
Java object	provides a mechanism that is similar to the Java Native Interface (JNI) for instantiating Java classes and accessing fields and methods on the resultant objects. For more information, see “Java Object Language Elements” in <i>SAS Language Reference: Dictionary</i> .
logger and appender objects	enable you to record logging events and write these events to the appropriate destination. For more information, see “Logger and Appender Object Language Elements” in <i>SAS Logging: Configuration and Programming Reference</i> .

The DATA step Component Interface enables you to create and manipulate these component objects using statements, attributes, operators, and methods. You use the DATA step object dot notation to access the component object’s attributes and methods. For detailed information about dot notation and the DATA step objects’ statements, attributes, methods, and operators, see the Dictionary of Component Language Elements in *SAS Language Reference: Dictionary*.

Note: The DATA step component object statement, attributes, methods, and operators are limited to those defined for these objects. You cannot use the SAS Component Language functionality with these predefined DATA step objects. \triangle

Using the Hash Object

Why Use the Hash Object?

The hash object provides an efficient, convenient mechanism for quick data storage and retrieval. The hash object stores and retrieves data based on lookup keys.

To use the DATA step Component Object Interface, follow these steps:

- 1 Declare the hash object.
- 2 Create an instance of (*instantiate*) the hash object.
- 3 Initialize lookup keys and data.

After you declare and instantiate a hash object, you can perform many tasks, including these:

- Store and retrieve data
- Maintain key summaries
- Replace and remove data
- Compare hash objects

- Output a data set that contains the data in the hash object

For example, suppose that you have a large data set that contains numeric lab results that correspond to a unique patient number and weight and a small data set that contains patient numbers (a subset of those in the large data set). You can load the large data set into a hash object using the unique patient number as the key and the weight values as the data. You can then iterate over the small data set using the patient number to look up the current patient in the hash object whose weight is over a certain value and output that data to a different data set.

Depending on the number of lookup keys and the size of the data set, the hash object lookup can be significantly faster than a standard format lookup.

Declaring and Instantiating a Hash Object

You declare a hash object using the DECLARE statement. After you declare the new hash object, use the `_NEW_` operator to instantiate the object.

```
declare hash myhash;
myhash = _new_ hash();
```

The DECLARE statement tells the compiler that the object reference MYHASH is of type hash. At this point, you have declared only the object reference MYHASH. It has the potential to hold a component object of type hash. You should declare the hash object only once. The `_NEW_` operator creates an instance of the hash object and assigns it to the object reference MYHASH.

As an alternative to the two-step process of using the DECLARE statement and the `_NEW_` operator to declare and instantiate a component object, you can use the DECLARE statement to declare and instantiate the component object in one step.

```
declare hash myhash();
```

The above statement is equivalent to the following code:

```
declare hash myhash;
myhash = _new_ hash();
```

For more information about the DECLARE statement and the `_NEW_` operator, see *SAS Language Reference: Dictionary*.

Initializing Hash Object Data Using a Constructor

When you create a hash object, you might want to provide initialization data. A *constructor* is a method that you can use to instantiate a hash object and initialize the hash object data.

The hash object constructor can have either of the following formats:

- `declare hash object_name(argument_tag-1 : value-1 <, ...argument_tag-n: value-n>);`
- `object_name = _new_ hash(argument_tag-1: value-1 <, ...argument_tag-n: value-n>);`

These are the valid hash object argument tags:

`dataset: 'dataset_name'`

is the name of a SAS data set to load into the hash object.

The name of the SAS data set can be a literal or a character variable. The data set name must be enclosed in single or double quotation marks. Macro variables must be in double quotation marks.

When declaring a hash object, you can use SAS data set options in the *dataset* argument tag. For a list of SAS data set options, see “Data Set Options by Category” in *SAS Language Reference: Dictionary*. For more information on using data set options when declaring a hash object, see the DECLARE Statement in *SAS Language Reference: Dictionary*.

Note: If the data set contains duplicate keys, by default, the first instance is stored in the hash object; subsequent instances are ignored. To store the last instance in the hash object, use the DUPLICATE argument tag. The DUPLICATE argument tag also writes an error to the SAS log if there is a duplicate key. Δ

duplicate: *'option'*

determines whether to ignore duplicate keys when loading a data set into the hash object. The default is to store the first key and ignore all subsequent duplicates. Options can be one of the following values:

'replace' | 'r'

stores the last duplicate key record.

'error' | 'e'

reports an error to the log if a duplicate key is found.

The following example using the REPLACE option stores **brown** for the key 620 and **blue** for the key 531. If you use the default, **green** would be stored for 620 and **yellow** would be stored for 531.

```
data table;
  input key data $;
  datalines;
  531 yellow
  620 green
  531 blue
  908 orange
  620 brown
  143 purple
run;

data _null_;
length key 8 data $ 8;
if _n_ = 1 then do;
  declare hash myhash(dataset: "table", duplicate: "r");
  rc = myhash.definekey('key');
  rc = myhash.definedata('data');
  myhash.definedone();
  call missing(key,data);
end;

rc = myhash.output(dataset:"otable");
run;

proc print data=otable;
run;
```

hashexp: *n*

is the hash object's internal table size, where the size of the hash table is 2^n .

The value of HASHEXP is used as a power-of-two exponent to create the hash table size. For example, a value of 4 for HASHEXP equates to a hash table size of 2^4 , or 16. The maximum value for HASHEXP is 20, which equates to a hash table size of 2^{20} or 65536.

The hash table size is not equal to the number of items that can be stored. Think of the hash table as an array of containers. A hash table size of 16 would have 16 containers. Each container can hold an infinite number of items. The efficiency of the hash tables lies in the ability of the hash function to map items to and retrieve items from the containers.

In order to maximize the efficiency of the hash object lookup routines, you should set the hash table size according to the amount of data in the hash object. Try different HASHEXP values until you get the best result. For example, if the hash object contains one million items, a hash table size of 16 (HASHEXP = 4) would not be very efficient. A hash table size of 512 or 1024 (HASHEXP = 9 or 10) would result in better performance.

Default: 8, which equates to a hash table size of 2^8 or 256.

ordered: *'option'*

specifies whether or how the data is returned in key-value order if you use the hash object with a hash iterator object or if you use the hash object OUTPUT method.

option can be one of the following values:

'ascending' | 'a' Data is returned in ascending key-value order. Specifying '**ascending**' is the same as specifying '**yes**'.

'descending' | 'd' Data is returned in descending key-value order.

'YES' | 'Y' Data is returned in ascending key-value order. Specifying '**yes**' is the same as specifying '**ascending**'.

'NO' | 'N' Data is returned in an undefined order.

Default: NO

You can enclose the argument value in double quotation marks, instead of single.

multidata: *'option'*

specifies whether multiple data items are allowed for each key.

option can be one of the following values:

'YES' | 'Y' Multiple data items are allowed for each key.

'NO' | 'N' Only one data item is allowed for each key.

Default: NO

See: "Non-Unique Key and Data Pairs" on page 397.

You can enclose the argument value in double quotation marks, instead of single.

suminc: *'variable-name'*

maintains a summary count of hash object keys. The SUMINC argument value is given a DATA step variable, which holds the sum increment, that is, how much to add to the key summary for each reference to the key. The SUMINC value can be greater than, less than, or equal to 0. For example, a key summary changes using the current value of the DATA step variable.

```
dcl hash myhash(suminc: 'count');
```

For more information, see "Maintaining Key Summaries" on page 399.

For more information on the DECLARE statement and the `_NEW_` operator, see the *SAS Language Reference: Dictionary*.

Defining Keys and Data

The hash object uses lookup keys to store and retrieve data. The keys and the data are DATA step variables that you use to initialize the hash object by using dot notation

method calls. A key is defined by passing the key variable name to the DEFINEKEY method. Data is defined by passing the data variable name to the DEFINEDATA method. After you have defined all key and data variables, the DEFINEDONE method is called. Keys and data can consist of any number of character or numeric DATA step variables.

For example, the following code initializes a character key and a character data variable.

```
length d $20;
length k $20;

if _N_ = 1 then do;
  declare hash h;
  rc = h.defineKey('k');
  rc = h.defineData('d');
  rc = h.defineDone();
end;
```

You can have multiple key and data variables, but the entire key must be unique. You can store more than one data item with a particular key. For example, you could modify the previous example to store auxiliary numeric values with the character key and data. In this example, each key and each data item consists of a character value and a numeric value.

```
length d1 8;
length d2 $20;
length k1 $20;
length k2 8;

if _N_ = 1 then do;
  declare hash h;
  rc = h.defineKey('k1', 'k2');
  rc = h.defineData('d1', 'd2');
  rc = h.defineDone();
end;
```

For more information about the DEFINEDATA, DEFINEDONE, and the DEFINEKEY methods, see the *SAS Language Reference: Dictionary*.

Note: The hash object does not assign values to key variables (for example, **h.find(key: 'abc')**), and the SAS compiler cannot detect the data variable assignments that are performed by the hash object and the hash iterator. Therefore, if no assignment to a key or data variable appears in the program, SAS will issue a note stating that the variable is uninitialized. To avoid receiving these notes, you can perform one of the following actions:

- Set the NONOTES system option.
- Provide an initial assignment statement (typically to a missing value) for each key and data variable.
- Use the CALL MISSING routine with all the key and data variables as parameters. Here is an example.

```
length d $20;
length k $20;

if _N_ = 1 then do;
  declare hash h;
  rc = h.defineKey('k');
```

```

rc = h.defineData('d');
rc = h.defineDone();
call missing(k, d);
end;

```

Δ

Non-Unique Key and Data Pairs

By default, all of the keys in a hash object are unique. This means one set of data variables exists for each key. In some situations, you might want to have duplicate keys in the hash object, that is, associate more than one set of data variables with a key.

For example, assume that the key is a patient ID and the data is a visit date. If the patient were to visit multiple times, multiple visit dates would be associated with the patient ID. When you create a hash object with the MULTIDATA:“YES” argument tag, multiple sets of the data variables are associated with the key.

If the data set contains duplicate keys, by default, the first instance is stored in the hash object and subsequent instances are ignored. To store the last instance in the hash object, use the DUPLICATE argument tag. The DUPLICATE argument tag also writes an error to the SAS log if there is a duplicate key.

However, the hash object allows storage of multiple values for each key if you use the MULTIDATA argument tag in the DECLARE statement or _NEW_ operator. The hash object keeps the multiple values in a list that is associated with the key. This list can be traversed and manipulated by using several methods such as HAS_NEXT or FIND_NEXT.

To traverse a multiple data item list, you must know the *current list item*. Start by calling the FIND method for a given key. The FIND method sets the current list item. Then to determine whether the key has multiple data values, call the HAS_NEXT method. After you have determined that the key has another data value, you can retrieve that value with the FIND_NEXT method. The FIND_NEXT method sets the current list item to the next item in the list and sets the corresponding data variable or variables for that item.

In addition to moving forward through the list for a given key, you can loop backwards through the list by using the HAS_PREV and FIND_PREV methods in a similar manner.

Note: For SAS 9.2 Phase 2 and later, the items in a multiple data item list are maintained in the order in which you insert them. Δ

For more information about these and other methods associated with non-unique key and data pairs, see “Hash and Hash Iterator Object Language Elements” in *SAS Language Reference: Dictionary*.

Storing and Retrieving Data

After you initialize the hash object’s key and data variables, you can store data in the hash object using the ADD method, or you can use the *dataset* argument tag to load a data set into the hash object. If you use the *dataset* argument tag, and if the data set contains more than one observation with the same value of the key, by default, SAS keeps the first observation in the hash table and ignores subsequent observations. To store the last instance in the hash object or to send an error to the log if there is a duplicate key, use the DUPLICATE argument tag. To allow duplicate values for each key, use the MULTIDATA argument tag.

You can then use the FIND method to search and retrieve data from the hash object if one data value exists for each key. Use the FIND_NEXT and FIND_PREV methods to search and retrieve data if multiple data items exist for each key.

For more information about the ADD, FIND, FIND_NEXT, and FIND_PREV methods, see the *SAS Language Reference: Dictionary*.

You can consolidate a FIND method and ADD method using the REF method. In the following example, you can reduce the amount of code from this:

```
rc = h.find();
  if (rc != 0) then
    rc = h.add();
```

to a single method call:

```
rc = h.ref();
```

For more information about the REF method, see the *SAS Language Reference: Dictionary*.

Note: You can also use the hash iterator object to retrieve the hash object data, one data item at a time, in forward and reverse order. For more information, see “Using the Hash Iterator Object” on page 406. Δ

Example 1: Using the ADD and FIND Methods to Store and Retrieve Data

The following example uses the ADD method to store the data in the hash object and associate the data with the key. The FIND method is then used to retrieve the data that is associated with the key value **Homer**.

```
data _null_;
  length d $20;
  length k $20;

  /* Declare the hash object and key and data variables */
  if _N_ = 1 then do;
    declare hash h( );
    rc = h.defineKey('k');
    rc = h.defineData('d');
    rc = h.defineDone();
  end;

  /* Define constant value for key and data */
  k = 'Homer';
  d = 'Odyssey';
  /* Use the ADD method to add the key and data to the hash object */
  rc = h.add();
  if (rc ne 0) then
    put 'Add failed.';

  /* Define constant value for key and data */
  k = 'Joyce';
  d = 'Ulysses';
  /* Use the ADD method to add the key and data to the hash object */
  rc = h.add();
  if (rc ne 0) then
    put 'Add failed.';

  k = 'Homer';
```

```

/* Use the FIND method to retrieve the data associated with 'Homer' key */
rc = h.find();
if (rc = 0) then
    put d=;
else
    put 'Key Homer not found.';
run;

```

The FIND method assigns the data value **Odyssey**, which is associated with the key value **Homer**, to the variable D.

Example 2: Loading a Data Set and Using the FIND Method to Retrieve Data

Assume the data set **SMALL** contains two numeric variables K (key) and S (data) and another data set, **LARGE**, contains a corresponding key variable K. The following code loads the **SMALL** data set into the hash object, and then searches the hash object for key matches on the variable K from the **LARGE** data set.

```

data match;
    length k 8;
    length s 8;
    if _N_ = 1 then do;
        /* load SMALL data set into the hash object */
        declare hash h(dataset: "work.small");
        /* define SMALL data set variable K as key and S as value */
        h.defineKey('k');
        h.defineData('s');
        h.defineDone();
        /* avoid uninitialized variable notes */
        call missing(k, s);
    end;

    /* use the SET statement to iterate over the LARGE data set using */
    /* keys in the LARGE data set to match keys in the hash object */
    set large;
    rc = h.find();
    if (rc = 0) then output;
run;

```

The *dataset* argument tag specifies the **SMALL** data set whose keys and data will be read and loaded by the hash object during the **DEFINEDONE** method. The **FIND** method is then used to retrieve the data.

Maintaining Key Summaries

You can maintain a summary count for a hash object key by using the **SUMINC** argument tag. This argument instructs the hash object to allocate internal storage in each record to store a summary value in the record each time that the record is used by a **FIND**, **CHECK**, or **REF** method. The **SUMINC** value is also used to maintain a summary count of hash object keys after a **FIND**, **CHECK**, or **REF** method. The **SUMINC** argument tag is given a **DATA** step variable, which holds the sum increment, that is, how much to add to the key summary for each reference to the key. The **SUMINC** value can be greater than, less than, or equal to 0.

The **SUMINC** value is also used to initialize the summary on an **ADD** method. Each time the **ADD** method occurs, the key to the **SUMINC** value is initialized.

In the following example, the initial ADD method sets the summary count for K=99 to 1 before the ADD. Then each time a new COUNT value is given, the following FIND method adds the value to the key summary. In this example, one data value exists for each key. The SUM method retrieves the current value of the key summary and the value is stored in the DATA step variable TOTAL. If multiple items exist for each key, the SUMDUP method would retrieve the current value of the key summary.

```
data _null_;
  length k count 8;
  length total 8;
  dcl hash myhash(suminc: 'count');
  myhash.defineKey('k');
  myhash.defineDone();

  k = 99;
  count = 1;
  myhash.add();

  /* COUNT is given the value 2.5 and the */
  /* FIND sets the summary to 3.5*/
  count = 2.5;
  myhash.find();

  /* The COUNT of 3 is added to the FIND and */
  /* sets the summary to 6.5. */
  count = 3;
  myhash.find();

  /* The COUNT of -1 sets the summary to 5.5. */
  count = -1;
  myhash.find();

  /* The SUM method gives the current value of */
  /* the key summary to the variable TOTAL. */
  myhash.sum(sum: total);

  /* The PUT statement prints total=5.5 in the log. */
  put total=;
  run;
```

In this example, a summary is maintained for each key value K=99 and K=100:

```
k = 99;
count = 1;
myhash.add();
/* key=99 summary is now 1 */

k = 100;
myhash.add();
/* key=100 summary is now 1 */

k = 99;
myhash.find();
/* key=99 summary is now 2 */
```

```

count = 2;
myhash.find();
/* key=99 summary is now 4 */

k = 100;
myhash.find();
/* key=100 summary is now 3 */

myhash.sum(sum: total);
put 'total for key 100 = 'total;

k = 99;
myhash.sum(sum:total);
put 'total for key 99 = ' total;

```

The first PUT statement prints the summary for K=100:

```
total for key 100 = 3
```

And the second PUT statement prints the summary for K=99:

```
total for key 99 = 4
```

You can use key summaries in conjunction with the *dataset* argument tag. As the data set is read into the hash object using the DEFINEDONE method, all key summaries are set to the SUMINC value and all subsequent FIND, CHECK, or ADD methods change the corresponding key summaries.

```
declare hash myhash(suminc: "keycount", dataset: "work.mydata");
```

You can use key summaries for counting the number of occurrences of given keys. In the following example, the data set MYDATA is loaded into a hash object and uses key summaries to keep count of the number of occurrences for each key in the data set KEYS. (The SUMINC variable is not set to a value, so the default initial value of zero is used.)

```

data mydata;
  input key;
datalines;
1
2
3
4
5
;
run;

```

```

data keys;
  input key;
datalines;
1
2
1
3
5
2

```

```

3
2
4
1
5
1
;
run;

data count;
  length total key 8;
  keep key total;

  declare hash myhash(suminc: "count", dataset:"mydata");
  myhash.defineKey('key');
  myhash.defineDone();
  count = 1;

  do while (not done);
    set keys end=done;
    rc = myhash.find();
  end;

  done = 0;
  do while (not done);
    set mydata end=done;
    rc = myhash.sum(sum: total);
    output;
  end;
  stop;
run;

```

Here is the output for the resulting data set.

Obs	total	key
1	4	1
2	3	2
3	2	3
4	1	4
5	2	5

For more information about the SUM and SUMDUP methods see the *SAS Language Reference: Dictionary*.

Replacing and Removing Data in the Hash Object

You can remove or replace data that is stored in the hash object using the following methods:

- Use the REMOVE method to remove all data items.
- Use the REPLACE method to replace all data items.
- Use the REMOVEDUP method to remove only the current data item.
- Use the REPLACEDUP method to replace only the current data item.

In the following example, the REPLACE method replaces the data **Odyssey** with **Iliad**, and the REMOVE method deletes the entire data entry associated with the **Joyce** key from the hash object.

```

data _null_;
length d $20;
length k $20;

/* Declare the hash object and key and data variables */
if _N_ = 1 then do;
  declare hash h;
  rc = h.defineKey('k');
  rc = h.defineData('d');
  rc = h.defineDone();
end;

/* Define constant value for key and data */
k = 'Joyce';
d = 'Ulysses';
/* Use the ADD method to add the key and data to the hash object */
rc = h.add();
if (rc ne 0) then
  put 'Add failed.';

/* Define constant value for key and data */
k = 'Homer';
d = 'Odyssey';
/* Use the ADD method to add the key and data to the hash object */
rc = h.add();
if (rc ne 0) then
  put 'Add failed.';

/* Use the REPLACE method to replace 'Odyssey' with 'Iliad' */
k = 'Homer';
d = 'Iliad';
rc = h.replace();
if (rc = 0) then
  put d;
else
  put 'Replace not successful.';

/* Use the REMOVE method to remove the 'Joyce' key and data */
k = 'Joyce';
rc = h.remove();
if (rc = 0) then
  put k 'removed from hash object';
else
  put 'Deletion not successful.';

run;

```

The following lines are written to the SAS log.

```

d=Iliad
Joyce removed from hash object

```

Note: If an associated hash iterator is pointing to the key, the REMOVE method will not remove the key or data from the hash object. An error message is issued to the log. Δ

For more information on the REMOVE, REMOVEDUP, REPLACE, and REPLACEDUP methods, see the *SAS Language Reference: Dictionary*.

Saving Hash Object Data in a Data Set

You can create a data set that contains the data in a specified hash object by using the OUTPUT method. In the following example, two keys and data are added to the hash object and then output to the WORK.OUT data set.

```
data test;
  length d1 8;
  length d2 $20;
  length k1 $20;
  length k2 8;

  /* Declare the hash object and two key and data variables */
  if _N_ = 1 then do;
    declare hash h;
    rc = h.defineKey('k1', 'k2');
    rc = h.defineData('d1', 'd2');
    rc = h.defineDone();
  end;

  /* Define constant value for key and data */
  k1 = 'Joyce';
  k2 = 1001;
  d1 = 3;
  d2 = 'Ulysses';
  rc = h.add();

  /* Define constant value for key and data */
  k1 = 'Homer';
  k2 = 1002;
  d1 = 5;
  d2 = 'Odyssey';
  rc = h.add();

  /* Use the OUTPUT method to save the hash object data to the OUT data set */
  rc = h.output(dataset: "work.out");
run;

proc print data=work.out;
run;
```

The following output shows the report that PROC PRINT generates.

Output 23.1 Data Set Created from the Hash Object

The SAS System			1
Obs	d1	d2	
1	5	Odyssey	
2	3	Ulysses	

Note that the hash object keys are not stored as part of the output data set. If you want to include the keys in the output data set, you must define the keys as data in the DEFINEDATA method. In the previous example, the DEFINEDATA method would be written this way:

```
rc = h.defineData('k1', 'k2', 'd1', 'd2');
```

For more information on the OUTPUT method, see *SAS Language Reference: Dictionary*.

Comparing Hash Objects

You can compare one hash object to another by using the EQUALS method. In the following example, two hash objects are being compared. Note that the EQUALS method has two argument tags. The HASH argument tag is the name of the second hash object and the RESULTS argument tag is a numeric variable name that will hold the result of the comparison (1 if equal and zero if not equal).

```
length eq k 8;

declare hash myhash1();
myhash1.defineKey('k');
myhash1.defineDone();

declare hash myhash2();
myhash2.defineKey('k');
myhash2.defineDone();

rc = myhash1.equals(hash: 'myhash2', result: eq);
```

For more information about the EQUALS method, see *SAS Language Reference: Dictionary*.

Using Hash Object Attributes

You can use the DATA Step Component Interface to retrieve information from a hash object using an attribute. Use the following syntax for an attribute:

```
attribute_value=obj.attribute_name;
```

There are two attributes available to use with hash objects. NUM_ITEMS returns the number of items in a hash object and ITEM_SIZE returns the size (in bytes) of an item. The following example retrieves the number of items in a hash object:

```
n = myhash.num_items;
```

The following example retrieves the size of an item in a hash object:

```
s = myhash.item_size;
```

You can obtain an idea of how much memory the hash object is using with the `ITEM_SIZE` and `NUM_ITEMS` attributes. The `ITEM_SIZE` attribute does not reflect the initial overhead that the hash object requires, nor does it take into account any necessary internal alignments. Therefore, the use of `ITEM_SIZE` will not provide exact memory usage, but it will give a good approximation.

For more information about the `NUM_ITEMS` and `ITEM_SIZE` attributes, see the *SAS Language Reference: Dictionary*.

Using the Hash Iterator Object

About the Hash Iterator Object

Use the hash iterator object to store and search data based on lookup keys. The hash iterator object enables you to retrieve the hash object data in either forward or reverse key order.

Declaring and Instantiating a Hash Iterator Object

You declare a hash iterator object by using the `DECLARE` statement. After you declare the new hash iterator object, use the `_NEW_` operator to instantiate the object, using the hash object name as an argument tag. For example:

```
declare hiter myiter;
myiter = _new_ hiter('h');
```

the `DECLARE` statement tells the compiler that the object reference `MYITER` is of type hash iterator. At this point, you have declared only the object reference `MYITER`. It has the potential to hold a component object of type hash iterator. You should declare the hash iterator object only once. The `_NEW_` operator creates an instance of the hash iterator object and assigns it to the object reference `MYITER`. The hash object, `H`, is passed as a constructor argument. The hash object, not the hash object variable, is specifically assigned to the hash iterator.

As an alternative to the two-step process of using the `DECLARE` statement and the `_NEW_` operator to declare and instantiate a component object, you can declare and instantiate a hash iterator object in one step by using the `DECLARE` statement as a constructor method. The syntax is as follows:

```
declare hiter object_name(hash_object_name);
```

In the above example, the hash object name must be enclosed in single or double quotation marks.

For example:

```
declare hiter myiter('h');
```

The previous statement is equivalent to these:

```
declare hiter myiter;
myiter = _new_ hiter('h');
```

Note: You must declare and instantiate a hash object before you create a hash iterator object. For more information, see “Declaring and Instantiating a Hash Object” on page 393. Δ

For example:

```
if _N_ = 1 then do;
  length key $10;
  declare hash myhash(dataset:"work.x", ordered: 'yes');
  declare hiter myiter('myhash');
  myhash.defineKey('key');
  myhash.defineDone();
end;
```

This code creates an instance of a hash iterator object with the variable name MYITER. The hash object, MYHASH, is passed as the constructor argument. Because the hash object was created with the ORDERED argument tag set to 'yes', the data will be returned in ascending key-value order.

For more information about the DECLARE statement and the _NEW_ operator, see the *SAS Language Reference: Dictionary*.

Example: Retrieving Hash Object Data by Using the Hash Iterator

Using the data set ASTRO that contains astronomical data, the following code creates the data set that contains Messier objects (OBJ) whose right-ascension (RA) values are greater than 12. The FIRST and NEXT methods are used to retrieve the data in ascending order. For more information about the FIRST and NEXT methods, see the *SAS Language Reference: Dictionary*.

```
data astro;
  input obj $1-4 ra $6-12 dec $14-19;
  datalines;
M31 00 42.7 +41 16
M71 19 53.8 +18 47
M51 13 29.9 +47 12
M98 12 13.8 +14 54
M13 16 41.7 +36 28
M39 21 32.2 +48 26
M81 09 55.6 +69 04
M100 12 22.9 +15 49
M41 06 46.0 -20 44
M44 08 40.1 +19 59
M10 16 57.1 -04 06
M57 18 53.6 +33 02
  M3 13 42.2 +28 23
M22 18 36.4 -23 54
M23 17 56.8 -19 01
M49 12 29.8 +08 00
M68 12 39.5 -26 45
M17 18 20.8 -16 11
M14 17 37.6 -03 15
M29 20 23.9 +38 32
M34 02 42.0 +42 47
M82 09 55.8 +69 41
M59 12 42.0 +11 39
M74 01 36.7 +15 47
```

```

M25 18 31.6 -19 15
;
run;

data out;
  if _N_ = 1 then do;
    length obj $10;
    length ra $10;
    length dec $10;
    /* Read ASTRO data set and store in asc order in hash obj */
    declare hash h(dataset:"work.astro", ordered: 'yes');
    /* Define variables RA and OBJ as key and data for hash object */
    declare hiter iter('h');
    h.defineKey('ra');
    h.defineData('ra', 'obj');
    h.defineDone();
    /* Avoid uninitialized variable notes */
    call missing(obj, ra, dec);
  end;
  /* Retrieve RA values in ascending order */
  rc = iter.first();
  do while (rc = 0);
    /* Find hash object keys greater than 12 and output data */
    if ra GE '12' then
      output;
    rc = iter.next();
  end;
run;

proc print data=work.out;
  var ra obj;
  title 'Messier Objects Greater than 12 Sorted by Right Ascension Values';
run;

```

The following output shows the report that PROC PRINT generates.

Output 23.2 Messier Objects Greater than 12, Sorted by Right Ascension Values

Messier Objects Greater than 12 Sorted by Right Ascension Values			
Obs	ra		obj
1	12	13.8	M98
2	12	22.9	M100
3	12	29.8	M49
4	12	39.5	M68
5	12	42.0	M59
6	13	29.9	M51
7	13	42.2	M3
8	16	41.7	M13
9	16	57.1	M10
10	17	37.6	M14
11	17	56.8	M23
12	18	20.8	M17
13	18	31.6	M25
14	18	36.4	M22
15	18	53.6	M57
16	19	53.8	M71
17	20	23.9	M29
18	21	32.2	M39

Using the Java Object

About the Java Object

The Java object provides a mechanism that is similar to the Java Native Interface (JNI) for instantiating Java classes and accessing fields and methods on the resultant objects. You can create hybrid applications that contain both Java and DATA step code.

CLASSPATH and Java Options

In previous versions of SAS, Java classes were found using the JREOPTIONS system option.

In SAS 9.2, you must set the CLASSPATH environment variable so the Java object can find your Java classes. The Java object represents an instance of a Java class that is found in the current Java classpath. Any class that you use must appear in the classpath. If the class is in a .jar file, then the .jar filename must appear in the classpath.

How you set the CLASSPATH environment variable depends on your operating environment. For most operating systems, you can set the CLASSPATH environment variable either locally (for use only in your SAS session) or globally. Table 23.1 shows methods and examples for different operating environments. For more information, see the SAS documentation for your operating environment.

Table 23.1 Setting the CLASSPATH Environment Variable in Different Operating Environments

Operating Environment	Method	Example
Windows		
Globally	Windows System Environment Variable in Control Panel	Control Panel ► System ► Advanced ► Environment Variables (Windows XP Classic view)
	SAS configuration file	set classpath c:\HelloWorld.jar
Locally	SAS command line	-set classpath c:\HelloWorld.jar
UNIX		
Globally	SAS configuration file	set classpath ~/HelloWorld.jar
Locally	EXPORT command ¹	export classpath=~/HelloWorld.jar;
z/OS		
Globally	TKMSENV data set	set TKJNI_OPT_CLASSPATH=/u/userid/java: /u/userid/java/test.jar: asis
Locally	Not available	
VMS		
Globally	Command line ² <i>detach_template.com</i> script that gets generated in sas\$root:[misc.base] at installation	\$ define java\$classpath disk:[subdir] abc.jar, disk:[subdir2]def.jar define java\$classpath disk:[subdir] abc.jar, disk:[subdir2]def.jar
Locally	Not available	

1 The syntax depends on the shell.

2 The command line should be defined before you invoke SAS so the process that the JVM actually runs in gets the definition as well.

Restrictions and Requirements for Using the Java Object

The following restrictions and requirements apply when using the Java object:

- The Java object is designed to call Java methods from SAS. The Java object is not intended to extend the SAS library of functions. Calling PROC FCMP functions is much more efficient for fast in-process extensions to the DATA step, especially when large data sets are involved. Using the Java object to perform this kind of processing with large data sets takes significantly more time.
- The only Java Runtime Environments (JREs) that are supported by SAS are those that are explicitly required during the installation of the SAS software.
- The only Java options that are supported by SAS are those that are set when SAS is installed.
- Ensure that your Java application runs correctly before using it with the Java object.

Declaring and Instantiating a Java Object

You declare a Java object by using the DECLARE statement. After you declare the new Java object, use the `_NEW_` operator to instantiate the object, using the Java object name as an argument tag.

```
declare javaobj j;
j = _new_ javaobj("somejavaclass");
```

In this example, the DECLARE statement tells the compiler that the object reference J is of type java—that is, the instance of the Java object is stored in the variable J. At this point, you have declared only the object reference J. It has the potential to hold a component object of type java. You should declare the Java object only once. The `_NEW_` operator creates an instance of the Java object and assigns it to the object reference J. The Java class name, SOMEJAVACLASS, is passed as a constructor argument, which is the first-and-only argument that is required for the Java object constructor. All other arguments are constructor arguments to the Java class itself.

As an alternative to the two-step process of using the DECLARE statement and the `_NEW_` operator to declare and instantiate a Java object, you can declare and instantiate a Java object in one step by using the DECLARE statement as a constructor method. The syntax is as follows:

```
DECLARE JAVAOBJ object-name("java-class", <argument-1, ... argument-n>);
```

For more information, see the DECLARE Statement and the `_NEW_` Operator for the Java object in *SAS Language Reference: Dictionary*.

Accessing Object Fields

Once you instantiate a Java object, you can access and modify its public and class fields in a DATA step through method calls on the Java object. Public fields are non-static and declared as public in the Java class. Class fields are static and accessed from Java classes.

Method calls to access object fields have one of these forms, depending on whether you are accessing non-static or static fields:

```
GETtypeFIELD("field-name", value);
```

```
GETSTATICtypeFIELD("field-name", value);
```

Method calls to modify object fields have one of these forms, depending on whether you access static or non-static fields:

```
SETtypeFIELD("field-name", value);
```

```
SETSTATICtypeFIELD("field-name", value);
```

Note: The *type* argument represents a Java data type. For more information about how Java data types relate to SAS data types, see “Type Issues” on page 412. The *field-name* argument specifies the type for the Java field, and *value* specifies the value that is returned or set by the method. Δ

For more information and examples, see “Java Object Language Elements” in *SAS Language Reference: Dictionary*.

Accessing Object Methods

Once you instantiate a Java object, you can access its public and class methods in a DATA step through method calls on the Java object. Public methods are non-static and declared as public in the Java class. Class methods are static and accessed from Java classes.

Method calls to access Java methods have one of these forms, depending on whether you are accessing non-static or static methods:

```
object.CALLtypeMETHOD ("method-name", <method-argument-1 ...,
method-argument-n>, <return value>);
```

```
object.CALLSTATICtypeMETHOD ("method-name", <method-argument-1 ...,
method-argument-n>, <return value>);
```

Note: The *type* argument represents a Java data type. For more information about how Java data types relate to SAS data types, see “Type Issues” on page 412. Δ

For more information and examples, see “Java Object Language Elements” in *SAS Language Reference: Dictionary*.

Type Issues

The Java type set is a superset of the SAS data types. Java has data types such as BYTE, SHORT, and CHAR in addition to the standard numeric and character values. SAS has only two data types: numeric and character.

Table 23.2 describes how Java data types are mapped to SAS data types when using the Java object method calls.

Table 23.2 How Java Data Types Map to SAS Data Types

Java Data Type	SAS Data Type
BOOLEAN	numeric
BYTE	numeric
CHAR	numeric
DOUBLE	numeric
FLOAT	numeric
INT	numeric
LONG	numeric
SHORT	numeric
STRING	character ¹

¹ Java string data types are mapped to SAS character data types as UTF-8 strings.

Other than STRING, it is not possible to return objects from Java classes to the DATA step. However, it is possible to pass objects to Java methods. For more information, see “Passing Java Object Arguments” on page 415.

Some Java methods that return objects can be handled by creating wrapper classes to convert the object values. In the following example, the Java hash table returns object values, but you can still use the hash table from the DATA step by creating

simple Java wrapper classes to handle the type conversions. Then you can access the **dhash** and **shash** classes from the DATA step.

```

/* Java code */
import java.util.*;

public class dhash
{
    private Hashtable table;

    public dhash()
    {
        table = new Hashtable ();
    }

    public void put(double key, double value)
    {
        table.put(new Double(key), new Double(value));
    }

    public double get(double key)
    {
        Double ret = table.get(new Double(key));
        return ret.doubleValue();
    }
}

import java.util.*;

public class shash
{
    private Hashtable table;

    public shash()
    {
        table = new Hashtable ();
    }

    public void put(double key, String value)
    {
        table.put(new Double(key), value);
    }

    public String get(double key)
    {
        return table.get(new Double(key));
    }
}

/* DATA step code */
data _null_;
    dcl javaobj sh('shash');
    dcl javaobj dh('dhash');

```

```

length d 8;
length s $20;

do i = 1 to 10;
    dh.callvoidmethod('vput', i, i * 2);
end;

do i = 1 to 10;
    sh.callvoidmethod('put', i, 'abc' || left(trim(i))); end;

do i = 1 to 10;
    dh.calldoublemethod('get', i, d);
    sh.callstringmethod('get', i, s);
    put d= s=;
end;
run;

```

The following lines are written to the SAS log:

```

d=2 s=abc1
d=4 s=abc2
d=6 s=abc3
d=8 s=abc4
d=10 s=abc5
d=12 s=abc6
d=14 s=abc7
d=16 s=abc8
d=18 s=abc9
d=20 s=abc10

```

Java Objects and Arrays

You can pass DATA step arrays to Java objects.

In the following example, the arrays **d** and **s** are passed to the Java object **j**.

```

/* Java code */
import java.util.*;
import java.lang.*;
class jtest
{
    public void dbl(double args[])
    {
        for(int i = 0; i < args.length; i++)
            System.out.println(args[i]);
    }

    public void str(String args[])
    {
        for(int i = 0; i < args.length; i++)
            System.out.println(args[i]);
    }
}

```

```

/* DATA Step code */
data _null_;
  dcl javaobj j("jtest");
  array s[3] $20 ("abc", "def", "ghi");
  array d[10] (1:10);
  j.callVoidMethod("dbl", d);
  j.callVoidMethod("str", s);
run;

```

The following lines are written to the SAS log:

```

1.0
2.0
3.0
4.0
5.0
6.0
7.0
8.0
9.0
10.0
abc
def
ghi

```

Only one-dimensional array parameters are supported. However, it is possible to pass multidimensional array arguments by taking advantage of the fact that the arrays are passed in row-major order. You must handle the dimensional indexing manually in the Java code—that is, you must declare a one-dimensional array parameter and index to the subarrays accordingly.

Passing Java Object Arguments

While it is not possible to return objects from Java classes to the DATA step, it is possible to pass objects, as well as strings, to Java class methods.

For example, suppose you have the following wrapper classes for `java/util/Vector` and its iterator:

```

/* Java code */
import java.util.*;

class mVector extends Vector
{
  public mVector()
  {
    super();
  }

  public mVector(double d)
  {
    super((int)d);
  }

  public void addElement(String s)

```

```

    {
        addElement((Object)s);
    }
}

import java.util.*;
public class mIterator
{
    protected mVector m_v;
    protected Iterator iter;

    public mIterator(mVector v)
    {
        m_v = v;
        iter = v.iterator();
    }

    public boolean hasNext()
    {
        return iter.hasNext();
    }

    public String next()
    {
        String ret = null;
        ret = (String)iter.next();
        return ret;
    }
}

```

These wrapper classes are useful for performing type conversions (for example, the **mVector** constructor takes a **DOUBLE** argument). Overloading the constructor is necessary because **java/util/Vector**'s constructor takes an integer value, but the **DATA** step has no integer type.

The following **DATA** step program uses these classes. The program creates and fills a vector, passes the vector to the iterator's constructor, and then lists all the values in the vector. Note that you must create the iterator after the vector is filled. The iterator keeps a copy of the vector's modification count at the time of creation, and this count must stay in synchronization with the vector's current modification count. The code would throw an exception if the iterator were created before the vector was filled.

```

/* DATA step code */
data _null_;
    length b 8;
    length val $200;
    dcl javaobj v("mVector");

    v.callVoidMethod("addElement", "abc");
    v.callVoidMethod("addElement", "def");
    v.callVoidMethod("addElement", "ghi");
    dcl javaobj iter("mIterator", v);

    iter.callBooleanMethod("hasNext", b);
    do while(b);

```

```

        iter.callStringMethod("next", val);
        put val=;
        iter.callBooleanMethod("hasNext", b);
    end;

    m.delete();
    v.delete();
    iter.delete();
run;

```

The following lines are written to the SAS log:

```

val=abc
val=def
val=ghi

```

One current limitation to passing objects is that the JNI method lookup routine does not perform a full class lookup based on a given signature. This means you could not change the **mIterator** constructor to take a **Vector** as shown in the following code:

```

/* Java code */
public mIterator(Vector v)
{
    m_v = v;
    iter = v.iterator();
}

```

Even though **mVector** is a subclass of **Vector**, the method lookup routine will not find the constructor. Currently, the only solution is to manage the types in Java by adding new methods or creating wrapper classes.

Exceptions

Java exceptions are handled through the **EXCEPTIONCHECK**, **EXCEPTIONCLEAR**, and **EXCEPTIONDESCRIBE** methods.

The **EXCEPTIONCHECK** method is used to determine whether an exception occurred during a method call. If you call a method that can throw an exception, it is strongly recommended that you check for an exception after the call. If an exception is thrown, you should take appropriate action and then clear the exception by using the **EXCEPTIONCLEAR** method.

The **EXCEPTIONDESCRIBE** method is used to turn exception debug logging on or off. If exception debug logging is on, exception information is printed to the JVM standard output. By default, JVM standard output is redirected to the SAS log. Exception debugging is off by default.

For more information, see the **EXCEPTIONCHECK** method, the **EXCEPTIONCLEAR** method, and the **EXCEPTIONDESCRIBE** method in *SAS Language Reference: Dictionary*.

Java Standard Output

Output from statements in Java that are directed to standard output such as the following are sent to the SAS log by default.

```
System.out.println("hello");
```

The Java output that is directed to the SAS log is flushed when the DATA step ends. This flushing causes the Java output to appear after any output that was generated while the DATA step was running. Use the FLUSHJAVAOUTPUT method to synchronize the output so that it appears in the order of execution.

Java Object Examples

Example 1: Calling a Simple Java Method

This Java class creates a simple method that sums three numbers.

```
/* Java code */
class MyClass
{
    double compute(double x, double y, double z)
    {
        return (x + y + z);
    }
}

/* DATA step code
data _null_;
    dcl javaobj j("MyClass");

    rc = j.callDoubleMethod("compute", 1, 2, 3, r);

    put rc= r=;
run;
```

The following lines are written to the SAS log:

```
rc=0 rc=6
```

Example 2: Creating a User Interface

In addition to providing a Java component access mechanism, you can use the Java object to create a simple Java user interface.

This Java class creates a simple user interface with several buttons. The user interface also maintains a queue of values that represent the sequence of button choices that are entered by a user.

```
/* Java code */
import java.awt.*;
import java.util.*;
import java.awt.event.*;

class colorsUI extends Frame
{
```

```

private Button red;
private Button blue;
private Button green;
private Button quit;
private Vector list;
private boolean d;
private colorsButtonListener cbl;

public colorsUI()
{
    d = false;
    list = new Vector();
    cbl = new colorsButtonListener();

    setBackground(Color.lightGray);
    setSize(320,100);
    setTitle("New Frame");
    setVisible(true);
    setLayout(new FlowLayout(FlowLayout.CENTER, 10, 15));
    addWindowListener(new colorsUIListener());

    red = new Button("Red");
    red.setBackground(Color.red);
    red.addActionListener(cbl);

    blue = new Button("Blue");
    blue.setBackground(Color.blue);
    blue.addActionListener(cbl);

    green = new Button("Green");
    green.setBackground(Color.green);
    green.addActionListener(cbl);

    quit = new Button("Quit");
    quit.setBackground(Color.yellow);
    quit.addActionListener(cbl);

    this.add(red);
    this.add(blue);
    this.add(green);
    this.add(quit);

    show();
}

public synchronized void enqueue(Object o)
{
    synchronized(list)
    {
        list.addElement(o);
        notify();
    }
}

```

```

public synchronized Object dequeue()
{
    try
    {
        while(list.isEmpty())
            wait();

        if (d)
            return null;

        synchronized(list)
        {
            Object ret = list.elementAt(0);
            list.removeElementAt(0);
            return ret;
        }
    }
    catch(Exception e)
    {
        return null;
    }
}

public String getNext()
{
    return (String)dequeue();
}

public boolean done()
{
    return d;
}

class colorsButtonListener implements ActionListener
{
    public void actionPerformed(ActionEvent e)
    {
        Button b;
        String l;
        b = (Button)e.getSource();
        l = b.getLabel();
        if ( l.equals("Quit") )
        {
            d = true;
            hide();
            l = "";
        }
        enqueue(l);
    }
}

class colorsUIListener extends WindowAdapter
{
    public void windowClosing(WindowEvent e)

```

```

        {
            Window w;
            w = e.getWindow();
            d = true;
            enqueue("");
            w.hide();
        }
    }

    public static void main(String s[])
    {
        colorsUI cui;
        cui = new colorsUI();
    }
}

/* DATA step code */
data colors;
    length s $10;
    length done 8;
    drop done;

    if (_n_ = 1) then do;
        /* Declare and instantiate colors object (from colorsUI.class) */
        dcl javaobj j("colorsUI");
    end;

    /*
     * colorsUI.class will display a simple UI and maintain a
     * queue to hold color choices.
     */


    /* Loop until user hits quit button */
    do while (1);
        j.callBooleanMethod("done", done);
        if (done) then
            leave;
        else do;
            /* Get next color back from queue */
            j.callStringMethod("getNext", s);
            if s ne "" then
                output;
            end;
        end;
    end;
run;
proc print data=colors;
run;
quit;

```

In the DATA step code, the **colorsUI** class is instantiated and the user interface is displayed. You enter a loop that is terminated when you click **Quit**. This action is communicated to the DATA step through the Done variable. While looping, the DATA

step retrieves the values from the Java class's queue and writes the values successively to the output data set.

Display 23.1 User Interface Created by the Java Object



Example 3: Creating a Custom Class Loader

You might not want to put all your Java classes in the classpath. You can write your own class loader to find the classes and load them. The following example illustrates how you can create a custom class loader.

In this example, you create a class, **x**, which resides in a folder or directory, **y**. You call the methods in this class by using the Java object with the classpath that includes the **y** folder.

```

/* Java code */
package com.sas;

public class x
{
    public void m()
    {
        System.out.println("method m in y folder");
    }

    public void m2()
    {
        System.out.println("method m2 in y folder");
    }
}

/* DATA step code */
data _null_;
    dcl javaobj j('com/sas/x');

```

```

        j.callvoidmethod('m');
        j.callvoidmethod('m2');
run;

```

The following lines are written to the SAS log.

```

method m in y folder
method m2 in y folder

```

Suppose you have another class, **x**, that is stored in a different folder, **z**.

```

/* Java code */
package com.sas;

public class z
{
    public void m()
    {
        System.out.println("method m in y folder");
    }

    public void m2()
    {
        System.out.println("method m2 in y folder");
    }
}

```

You can call methods in this class instead of the class in folder **y** by changing the classpath, but this requires restarting SAS. The following method allows for more dynamic control of how classes are loaded.

To create a custom class loader, first you create an interface that contains all the methods you will call through the Java object—in this program, **m** and **m2**.

```

/* Java code */
public interface apiInterface
{
    public void m();
    public void m2();
}

```

Then you create a class for the actual implementation.

```

/* Java code */
import com.sas.x;

public class apiImpl implements apiInterface
{
    private x x;

    public apiImpl()
    {
        x = new x();
    }
}

```

```

    public void m()
    {
        x.m();
    }

    public void m2()
    {
        x.m2();
    }
}

```

These methods are called by delegating to the Java object instance class. Note that the code to create the **apiClassLoader** custom class loader is provided later in this section.

```

/* Java code */
public class api
{
    /* Load classes from the z folder */
    static ClassLoader customLoader = new apiClassLoader("C:\\z");
    static String API_IMPL = "apiImpl";
    apiInterface cp = null;

    public api()
    {
        cp = load();
    }

    public void m()
    {
        cp.m();
    }

    public void m2()
    {
        cp.m2();
    }

    private static apiInterface load()
    {
        try
        {
            Class aClass = customLoader.loadClass(API_IMPL);
            return (apiInterface) aClass.newInstance();
        }
        catch (Exception e)
        {
            e.printStackTrace();
            return null;
        }
    }
}

```

The following DATA step program calls these methods by delegating through the **api** Java object instance class. The Java object instantiates the **api** class, which creates a

custom class loader to load classes from the **z** folder. The **api** class calls the custom loader and returns an instance of the **apiImpl** interface implementation class to the Java object. When methods are called through the Java object, the **api** class delegates them to the implementation class.

```
/* DATA step code */
data _null_;
  dcl javaobj j('api');
  j.callvoidmethod('m');
  j.callvoidmethod('m2');
run;
```

The following lines are written to the SAS log:

```
method m is z folder
method m2 in z folder
```

In the previous Java code, you could also use **.jar** files to augment the classpath in the **ClassLoader** constructor.

```
static ClassLoader customLoader = new apiClassLoader("C:\\z;C:\\temp\\some.jar");
```

In this case, the Java code for the custom class loader is as follows. This code for this class loader can be added to or modified as needed.

```
import java.io.*;
import java.util.*;
import java.util.jar.*;
import java.util.zip.*;

public class apiClassLoader extends ClassLoader
{
  //class repository where findClass performs its search
  private List classRepository;

  public apiClassLoader(String loadPath)
  {
    super(apiClassLoader.class.getClassLoader());
    initLoader(loadPath);
  }

  public apiClassLoader(ClassLoader parent,String loadPath)
  {
    super(parent);
    initLoader(loadPath);
  }

  /**
   * This method will look for the class in the class repository. If
   * the method cannot find the class, the method will delegate to its parent
   * class loader.
   *
   * @param className A String specifying the class to be loaded
   * @return A Class object loaded by the apiClassLoader
   * @throws ClassNotFoundException if the method is unable to load the class
   */
}
```

```

*/
public Class loadClass(String name) throws ClassNotFoundException
{
    // Check if the class is already loaded
    Class loadedClass = findLoadedClass(name);

    // Search for class in local repository before delegating
    if (loadedClass == null)
    {
        loadedClass = myFindClass(name);
    }

    // If class not found, delegate to parent
    if (loadedClass == null)
    {
        loadedClass = this.getClass().getClassLoader().loadClass(name);
    }
    return loadedClass;
}

private Class myFindClass(String className) throws ClassNotFoundException
{
    byte[] classBytes = loadFromCustomRepository(className);
    if(classBytes != null)
    {
        return defineClass(className,classBytes,0,classBytes.length);
    }
    return null;
}

/**
 * This method loads binary class file data from the classRepository.
 */
private byte[] loadFromCustomRepository(String classFileName)
throws ClassNotFoundException
{
    Iterator dirs = classRepository.iterator();
    byte[] classBytes = null;
    while (dirs.hasNext())
    {
        String dir = (String) dirs.next();

        if (dir.endsWith(".jar"))
        {
            // Look for class in jar

            String jclassFileName = classFileName;

            jclassFileName = jclassFileName.replace('.', '/');
            jclassFileName += ".class";

            try
            {
                JarFile j = new JarFile(dir);

```

```

for (Enumeration e = j.entries(); e.hasMoreElements() ;)
{
    Object n = e.nextElement();

    if (jclassFileName.equals(n.toString()))
    {
        ZipEntry zipEntry = j.getEntry(jclassFileName);
        if (zipEntry == null)
        {
            return null;
        }
        else
        {
            // read file
            InputStream is = j.getInputStream(zipEntry);
            classBytes = new byte[is.available()];
            is.read(classBytes);
            break;
        }
    }
}
catch (Exception e)
{
    System.out.println("jar file exception");
    return null;
}
}
else
{
    // Look for class in directory
    String fclassFileName = classFileName;


    fclassFileName = fclassFileName.replace('.', File.separatorChar);
    fclassFileName += ".class";

    try
    {
        File file = new File(dir, fclassFileName);
        if(file.exists()) {
            //read file
            InputStream is = new FileInputStream(file);
            classBytes = new byte[is.available()];
            is.read(classBytes);
            break;
        }
    }
    catch(IOException ex)
    {
        System.out.println("IOException raised while reading class file data");
        ex.printStackTrace();
        return null;
    }
}
}

```

```
    }
    return classBytes;
}

private void initLoader(String loadPath)
{
    /*
     * loadPath is passed in as a string of directories/jar files
     * separated by the File.pathSeparator
     */
    classRepository = new ArrayList();
    if((loadPath != null) && !(loadPath.equals("")))
    {
        StringTokenizer tokenizer =
            new StringTokenizer(loadPath,File.pathSeparator);
        while(tokenizer.hasMoreTokens())
        {
            classRepository.add(tokenizer.nextToken());
        }
    }
}
}
```



CHAPTER

24

Array Processing

<i>Definitions for Array Processing</i>	429
<i>A Conceptual View of Arrays</i>	430
<i>One-Dimensional Array</i>	430
<i>Two-Dimensional Array</i>	431
<i>Syntax for Defining and Referencing an Array</i>	431
<i>Processing Simple Arrays</i>	432
<i>Grouping Variables in a Simple Array</i>	432
<i>Using a DO Loop to Repeat an Action</i>	433
<i>Using a DO Loop to Process Selected Elements in an Array</i>	433
<i>Selecting the Current Variable</i>	434
<i>Defining the Number of Elements in an Array</i>	435
<i>Rules for Referencing Arrays</i>	435
<i>Variations on Basic Array Processing</i>	436
<i>Determining the Number of Elements in an Array Efficiently</i>	436
<i>DO WHILE and DO UNTIL Expressions</i>	436
<i>Using Variable Lists to Define an Array Quickly</i>	436
<i>Multidimensional Arrays: Creating and Processing</i>	437
<i>Grouping Variables in a Multidimensional Array</i>	437
<i>Using Nested DO Loops</i>	438
<i>Specifying Array Bounds</i>	439
<i>Identifying Upper and Lower Bounds</i>	439
<i>Determining Array Bounds: LBOUND and HBOUND Functions</i>	440
<i>When to Use the HBOUND Function instead of the DIM Function</i>	440
<i>Specifying Bounds in a Two-Dimensional Array</i>	441
<i>Examples of Array Processing</i>	441
<i>Example 1: Using Character Variables in an Array</i>	441
<i>Example 2: Assigning Initial Values to the Elements of an Array</i>	442
<i>Example 3: Creating an Array for Temporary Use in the Current DATA Step</i>	443
<i>Example 4: Performing an Action on All Numeric Variables</i>	444

Definitions for Array Processing

array

is a temporary grouping of SAS variables that are arranged in a particular order and identified by an *array-name*. The array exists only for the duration of the current DATA step. The array-name distinguishes it from any other arrays in the same DATA step; it is not a variable.

Note: Arrays in SAS are different from those in many other programming languages. In SAS, an array is not a data structure but is just a convenient way of temporarily identifying a group of variables. Δ

array processing

is a method that enables you to perform the same tasks for a series of related variables.

array reference

is a method to reference the elements of an array.

one-dimensional array

is a simple grouping of variables that, when processed, results in output that can be represented in simple row format.

multidimensional array

is a more complex grouping of variables that, when processed, results in output that could have two or more dimensions, such as columns and rows.

Basic array processing involves the following steps:


- grouping variables into arrays
- selecting a current variable for an action
- repeating an action

A Conceptual View of Arrays

One-Dimensional Array

The following figure is a conceptual representation of two one-dimensional arrays, MISC and MDAY.

Figure 24.1 One-Dimensional Array



MISC contains eight elements, the variables MISC1 through MISC8. To reference the data in these variables, use the form MISC(*n*), where *n* is the element number in the array. For example, MISC(6) is the sixth element in the array.

MDAY contains seven elements, the variables MDAY1 through MDAY7. MDAY(3) is the third element in the array.

Two-Dimensional Array

The following figure is a conceptual representation of the two-dimensional array EXPENSES.

Figure 24.2 Example of a Two-Dimensional Array

Expense Categories	First Dimension	Second Dimension							Total
		Days of the Week							
		1	2	3	4	5	6	7	8
Hotel	1	hotel1	hotel2	hotel3	hotel4	hotel5	hotel6	hotel7	hotel8
Phone	2	phone1	phone2	phone3	phone4	phone5	phone6	phone7	phone8
Pers. Auto	3	peraut1	peraut2	peraut3	peraut4	peraut5	peraut6	peraut7	peraut8
Rental Car	4	carrnt1	carrnt2	carrnt3	carrnt4	carrnt5	carrnt6	carrnt7	carrnt8
Airfare	5	airlin1	airlin2	airlin3	airlin4	airlin5	airlin6	airlin7	airlin8
Dues	6	dues1	dues2	dues3	dues4	dues5	dues6	dues7	dues8
Registration Fees	7	regfee1	regfee2	regfee3	regfee4	regfee5	regfee6	regfee7	regfee8
Other	8	other1	other2	other3	other4	other5	other6	other7	other8
Tips (non-meal)	9	tips1	tips2	tips3	tips4	tips5	tips6	tips7	tips8
Meals	10	meals1	meals2	meals3	meals4	meals5	meals6	meals7	meals8

The EXPENSES array contains ten groups of eight variables each. The ten groups (expense categories) comprise the first dimension of the array, and the eight variables (days of the week) comprise the second dimension. To reference the data in the array variables, use the form EXPENSES(m,n), where m is the element number in the first dimension of the array, and n is the element number in the second dimension of the array. EXPENSES(6,4) references the value of dues for the fourth day (the variable is DUES4).

Syntax for Defining and Referencing an Array

To define a simple or a multidimensional array, use the ARRAY statement. The ARRAY statement has the following form:

```
ARRAY array-name {number-of-elements} <$> <length> <array-elements>
<(initial-value-list)>;
```

where

array-name

is a SAS name that identifies the group of variables.

number-of-elements

is the number of variables in the group. You must enclose this value in either parentheses (), braces { }, or brackets [].

\$

specifies that the elements in the array are character elements.

length

specifies the length of the elements in the array that have not been previously assigned a length.

array-elements

is a list of the names of the variables in the group. All variables that are defined in a given array must be of the same type, either all character or all numeric.

initial-value-list

is a list of the initial values for the corresponding elements in the array.

For complete information about the ARRAY statement, see *SAS Language Reference: Dictionary*.

To reference an array that was previously defined in the same DATA step, use an Array Reference statement. An array reference has the following form:

```
array-name {subscript}
```

where

array-name

is the name of an array that was previously defined with an ARRAY statement in the same DATA step.

subscript

specifies the subscript, which can be a numeric constant, the name of a variable whose value is the number, a SAS numeric expression, or an asterisk (*).

Note: Subscripts in SAS are 1-based by default, and not 0-based as they are in some other programming languages. Δ

For complete information about the Array Reference statement, see *SAS Language Reference: Dictionary*.

Processing Simple Arrays

Grouping Variables in a Simple Array

The following ARRAY statement creates an array named BOOKS that contains the three variables Reference, Usage, and Introduction:

```
array books{3} Reference Usage Introduction;
```

When you define an array, SAS assigns each array element an *array reference* with the form *array-name*{*subscript*}, where *subscript* is the position of the variable in the list. The following table lists the array reference assignments for the previous ARRAY statement:

Variable	Array Reference
Reference	books{1}
Usage	books{2}
Introduction	books{3}

Later in the DATA step, when you want to process the variables in the array, you can refer to a variable by either its name or its array reference. For example, the names Reference and books{1} are equivalent.

Using a DO Loop to Repeat an Action

To perform the same action several times, use an iterative DO loop. A simple iterative DO loop that processes an array has the following form:

```
DO index-variable=1 TO number-of-elements-in-array;
    ... more SAS statements ...
END;
```

The loop is processed repeatedly (iterates) according to the instructions in the iterative DO statement. The iterative DO statement contains an *index-variable* whose name you specify and whose value changes at each iteration of the loop.

To execute the loop as many times as there are variables in the array, specify that the values of *index-variable* are 1 TO *number-of-elements-in-array*. SAS increases the value of *index-variable* by 1 before each new iteration of the loop. When the value exceeds the *number-of-elements-in-array*, SAS stops processing the loop. By default, SAS automatically includes *index-variable* in the output data set. Use a DROP statement or the DROP= data set option to prevent the index variable from being written to your output data set.

An iterative DO loop that executes three times and has an index variable named count has the following form:

```
do count=1 to 3;
    ... more SAS statements ...
end;
```

The first time the loop processes, the value of count is 1; the second time, 2; and the third time, 3. At the beginning of the fourth iteration, the value of count is 4, which exceeds the specified range and causes SAS to stop processing the loop.

Using a DO Loop to Process Selected Elements in an Array

To process particular elements of an array, specify those elements as the range of the iterative DO statement. For example, the following statement creates an array DAYS that contains seven elements:

```
array days{7} D1-D7;
```

The following DO statements process selected elements of the array DAYS:

<code>do i=2 to 4;</code>	processes elements 2 through 4
<code>do i=1 to 7 by 2;</code>	processes elements 1, 3, 5, and 7
<code>do i=3,5;</code>	processes elements 3 and 5

Selecting the Current Variable

You must tell SAS which variable in the array to use in each iteration of the loop. Recall that you identify variables in an array by their array references and that you use a variable name, a number, or an expression as the subscript of the reference. Therefore, you can write programming statements so that the index variable of the DO loop is the subscript of the array reference (for example, *array-name{index-variable}*). When the value of the index variable changes, the subscript of the array reference (and therefore the variable that is referenced) also changes.

The following example uses the index variable count as the subscript of array references inside a DO loop:

```
array books{3} Reference Usage Introduction;
do count=1 to 3;
  if books{count}=. then books{count}=0;
end;
```

When the value of count is 1, SAS reads the array reference as books{1} and processes the IF-THEN statement on books{1}, which is the variable Reference. When count is 2, SAS processes the statement on books{2}, which is the variable Usage. When count is 3, SAS processes the statement on books{3}, which is the variable Introduction.

The statements in the example tell SAS to

- perform the actions in the loop three times
- replace the array subscript count with the current value of count for each iteration of the IF-THEN statement
- locate the variable with that array reference and process the IF-THEN statement on it
- replace missing values with zero if the condition is true.

The following DATA step defines the array BOOK and processes it with a DO loop.

```
options nodate pageno=1 linesize=80 pagesize=60;

data changed(drop=count);
  input Reference Usage Introduction;
  array book{3} Reference Usage Introduction;
  do count=1 to 3;
    if book{count}=. then book{count}=0;
  end;
  datalines;
45 63 113
. 75 150
62 . 98
;
```

```
proc print data=changed;
  title 'Number of Books Sold';
run;
```

The following output shows the CHANGED data set.

Output 24.1 Using an Array Statement to Process Missing Data Values

Number of Books Sold				1
Obs	Reference	Usage	Introduction	
1	45	63	113	
2	0	75	150	
3	62	0	98	

Defining the Number of Elements in an Array

When you define the number of elements in an array, you can either use an asterisk enclosed by braces ({*}), brackets ([*]), or parentheses ((*)) to count the number of elements or to specify the number of elements. You must list each array element if you use the asterisk to designate the number of elements. In the following example, the array C1TEMP references five variables with temperature measures.

```
array c1temp{*} c1t1 c1t2 c1t3 c1t4 c1t5;
```

If you specify the number of elements explicitly, you can omit the names of the variables or array elements in the ARRAY statement. SAS then creates variable names by concatenating the array name with the numbers 1, 2, 3, and so on. If a variable name in the series already exists, SAS uses that variable instead of creating a new one. In the following example, the array c1t references five variables: c1t1, c1t2, c1t3, c1t4, and c1t5.

```
array c1t{5};
```

Rules for Referencing Arrays

Before you make any references to an array, an ARRAY statement must appear in the same DATA step that you used to create the array. Once you have created the array, you can

- use an array reference anywhere that you can write a SAS expression
- use an array reference as the arguments of some SAS functions
- use a subscript enclosed in braces, brackets, or parentheses to reference an array
- use the special array subscript asterisk (*) to refer to all variables in an array in an INPUT or PUT statement or in the argument of a function.

Note: You cannot use the asterisk with `_TEMPORARY_` arrays. Δ

An array definition is in effect only for the duration of the DATA step. If you want to use the same array in several DATA steps, you must redefine the array in each step. You can, however, redefine the array with the same variables in a later DATA step by using a macro variable. A macro variable is useful for storing the variable names you need, as shown in this example:

```
%let list=NC SC GA VA;
```

```

data one;
  array state(*) &list;
  ... more SAS statements ...
run;

data two;
  array state(*) &list;
  ... more SAS statements ...
run;

```

Variations on Basic Array Processing

Determining the Number of Elements in an Array Efficiently

The DIM function in the iterative DO statement returns the number of elements in a one-dimensional array or the number of elements in a specified dimension of a multidimensional array, when the lower bound of the dimension is 1. Use the DIM function to avoid changing the upper bound of an iterative DO group each time you change the number of elements in the array.

The form of the DIM function is as follows:

DIM n (array-name)

where n is the specified dimension that has a default value of 1.

You can also use the DIM function when you specify the number of elements in the array with an asterisk. Here are some examples of the DIM function:

- `do i=1 to dim(days);`
- `do i=1 to dim4(days) by 2;`

DO WHILE and DO UNTIL Expressions

Arrays are often processed in iterative DO loops that use the array reference in a DO WHILE or DO UNTIL expression. In this example, the iterative DO loop processes the elements of the array named TREND.

```

data test;
  array trend{5} x1-x5;
  input x1-x5 y;
  do i=1 to 5 while(trend{i}<y);
  ... more SAS statements ...
  end;
  datalines;
... data lines ...
;

```

Using Variable Lists to Define an Array Quickly

SAS reserves the following three names for use as variable list names:

```

_CHARACTER_
_NUMERIC_
_ALL_

```

You can use these variable list names to reference variables that have been previously defined in the same DATA step. The `_CHARACTER_` variable lists character values only. The `_NUMERIC_` variable lists numeric values only. The `_ALL_` variable lists either all character or all numeric values, depending on how you previously defined the variables.

For example, the following INPUT statement reads in variables X1 through X3 as character values using the \$8. informat, and variables X4 through X5 as numeric variables. The following ARRAY statement uses the variable list `_CHARACTER_` to include only the character variables in the array. The asterisk indicates that SAS will determine the subscript by counting the variables in the array.

```
input (x1-x3) ($8.) x4-x5;
array item {*} _character_;
```

You can use the `_NUMERIC_` variable in your program if, for example, you need to convert currency. In this application, you do not need to know the variable names. You need only to convert all values to the new currency.

For more information about variable lists, see the ARRAY statement in *SAS Language Reference: Dictionary*.

Multidimensional Arrays: Creating and Processing

Grouping Variables in a Multidimensional Array

To create a multidimensional array, place the number of elements in each dimension after the array name in the form $\{n, \dots\}$ where n is required for each dimension of a multidimensional array.

From right to left, the rightmost dimension represents columns; the next dimension represents rows. Each position farther left represents a higher dimension. The following ARRAY statement defines a two-dimensional array with two rows and five columns. The array contains ten variables: five temperature measures (t1 through t5) from two cities (c1 and c2):

```
array temprg{2,5} c1t1-c1t5 c2t1-c2t5;
```

SAS places variables into a multidimensional array by filling all rows in order, beginning at the upper-left corner of the array (known as row-major order). You can think of the variables as having the following arrangement:

```
c1t1 c1t2 c1t3 c1t4 c1t5
c2t1 c2t2 c2t3 c2t4 c2t5
```

To refer to the elements of the array later with an array reference, you can use the array name and subscripts. The following table lists some of the array references for the previous example:

Variable	Array Reference
c1t1	temprg[1,1]
c1t2	temprg[1,2]

Variable	Array Reference
c2t2	temprg{2,2}
c2t5	temprg{2,5}

Using Nested DO Loops

Multidimensional arrays are usually processed inside nested DO loops. As an example, the following is one form that processes a two-dimensional array:

```
DO index-variable-1=1 TO number-of-rows;
  DO index-variable-2=1 TO number-of-columns;
    ... more SAS statements ...
  END;
END;
```

An array reference can use two or more index variables as the subscript to refer to two or more dimensions of an array. Use the following form:

```
array-name {index-variable-1, ...,index-variable-n}
```

The following example creates an array that contains ten variables- five temperature measures (t1 through t5) from two cities (c1 and c2). The DATA step contains two DO loops.

- The outer DO loop (DO I=1 TO 2) processes the inner DO loop twice.
- The inner DO loop (DO J=1 TO 5) applies the ROUND function to all the variables in one row.

For each iteration of the DO loops, SAS substitutes the value of the array element corresponding to the current values of I and J.

```
options nodate pageno=1 linesize=80 pagesize=60;

data temps;
  array temprg{2,5} c1t1-c1t5 c2t1-c2t5;
  input c1t1-c1t5 /
        c2t1-c2t5;
  do i=1 to 2;
    do j=1 to 5;
      temprg{i,j}=round(temprg{i,j});
    end;
  end;
  datalines;
89.5 65.4 75.3 77.7 89.3
73.7 87.3 89.9 98.2 35.6
75.8 82.1 98.2 93.5 67.7
101.3 86.5 59.2 35.6 75.7
;

proc print data=temps;
  title 'Temperature Measures for Two Cities';
run;
```

The following data set TEMPS contains the values of the variables rounded to the nearest whole number.

Output 24.2 Using a Multidimensional Array

Temperature Measures for Two Cities											1	
Obs	c1t1	c1t2	c1t3	c1t4	c1t5	c2t1	c2t2	c2t3	c2t4	c2t5	i	j
1	90	65	75	78	89	74	87	90	98	36	3	6
2	76	82	98	94	68	101	87	59	36	76	3	6

The previous example can also use the DIM function to produce the same result:

```
do i=1 to dim1(temprg);
  do j=1 to dim2(temprg);
    temprg{i,j}=round(temprg{i,j});
  end;
end;
```

The value of DIM1(TEMPRG) is 2; the value of DIM2(TEMPRG) is 5.

Specifying Array Bounds

Identifying Upper and Lower Bounds

Typically in an ARRAY statement, the subscript in each dimension of the array ranges from 1 to n , where n is the number of elements in that dimension. Thus, 1 is the lower bound and n is the upper bound of that dimension of the array. For example, in the following array, the lower bound is 1 and the upper bound is 4:

```
array new{4} Jackson Poulenc Andrew Parson;
```

In the following ARRAY statement, the bounds of the first dimension are 1 and 2 and those of the second dimension are 1 and 5:

```
array test{2,5} test1-test10;
```

Bounded array dimensions have the following form:

```
{<lower-1:>upper-1<,...<lower-n:>upper-n<>}
```

Therefore, you can also write the previous ARRAY statements as follows:

```
array new{1:4} Jackson Poulenc Andrew Parson;
array test{1:2,1:5} test1-test10;
```

For most arrays, 1 is a convenient lower bound, so you do not need to specify the lower bound. However, specifying both the lower and the upper bounds is useful when the array dimensions have beginning points other than 1.

In the following example, ten variables are named Year76 through Year85. The following ARRAY statements place the variables into two arrays named FIRST and SECOND:

```
array first{10} Year76-Year85;
array second{76:85} Year76-Year85;
```

In the first ARRAY statement, the element first{4} is variable Year79, first{7} is Year82, and so on. In the second ARRAY statement, element second{79} is Year79 and second{82} is Year82.

To process the array names SECOND in a DO group, make sure that the range of the DO loop matches the range of the array as follows:

```
do i=76 to 85;
    if second{i}=9 then second{i}=.;
end;
```

Determining Array Bounds: LBOUND and HBOUND Functions

You can use the LBOUND and HBOUND functions to determine array bounds. The LBOUND function returns the lower bound of a one-dimensional array or the lower bound of a specified dimension of a multidimensional array. The HBOUND function returns the upper bound of a one-dimensional array or the upper bound of a specified dimension of a multidimensional array.

The form of the LBOUND and HBOUND functions is as follows:

LBOUND $_n$ (array-name)

HBOUND $_n$ (array-name)

where

n

is the specified dimension and has a default value of 1.

You can use the LBOUND and HBOUND functions to specify the starting and ending values of the iterative DO loop to process the elements of the array named SECOND:

```
do i=lbound{second} to hbound{second};
    if second{i}=9 then second{i}=.;
end;
```

In this example, the index variable in the iterative DO statement ranges from 76 to 85.

When to Use the HBOUND Function instead of the DIM Function

The following ARRAY statement defines an array containing a total of five elements, a lower bound of 72, and an upper bound of 76. It represents the calendar years 1972 through 1976:

```
array years{72:76} first second third fourth fifth;
```

To process the array named YEARS in an iterative DO loop, make sure that the range of the DO loop matches the range of the array as follows:

```
do i=lbound(years) to hbound(years);
    if years{i}=99 then years{i}=.;
end;
```

The value of LBOUND(YEARS) is 72; the value of HBOUND(YEARS) is 76.

For this example, the DIM function would return a value of 5, the total count of elements in the array YEARS. Therefore, if you used the DIM function instead of the HBOUND function for the upper bound of the array, the statements inside the DO loop would not have executed.

Specifying Bounds in a Two-Dimensional Array

The following list contains 40 variables named X60 through X99. They represent the years 1960 through 1999.

```
X60  X61  X62  X63  X64  X65  X66  X67  X68  X69
X70  X71  X72  X73  X74  X75  X76  X77  X78  X79
X80  X81  X82  X83  X84  X85  X86  X87  X88  X89
X90  X91  X92  X93  X94  X95  X96  X97  X98  X99
```

The following ARRAY statement arranges the variables in an array by decades. The rows range from 6 through 9, and the columns range from 0 through 9.

```
array X{6:9,0:9} X60-X99;
```

In array X, variable X63 is element X{6,3} and variable X89 is element X{8,9}. To process array X with iterative DO loops, use one of these methods:

Method 1:

```
do i=6 to 9;
  do j=0 to 9;
    if X{i,j}=0 then X{i,j}=.;
  end;
end;
```

Method 2:

```
do i=lbound1(X) to hbound1(X);
  do j=lbound2(X) to hbound2(X);
    if X{i,j}=0 then X{i,j}=.;
  end;
end;
```

Both examples change all values of 0 in variables X60 through X99 to missing. The first example sets the range of the DO groups explicitly, and the second example uses the LBOUND and HBOUND functions to return the bounds of each dimension of the array.

Examples of Array Processing

Example 1: Using Character Variables in an Array

You can specify character variables and their lengths in ARRAY statements. The following example groups variables into two arrays, NAMES and CAPITALS. The dollar sign (\$) tells SAS to create the elements as character variables. If the variables have already been declared as character variables, a dollar sign in the array is not necessary. The INPUT statement reads all the variables in array NAMES.

The statement inside the DO loop uses the UPCASE function to change the values of the variables in array NAMES to uppercase and then store the uppercase values in the variables in the CAPITALS array.

```
options nodate pageno=1 linesize=80 pagesize=60;

data text;
```

```

array names{*} $ n1-n10;
array capitals{*} $ c1-c10;
input names{*};
  do i=1 to 10;
    capitals{i}=upcase(names{i});
  end;
datalines;
smithers michaelson gonzalez hurth frank bleigh
rounder joseph peters sam
;

proc print data=text;
  title 'Names Changed from Lowercase to Uppercase';
run;

```

The following output shows the TEXT data set.

Output 24.3 Using Character Variables in an Array

Names Changed from Lowercase to Uppercase											1
Obs	n1	n2	n3	n4	n5	n6	n7	n8	n9	n10	
1	smithers	michaelson	gonzalez	hurth	frank	bleigh	rounder	joseph	peters	sam	
Obs	c1	c2	c3	c4	c5	c6	c7	c8	c9	c10	i
1	SMITHERS	MICHAELSON	GONZALEZ	HURTH	FRANK	BLEIGH	ROUNDER	JOSEPH	PETERS	SAM	11

Example 2: Assigning Initial Values to the Elements of an Array

This example creates variables in the array TEST and assigns them the initial values 90, 80, and 70. It reads values into another array named SCORE and compares each element of SCORE to the corresponding element of TEST. If the value of the element in SCORE is greater than or equal to the value of the element in TEST, the variable NewScore is assigned the value in the element SCORE, and the OUTPUT statement writes the observation to the SAS data set.

The INPUT statement reads a value for the variable named ID and then reads values for all the variables in the SCORE array.

```

options nodate pageno=1 linesize=80 pagesize=60;

data score1(drop=i);
  array test{3} t1-t3 (90 80 70);
  array score{3} s1-s3;
  input id score{*};
  do i=1 to 3;
    if score{i}>=test{i} then
      do;
        NewScore=score{i};
        output;
      end;
  end;
datalines;
1234 99 60 82

```

```

5678 80 85 75
;

proc print noobs data=score1;
  title 'Data Set SCORE1';
run;

```

The following output shows the SCORE1 data set.

Output 24.4 Assigning Initial Values to the Elements of an Array

Data Set SCORE1								1
t1	t2	t3	s1	s2	s3	id	New Score	
90	80	70	99	60	82	1234	99	
90	80	70	99	60	82	1234	82	
90	80	70	80	85	75	5678	85	
90	80	70	80	85	75	5678	75	

Example 3: Creating an Array for Temporary Use in the Current DATA Step

When elements of an array are constants that are needed only for the duration of the DATA step, you can omit variables from an array group and instead use temporary array elements. You refer to temporary data elements by the array name and dimension. Although they behave like variables, temporary array elements do not have names, and they do not appear in the output data set. Temporary array elements are automatically retained, instead of being reset to missing at the beginning of the next iteration of the DATA step.

To create a temporary array, use the `_TEMPORARY_` argument. The following example creates a temporary array named TEST:

```

options nodate pageno=1 linesize=80 pagesize=60;

data score2(drop=i);
  array test{3} _temporary_ (90 80 70);
  array score{3} s1-s3;
  input id score{*};
  do i=1 to 3;
    if score{i}>=test{i} then
      do;
        NewScore=score{i};
        output;
      end;
    end;
  datalines;
1234 99 60 82
5678 80 85 75
;

proc print noobs data=score2;
  title 'Data Set SCORE2';
run;

```

The following output shows the SCORE2 data set.

Output 24.5 Using `_TEMPORARY_` Arrays

Data Set SCORE2					1
s1	s2	s3	id	New Score	
99	60	82	1234	99	
99	60	82	1234	82	
80	85	75	5678	85	
80	85	75	5678	75	

Example 4: Performing an Action on All Numeric Variables

This example multiplies all the numeric variables in array TEST by 3.

```
options nodate pageno=1 linesize=80 pagesize=60;
```

```
data sales;
  infile datalines;
  input Value1 Value2 Value3 Value4;
  datalines;
11 56 58 61
22 51 57 61
22 49 53 58
;


data convert(drop=i);
  set sales;
  array test{*} _numeric_;
  do i=1 to dim(test);
    test{i} = (test{i}*3);
  end;
run;

proc print data=convert;
  title 'Data Set CONVERT';
run;
```

The following output shows the CONVERT data set.


Output 24.6 Output From Using a `_NUMERIC_` Variable List

Data Set CONVERT					1
Obs	Value1	Value2	Value3	Value4	
1	33	168	174	183	
2	66	153	171	183	
3	66	147	159	174	



SAS Files Concepts

- Chapter 25*.....**SAS Libraries** 447
- Chapter 26*.....**SAS Data Sets** 461
- Chapter 27*.....**SAS Data Files** 475
- Chapter 28*.....**SAS Views** 529
- Chapter 29*.....**Stored Compiled DATA Step Programs** 539
- Chapter 30*.....**DICTIONARY Tables** 549
- Chapter 31*.....**SAS Catalogs** 553
- Chapter 32*.....**About SAS/ACCESS Software** 561
- Chapter 33*.....**Processing Data Using Cross-Environment Data Access (CEDA)** 567
- Chapter 34*.....**SAS 9.2 Compatibility with SAS Files from Earlier Releases** 577
- Chapter 35*.....**File Protection** 581
- Chapter 36*.....**SAS Engines** 589
- Chapter 37*.....**SAS File Management** 599
- Chapter 38*.....**External Files** 605



CHAPTER 25

SAS Libraries

<i>Definition of a SAS Library</i>	447
<i>Library Engines</i>	449
<i>Library Names</i>	449
<i>Physical Names and Logical Names (Librefs)</i>	449
<i>Assigning Librefs</i>	450
<i>Associating and Clearing Logical Names (Librefs) with the LIBNAME Statement</i>	450
<i>Reserved Librefs</i>	451
<i>Accessing Remote SAS Libraries on SAS/CONNECT, SAS/SHARE, and WebDAV Servers</i>	451
<i>Remote Library Access for SAS/CONNECT and SAS/SHARE</i>	451
<i>Remote Library Access for WebDAV Servers</i>	452
<i>Library Concatenation</i>	452
<i>Definition of Library Concatenation</i>	452
<i>How SAS Concatenates Library Members</i>	452
<i>Rules for Library Concatenation</i>	453
<i>Permanent and Temporary Libraries</i>	454
<i>SAS System Libraries</i>	455
<i>Introduction to SAS System Libraries</i>	455
<i>WORK Library</i>	455
<i>Definition of WORK Library</i>	455
<i>Using the WORK Library</i>	455
<i>Relation to the USER Library</i>	455
<i>USER Library</i>	456
<i>Definition of USER Library</i>	456
<i>Ways to Assign the USER Libref</i>	456
<i>Relation to WORK Library</i>	456
<i>SASHELP Library</i>	457
<i>SASUSER Library</i>	457
<i>Sequential Data Libraries</i>	457
<i>Tools for Managing Libraries</i>	458
<i>SAS Utilities</i>	458
<i>Library Directories</i>	459
<i>Accessing Permanent SAS Files without a Libref</i>	459
<i>Operating Environment Commands</i>	460

Definition of a SAS Library

The logical concept of a SAS library remains constant, regardless of the operating environment. In any operating environment where SAS can be installed, the structure for organizing, locating, and managing SAS files is the same.


At the operating environment level, however, a SAS library has different physical implementations. Most SAS libraries implement the storage of files in a manner similar to the way the operating environment stores and accesses files.

For example, in directory-based operating environments, a SAS library is a group of SAS files that are stored in the same directory and accessed by the same engine. Other files can be stored in the directory, but only the files with file extensions that are assigned by SAS are recognized as part of the SAS library. Under OS/390 or z/OS, a SAS library can be implemented as either a bound library in a traditional OS data set or as a directory under UNIX System Services.

SAS files can be any of the following file types:

- SAS data set (SAS data file or SAS view)
- SAS catalog
- stored compiled SAS program
- SAS utility file
- access descriptors
- multi-dimensional database files such as MDDB, FDB, and DMDB files
- item store files

Figure 25.1 Types of Files in a SAS Library



Each SAS file, in turn, stores information in smaller units that are characteristic of the *SAS file type*. For example, SAS data sets store information as variables and observations, while SAS catalogs store information in units called *entries*. SAS determines the type of a file from the context of the SAS program in which the file is created or specified; therefore, a library can contain files with the same name but with different member types.

SAS libraries can contain files that you create, or they can be one of several special libraries that SAS provides for convenience, support, and customizing capability such as the WORK library. SAS does not limit the number of SAS files you can store in a SAS library.

Library Engines

Each SAS library is associated with a library engine. SAS library engines are software components that form the interface between SAS and the SAS library. It is the SAS library engine that locates files in a SAS library and renders the file contents to SAS in a form that it can recognize. Library engines perform such tasks as:

- reading and writing data
- listing the files in the library
- deleting and renaming files.

SAS has a *Multi Engine Architecture* in order to read to and write from files in different formats. Each SAS engine has specific processing characteristics, such as the ability to

- process a SAS file generated by an older version of SAS
- read database files created by other software programs
- store and access files on disk or tape
- determine how variables and observations are placed in a file
- place data into memory from its physical location
- transport SAS files between operating environments.

You generally are not aware of the particular type of engine that is processing data at any given time. If you issue an instruction that is not supported by the engine, an error message is displayed in the log. When needed, you can select a specific engine to perform a task. But usually, you do not have to specify an engine, because SAS automatically selects the appropriate one.

More than one engine might be involved in processing a DATA step; for example, one engine might be used to input data, and another engine might be used to write observations to the output data set.

For more information on library engines, including a list of engines available in Base SAS, see “About Library Engines” on page 593.

Library Names

Physical Names and Logical Names (Librefs)

Before you can use a SAS library, you must tell SAS where it is. SAS recognizes SAS libraries based on either operating environment naming conventions or SAS naming conventions. There are two ways to define SAS libraries.

- a physical location name that the operating environment recognizes
- a logical name (libref) that you assign using the LIBNAME statement, LIBNAME function, or the New Library window.

The physical location name of the SAS library is a name that identifies your SAS files to the operating environment. The physical location name must conform to the naming conventions of your operating environment. The physical location name fully identifies the directory, or operating environment data set that contains the SAS library.

The logical name, or *libref*, is the way you identify a group of files to SAS. A libref is a name that you associate with the physical location of the SAS library.

Assigning Librefs

Librefs can be assigned using the following methods:

- LIBNAME statement
- LIBNAME function
- New Library window that is available in your toolbar
- operating environment commands

Once the libref is assigned, you can read, create, or update files in a SAS library. A libref is valid only for the current SAS session, unless it is assigned using the New Library window with the **Enable at startup** box checked.

A libref can have a maximum length of eight characters. You can use the LIBREF function to verify that a libref has been assigned. Librefs can be referenced repeatedly within a SAS session. SAS does not limit the number of librefs you can assign during a session; however, your operating environment or site might set limitations. If you are running in batch mode, the library must exist before you can allocate or assign it. In interactive mode, you might be allowed to create it if it does not already exist.

Operating Environment Information: Here are examples of the LIBNAME statement for different operating environments. The rules for assigning and using librefs differ across operating environments. See the SAS documentation for your operating environment for specific information. \triangle

Table 25.1 Syntax for Assigning a Libref

Operating Environment	Examples
DOS, Windows	<code>libname mylibref 'c:\root\mystuff\sasstuff\work';</code>
UNIX	<code>libname mylibref '/u/mystuff/sastuff/work';</code>
z/OS	<code>libname mylibref 'userid.mystuff.sastuff.work';</code> <code>libname mylibref '/mystuff/sastuff/work';</code>
OpenVMS for Integrity servers	<code>libname mylibref 'filename filetype filemode';</code>

You can also access files without using a libref. See “Accessing Permanent SAS Files without a Libref” on page 459.

Associating and Clearing Logical Names (Librefs) with the LIBNAME Statement

You can assign or clear a physical name with a libref using the LIBNAME statement or the LIBNAME function, which are described in the *SAS Language Reference: Dictionary*.

Operating Environment Information: For some operating environments, you can use operating environment commands to associate a libref with a SAS library. When using operating environment commands to assign librefs to a SAS library, the association might persist beyond the SAS session in which the libref was created. For some operating environments you can use only the LIBNAME statement or function. See the SAS documentation for your operating environment for more information on assigning librefs. \triangle

The most common form of the LIBNAME statement is used in this example to associate the libref ANNUAL with the physical name of the SAS library.

```
libname annual 'SAS-library';
```

If you use the LIBNAME statement to assign the libref, SAS clears (deassigns) the libref automatically at the end of each SAS session. If you want to clear the libref ANNUAL before the end of the session, you can issue the following form of the LIBNAME statement:

```
libname annual clear;
```

SAS also provides a New Library window to assign or clear librefs and SAS Explorer to view, add, or delete SAS libraries. You can select the New Library or the SAS Explorer icon from the Toolbar.

Reserved Librefs

SAS reserves a few names for special uses. You should not use SASHELP, SASUSER or SASWORK as librefs, except as intended. The purpose and content of these libraries are discussed in “Permanent and Temporary Libraries” on page 454.

Operating Environment Information: There are other librefs reserved for SAS under some operating environments. In addition, your operating environment might have reserved certain words that cannot be used as librefs. See the SAS documentation for your operating environment for more information. △

Accessing Remote SAS Libraries on SAS/CONNECT, SAS/SHARE, and WebDAV Servers

Remote Library Access for SAS/CONNECT and SAS/SHARE

You can use a LIBNAME statement to read, write, and update server (remote) data as if it were stored on the client’s disk. SAS processes the data in client memory, which gets overwritten in subsequent client requests for server data.

The LIBNAME statement can be used to access SAS data sets across computers that have different architectures. The LIBNAME statement also provides read-only access to some SAS catalog entry types across computers that have different architectures.

The LIBNAME statement provides access to remote server data by associating a SAS library reference (libref) with a permanent SAS library.

SAS/CONNECT Example:

The SAS/CONNECT client creates a LIBNAME statement to access a server library that is located on a SAS/CONNECT server. The client creates the new libref, REPORTS.

```
signon rempc;
libname reports 'd:\prod\reports' server=rempc;
```

The SAS/CONNECT client signs on to the SAS/CONNECT server named REMPC. A server library is assigned to the client session. The value for SERVER= is the same as the server session ID that is used in the SIGNON statement.

For more information about SAS/CONNECT, see *SAS/CONNECT User’s Guide*.

SAS/SHARE Example:

The SAS/SHARE client uses a LIBNAME statement to access a server library via the existing libref, SALES, which was pre-defined at the SAS/SHARE server for client access.

```
libname sales server=server1;
```

For more information about SAS/SHARE, see *SAS/SHARE User's Guide*.

Remote Library Access for WebDAV Servers

WebDAV (Web Distributed Authoring and Versioning) is a protocol that enhances the HTTP protocol. It provides a standard infrastructure for collaborative authoring across the Internet. WebDAV enables you to edit Web documents, stores versions for later retrieval, and provides a locking mechanism to prevent overwriting. SAS supports the WebDAV protocol under the UNIX and Windows operating environments.

You use a LIBNAME statement to access WebDAV servers, as shown in the following example:

```
libname davdata v9 "http://www.webserver.com/users/mydir/datadir"
      webdav user="mydir" pw="12345";
```

When you access files on a WebDAV server, SAS pulls the file from the server to your local disk for processing. The files are temporarily stored in the SAS WORK directory, unless you use the LOCALCACHE= option in the LIBNAME statement, which specifies a different directory for temporary storage. When you finish updating the file, SAS pushes the file back to the WebDAV server for storage and removes the file from the local disk.

For more information, see the “LIBNAME Statement for WebDAV Server Access” in *SAS Language Reference: Dictionary*.

Library Concatenation

Definition of Library Concatenation

Concatenation is the logical combining of two or more libraries. Concatenation enables you to access the SAS data sets in several libraries with one libref.

You can concatenate two or more libraries by specifying their librefs or physical names in the LIBNAME statement or function.

Physical names must be enclosed in single or double quotation marks in a LIBNAME statement. Otherwise SAS looks for a previously assigned libref with the same name.

In the following examples, summer, winter, spring, fall, and annual are previously defined librefs:

```
libname annual (summer winter spring fall);
```

```
libname annual ('path1' 'path2' 'path3');
```

```
libname annual ('path' winter spring fall);
```

```
libname total (annual 'path');
```

How SAS Concatenates Library Members

When there are members of the same name in more than one library, the first occurrence of the member is used for input and update processes. Output will always go to the first library.

This example contains three SAS libraries, and each library contains two SAS data files:

```
LIB1          APPLES and PEARS
LIB2          APPLES and ORANGES
LIB3          ORANGES and PLUMS
```

The LIBNAME statement concatenates LIB1, LIB2, and LIB3:

```
libname fruit (lib1 lib2 lib3);
```

The concatenated library FRUIT has the following members:

```
APPLES
PEARS
ORANGES
PLUMS
```

Note: Output will always go to the first library. For example, the following statement writes to the first library in the concatenation, LIB1:

```
data fruit.oranges;
```

Δ

Note that in this example, if the file APPLES in LIB1 was different from the file APPLES in LIB2, and if an update to APPLES was specified, it will be updated only in LIB1 because that is the first occurrence of the member APPLES.

For complete documentation on library concatenation, see the LIBNAME statement or function in *SAS Language Reference: Dictionary*.

Operating Environment Information: For more information about how specific operating environments handle concatenation, see the SAS documentation for your operating environment. Δ

Rules for Library Concatenation

After you create a library concatenation, you can specify the libref in any context that accepts a simple (nonconcatenated) libref. These rules determine how SAS files (that is, members of SAS libraries) are located among the concatenated libraries:

- If you specify any options or an engine, the options apply only to the libraries that you specified with the physical name, not to any library that you specified with a libref.
- When a SAS file is opened for input or update, the concatenated libraries are searched and the first occurrence of the specified file is used.
- When a SAS file is opened for output, it is created in the first library that is listed in the concatenation.

Note: A new SAS file is created in the first library even if there is a file with the same name in another part of the concatenation. Δ

- When you delete or rename a SAS file, only the first occurrence of the file is affected.
- Any time a list of SAS files is displayed, only one occurrence of a filename is shown, even if the name occurs multiple times in the concatenation. For example,

if library ONE contains A.DATA and library TWO contains A.DATA, only A.DATA from library ONE is listed because it is the first occurrence of the filename.

In addition, a SAS file that is logically connected to another file (such as an index to a data file) is listed only if the parent file is the first (or only) occurrence of the filename. For example, if library ONE contains A.DATA and library TWO contains A.DATA and A.INDEX, only A.DATA from library ONE is listed. A.DATA and A.INDEX from library TWO are not listed.

- If any library in the concatenation is sequential, then the concatenated library is considered sequential by applications that require random access. The DATASETS procedure, for example, cannot process sequential libraries, and therefore cannot process a concatenated library that contains one or more sequential libraries.
- The attributes of the first library that is specified determine the attributes of the concatenation. For example, if the first SAS library that is listed is “read only,” then the entire concatenated library is “read only.”
- Once a libref has been assigned in a concatenation, any changes made to the libref will not affect the concatenation.
- Changing a data set name to an existing name in the concatenation will fail.

Permanent and Temporary Libraries

SAS libraries are generally stored as permanent data libraries; however, SAS provides a temporary or scratch library where you can store files for the duration of a SAS session or job.

A *permanent SAS library* is one that resides on the external storage medium of your computer and is not deleted when the SAS session terminates. Permanent SAS libraries are stored until you delete them. The library is available for processing in subsequent SAS sessions. When working with files in a permanent SAS library, you generally specify a libref as the first part of a two-level SAS filename. The libref tells SAS where to find or store the file.

Note: You can also skip using a libref and point directly to the file you want to use, using syntax that your operating system understands. An example of this in the Windows environment is

```
data 'C:\root\sasfiles\myfile.ext';
```

Operating Environment Information: Files are specified differently in various operating environments. See the SAS documentation for your operating environment for more information. \triangle

\triangle

A *temporary SAS library* is one that exists only for the current SAS session or job. SAS files that are created during the session or job are held in a special work space that might or might not be an external storage medium. This work space is generally assigned the default libref WORK. Files in the temporary WORK library can be used in any DATA step or SAS procedure during the SAS session, but they are typically not available for subsequent SAS sessions. Normally, you specify that data sets be stored in or retrieved from this library by specifying a one-level name. Files held in the WORK library are deleted at the end of the SAS session if it ends normally.

There are a number of SAS system options that enable you to customize the way you name and work with your permanent and temporary SAS libraries. See the USER=, WORK=, WORKINIT, and WORKTERM system options in *SAS Language Reference: Dictionary* for more information.

SAS System Libraries

Introduction to SAS System Libraries

Four special SAS-supplied libraries provide convenience, support, and customization capability:

- WORK
- USER
- SASHELP
- SASUSER

WORK Library

Definition of WORK Library

The WORK library is the temporary (scratch) library that is automatically defined by SAS at the beginning of each SAS session. The WORK library stores two types of temporary files: those that you create and those that are created internally by SAS as part of normal processing. Typically, the WORK library is deleted at the end of each SAS session if the session terminates normally.

Using the WORK Library

To store or retrieve SAS files in the WORK library, specify a one-level name in your SAS program statements. The libref WORK is automatically assigned to these files as a system default unless you have assigned the USER libref. The following examples contain valid names for SAS data sets stored in the WORK library:

- `data test2;`
- `data work.test2;`
- `proc contents data=testdata;`
- `proc contents data=work.testdata;`

Operating Environment Information: The WORK library is implemented differently in various operating environments. See the SAS documentation for your operating environment for more information. △

Relation to the USER Library

While the WORK library is designed to hold temporary files used during the current SAS session, the USER library is designed to hold files after the SAS session is over. If you associate the libref USER with a SAS library, use a one-level name to create and access files that are not deleted at the end of your SAS session. When SAS encounters a one-level filename, it looks first in the USER library, if it has been defined, and then it looks in WORK. If you want to place a file in the USER library, so that it is not deleted

after your SAS session is over, any single-level file goes there by default. At that point, if you want to create a temporary file in WORK, you must use a two-level name, such as WORK.NAME.

USER Library

Definition of USER Library

The USER library enables you to read, create, and write to files in a SAS library other than WORK without specifying a libref as part of the SAS filename. Once you associate the libref USER with a SAS library, SAS stores any file with a one-level name in that library. Unlike the WORK library, files stored in this library are not deleted by SAS when the session terminates.

Ways to Assign the USER Libref

You can assign the USER libref using one of the following methods:

- LIBNAME statement
- LIBNAME function
- USER= system option
- operating environment command

In this example, the LIBNAME statement is used with a DATA step, which stores the data set REGION in a permanent SAS library.

```
libname user 'SAS-library';
data region;
... more DATA step statements ...
run;
```

In this example, the LIBNAME function assigns the USER libref:

```
data _null_;
  x=libname ('user', 'SAS-library');
run;
```

When assigning a libref using the USER= system option, you must first assign a libref to a SAS library, then use the USER= system option to specify that library as the default for one-level names. In this example, the DATA step stores the data set PROCHLOR in the SAS library TESTLIB.

```
libname testlib 'SAS-library';
options user=testlib;
data prochlor;
... more DATA step statements ...
run;
```

Operating Environment Information: The methods and results of assigning the USER libref vary slightly from one operating environment to another. See the SAS documentation for your operating environment for more information. \triangle

Relation to WORK Library

The USER libref overrides the default libref WORK for one-level names. When you refer to a file by a one-level name, SAS looks first for the libref USER. If USER is

assigned to a SAS library, files with one-level names are stored there. If you have not assigned the libref USER to a library, the files with one-level names are stored in the temporary library WORK. To refer to SAS files in the WORK library while the USER libref is assigned, you must specify a two-level name with WORK as the libref. Data files that SAS creates internally still go to the WORK library.

SASHELP Library

Each SAS site receives the SASHELP library, which contains a group of catalogs and other files containing information that is used to control various aspects of your SAS session. The defaults stored in this library are for everyone using SAS at your installation. Your personal settings are stored in the SASUSER library, which is discussed later in this section.

If SAS products other than Base SAS are installed at your site, the SASHELP library contains catalogs that are used by those products. In many instances, the defaults in this library are tailored to your site by your on-site SAS support personnel. You can list the catalogs stored at your site by using one of the file management utilities discussed later in this section.

SASUSER Library

The SASUSER library contains SAS catalogs that enable you to tailor features of SAS 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. For example, in Base SAS, you can store your own defaults for function key settings or window attributes in a personal profile catalog named SASUSER.PROFILE.

SAS assigns the SASUSER library during system initialization, according to the information supplied by the SASUSER system option.

A system option called RSASUSER enables the system administrator to control the mode of access to the SASUSER library at installations that have one SASUSER library for all users and that want to prevent users from modifying it.

Operating Environment Information: In most operating environments, the SASUSER library is created if it does not already exist. However, the SASUSER library is implemented differently in various operating environments. See the SAS documentation for your operating environment for more information. △

Sequential Data Libraries

SAS provides a number of features and procedures for reading from and writing to files that are stored on sequential format devices, either disk or tape. Before you store SAS libraries in sequential format, you should consider the following:

- You cannot use random access methods with sequential SAS data sets.
- You can access only one of the SAS files in a sequential library, or only one of the SAS files on a tape, at any point in a SAS job.

For example, you cannot read two or more SAS data sets in the same library or on the same tape at the same time in a single DATA step. However, you can access

- two or more SAS files in different sequential libraries, or on different tapes at the same time, if there are enough tape drives available

- a SAS file during one DATA or PROC step, then access another SAS file in the same sequential library or on the same tape during a later DATA or PROC step.

Also, when you have more than one SAS data set on a tape or in a sequential library in the same DATA or PROC step, one SAS data set file might be opened during the compilation phase, and the additional SAS data sets are opened during the execution phase. For more information, see the SET statement OPEN= option in *SAS Language Reference: Dictionary*.

- For some operating environments, you can read from or write to SAS data sets only during a DATA or PROC step. However, you can always use the COPY procedure to transfer all members of a SAS library to tape for storage and backup purposes.
- Considerations specific to your site can affect your use of tape. For example, it might be necessary to manually mount a tape before the SAS libraries become available. Consult your operations staff if you are not familiar with using tape storage at your location.

For information on sequential engines, see Chapter 36, “SAS Engines,” on page 589.

Operating Environment Information: The details for storing and accessing SAS files in sequential format vary with the operating environment. See the SAS documentation for your operating environment for more information. Δ

Tools for Managing Libraries

SAS Utilities

The SAS utilities that are available for SAS file management enable you to work with more than one SAS file at a time, as long as the files belong to the same library. The advantage of learning and using SAS Explorer, functions, options, and procedures is that they automatically copy, rename, or delete any index files or integrity constraints, audit trails, backups, and generation data sets that are associated with your SAS data files. Another advantage is that SAS utility procedures work on any operating environment at any level.

There are several SAS window options, functions, and procedures available for performing file management tasks. You can use the following features alone or in combination, depending on what works best for you. See “Choosing the Right Procedure” in *Base SAS Procedures Guide* for detailed information on SAS utility procedures. The SAS windowing environment and how to use it for managing SAS files is discussed in Chapter 17, “Introduction to the SAS Windowing Environment,” on page 247 and Chapter 18, “Managing Your Data in the SAS Windowing Environment,” on page 263 as well as in the online Help.

CATALOG procedure

provides catalog management utilities with the COPY, CONTENTS, and APPEND procedures.

DATASETS procedure

provides all library management functions for all member types except catalogs. If your site does not use the SAS Explorer, or if SAS executes in batch or interactive line mode, using this procedure can save you time and resources.

SAS Explorer

includes windows that enable you to perform most file management tasks without submitting SAS program statements. Type LIBNAME, CATALOG, or DIR in the Toolbar window to use SAS Explorer, or select the Explorer icon from the Toolbar menu.

DETAILS system option

Sets the default display for file information when using the CONTENTS, or DATASETS procedure. When enabled, DETAILS provides additional information about files, depending on which procedure or window you use.

Library Directories

SAS Explorer and SAS procedures enable you to obtain a list, or *directory*, of the members in a SAS library. Each directory contains the name of each member and its member type. For the member type DATA, the directory indicates whether an index, audit trail, backup, or generation data set is associated with the data set. The directory also describes some attributes of the library, but the amount and nature of this information vary with the operating environment.

Note: SAS libraries can also contain various SAS utility files. These files are not listed in the library directory and are for internal processing. △

Accessing Permanent SAS Files without a Libref

SAS provides another method of accessing files in addition to assigning a libref with the LIBNAME statement or using the New Library window. To use this method, enclose the filename, or the filename and path, in single quotation marks.

For example, in a directory based system, if you want to create a data set named MYDATA in your default directory, that is, in the directory that you are running SAS in, you can write the following line of code:

```
data 'mydata';
```

SAS creates the data set and remembers its location for the duration of the SAS session.

If you omit the single quotation marks, SAS creates the data set MYDATA in the temporary WORK subdirectory, named WORK.MYDATA:

```
data mydata;
```

If you want to create a data set named MYDATA in a library other than the directory in which you are running SAS, enclose the entire path in quotation marks, following the naming conventions of your operating environment. For example, the following DATA step creates a data set named FOO in the directory C:\sasrun\mydata.

```
data 'c:\sasrun\mydata\foo';
```

This method of accessing files works on all operating environments and in most contexts where a *libref.data-set-name* is accepted as a SAS data set. Most data set options can be specified with a quoted name.

You cannot use quoted names for

- SAS catalogs
- MDDb and FDB references
- contexts that do not accept a libref, such as the SELECT statement of PROC COPY and most PROC DATASETS statements

- PROC SQL
- DATA step stored programs, or SAS views
- SAS Component Language (SCL) open function


Operating Environment	Examples
DOS, Windows	<code>data 'c:\root\mystuff\sasstuff\work\myfile';</code>
UNIX	<code>data '/u/root/mystuff/sastuff/work/myfile';</code>
z/OS	<code>data 'user489.mystuff.saslib(member1)';</code> <code>/* bound SAS library */</code> <code>data '/mystuff/sasstuff/work/myfile';</code> <code>/* UNIX file system library */</code>
OpenVMS	<code>data 'filename filetype filemode';</code>

Operating Environment Commands

You can use operating environment commands to copy, rename, and delete the operating environment file or files that make up a SAS library. However, to maintain the integrity of your files, you must know how the SAS library model is implemented in your operating environment. For example, in some operating environments, SAS data sets and their associated indexes can be copied, deleted, or renamed as separate files. If you rename the file containing the SAS data set, but not its index, the data set will be marked as damaged.

CAUTION:

Using operating environment commands can damage files. You can avoid problems by always using SAS utilities to manage SAS files. Δ



CHAPTER 26

SAS Data Sets

<i>Definition of a SAS Data Set</i>	461
<i>Descriptor Information for a SAS Data Set</i>	462
<i>Data Set Names</i>	463
<i>Where to Use Data Set Names</i>	463
<i>How and When SAS Data Set Names Are Assigned</i>	463
<i>Parts of a Data Set Name</i>	463
<i>Two-level SAS Data Set Names</i>	464
<i>One-level SAS Data Set Names</i>	464
<i>Data Set Lists</i>	464
<i>Special SAS Data Sets</i>	465
<i>Null Data Sets</i>	465
<i>Default Data Sets</i>	466
<i>Automatic Naming Convention</i>	466
<i>Sorted Data Sets</i>	466
<i>The Sort Indicator</i>	466
<i>Example 1: Using No Sorting</i>	467
<i>Example 2: Using the SORTEDBY= Data Set Option</i>	468
<i>Example 3: Using the SORT Procedure</i>	469
<i>How SAS Uses the Sort Indicator to Improve Performance</i>	471
<i>Validating That a Data Set Is Sorted</i>	472
<i>Tools for Managing Data Sets</i>	472
<i>Viewing and Editing SAS Data Sets</i>	473

Definition of a SAS Data Set

A *SAS data set* is a SAS file stored in a SAS library that SAS creates and processes. A SAS data set contains data values that are organized as a table of observations (rows) and variables (columns) that can be processed by SAS software. A SAS data set also contains descriptor information such as the data types and lengths of the variables, as well as which engine was used to create the data.

A SAS data set can be one of the following:

SAS data file	contains both the data and the descriptor information. SAS data files have a member type of DATA. For specific information, see Chapter 27, “SAS Data Files,” on page 475.
SAS view	is a virtual data set that points to data from other sources. SAS views have a member type of VIEW. For specific information, see Chapter 28, “SAS Views,” on page 529.

Note: The term SAS data set is used when a SAS view and a SAS data file can be used in the same manner. △


Descriptor Information for a SAS Data Set

The descriptor information for a SAS data set makes the file self-documenting; that is, each data set can supply the attributes of the data set and of its variables. Once the data is in the form of a SAS data set, you do not have to specify the attributes of the data set or the variables in your program statements. SAS obtains the information directly from the data set.

Descriptor information includes the number of observations, the observation length, the date that the data set was last modified, and other facts. Descriptor information for individual variables includes attributes such as name, type, length, format, label, and whether the variable is indexed.

The following figure illustrates the logical components of a SAS data set:

Figure 26.1 Logical Components of a SAS Data Set



The following items correspond to the numbers in the figure above:

- 1 A SAS view (member type VIEW) contains descriptor information and uses data values from one or more data sets.
- 2 A SAS data file (member type DATA) contains descriptor information and data values. SAS data sets can be a member type DATA (SAS data file) or VIEW (SAS view).
- 3 An *index* is a separate file that you can create for a SAS data file in order to provide direct access to specific observations. The index file has the same name as its data file and a member type of INDEX. Indexes can provide faster access to specific observations, particularly when you have a large data set.

Data Set Names

Where to Use Data Set Names

You can use SAS data sets as input for DATA or PROC steps by specifying the name of the data set in the following:

- a SET statement
- a MERGE statement
- an UPDATE statement
- a MODIFY statement
- the DATA= option of a SAS procedure
- the OPEN function

How and When SAS Data Set Names Are Assigned

You name SAS data sets when you create them. Output data sets that you create in a DATA step are named in the DATA statement. SAS data sets that you create in a procedure step are usually given a name in the procedure statement or an OUTPUT statement. If you do not specify a name for an output data set, SAS assigns a default name.

If you are creating SAS views, you assign the data set name using one of the following:

- the SQL procedure
- the ACCESS procedure
- the VIEW= option in the DATA statement

Note: Because you can specify both SAS data files and SAS views in the same program statements but cannot specify the member type, SAS cannot determine from the program statement which one you want to process. This is why SAS prevents you from giving the same name to SAS views and SAS data sets in the same library. △

Parts of a Data Set Name

The complete name of every SAS data set has three elements. You assign the first two; SAS supplies the third. The form for SAS data set names is as follows:

libref.SAS-data-set.member-type

The elements of a SAS data set name include the following:

libref

is the logical name that is associated with the physical location of the SAS library.

SAS-data-set

is the data set name, which can be up to 32 bytes long for the Base SAS engine starting in Version 7. Earlier SAS versions are limited to 8-byte names.

member-type

is assigned by SAS. The member type is DATA for SAS data files and VIEW for SAS views.

When you refer to SAS data sets in your program statements, use a one- or two-level name. You can use a one-level name when the data set is in the temporary library WORK. In addition, if the reserved libref USER is assigned, you can use a one-level name when the data set is in the permanent library USER. Use a two-level name when the data set is in some other permanent library you have established. A two-level name consists of both the libref and the data set name. A one-level name consists of just the data set name.

Two-level SAS Data Set Names

The form most commonly used to create, read, or write to SAS data sets in permanent SAS libraries is the two-level name as shown here:

```
libref.SAS-data-set
```

When you create a new SAS data set, the libref indicates where it is to be stored. When you reference an existing data set, the libref tells SAS where to find it. The following examples show the use of two-level names in SAS statements:

```
data revenue.sales;

proc sort data=revenue.sales;
```

One-level SAS Data Set Names

You can omit the libref, and refer to data sets with a one-level name in the following form:

```
SAS-data-set
```

Data sets with one-level names are automatically assigned to one of two SAS libraries: WORK or USER. Most commonly, they are assigned to the temporary library WORK and are deleted at the end of a SAS job or session. If you have associated the libref USER with a SAS library or used the USER= system option to set the USER library, data sets with one-level names are stored in that library. See Chapter 25, “SAS Libraries,” on page 447 for more information on using the USER and WORK libraries. The following examples show how one-level names are used in SAS statements.

```
/* create perm data set in location of USER=option*/
options user='c:\temp'
data test3;

/* create perm data set in current directory */data 'test3';

/* create a temp data set in WORK directory if USER= is not specified*/
data stratifiedsample1;
```

Data Set Lists

In the DATASETS procedure and the DATA step MERGE and SET statements, data set lists provide a quick way to reference existing groups of data sets. These data set lists can be numbered range lists or colon (name prefix) lists.

- Numbered range lists require you to have a series of data sets with the same name, except for the last character or characters, which are consecutive numbers.

In a numbered range list, you can begin with any number and end with any number. For example, the following two lists refer to the same data sets:

```
sales1 sales2 sales3 sales4
```

```
sales1--sales4
```

Note: If the numeric suffix of the first data set name contains leading zeros, the number of digits in the numeric suffix of the first and last data set name must be greater than or equal to the number of digits in the first data set name; otherwise, an error will occur. For example, the data set lists *sales001–sales99* and *sales01–sales9* will cause an error to occur. The data set list *sales001–sales999* is valid. If the numeric suffix of the first data set name does not contain leading zeros, then the number of digits in the numeric suffix of the first and last data set names does not have to be equal. For example, the data set list *sales1–sales999* is valid. △

- Colon (name prefix) lists require you to have a series of data sets with the same starting character or characters. For example, the following two lists refer to the same data sets:

```
abc:
```

```
abc1 abc2 abcr abcx
```

In the DATASETS procedure, data set lists can be used with the following statements:

- COPY SELECT
- COPY EXCLUDE
- DELETE
- REPAIR
- REBUILD
- the variables that are specified in MODIFY SORTEDBY

For more information about how data set lists can be used in the DATA step, see the MERGE and SET statements in *SAS Language Reference: Dictionary*.

Special SAS Data Sets

Null Data Sets

If you want to execute a DATA step but do not want to create a SAS data set, you can specify the keyword `_NULL_` as the data set name. The following statement begins a DATA step that does not create a data set:

```
data _null_;
```

Using `_NULL_` causes SAS to execute the DATA step as if it were creating a new data set, but no observations or variables are written to an output data set. This process can be a more efficient use of computer resources if you are using the DATA step for some function, such as report writing, for which the output of the DATA step does not need to be stored as a SAS data set.

Default Data Sets

SAS keeps track of the most recently created SAS data set through the reserved name `_LAST_`. When you execute a DATA or PROC step without specifying an input data set, by default, SAS uses the `_LAST_` data set. Some functions use the `_LAST_` default as well.

The `_LAST_` system option enables you to designate a data set as the `_LAST_` data set. The name you specify is used as the default data set until you create a new data set. You can use the `_LAST_` system option when you want to use an existing permanent data set for a SAS job that contains a number of procedure steps. Issuing the `_LAST_` system option enables you to avoid specifying the SAS data set name in each procedure statement. The following OPTIONS statement specifies a default SAS data set:

```
options _last_=schedule.january;
```

Automatic Naming Convention

If you do not specify a SAS data set name or the reserved name `_NULL_` in a DATA statement, SAS automatically creates data sets with the names DATA1, DATA2, and so on, to successive data sets in the WORK or USER library. This feature is referred to as the `DATA n` naming convention. The following statement produces a SAS data set using the `DATA n` naming convention:

```
data;
```

Sorted Data Sets

The Sort Indicator

After a data set is sorted, a *sort indicator* is added to the data set descriptor information. The sort indicator is updated without a permanent sort of the data set by using the SORTEDBY= data set option. The **Sortedby** and **Validated** sort information is updated when the SORTEDBY= data set option is used.

The sort indicator contains some or all of the following sort information of a SAS data set:

- how the data set is sorted by which variable or variables
- whether the order for a variable is descending or ascending
- the character set used for character variables
- the collating sequence used for ordering character data
- collation rules if the data set is sorted linguistically
- whether there is only one observation for any given BY group (use of NODUPKEY option)
- whether there are no adjacent duplicate observations (use of NODUPREC option)
- whether the data set is validated

The sort indicator is set when a data set is sorted by a SORT procedure, an SQL procedure with an ORDER BY clause, a DATASETS procedure MODIFY statement, or a SORTEDBY= data set option. If the SORT or SQL procedures were used to sort the

data set, which is being sorted by SAS, the CONTENTS procedure output indicates that the **validated** sort information is YES. If the SORTEDBY= data set option was used to sort the data set, which is being sorted by the user, the CONTENTS procedure output indicates the **validated** sort information is set to NO and the **Sortedby** sort information is updated with the variable or variables specified by the data set option.

Data sets can be sorted outside of SAS. In that case, you might use the SORTEDBY= data set option or the DATASETS procedure MODIFY statement to add the sort order to the sort indicator. In this case, they are not validated. For more information, see “Validating That a Data Set Is Sorted” on page 472.

To view the sort indicator information, use the CONTENTS procedure or the CONTENTS statement in the DATASETS procedure. The following three examples show the sort indicator information in the CONTENTS procedure output.

Example 1: Using No Sorting

The first example is a data set created without any type of sort:

```
options yearcutoff=1920;

data sasuser.sortttest1;
  input priority 1. +1 indate date7.
        +1 office $ code $;
  format indate date7.;
  datalines;
1 03may07 CH J8U
1 21mar07 LA M91
1 01dec07 FW L6R
1 27feb06 FW Q2A
2 15jan07 FW I9U
2 09jul07 CH P3Q
3 08apr06 CH H5T
3 31jan06 FW D2W
;
proc contents data=sasuser.sortttest1;
run;
```

Note that the CONTENTS procedure output indicates there was no sort. SAS did not sort the data set, and the user did not specify that the data is sorted.

Output 26.1 Contents of SORTTEST1 Data Set —No Sorting

The SAS System		12:47 Wednesday, October 3, 2007		1	
The CONTENTS Procedure					
Data Set Name	SASUSER.SORTTEST1	Observations	8		
Member Type	DATA	Variables	4		
Engine	V9	Indexes	0		
Created	Wed, Oct 03, 2007 12:55:59 PM	Observation Length	32		
Last Modified	Wed, Oct 03, 2007 12:55:59 PM	Deleted Observations	0		
Protection		Compressed	NO		
Data Set Type		Sorted	NO		
Label					
Data Representation	WINDOWS_32				
Encoding	wlatin1 Western (Windows)				
Engine/Host Dependent Information					
Data Set Page Size	4096				
Number of Data Set Pages	1				
First Data Page	1				
Max Obs per Page	126				
Obs in First Data Page	8				
Number of Data Set Repairs	0				
Filename	C:\Documents and Settings\suholm\sasuser\sorttest1.sas7bdat				
Release Created	9.0201B0				
Host Created	XP_PRO				
Alphabetic List of Variables and Attributes					
	#	Variable	Type	Len	Format
	4	code	Char	8	
	2	indate	Num	8	DATE7.
	3	office	Char	8	
	1	priority	Num	8	

Example 2: Using the SORTEDBY= Data Set Option

In the second example, the data set is created using the SORTEDBY= data set option in the DATA statement.

```
options yearcutoff=1920;

data sasuser.sorttest1 (
    input priority 1. +1 indate date7.
        +1 office $ code $;
    format indate date7.;
    datalines;
1 03may01 CH J8U
1 21mar01 LA M91
1 01dec00 FW L6R
1 27feb99 FW Q2A
2 15jan00 FW I9U
2 09jul99 CH P3Q
3 08apr99 CH H5T
3 31jan99 FW D2W
;
proc contents data=sasuser.sorttest1;
run;
```

Note that the CONTENTS procedure output shows that the data set is sorted. Therefore, a **Sort Information** section containing sort indicator information is created. In the **Sort Information** section, the **Sortedby** information indicates the data set is sorted by the PRIORITY variable and is in descending order by the INDATE variable. The data set is sorted using the SORTEDBY= data set option, so the **Validated** information is NO. The **Character Set** information for the data set is ANSI.

Output 26.2 Contents of SORTTEST1 Data Set – Sorted

```

The SAS System      12:47 Wednesday, October 3, 2007   2

The CONTENTS Procedure

Data Set Name      SASUSER.SORTTEST1      Observations      8
Member Type       DATA                Variables         4
Engine            V9                    Indexes           0
Created           Wed, Oct 03, 2007 01:02:55 PM  Observation Length 32
Last Modified     Wed, Oct 03, 2007 01:02:55 PM  Deleted Observations 0
Protection        Compressed           NO
Data Set Type     Sorted                YES
Label
Data Representation WINDOWS_32
Encoding          wlatin1 Western (Windows)

Engine/Host Dependent Information

Data Set Page Size      4096
Number of Data Set Pages 1
First Data Page        1
Max Obs per Page       126
Obs in First Data Page 8
Number of Data Set Repairs 0
Filename                C:\Documents and Settings\suholm\sasuser\sorttest1.sas7bdat
Release Created         9.0201B0
Host Created            XP_PRO

Alphabetic List of Variables and Attributes

#   Variable   Type   Len   Format
4   code       Char   8
2   indate    Num    8   DATE7.
3   office    Char   8
1   priority  Num    8

Sort Information

Sortedby      priority DESCENDING indate
Validated     NO
Character Set ANSI

```

Example 3: Using the SORT Procedure

In the third example, the data set is sorted using the SORT procedure.

```

options yearcutoff=1920;

data sasuser.sorttest1;
  input priority 1. +1 indate date7.
        +1 office $ code $;

```

```
format indate date7.;
datalines;
1 03may01 CH J8U
1 21mar01 LA M91
1 01dec00 FW L6R
1 27feb99 FW Q2A
2 15jan00 FW I9U
2 09jul99 CH P3Q
3 08apr99 CH H5T
3 31jan99 FW D2W
;
proc sort data=sasuser.sortttest1; by priority descending indate;run;

proc contents data=sasuser.sortttest1;
run;
```

Note that the CONTENTS procedure output shows that the data set is sorted. Therefore, a **Sort Information** section containing sort indicator information is created. In the **Sort Information** section, the **Sortedby** information indicates the data set is sorted by the PRIORITY variable and is in descending order by the INDATE variable. The data set is sorted using the SORT procedure, so the **Validated** information is YES. The **Character Set** information for the data set is ANSI.

Output 26.3 Contents of SORTTEST1 Data Set – Validated Sort

```

The SAS System      12:47 Wednesday, October 3, 2007   4

The CONTENTS Procedure

Data Set Name      SASUSER.SORTTEST1      Observations      8
Member Type       DATA              Variables         4
Engine            V9                  Indexes          0
Created           Wed, Oct 03, 2007 01:08:35 PM  Observation Length 32
Last Modified     Wed, Oct 03, 2007 01:08:35 PM  Deleted Observations 0
Protection        Compressed          NO
Data Set Type     Sorted              YES
Label
Data Representation WINDOWS_32
Encoding          wlatin1 Western (Windows)

Engine/Host Dependent Information

Data Set Page Size      4096
Number of Data Set Pages 1
First Data Page        1
Max Obs per Page       126
Obs in First Data Page 8
Number of Data Set Repairs 0
Filename                C:\Documents and Settings\suholm\sasuser\sorttest1.sas7bdat
Release Created         9.0201B0
Host Created            XP_PRO

Alphabetic List of Variables and Attributes

#   Variable   Type   Len   Format
4   code       Char   8
2   indate    Num    8   DATE7.
3   office    Char   8
1   priority  Num    8

Sort Information

Sortedby  priority DESCENDING indate
Validated  YES
Character Set ANSI

```

How SAS Uses the Sort Indicator to Improve Performance

The sort information provided by the sort indicator is used internally for performance improvements. There are several ways to improve performance using the sort indicator:

- SAS uses the sort indicator to validate whether there was a previous sort. If there was a previous sort from a SORT procedure or an SQL procedure with an ORDER BY clause, then SAS does not perform another sort.
 - The SORT procedure sets the sort indicator when a sort occurs. The SORT procedure checks for the sort indicator before it sorts a data set so that data is not resorted unnecessarily. For more information, see the SORT Procedure in *Base SAS Procedures Guide*.
 - The SQL procedure uses the sort indicator to process queries more efficiently and to determine whether an internal sort is necessary before performing a join. For more information, see the SQL Procedure in *Base SAS Procedures Guide*.

- When using the sort indicator during index creation, SAS determines whether the data is already sorted by the key variable or variables in ascending order by checking the sort indicator in the data file. If the values in the sort indicator are in ascending order, SAS does not sort the values for the index file. For more information, see “Overview of Creating Indexes” on page 513.
- When processing a WHERE expression without an index, SAS first checks the sort indicator. If the **validated** sort information is YES, SAS stops reading the file once there are no more values that satisfy the WHERE expression.
- If an index is selected for WHERE expression processing, the sort indicator for that data set is changed to reflect the order that is specified by the index.
- For BY-group processing, if the data set is already sorted by the BY variable, SAS does not use the index, even if the data set is indexed on the BY variable.
- If the **validated** sort information is set to YES, SAS does not need to perform another sort.

Validating That a Data Set Is Sorted

Any SAS procedure that requires data to be sorted as part of the process checks the sort indicator information. The sort indicator is set when a data set is sorted by a SORT procedure, an SQL procedure with an ORDER BY clause, a DATASETS procedure MODIFY statement, or a SORTEDBY= data set option. If the SORT or SQL procedures were used to sort the data set, the CONTENTS procedure output indicates the **validated** sort information is YES. If the SORTEDBY= data set option was used to sort the data set, the CONTENTS procedure output indicates the **validated** sort information is NO. For examples of the CONTENTS procedure output, see “Example 1: Using No Sorting” on page 467, “Example 2: Using the SORTEDBY= Data Set Option” on page 468, and “Example 3: Using the SORT Procedure” on page 469.

You can use the SORTVALIDATE system option to specify whether the SORT procedure validates that a data set is sorted correctly when the data set sort indicator shows that a user has specified that the data set is sorted. The user can specify a sort order by using the SORTEDBY= data set option in a DATA statement or by using the SORTEDBY= option in the DATASETS procedure MODIFY statement. When the sort indicator is set by a user, SAS cannot be absolutely certain that a data set is sorted according to the variables in the BY statement.

If the SORTVALIDATE system option is set and the data set sort indicator was set by a user, the SORT procedure performs a sequence check on each observation to ensure that the data set is sorted according to the variables in the BY statement. If the data set is not sorted correctly, SAS sorts the data set.

At the end of a successful sequence check or at the end of a sort, the SORT procedure sets the sort indicator **validated** sort information to YES. If a sort is performed, the SORT procedure also updates the sort indicator **Sortedby** sort information to the variables in the BY statement. If an output data set is specified, the sort indicator **validated** sort information in the output data set is set to YES. If no sort is necessary, the data set is copied to the output data set. For more information about validated data sets, see the SORTVALIDATE= system option in *SAS Language Reference: Dictionary*.

Tools for Managing Data Sets

To copy, rename, delete, or obtain information about the contents of SAS data sets, use the same windows, procedures, functions, and options you do for SAS libraries. For a list of those windows and procedures, see Chapter 25, “SAS Libraries,” on page 447.

There are also functions available that enable you to work with your SAS data set. See each individual function for more complete information.

Viewing and Editing SAS Data Sets

The VIEWTABLE window enables you to browse, edit, or create data sets. This window provides two viewing modes:

Table View uses a tabular format to display multiple observations in the data set.

Form View displays data one observation at a time in a form layout.


You can customize your view of a data set, for example, by sorting your data, changing the color and fonts of columns, displaying variable labels instead of variable names, or removing or adding variables. You can also load an existing DATAFORM catalog entry in order to apply a previously defined variable, data set, and viewer attributes.

To view a data set, select the following: **Tools ► Table Editor**. This action brings up VIEWTABLE or FSVIEW (z/OS). You can also double-click on the data set in the Explorer window.

The following SAS files are supported within the VIEWTABLE window:

- SAS data files
- SAS views
- MDDB files

For more information, see the SAS System Help for VIEWTABLE in Base SAS.



CHAPTER

27

SAS Data Files

<i>Definition of a SAS Data File</i>	477
<i>Differences between Data Files and SAS Views</i>	477
<i>Understanding an Audit Trail</i>	479
<i>Definition of an Audit Trail</i>	479
<i>Audit Trail Description</i>	479
<i>Operation in a Shared Environment</i>	481
<i>Performance Implications</i>	481
<i>Preservation by Other Operations</i>	481
<i>Programming Considerations</i>	481
<i>Other Considerations</i>	481
<i>Initiating an Audit Trail</i>	482
<i>Controlling the Audit Trail</i>	482
<i>Reading and Determining the Status of the Audit Trail</i>	482
<i>Audit Trails and CEDA Processing</i>	483
<i>Examples of Using Audit Trails</i>	484
<i>Example of Initiating an Audit Trail</i>	484
<i>Example of a Data File Update</i>	485
<i>Example of Using the Audit Trail to Capture Rejected Observations</i>	485
<i>Understanding Generation Data Sets</i>	487
<i>Definition of Generation Data Set</i>	487
<i>Terminology for Generation Data Sets</i>	487
<i>Invoking Generation Data Sets</i>	488
<i>Understanding How a Generation Group Is Maintained</i>	488
<i>Processing Specific Versions of a Generation Group</i>	490
<i>Managing Generation Groups</i>	491
<i>Introduction</i>	491
<i>Displaying Data Set Information</i>	491
<i>Copying Generation Groups</i>	491
<i>Appending Generation Groups</i>	491
<i>Modifying the Number of Versions</i>	491
<i>Deleting Versions in a Generation Group</i>	492
<i>Renaming Versions in a Generation Group</i>	493
<i>Using Passwords in a Generation Group</i>	493
<i>Understanding Integrity Constraints</i>	493
<i>Definition of Integrity Constraints</i>	493
<i>General Integrity Constraints</i>	493
<i>Referential Integrity Constraints</i>	494
<i>Overlapping Primary Key and Foreign Key Constraints</i>	495
<i>Preservation of Integrity Constraints</i>	495
<i>Indexes and Integrity Constraints</i>	496
<i>Locking Integrity Constraints</i>	497

<i>Specifying Integrity Constraints</i>	497
<i>Specifying Physical Location for Inter-Libref Referential Integrity Constraints When Sharing Disk Space</i>	497
<i>Listing Integrity Constraints</i>	498
<i>Rejected Observations</i>	498
<i>Integrity Constraints and CEDA Processing</i>	498
<i>Examples</i>	499
<i>Creating Integrity Constraints with the DATASETS Procedure</i>	499
<i>Creating Integrity Constraints with the SQL Procedure</i>	499
<i>Creating Integrity Constraints by Using SCL</i>	500
<i>Removing Integrity Constraints</i>	503
<i>Reactivating an Inactive Integrity Constraint</i>	504
<i>Defining Overlapping Primary Key and Foreign Key Constraints</i>	504
<i>Understanding SAS Indexes</i>	505
<i>Definition of SAS Indexes</i>	505
<i>Benefits of an Index</i>	506
<i>The Index File</i>	506
<i>Types of Indexes</i>	507
<i>Simple Index</i>	507
<i>Composite Index</i>	508
<i>Unique Values</i>	508
<i>Missing Values</i>	509
<i>Deciding Whether to Create an Index</i>	509
<i>Costs of an Index</i>	509
<i>CPU Cost</i>	509
<i>I/O Cost</i>	510
<i>Buffer Requirements</i>	510
<i>Disk Space Requirements</i>	511
<i>Guidelines for Creating Indexes</i>	511
<i>Data File Considerations</i>	511
<i>Index Use Considerations</i>	512
<i>Key Variable Candidates</i>	512
<i>Creating an Index</i>	513
<i>Overview of Creating Indexes</i>	513
<i>Using the DATASETS Procedure</i>	514
<i>Using the INDEX= Data Set Option</i>	514
<i>Using the SQL Procedure</i>	514
<i>Using Other SAS Products</i>	514
<i>Using an Index for WHERE Processing</i>	515
<i>Identifying Available Index or Indexes</i>	515
<i>Compound Optimization</i>	516
<i>Estimating the Number of Qualified Observations</i>	518
<i>Comparing Resource Usage</i>	519
<i>Controlling WHERE Processing Index Usage with Data Set Options</i>	519
<i>Displaying Index Usage Information in the SAS Log</i>	520
<i>Using an Index with SAS Views</i>	520
<i>Using an Index for BY Processing</i>	521
<i>Using an Index for Both WHERE and BY Processing</i>	522
<i>Specifying an Index with the KEY= Option for SET and MODIFY Statements</i>	522
<i>Taking Advantage of an Index</i>	523
<i>Maintaining Indexes</i>	523
<i>Displaying Data File Information</i>	523
<i>Copying an Indexed Data File</i>	524
<i>Updating an Indexed Data File</i>	525

<i>Sorting an Indexed Data File</i>	525
<i>Adding Observations to an Indexed Data File</i>	525
<i>Multiple Occurrences of Values</i>	525
<i>Appending Data to an Indexed Data File</i>	526
<i>Recovering a Damaged Index</i>	526
<i>Indexes and Integrity Constraints</i>	526
<i>Indexes and CEDA Processing</i>	526
<i>Compressing Data Files</i>	526
<i>Definition of Compression</i>	526
<i>Requesting Compression</i>	527
<i>Disabling a Compression Request</i>	527

Definition of a SAS Data File

A *SAS data file* is a type of SAS data set that contains both the data values and the descriptor information. SAS data files have the member type of DATA. There are two general types of SAS data files:

native SAS data file

stores the data values and descriptor information in a file that is formatted by SAS.

interface SAS data file

stores the data in a file that was formatted by software other than SAS. SAS provides engines for reading and writing data from files that were formatted by software such as ORACLE, DB2, Sybase, ODBC, BMDP, SPSS, and OSIRIS. These files are interface SAS data files, and when their data values are accessed through an engine, SAS recognizes them as SAS data sets.

Note: The availability of engines that can access different types of interface data files is determined by your site licensing agreement. See your system administrator to determine which engines are available. For more information about SAS Multi Engine Architecture, see Chapter 36, “SAS Engines,” on page 589. Δ

Differences between Data Files and SAS Views

While the terms “SAS data files” and “SAS views” can often be used interchangeably, here are a few differences to consider:

- *The main difference is where the values are stored.* A SAS data file is a type of SAS data set that contains both descriptor information about the data and the data values themselves. SAS views contain only descriptor information and instructions for retrieving data that is stored elsewhere. Once the data is retrieved by SAS, it can be manipulated in a DATA step.
- *A data file is static; a SAS view is dynamic.* When you reference a data file in a later PROC step, you see the data values as they were when the data file was created or last updated. When you reference a SAS view in a PROC step, the view executes and provides you with an image of the data values as they currently exist, not as they existed when the view was defined.
- *SAS data files can be created on tape or on any other storage medium.*


SAS views cannot be stored on tape. Because of their dynamic nature, SAS views must derive their information from data files on random-access storage

devices, such as disk drives. SAS views cannot derive their information from files stored on sequentially accessed storage devices, such as tape drives.

- *SAS views are read only.* You cannot write to a SAS view, but some SQL views can be updated.
- *SAS data files can have an audit trail.* The audit trail is an optional SAS file that logs modifications to a SAS data file. Each time an observation is added, deleted, or updated, information is written to the audit trail about who made the modification, what was modified, and when.
- *SAS data files can have generations.* Generations provide the ability to keep multiple copies of a SAS data file. The multiple copies represent versions of the same data file, which is archived each time it is replaced.
- *SAS data files can have integrity constraints.* When you update a SAS data file, you can ensure that the data conforms to certain standards by using integrity constraints. With SAS views, you can assign integrity constraints to the data files that the views reference.
- *SAS data files can be indexed.* Indexing might enable SAS to find data in a SAS data file more quickly. SAS views cannot be indexed.
- *SAS data files can be encrypted.* Encryption provides an extra layer of security to physical files. SAS views cannot be encrypted.
- *SAS data files can be compressed.* Compression makes it possible to store physical files in less space. SAS views cannot be compressed.

The following figure illustrates native and interface SAS data files and their relationship to SAS views.

Figure 27.1 Types of SAS Data Sets



Understanding an Audit Trail

Definition of an Audit Trail

The *audit trail* is an optional SAS file that you can create in order to log modifications to a SAS data file. Each time an observation is added, deleted, or updated, information is written to the audit trail about who made the modification, what was modified, and when.

Many businesses and organizations require an audit trail for security reasons. The audit trail maintains historical information about the data, which gives you the opportunity to develop usage statistics and patterns. The historical information enables you to track individual pieces of data from the moment they enter the data file to the time they leave.

The audit trail is also the only facility in SAS that stores observations from failed append operations and that were rejected by integrity constraints. (The integrity constraints feature is described in “Understanding Integrity Constraints” on page 493.) The audit trail enables you to write a DATA step to extract the failed or rejected observations, use information describing why the observations failed to correct them, and then reapply the observations to the data file.

Audit Trail Description

The audit trail is created by the default Base SAS engine and has the same libref and member name as the data file, but has a data set type of AUDIT. It replicates the variables in the data file and additionally stores two types of audit variables:

- `_AT*` variables, which automatically store modification data
- user variables, which are optional variables you can define to collect modification data

The `_AT*` variables are described in the following table.

Table 27.1 `_AT*` Variables

<code>_AT*_ Variable</code>	Description
<code>_ATDATETIME_</code>	Stores the date and time of a modification
<code>_ATUSERID_</code>	Stores the logon user ID that is associated with a modification
<code>_ATOBSNO_</code>	Stores the observation number that is affected by the modification, except when <code>REUSE=YES</code> (because the observation number is always 0)
<code>_ATRETURNCODE_</code>	Stores the event return code
<code>_ATMESSAGE_</code>	Stores the SAS log message at the time of the modification
<code>_ATOPCODE_</code>	Stores a code that describes the type of modification

The `_ATOPCODE_` values are listed in the following table.

Table 27.2 _ATOPCODE_ Values

Code	Modification
AL	Auditing is resumed
AS	Auditing is suspended
DA	Added data record image
DD	Deleted data record image
DR	Before-update record image
DW	After-update record image
EA	Observation add failed
ED	Observation delete failed
EU	Observation update failed

The type of entries stored in the audit trail, along with their corresponding `_ATOPCODE_` values, are determined by the options specified in the LOG statement when the audit trail is initiated. Note that if the LOG statement is omitted when the audit trail is initiated, the default behavior is to log all images.

- The A operation codes are controlled by the `ADMIN_IMAGE` option.
- The DR operation code is controlled by the `BEFORE_IMAGE` option.
- All other D operation codes are controlled with the `DATA_IMAGE` option.
- The E operation codes are controlled by the `ERROR_IMAGE` option.

The user variable is a variable that associates data values with the data file without making them part of the data file. That is, the data values are stored in the audit file, but you update them in the data file like any other variable. You might want to define a user variable to enable end users to enter a reason for each update.

User variables are defined at audit trail initiation with the `USER_VAR` statement. For example, the following code initiates an audit trail and creates a user variable `REASON_CODE` for data file `MYLIB.SALES`:

```
proc datasets lib=mylib;
  audit sales;
  initiate;
  user_var reason_code $ 20;
run;
```

After the audit trail is initiated, SAS retrieves the user variables from the audit trail and displays them when the data file is opened for update. You can enter data values for the user variables as you would for any data variable. The data values are saved to the audit trail as each observation is saved. (In applications that save observations as you scroll through them, it might appear that the data values have disappeared.) The user variables are not available when the data file is opened for browsing or printing. However, to rename a user variable or modify its attributes, you modify the data file, not the audit file. The following example uses `PROC DATASETS` to rename the user variable:

```
proc datasets lib=mylib;
  modify sales;
  rename reason_code = Reason;
run;
quit;
```

You must also define attributes such as format and informat in the data file with PROC DATASETS. If you define user variables, you must store values in them for the variables to be meaningful.

A data file can have one audit file, and the audit file must reside in the same SAS library as the data file.

Operation in a Shared Environment

The audit trail operates similarly in local and remote environments. The only difference for applications and users networking with SAS/CONNECT and SAS/SHARE is that the audit trail logs events when the observation is written to permanent storage; that is, when the data is written to the remote SAS session or server. Therefore, the time that the transaction is logged might be different from the user's SAS session.

Performance Implications

Because each update to the data file is also written to the audit file, the audit trail can negatively impact system performance. You might want to consider suspending the audit trail for large, regularly scheduled batch updates. Note that the audit variables are unavailable when the audit trail is suspended.

Preservation by Other Operations

The audit trail is not recommended for data files that are copied, moved, sorted in place, replaced, or transferred to another operating environment, because those operations do not preserve the audit trail. In a copy operation on the same host, you can preserve the data file and audit trail by renaming them using the generation data sets feature; however, logging stops because neither the auditing process nor the generation data sets feature saves the source program that caused the replacement. For more information about generation data sets, see “Understanding Generation Data Sets” on page 487.

Programming Considerations

For data files whose audit file contains user variables, the variable list is different when browsing and updating the data file. The user variables are selected for update but not for browsing. You should be aware of this difference when you are developing your own full-screen applications.

Other Considerations

Data values that are entered for user variables are not stored in the audit trail for delete operations.

If the audit file becomes damaged, you cannot process the data file until you terminate the audit trail. Then you can initiate a new audit trail or process the data file without one. To terminate the audit trail for a generation data set, use the GENNUM= data set option in the AUDIT statement. You cannot initiate an audit trail for a generation data set.

In indexed data sets, the fast-append feature can cause some observations to be written to the audit trail twice, first with a DA operation code and then with an EA operation code. The observations with EA represent the observations rejected by index restrictions. For more information, see “Appending to an Indexed Data Set” in the PROC DATASETS APPEND statement documentation in *Base SAS Procedures Guide*.

Initiating an Audit Trail

You initiate an audit trail in the DATASETS procedure with the AUDIT statement. Consult the PROC DATASETS AUDIT statement in *Base SAS Procedures Guide* for syntax information.

The audit file uses the SAS password assigned to its associated data file; therefore, it is recommended that the data file have an ALTER password. An ALTER-level password restricts read and edit access to SAS files. If a password other than ALTER is used, or no password is used, the software generates a warning message that the files are not protected from accidental update or deletion.

Controlling the Audit Trail

Once active, you can suspend and resume logging, and terminate (delete) the audit trail. The syntax for controlling the audit trail is described in the PROC DATASETS AUDIT statement documentation. Note that replacing the associated data file also deletes the audit trail.

Reading and Determining the Status of the Audit Trail

The audit trail is read-only. You can read the audit trail with any component of SAS that reads a data set. To refer to the audit trail, use the TYPE= data set option. For example, to view the contents of the audit trail, issue the following statement. Note that the parentheses around the TYPE= option are required.

```
proc contents data=mylib.sales (type=audit);  
run;
```

The PROC CONTENTS output is shown below. Notice that the listing contains all of the variables from the corresponding data file, the _AT*_ variables, and the user variable.

Output 27.1 PROC CONTENTS Listing for Data File MYLIB.SALES

The CONTENTS Procedure				
Data Set Name	MYLIB.SALES.AUDIT	Observations	0	
Member Type	AUDIT	Variables	10	
Engine	V9	Indexes	0	
Created	12:34 Wednesday, January 24, 2007	Observation Length	111	
Last Modified	12:34 Wednesday, January 24, 2007	Deleted Observations	0	
Protection		Compressed	NO	
Data Set Type	AUDIT	Sorted	NO	
Label				
Data Representation	WINDOWS			
Encoding	Default			
Engine/Host Dependent Information				
Data Set Page Size	4096			
Number of Data Set Pages	1			
First Data Page	1			
Max Obs per Page	33			
Obs in First Data Page	0			
Number of Data Set Repairs	0			
File Name	C:\My Documents\myfiles\sales.sas7baud			
Release Created	9.2000A0			
Host Created	XP_PRO			
Alphabetic List of Variables and Attributes				
#	Variable	Type	Len	Format
5	_ATDATETIME_	Num	8	DATETIME19.
10	_ATMESSAGE_	Char	8	
6	_ATOBSNO_	Num	8	
9	_ATOPCODE_	Char	2	
7	_ATRETURNCODE_	Num	8	
8	_ATUSERID_	Char	32	
2	invoice	Num	8	
1	product	Char	9	
4	reason_code	Char	20	
3	renewal	Num	8	

You can also use your favorite reporting tool, such as PROC REPORT or PROC TABULATE, on the audit trail.

Audit Trails and CEDA Processing

When a SAS data file requires processing with CEDA, audit trails are not supported. For example, if you transfer a SAS data file with an initiated audit trail from one operating environment such as Windows to a different operating environment such as UNIX, CEDA translates the file for you, but the audit trail is not available. For information on CEDA processing, see Chapter 33, “Processing Data Using Cross-Environment Data Access (CEDA),” on page 567.

The MIGRATE procedure preserves audit trails when migrating data files. For more information, see the MIGRATE Procedure in *Base SAS Procedures Guide*. The CPORT and CIMPORT procedures preserve audit trails when transporting SAS data files from one operating environment to another operating environment. The CPORT procedure makes a copy of the data file in a transportable format, and the CIMPORT procedure reads the transport file and creates a new host-specific copy of the data file. For more

information, see the CPORT Procedure and CIMPORT Procedure in *Base SAS Procedures Guide*.

CAUTION:

If your data files contain audit trails, do not use your operating environment commands to copy, move, or delete your data files. Δ

Examples of Using Audit Trails

Example of Initiating an Audit Trail

The following example shows the data and code that are used to create and initiate an audit trail for the data file MYLIB.SALES that is used in earlier examples in this section. MYLIB.SALES stores fictional invoice and renewal figures for SAS products. The audit trail records all events and stores one user variable, REASON_CODE, for users to enter a reason for the update.

Subsequent examples illustrate the effect of a data file update on the audit trail and how to use audit variables to capture observations that are rejected by integrity constraints. The system option LINESIZE is set in advance for the integrity constraints example. A large LINESIZE value is recommended to display the content of the `_ATMESSAGE_` variable. The output examples have been modified to fit on the page.

```
options linesize=250;
/*-----*/
/* Create SALES data set.          */
/*-----*/

data mylib.sales;
  length product $9;
  input product invoice renewal;
datalines;
FSP          1270.00      570
SAS          1650.00      850
STAT         570.00       0
STAT         970.82      600
OR           239.36       0
SAS          7478.71     1100
SAS           800.00      800
;

/*-----*/
/* Create an audit trail with a     */
/* user variable.                  */
/*-----*/

proc datasets lib=mylib nolist;
  audit sales;
  initiate;
  user_var reason_code $ 20;
run;
```

Example of a Data File Update

The following example inserts an observation into MYLIB.SALES.DATA and prints the update data in the MYLIB.SALES.AUDIT.

```

/*-----*/
/* Do an update. */
/*-----*/
proc sql;
  insert into mylib.sales
    set product = 'AUDIT',
       invoice = 2000,
       renewal = 970,
       reason_code = "Add new product";
quit;

/*-----*/
/* Print the audit trail. */
/*-----*/
proc sql;
  select product,
         reason_code,
         _atopcode_,
         _atuserid_ format=$6.,
         _atdatetime_
  from mylib.sales(type=audit);
quit;

```

Output 27.2 Updated Data in MYLIB.SALES.AUDIT

The SAS System				
product	reason_code	_ATOPCODE_	_ATUSERID_	_ATDATETIME_
AUDIT	Add new product	DA	xxxxxx	31OCT2007:11:24:32

Example of Using the Audit Trail to Capture Rejected Observations

The following example adds integrity constraints to MYLIB.SALES.DATA and records observations that are rejected as a result of the integrity constraints in MYLIB.SALES.AUDIT. For more information about integrity constraints, see “Understanding Integrity Constraints” on page 493.

```

/*-----*/
/* Create integrity constraints. */
/*-----*/
proc datasets lib=mylib;
  modify sales;
  ic create null_renewal = not null (invoice)
    message = "Invoice must have a value.";
  ic create invoice_amt = check (where=((invoice > 0) and
    (renewal <= invoice)))
    message = "Invoice and/or renewal are invalid.";
run;

/*-----*/

```

```

/* Do some updates. */
/*-----*/
proc sql; /* this update works */
  update mylib.sales
    set invoice = invoice * .9,
    reason_code = "10% price cut"
  where renewal > 800;

proc sql; /* this update fails */
  insert into mylib.sales
    set product = 'AUDIT',
    renewal = 970,
    reason_code = "Add new product";

proc sql; /* this update works */
  insert into mylib.sales
    set product = 'AUDIT',
    invoice = 10000,
    renewal = 970,
    reason_code = "Add new product";

proc sql; /* this update fails */
  insert into mylib.sales
    set product = 'AUDIT',
    invoice = 100,
    renewal = 970,
    reason_code = "Add new product";
quit;

/*-----*/
/* Print the audit trail. */
/*-----*/
proc print data=mylib.sales(type=audit);
  format _atuserid_ $6.;
  var product reason_code _atopcode_ _atuserid_ _atdatetime_;
title 'Contents of the Audit Trail';
run;

/*-----*/
/* Print the rejected records. */
/*-----*/
proc print data=mylib.sales(type=audit);
  where _atopcode_ eq "EA";
  format _atmessage_ $250.;
  var product invoice renewal _atmessage_ ;
title 'Rejected Records';
run;

```

Output 27.3 shows the contents of MYLIB.SALES.AUDIT after several updates of MYLIB.SALES.DATA were attempted. Integrity constraints were added to the file, then updates were attempted. Output 27.4 prints information about the rejected observations on the audit trail.

Output 27.3 Contents of MYLIB.SALES.AUDIT After an Update with Integrity Constraints

Contents of the Audit Trail					
Obs	product	reason_code	_ATOPCODE_	_ATUSERID_	_ATDATETIME_
1	AUDIT	Add new product	DA	xxxxxxx	22JAN2007:12:43:27
2	SAS		DR	xxxxxxx	22JAN2007:13:02:02
3	SAS	10% price cut	DW	xxxxxxx	22JAN2007:13:02:02
4	SAS		DR	xxxxxxx	22JAN2007:13:02:02
5	SAS	10% price cut	DW	xxxxxxx	22JAN2007:13:02:02
6	AUDIT		DR	xxxxxxx	22JAN2007:13:02:02
7	AUDIT	10% price cut	DW	xxxxxxx	22JAN2007:13:02:02
8	AUDIT	Add new product	EA	xxxxxxx	22JAN2007:13:02:02
9	AUDIT	Add new product	DA	xxxxxxx	22JAN2007:13:02:02
10	AUDIT	Add new product	EA	xxxxxxx	22JAN2007:13:02:02

Output 27.4 Rejected Records on the Audit Trail

Rejected Records				
Obs	product	invoice	renewal	_ATMESSAGE_
1	AUDIT	.	970	ERROR: Invoice must have a value. Add/Update failed for data set MYLIB.SALES because data value(s) do not comply with integrity constraint null_renewal.
2	AUDIT	100	970	ERROR: Invoice and/or renewal are invalid. Add/update failed for data set MYLIB.SALES because data value(s)a do not comply with integrity constraint invoice_amt.

Understanding Generation Data Sets

Definition of Generation Data Set

A *generation data set* is an archived version of a SAS data set that is stored as part of a generation group. A generation data set is created each time the file is replaced. Each generation data set in a generation group has the same root member name, but each has a different version number. The most recent version of the generation data set is called the base version.

You can request generations for a SAS data file only. You cannot request generations for a SAS view.

Note: Generation data sets provide historical versions of a data set; they do not track observation updates for an individual data set. To log each time an observation is added, deleted, or updated, see “Understanding an Audit Trail” on page 479. △

Terminology for Generation Data Sets

The following terms are relevant to generation data sets:

base version

is the most recently created version of a data set. Its name does not have the four-character suffix for the generation number.

generation group

is a group of data sets that represent a series of replacements to the original data set. The generation group consists of the base version and a set of historical versions.

generation number

is a monotonically increasing number that identifies one of the historical versions in a generation group. For example, the data set named AIR#272 has a generation number of 272.

GENMAX=

is an output data set option that requests generations for a data set and specifies the maximum number of versions (including the base version and all historical versions) to keep for a given data set. The default is GENMAX=0, which means that the generation data sets feature is not in effect.

GENNUM=

is an input data set option that references a specific version from a generation group. Positive numbers are absolute references to a historical version by its generation number. Negative numbers are a relative reference to historical versions. For example, GENNUM=-1 refers to the youngest version.

historical versions

are the older copies of the base version of a data set. Names for historical versions have a four-character suffix for the generation number, such as #003.

oldest version

is the oldest version in a generation group.

rolling over

specifies the process of the version number moving from 999 to 000. When the generation number reaches 999, its next value is 000.

youngest version

is the version that is chronologically closest to the base version.

Invoking Generation Data Sets

To invoke generation data sets and to specify the maximum number of versions to maintain, include the output data set option GENMAX= when creating or replacing a data set. For example, the following DATA step creates a new data set and requests that up to four copies be kept (one base version and three historical versions):

```
data a (genmax=4);
  x=1;
  output;
run;
```

Once the GENMAX= data set option is in effect, the data set member name is limited to 28 characters (rather than 32), because the last four characters are reserved for a version number. When the GENMAX= data set option is not in effect, the member name can be up to 32 characters. See the GENMAX= data set option in *SAS Language Reference: Dictionary*.

Understanding How a Generation Group Is Maintained

The first time a data set with generations in effect is replaced, SAS keeps the replaced data set and appends a four-character version number to its member name,

which includes # and a three-digit number. That is, for a data set named A, the replaced data set becomes A#001. When the data set is replaced for the second time, the replaced data set becomes A#002; that is, A#002 is the version that is chronologically closest to the base version. After three replacements, the result is:

- A base (current) version
- A#003 most recent (youngest) historical version
- A#002 second most recent historical version
- A#001 oldest historical version.

With GENMAX=4, a fourth replacement deletes the oldest version, which is A#001. As replacements occur, SAS always keeps four copies. For example, after ten replacements, the result is:

- A base (current) version
- A#010 most recent (youngest) historical version
- A#009 2nd most recent historical version
- A#008 oldest historical version

The limit for version numbers that SAS can append is #999. After 999 replacements, the youngest version is #999. After 1,000 replacements, SAS rolls over the youngest version number to #000. After 1,001 replacements, the youngest version number is #001. For example, using data set A with GENNUM=4, the results would be:

- 999 replacements
 - A (current)
 - A#999 (most recent)
 - A#998 (2nd most recent)
 - A#997 (oldest)
- 1,000 replacements
 - A (current)
 - A#000 (most recent)
 - A#999 (2nd most recent)
 - A#998 (oldest)
- 1,001 replacements
 - A (current)
 - A#001 (most recent)
 - A#000 (2nd most recent)
 - A#999 (oldest)

The following table shows how names are assigned to a generation group:

Table 27.3 Naming Generation Group Data Sets

Time	SAS Code	Data Set Names	GENNUM= Absolute Reference	GENNUM= Relative Reference	Explanation
1	data air (genmax=3);	AIR	1	0	The AIR data set is created, and three generations are requested
2	data air;	AIR AIR#001	2 1	0 -1	AIR is replaced. AIR from time 1 is renamed AIR#001.
3	data air;	AIR AIR#002 AIR#001	3 2 1	0 -1 -2	AIR is replaced. AIR from time 2 is renamed AIR#002.

Time	SAS Code	Data Set Names	GENNUM= Absolute Reference	GENNUM= Relative Reference	Explanation
4	data air;	AIR	4	0	AIR is replaced. AIR from time 3 is renamed AIR#003. AIR#001 from time 1, which is the oldest, is deleted.
		AIR#003	3	-1	
		AIR#002	2	-2	
5	data air (genmax=2);	AIR	5	0	AIR is replaced, and the number of generations is changed to two. AIR from time 4 is renamed AIR#004. The two oldest versions are deleted.
		AIR#004	4	-1	

Processing Specific Versions of a Generation Group

When a generation group exists, SAS processes the base version by default. For example, the following PRINT procedure prints the base version:

```
proc print data=a;
run;
```

To request a specific version from a generation group, use the GENNUM= input data set option. There are two methods that you can use:

- A positive integer (excluding zero) references a specific historical version number. For example, the following statement prints the historical version #003:

```
proc print data=a(gennum=3);
run;
```

Note: After 1,000 replacements, if you want historical version #000, specify GENNUM=1000. Δ

- A negative integer is a relative reference to a version in relation to the base version, from the youngest predecessor to the oldest. For example, GENNUM=-1 refers to the youngest version. The following statement prints the data set that is three versions previous to the base version:

```
proc print data=a(gennum=-3);
run;
```

Table 27.4 Requesting Specific Generation Data Sets

SAS statement	Result
<code>proc print data=air (gennum=0);</code> <code>proc print data=air;</code>	Prints the current (base) version of the AIR data set.
<code>proc print data=air (gennum=-2);</code>	Prints the version two generations back from the current version.
<code>proc print data=air (gennum=3);</code>	Prints the file AIR#003.
<code>proc print data=air (gennum=1000);</code>	After 1,000 replacements, prints the file AIR#000, which is the file that is created after AIR#999.

Managing Generation Groups

Introduction

The DATASETS procedure provides a variety of statements for managing generation groups. Note that for the DATASETS procedure, GENNUM= has the following additional functionality:

- For the PROC DATASETS and DELETE statements, GENNUM= supports the additional values ALL, HIST, and REVERT.
- For the CHANGE statement, GENNUM= supports the additional value ALL.
- For the CHANGE statement, specifying GENNUM=0 refers to all versions rather than just the base version.

Displaying Data Set Information

A variety of statements in the DATASETS procedure can process a specific historical version. For example, you can display data set version numbers for historical copies using the CONTENTS statement in PROC DATASETS:

```
proc datasets library=myfiles;
    contents data=test (gennum=2);
run;
```

Copying Generation Groups

You can use the COPY statement in the DATASETS procedure or the COPY procedure to copy a generation group. However, you cannot copy an individual version.

For example, the following DATASETS procedure uses the COPY statement to copy a generation group for data set MYGEN1 from library MYLIB1 to library MYLIB2.

```
libname mylib1 'SAS-library-1';
libname mylib2 'SAS-library-2';

proc datasets;
    copy in=mylib1 out=mylib2;
    select mygen1;
run;
```

Appending Generation Groups

You can use the GENNUM= data set option to append a specific historical version. For example, the following DATASETS procedure uses the APPEND statement to append a historical version of data set B to data set A. Note that by default, SAS uses the base version for the BASE= data set.

```
proc datasets;
    append base=a data=b(gennum=2);
run;
```

Modifying the Number of Versions

When you modify the attributes of a data set, you can increase or decrease the number of versions for an existing generation group.

For example, the following MODIFY statement in the DATASETS procedure changes the number of generations for data set MYLIB.AIR to 4:

```
libname mylib 'SAS-library';

proc datasets library=mylib;
    modify air(genmax=4);
run;
```

CAUTION:

You cannot increase the number of versions after a generation group has rolled over. To increase the number of versions for a generation group that has rolled over, save the generation group, and then create a new generation data set and specify the desired maximum number of versions to maintain. Δ

CAUTION:

If you decrease the number of versions, SAS deletes the oldest version or versions so as not to exceed the new maximum. For example, the following MODIFY statement decreases the number of versions for MYLIB.AIR from 4 to 0, which causes SAS to automatically delete the three historical versions:

```
proc datasets library=mylib;
    modify air (genmax=0);
run;
```

Δ

Deleting Versions in a Generation Group

When you delete data sets, you can specify a specific version or an entire generation group to delete. The following table shows the types of delete operations and their effects when you delete versions of a generation group.

The following examples assume that the base version of AIR and two historical versions (AIR#001 and AIR#002) exist for each command.

SAS statement in PROC DATASETS	Results
<code>delete air;</code>	Deletes the base version and shifts up historical versions.
<code>delete air(gennum=0);</code>	AIR#002 is renamed to AIR and becomes the new base version.
<code>delete air(gennum=2);</code>	Deletes historical version AIR#002.
<code>delete air(gennum=-2);</code>	Deletes the second youngest historical version (AIR#001).
<code>delete air(gennum=all);</code>	Deletes all data sets in the generation group, including the base version.
<code>delete air(gennum=hist);</code>	Deletes all data sets in the generation group, except the base version.

Note: Both an absolute reference and a relative reference refer to a specific version. A relative reference does not skip deleted versions. Therefore, when you are working with a generation group that includes one or more deleted versions, using a relative

reference results in an error if the referenced version has been deleted. For example, if you have the base version AIR and three historical versions (AIR#001, AIR#002, and AIR#003) and you delete AIR#002, the following statements return an error, because AIR#002 does not exist. SAS does not assume that you mean AIR#003:

```
proc print data=air (gennum= -2);
run;
```

Δ

Renaming Versions in a Generation Group

When you rename a data set, you can rename an entire generation group:

```
proc datasets;
  change a=newa;
run;
```

You can also rename a single version by including GENNUM=:

```
proc datasets;
  change a(gennum=2)=newa;
```

Note: For the CHANGE statement in PROC DATASETS, specifying GENNUM=0 refers to the entire generation group. Δ

Using Passwords in a Generation Group

Passwords for versions in a generation group are maintained as follows:

- If you assign a password to the base version, the password is maintained in subsequent historical versions. However, the password is not applied to any existing historical versions.
- If you assign a password to a historical version, the password applies to that individual data set only.

Understanding Integrity Constraints

Definition of Integrity Constraints

Integrity constraints are a set of data validation rules that you can specify in order to restrict the data values that can be stored for a variable in a SAS data file. Integrity constraints help you preserve the validity and consistency of your data. SAS enforces the integrity constraints when the values associated with a variable are added, updated, or deleted.

There are two categories of integrity constraints: general and referential.

General Integrity Constraints

General integrity constraints enable you to restrict the values of variables within a single file. There are four types of general constraints:

check	limits the data values of variables to a specific set, range, or list of values. Check constraints can also be used to ensure that the data
-------	---

	values in one variable within an observation are contingent on the data values of another variable in the same observation.
not null	requires that a variable contain a data value. Null (missing) values are not allowed.
unique	requires that the specified variable or variables contain unique data values. A null data value is allowed but is limited to a single instance.
primary key	requires that the specified variable or variables contain unique data values and that null data values are not allowed. Only one primary key can exist in a data file.

Note: A *primary key* is a general integrity constraint if it does not have any foreign key constraints referencing it. Δ

Referential Integrity Constraints

A referential integrity constraint is created when a primary key integrity constraint in one data file is referenced by a foreign key integrity constraint in another data file.

The *foreign key* constraint links the data values of one or more variables in the foreign key data file, to corresponding variables and values in the primary key data file. Data values in the foreign key data file must have a matching value in the primary key data file, or they must be null. When data is updated or deleted in the primary key data file, the modifications are controlled by a referential action that is defined as part of the foreign key constraint.

Separate referential actions can be defined for the update and delete operations. There are three types of referential actions:

restrict	prevents the data values of the primary key variables from being updated or deleted if there is a matching value in one of the foreign key data file's corresponding foreign key variables. The restrict type of action is the default action if one is not specified.
set null	enables the data values of the primary key variables to be updated or deleted, but matching data values in the foreign key data files are changed to null (missing) values.
cascade	enables the data values in the primary key variables to be updated, and additionally updates matching data values in the foreign key data files to the same value. The cascade type of action is supported only for update operations.

The requirements for establishing a referential relationship are as follows:

- The primary key and foreign key must reference the same number of variables, and the variables must be in the same order.
- The variables must be of the same type (character or numeric) and length.
- If the foreign key is being added to a data file that already contains data, the data values in the foreign key data file must either match existing values in the primary key data file, or the values must be null.

The foreign key data file can exist in the same SAS library as the referenced primary key data file (intra-libref), or in a different SAS library (inter-libref). However, if the library that contains the foreign key data file is temporary, the library that contains the primary key data file must be temporary as well. In addition, referential integrity constraints cannot be assigned to data files in concatenated libraries.

There is no limit to the number of foreign keys that can reference a primary key. However, additional foreign keys can adversely impact the performance of update and delete operations.

When a referential constraint exists, a primary key integrity constraint is not deleted until all of the foreign keys that reference it have been deleted. There are no restrictions on deleting foreign keys.

Overlapping Primary Key and Foreign Key Constraints

Variables in a SAS data file can be part of both a primary key (general integrity constraint) and a foreign key (referential integrity constraint). However, there are restrictions when you define a primary key and a foreign key constraint that use the same variables:

- The foreign key's update and delete referential actions must both be RESTRICT.
- When the same variables are used in a primary key and foreign key definition, the variables must be defined in a different order.

For an example, see “Defining Overlapping Primary Key and Foreign Key Constraints” on page 504.

Preservation of Integrity Constraints

These procedures preserve integrity constraints when their operation results in a copy of the original data file:

- in Base SAS software, the APPEND, COPY, CPORT, CIMPORT, MIGRATE, and SORT procedures
- in SAS/CONNECT software, the UPLOAD and DOWNLOAD procedures
- PROC APPEND
 - for an existing BASE= data file, integrity constraints in the BASE= file are preserved, but integrity constraints in the DATA= file that is being appended to the BASE= file are not preserved.
 - for a non-existent BASE= data file, general integrity constraints in the DATA= file that is being appended to the new BASE= file are preserved. Referential constraints in the DATA= file are not preserved.
- PROC SORT, and PROC UPLOAD, and PROC DOWNLOAD, when an OUT= data file is not specified
- the SAS Explorer window.

You can also use the CONSTRAINT option to control whether integrity constraints are preserved for the COPY, CPORT, CIMPORT, UPLOAD, and DOWNLOAD procedures.

General integrity constraints are preserved in an active state. The state in which referential constraints are preserved depends on whether the procedure causes the primary key and foreign key data files to be written to the same or different SAS libraries (intra-libref versus inter-libref integrity constraints). Intra-libref constraints are preserved in an active state. Inter-libref constraints are preserved in an inactive state; that is, the primary key portion of the integrity constraint is enforced as a general integrity constraint but the foreign key portion is inactive. You must use the DATASETS procedure statement IC REACTIVATE to reactivate the inactive foreign keys.

The following table summarizes the circumstances under which integrity constraints are preserved.

Table 27.5 Circumstances That Cause Integrity Constraints to Be Preserved

Procedure	Condition	Constraints That Are Preserved
APPEND	DATA= data set does not exist	General constraints Referential constraints are not affected
COPY	CONSTRAINT=yes	General constraints Intra-libref constraints are referential in an active state Inter-libref constraints are referential in an inactive state
CPORT/CIMPORT	CONSTRAINT=yes	General constraints Intra-libref constraints are referential in an active state Inter-libref constraints are referential in an inactive state
SORT	OUT= data set is not specified	General constraints Referential constraints are not affected
UPLOAD/DOWNLOAD	CONSTRAINT=yes and OUT= data set is not specified	General constraints Intra-libref constraints are referential in an active state Inter-libref constraints are referential in an inactive state
SAS Explorer window		General constraints

Indexes and Integrity Constraints

The unique, primary key, and foreign key integrity constraints store data values in an index file. If an index file already exists, it is used; otherwise, one is created. Consider the following points when you create or delete an integrity constraint:

- When a user-defined index exists, the index's attributes must be compatible with the integrity constraint in order for the integrity constraint to be created. For example, when you add a primary key integrity constraint, the existing index must have the UNIQUE attribute. When you add a foreign key integrity constraint, the index must *not* have the UNIQUE attribute.
- The unique integrity constraint has the same effect as the UNIQUE index attribute; therefore, when one is used, the other is not necessary.
- The NOMISS index attribute and the not-null integrity constraint have different effects. The integrity constraint prevents missing values from being written to the SAS data file and cannot be added to an existing data file that contains missing values. The index attribute allows missing data values in the data file but excludes them from the index.
- When any index is created, it is marked as being “owned” by the user, the integrity constraint, or both. A user cannot delete an index that is also owned by an integrity constraint and vice versa. If an index is owned by both, the index is deleted only after both the integrity constraint and the user have requested the index's deletion. A note in the log indicates when an index cannot be deleted.

Locking Integrity Constraints

Integrity constraints support both member-level and record-level locking. You can override the default locking level with the CNTLLEV= data set option. For more information, see the CNTLLEV= data set option in *SAS Language Reference: Dictionary*.

Specifying Integrity Constraints

You can create integrity constraints in the SQL procedure, the DATASETS procedure, or in SCL (SAS Component Language). The constraints can be specified when the data file is created or can be added to an existing data file. When you add integrity constraints to an existing file, SAS verifies that the existing data values conform to the constraints that are being added.

When you specify integrity constraints, you must specify a separate statement for each constraint. In addition, you must specify a separate statement for each variable to which you want to assign the not-null integrity constraint. When multiple variables are included in the specification for a primary key, foreign key, or a unique integrity constraint, a composite index is created and the integrity constraint enforces the combination of variable values. The relationship between SAS indexes and integrity constraints is described in “Indexes and Integrity Constraints” on page 496. For more information, see “Understanding SAS Indexes” on page 505.

When you add an integrity constraint in SCL, open the data set in utility mode. See “Creating Integrity Constraints by Using SCL” on page 500 for an example. Integrity constraints must be deleted in utility open mode. For detailed syntax information, see *SAS Component Language: Reference*.

When generation data sets are used, you must create the integrity constraints in each data set generation that includes protected variables.

CAUTION:

CHECK constraints in SAS 9.2 are not compatible with earlier releases of SAS. If you add a CHECK constraint to an existing SAS data set or create a SAS data set that includes a CHECK constraint, the data set cannot be accessed by a release before SAS 9.2. \triangle

Specifying Physical Location for Inter-Libref Referential Integrity Constraints When Sharing Disk Space

When you share disk space over a network and access referential integrity constraints in which the foreign key and primary key data files are in different SAS libraries, a standard should be established for the physical location of the shared files. A standard is required when you create the shared files so that network machines use the same physical name in order to access the files. If the physical names do not match, SAS cannot open the referenced foreign key or primary key SAS data file.

For example, a standard might be established that all shared files are placed on disk T: so that network machines use the same pathname in order to access the files.

Here is an example of a problem regarding files that were created without a standard. Suppose a primary key and a foreign key SAS data file were created on machine D4064 in different directories **C:\Public\pkey_directory** and **C:\Public\fkey_directory**. The pathnames are stored in the descriptor information of the SAS data files.

To access the primary key data file from a different machine such as F2760, the following LIBNAME statement would be executed:

```
libname pkds '\\D4064\Public\pkey_directory';
```

When the primary key data file is opened for update processing, SAS automatically tries to open the foreign key data file by using the foreign key data file's physical name that is stored in the primary key data file, which is `C:\Public\fkey_directory`. However, that directory does not exist on machine F2760, and so opening the foreign key data file fails.

Listing Integrity Constraints

PROC CONTENTS and PROC DATASETS report integrity constraint information without special options. In addition, you can print information about integrity constraints and indexes to a data set by using the OUT2= option. In PROC SQL, the DESCRIBE TABLE and DESCRIBE TABLE CONSTRAINTS statements report integrity constraint characteristics as part of the data file definition or alone, respectively. SCL provides the ICTYPE, ICVALUE, and ICDESCRIBE functions for getting information about integrity constraints. Refer to *Base SAS Procedures Guide* and *SAS Component Language: Reference* for more information.

Rejected Observations

You can customize the error message that is associated with an integrity constraint when you create the constraint by using the MESSAGE= and MSGTYPE= options. The MESSAGE= option enables you to prepend a user-defined message to the SAS error message associated with an integrity constraint. The MSGTYPE= option enables you to suppress the SAS portion of the message. For more information, see the PROC DATASETS, PROC SQL, and SCL documentation.

Rejected observations can be collected in a special file by using an audit trail.

Integrity Constraints and CEDA Processing

When a SAS data file requires processing with CEDA, integrity constraints are not supported. For example, if you transfer a SAS data file with defined integrity constraints from one operating environment such as Windows to a different operating environment such as UNIX, CEDA translates the file for you, but the integrity constraints are not available. For information about CEDA processing, see Chapter 33, "Processing Data Using Cross-Environment Data Access (CEDA)," on page 567.

The MIGRATE procedure preserves integrity constraints when migrating data files. For more information, see the MIGRATE Procedure in *Base SAS Procedures Guide*. The CPORT and CIMPORT procedures preserve integrity constraints when transporting SAS data files from one operating environment to another operating environment. The CPORT procedure makes a copy of the data file in a transportable format, and the CIMPORT procedure reads the transport file and creates a new host-specific copy of the data file. For more information, see the CPORT Procedure and CIMPORT Procedure in *Base SAS Procedures Guide*.

CAUTION:

If your data files contain integrity constraints, do not use your operating environment commands to copy, move, or delete your data files. △

Examples

Creating Integrity Constraints with the DATASETS Procedure

The following sample code creates integrity constraints by means of the DATASETS procedure. The data file TV_SURVEY checks the percentage of viewing time spent on networks, PBS, and other channels, with the following integrity constraints:

- The viewership percentage cannot exceed 100 percent.
- Only adults can participate in the survey.
- GENDER can be male or female.

```
data tv_survey(label='Validity checking');
  length idnum age 4 gender $1;
  input idnum gender age network pbs other;
datalines;
1 M 55 80 . 20
2 F 36 50 40 10
3 M 42 20 5 75
4 F 18 30 0 70
5 F 84 0 100 0
;

proc datasets nolist;
  modify tv_survey;
  ic create val_gender = check(where=(gender in ('M','F')))
    message = "Valid values for variable GENDER are
    either 'M' or 'F'.";
  ic create val_age = check(where=(age >= 18 and age <= 120))
    message = "An invalid AGE has been provided.";
  ic create val_new = check(where=(network <= 100));
  ic create val_pbs = check(where=(pbs <= 100));
  ic create val_ot = check(where=(other <= 100));
  ic create val_max = check(where=((network+pbs+other)<= 100));
quit;
```

Creating Integrity Constraints with the SQL Procedure

The following sample program creates integrity constraints by means of the SQL procedure. The data file PEOPLE lists employees and contains employment information. The data file SALARY contains salary and bonus information. The integrity constraints are as follows:

- The names of employees receiving bonuses must be found in the PEOPLE data file.
- The names identified in the primary key must be unique.
- GENDER can be male or female.
- Job status can be permanent, temporary, or terminated.

```
proc sql;
  create table people
```

```

(
  name      char(14),
  gender    char(6),
  hired     num,
  jobtype   char(1) not null,
  status    char(10),

  constraint prim_key primary key(name),
  constraint gender check(gender in ('male' 'female')),
  constraint status check(status in ('permanent'
                                     'temporary' 'terminated'))
);

create table salary
(
  name      char(14),
  salary    num not null,
  bonus     num,

  constraint for_key foreign key(name) references people
    on delete restrict on update set null
);
quit;

```

Creating Integrity Constraints by Using SCL

To add integrity constraints to a data file by using SCL, you must create and build an SCL catalog entry. The following sample program creates and compiles catalog entry EXAMPLE.IC_CAT.ALLICS.SCL.

```

INIT:
  put "Test SCL integrity constraint functions start.";
return;

MAIN:
  put "Opening WORK.ONE in utility mode.";
  dsid = open('work.one', 'V');/* Utility mode.*/
  if (dsid = 0) then
    do;
      _msg_=sysmsg();
      put _msg_=;
    end;
  else do;
    if (dsid > 0) then
      put "Successfully opened WORK.ONE in"
        "UTILITY mode.";
    end;

  put "Create a check integrity constraint named teen.";
  rc = iccreate(dsid, 'teen', 'check',
    '(age > 12) && (age < 20)');

  if (rc > 0) then
    do;
      put rc=;
    end;

```

```
    _msg_=sysmsg();
    put _msg_=;
end;
else do;
    put "Successfully created a check"
        "integrity constraint.";
end;

put "Create a not-null integrity constraint named nn.";
rc = icreate(dsid, 'nn', 'not-null', 'age');

if (rc > 0) then
    do;
        put rc=;
        _msg_=sysmsg();
        put _msg_=;
    end;
else do;
    put "Successfully created a not-null"
        "integrity constraint.";
end;

put "Create a unique integrity constraint named uq.";
rc = icreate(dsid, 'uq', 'unique', 'age');

if (rc > 0) then
    do;
        put rc=;
        _msg_=sysmsg();
        put _msg_=;
    end;
else do;
    put "Successfully created a unique"
        "integrity constraint.";
end;

put "Create a primary key integrity constraint named pk.";
rc = icreate(dsid, 'pk', 'Primary', 'name');

if (rc > 0) then
    do;
        put rc=;
        _msg_=sysmsg();
        put _msg_=;
    end;
else do;
    put "Successfully created a primary key"
        "integrity constraint.";
end;

put "Closing WORK.ONE.";
rc = close(dsid);
if (rc > 0) then
    do;
```

```

    put rc=;
    _msg_=sysmsg();
    put _msg_=;
end;

put "Opening WORK.TWO in utility mode.";
dsid2 = open('work.two', 'V');
/*Utility mode */
if (dsid2 = 0) then
do;
_msg_=sysmsg();
put _msg_=;
end;
else do;
if (dsid2 > 0) then
put "Successfully opened WORK.TWO in"
"UTILITY mode.";
end;

put "Create a foreign key integrity constraint named fk.";
rc = iccreate(dsid2, 'fk', 'foreign', 'name',
'work.one', 'null', 'restrict');

if (rc > 0) then
do;
put rc=;
_msg_=sysmsg();
put _msg_=;
end;
else do;
put "Successfully created a foreign key"
"integrity constraint.";
end;

put "Closing WORK.TWO.";
rc = close(dsid2);
if (rc > 0) then
do;
put rc=;
_msg_=sysmsg();
put _msg_=;
end;
return;

TERM:
put "End of test SCL integrity constraint"
"functions.";
return;

```

The previous code creates the SCL catalog entry. The following code creates two data files, ONE and TWO, and executes the SCL entry EXAMPLE.IC_CAT.ALLICS.SCL:

```

/* Submit to create data files. */

data one two;

```

```

    input name $ age;
datalines;
Morris 13
Elaine 14
Tina 15
;

    /* after compiling, run the SCL program */

proc display catalog= example.ic_cat.allics.scl;
run;

```

Removing Integrity Constraints

The following sample program segments remove integrity constraints. In the code that deletes a primary key integrity constraint, note that the foreign key integrity constraint is deleted first.

This program segment deletes integrity constraints using PROC SQL.

```

proc sql;
    alter table salary
        DROP CONSTRAINT for_key;
    alter table people
        DROP CONSTRAINT gender
        DROP CONSTRAINT _nm0001_
        DROP CONSTRAINT status
        DROP CONSTRAINT prim_key
    ;
quit;

```

This program segment removes integrity constraints using PROC DATASETS.

```

proc datasets nolist;
    modify tv_survey;
        ic delete val_max;
        ic delete val_gender;
        ic delete val_age;
run;
quit;

```

This program segment removes integrity constraints using SCL.

```

TERM:
    put "Opening WORK.TWO in utility mode.";
    dsid2 = open( 'work.two' , 'V' ); /* Utility mode.    */
    if (dsid2 = 0) then
        do;
            _msg_=sysmsg();
            put _msg_=;
        end;
    else do;
        if (dsid2 > 0) then
            put "Successfully opened WORK.TWO in Utility mode.";
        end;

rc = icdelete(dsid2, 'fk');
if (rc > 0) then

```

```

do;
  put rc=;
  _msg_=sysmsg();
end;
else
  do;
    put "Successfully deleted a foreign key integrity constraint.";
  end;
rc = close(dsid2);
return;

```

Reactivating an Inactive Integrity Constraint

The following program segment reactivates a foreign key integrity constraint that has been inactivated as a result of a COPY, CPORT, CIMPORT, UPLOAD, or DOWNLOAD procedure.

```

proc datasets;
  modify SAS-data-set;
    ic reactivate fkname references libref;
  run;
quit;

```

Defining Overlapping Primary Key and Foreign Key Constraints

The following code illustrates defining overlapping primary key and foreign key constraints:

```

data Singers1;
  input FirstName $ LastName $ Age;
  datalines;
Tom Jones 62
Kris Kristofferson 66
Willie Nelson 69
Barbra Streisand 60
Paul McCartney 60
Randy Travis 43
;
data Singers2;
  input FirstName $ LastName $ Style $;
  datalines;
Tom Jones Rock
Kris Kristofferson Country
Willie Nelson Country
Barbra Streisand Contemporary
Paul McCartney Rock
Randy Travis Country
;
proc datasets library=work nolist;
  modify Singers1;
    ic create primary key (FirstName LastName); ❶
  run;
  modify Singers2;
    ic create foreign key (FirstName LastName) references Singers1
      on delete restrict on update restrict; ❷
  run;

```

```

modify Singers2;
    ic create primary key (LastName FirstName); ❸
run;
modify Singers1;
    ic create foreign key (LastName FirstName) references Singers2
        on delete restrict on update restrict; ❹
run;

quit;

```

- ❶ Defines a primary key constraint for data set Singers1, for variables FIRSTNAME and LASTNAME.
- ❷ Defines a foreign key constraint for data set Singers2 for variables FIRSTNAME and LASTNAME that references the primary key defined in Step 1. Because the intention is to define a primary key using the same variables, the foreign key update and delete referential actions must both be RESTRICT.
- ❸ Defines a primary key constraint for data set Singers2 for variables LASTNAME and FIRSTNAME. Because those exact same variables are already defined as a foreign key, the order must be different.
- ❹ Defines a foreign key constraint for data set Singers1 for variables LASTNAME and FIRSTNAME that references the primary key defined in Step 3. Because those exact same variables are already defined as a primary key, the order must be different. Because a primary key is already defined using the same variables, the foreign key's update and delete referential actions must both be RESTRICT.

Understanding SAS Indexes

Definition of SAS Indexes

An *index* is an optional file that you can create for a SAS data file in order to provide direct access to specific observations. The index stores values in ascending value order for a specific variable or variables and includes information as to the location of those values within observations in the data file. In other words, an index enables you to locate an observation by value.

For example, suppose you want the observation with SSN (Social Security number) equal to 123-45-6789:

- Without an index, SAS accesses observations sequentially in the order in which they are stored in the data file. SAS reads each observation, looking for SSN=123-45-6789 until all observations are read.
- With an index on variable SSN, SAS accesses the observation directly. SAS satisfies the condition using the index and goes straight to the observation that contains the value without having to read each observation.

You can either create an index when you create a data file or create an index for an existing data file. The data file can be either compressed or uncompressed. For each data file, you can create one or multiple indexes. Once an index exists, SAS treats it as part of the data file. That is, if you add or delete observations or modify values, the index is automatically updated.

Benefits of an Index

In general, SAS can use an index to improve performance in the following situations:

- For WHERE processing, an index can provide faster and more efficient access to a subset of data. Note that to process a WHERE expression, SAS by default decides whether to use an index or to read the data file sequentially.
- For BY processing, an index returns observations in the index order, which is in ascending value order, without using the SORT procedure even when the data file is not stored in that order.

Note: If you use the SORT procedure, the index is not used. Δ

- For the SET and MODIFY statements, the KEY= option enables you to specify an index in a DATA step to retrieve particular observations in a data file.

In addition, an index can benefit other areas of SAS. In SCL (SAS Component Language), an index improves the performance of table lookup operations. For the SQL procedure, an index enables the software to process certain classes of queries more efficiently, for example, join queries. For the SAS/IML software, you can explicitly specify that an index be used for read, delete, list, or append operations.

Even though an index can reduce the time required to locate a set of observations, especially for a large data file, there are costs associated with creating, storing, and maintaining the index. When deciding whether to create an index, you must consider increased resource usage, along with the performance improvement.

Note: An index is never used for the subsetting IF statement in a DATA step, or for the FIND and SEARCH commands in the FSEDIT procedure. Δ

The Index File

The *index file* is a SAS file that has the same name as its associated data file, and that has a member type of INDEX. There is only one index file per data file; that is, all indexes for a data file are stored in a single file.

The index file might be a separate file, or be part of the data file, depending on the operating environment. In any case, the index file is stored in the same SAS library as its data file.

The index file consists of entries that are organized hierarchically and connected by pointers, all of which are maintained by SAS. The lowest level in the index file hierarchy consists of entries that represent each distinct value for an indexed variable, in ascending value order. Each entry contains this information:

- a distinct value
- one or more unique record identifiers (referred to as a *RID*) that identifies each observation containing the value. (Think of the RID as an internal observation number.)

That is, in an index file, each value is followed by one or more RIDs, which identify the observations in the data file that contains the value. (Multiple RIDs result from multiple occurrences of the same value.) For example, the following represents index file entries for the variable LASTNAME:

Value	RID
Avery	10
Brown	6, 22, 43
Craig	5, 50
Dunn	1

When an index is used to process a request, such as a WHERE expression, SAS performs a binary search on the index file and positions the index to the first entry that contains a qualified value. SAS then uses the value's RID to read the observation that contains the value. If a value has more than one RID (such as in the value for Brown in the previous example), SAS reads the observation that is pointed to by the next RID in the list. The result is that SAS can quickly locate the observations that are associated with a value or range of values.

For example, using an index to process the WHERE expression, SAS positions the index to the index entry for the first value greater than 20 and uses the value's RID or RIDs to read the observation or observations **where age > 20 and age < 35**;. SAS then moves sequentially through the index entries reading observations until it reaches the index entry for the value that is equal to or greater than 35.

SAS automatically keeps the index file balanced as updates are made, which means that it ensures a uniform cost to access any index entry, and all space that is occupied by deleted values is recovered and reused.

Types of Indexes

When you create an index, you designate which variable or variables to index. An indexed variable is called a *key variable*. You can create two types of indexes:

- A *simple index*, which consists of the values of one variable.
- A *composite index*, which consists of the values of more than one variable, with the values concatenated to form a single value.

In addition to deciding whether you want a simple index or a composite index, you can also limit an index (and its data file) to *unique values* and exclude from the index *missing values*.

Simple Index

The most common index is a simple index, which is an index of values for one key variable. The variable can be numeric or character. When you create a simple index, SAS assigns to the index the name of the key variable.

The following example shows the DATASETS procedure statements that are used to create two simple indexes for variables CLASS and MAJOR in data file COLLEGE.SURVEY:

```
proc datasets library=college;
  modify survey;
    index create class;
    index create major;
run;
```

To process a WHERE expression using an index, SAS uses only one index. When the WHERE expression has multiple conditions using multiple key variables, SAS

determines which condition qualifies the smallest subset. For example, suppose that COLLEGE.SURVEY contains the following data:

- 42,000 observations contain CLASS=97
- 6,000 observations contain MAJOR='Biology'
- 350 observations contain both CLASS=97 and MAJOR='Biology'.

With simple indexes on CLASS and MAJOR, SAS would select MAJOR to process the following WHERE expression:

```
where class=97 and major='Biology';
```

Composite Index

A composite index is an index of two or more key variables with their values concatenated to form a single value. The variables can be numeric, character, or a combination. An example is a composite index for the variables LASTNAME and FRSTNAME. A value for this index consists of the value for LASTNAME immediately followed by the value for FRSTNAME from the same observation. When you create a composite index, you must specify a unique index name.

The following example shows the DATASETS procedure statements that are used to create a composite index for the data file COLLEGE.MAILLIST, specifying two key variables: ZIPCODE and SCHOOLID.

```
proc datasets library=college;
  modify maillist;
  index create zipid=(zipcode schoolid);
run;
```

Often, only the first variable of a composite index is used. For example, for a composite index on ZIPCODE and SCHOOLID, the following WHERE expression can use the composite index for the variable ZIPCODE because it is the first key variable in the composite index:

```
where zipcode = 78753;
```

However, you can take advantage of all key variables in a composite index by the way you construct the WHERE expression, which is referred to as *compound optimization*. Compound optimization is the process of optimizing multiple WHERE expression conditions using a single composite index. If you issue the following WHERE expression, the composite index is used to find all occurrences where the Zipcode is 78753 and the school identification number is 55. In this way, all of the conditions are satisfied with a single search of the index:

```
where zipcode = 78753 and schoolid = 55;
```

When you are deciding whether to create a simple index or a composite index, consider how you will access the data. If you often access data for a single variable, a simple index will do. But if you frequently access data for multiple variables, a composite index could be beneficial.

Unique Values

Often it is important to require that values for a variable be unique, like Social Security number and employee number. You can declare unique values for a variable by creating an index for the variable and including the UNIQUE option. A unique index guarantees that values for one variable or the combination of a composite group of variables remain unique for every observation in the data file. If an update tries to add a duplicate value to that variable, the update is rejected.

The following example creates a simple index for the variable IDNUM and requires that all values for IDNUM be unique:

```
proc datasets library=college;
  modify student;
    index create idnum / unique;
run;
```

Missing Values

If a variable has a large number of missing values, it might be desirable to keep them from using space in the index. Therefore, when you create an index, you can include the NOMISS option to specify that missing values are not maintained by the index.

The following example creates a simple index for the variable RELIGION and specifies that the index does not maintain missing values for the variable:

```
proc datasets library=college;
  modify student;
    index create religion / nomiss;
run;
```

In contrast to the UNIQUE option, observations with missing values for the key variable can be added to the data file, even though the missing values are not added to the index.

SAS will not use an index that was created with the NOMISS option to process a BY statement or to process a WHERE expression that qualifies observations that contain missing values. If no missing values are present, SAS considers using the index in processing the BY statement or WHERE expression.

In the following example, the index AGE was created with the NOMISS option and observations exist that contain missing values for the variable AGE. In this case, SAS will not use the index:

```
proc print data=mydata.employee;
  where age < 35;
run;
```

Deciding Whether to Create an Index

Costs of an Index

An index exists to improve performance. However, an index conserves some resources at the expense of others. Therefore, you must consider costs associated with creating, using, and maintaining an index. The following topics provide information on resource usage and give you some guidelines for creating indexes.

When you are deciding whether to create an index, you must consider CPU cost, I/O cost, buffer requirements, and disk space requirements.

CPU Cost

Additional CPU time is necessary to create an index as well as to maintain the index when the data file is modified. That is, for an indexed data file, when a value is added, deleted, or modified, it must also be added, deleted, or modified in the appropriate index(es).

When SAS uses an index to read an observation from a data file, there is also increased CPU usage. The increased usage results from SAS using a more complicated process than is used when SAS retrieves data sequentially. Although CPU usage is

greater, you benefit from SAS reading only those observations that meet the conditions. Note that increased CPU usage is why using an index is more expensive when there is a larger number of observations that meet the conditions.

Note: To compare CPU usage with and without an index, for some operating environments, you can issue the STIMER or FULLSTIMER system options in order to write performance statistics to the SAS log. Δ

I/O Cost

Using an index to read observations from a data file can increase the number of I/O (input/output) requests compared to reading the data file sequentially. For example, processing a BY statement with an index might increase I/O count, but you save in not having to issue the SORT procedure. For WHERE processing, SAS considers I/O count when deciding whether to use an index.

To process a request using an index, the following occurs:

- 1 SAS does a binary search on the index file and positions the index to the first entry that contains a qualified value.
- 2 SAS uses the value's RID (identifier) to directly access the observation containing the value. SAS transfers the observation between external storage to a *buffer*, which is the memory into which data is read or from which data is written. The data is transferred in *pages*, which is the amount of data (the number of observations) that can be transferred for one I/O request; each data file has a specified page size.
- 3 SAS then continues the process until the WHERE expression is satisfied. Each time SAS accesses an observation, the data file page containing the observation must be read into memory if it is not already there. Therefore, if the observations are on multiple data file pages, an I/O operation is performed for each observation.

The result is that the more random the data, the more I/Os are required to use the index. If the data is ordered more like the index, which is in ascending value order, a smaller number of I/Os are required to access the data.

The number of buffers determines how many pages of data can simultaneously be in memory. Frequently, the larger the number of buffers, the smaller the number of I/Os that are required. For example, if the page size is 4096 bytes and one buffer is allocated, then one I/O transfers 4096 bytes of data (or one page). To reduce I/Os, you can increase the page size but you will need a larger buffer. To reduce the buffer size, you can decrease the page size but you will use more I/Os.

For information on data file characteristics like the data file page size and the number of data file pages, issue the CONTENTS procedure (or use the CONTENTS statement in the DATASETS procedure). With this information, you can determine the data file page size and experiment with different sizes. Note that the information that is available from PROC CONTENTS depends on the operating environment.

The BUFSIZE= data set option (or system option) sets the permanent page size for a data file when it is created. The page size is the amount of data that can be transferred for an I/O operation to one buffer. The BUFNO= data set option (or system option) specifies how many buffers to allocate for a data file and for the overall system for a given execution of SAS; that is, BUFNO= is not stored as a data set attribute.

Buffer Requirements

In addition to the resources that are used to create and maintain an index, SAS also requires additional memory for buffers when an index is actually used. Opening the data file opens the index file but none of the indexes. The buffers are not required unless SAS uses the index but they must be allocated in preparation for the index that is being used.

The number of buffers that are allocated depends on the number of levels in the index tree and in the data file open mode. If the data file is open for input, the maximum number of buffers is three; for update, the maximum number is four. (Note that these buffers are available for other uses; they are not dedicated to indexes.)

The `IBUFSIZE=` system option specifies the page size on disk for an index file when it is created. The default setting causes SAS to use the minimum optimal page size for the operating environment. Typically, you do not need to specify an index page size. However, there are situations that could require a different page size. For more information, see the `IBUFSIZE=` system option in *SAS Language Reference: Dictionary*.

The `IBUFNO=` system option specifies an optional number of extra buffers to be allocated when navigating an index file. SAS automatically allocates a minimal number of buffers. Typically, you do not need to specify extra buffers. However, using `IBUFNO=` to specify extra buffers could improve execution time by limiting the number of input/output operations that are required for a particular index file. The improvement in execution time, however, comes at the expense of increased memory consumption. For more information, see the `IBUFNO=` system option in *SAS Language Reference: Dictionary*.

Disk Space Requirements

Additional disk space is required to store the index file, which might show up as a separate file or appear to be part of the data file, depending on the operating environment.

For information on the index file size, issue the `CONTENTS` procedure (or the `CONTENTS` statement in the `DATASETS` procedure). Note that the available information from `PROC CONTENTS` depends on the operating environment.

Guidelines for Creating Indexes

Data File Considerations

- For a small data file, sequential processing is often just as efficient as index processing. *Do not create an index if the data file page count is less than three pages.* It would be faster to access the data sequentially. To see how many pages are in a data file, use the `CONTENTS` procedure (or use the `CONTENTS` statement in the `DATASETS` procedure). Note that the information that is available from `PROC CONTENTS` depends on the operating environment.
- Consider the cost of an index for a data file that is frequently changed. If you have a data file that changes often, the overhead associated with updating the index after each change can outweigh the processing advantages you gain from accessing the data with an index.
- Create an index when you intend to retrieve a small subset of observations from a large data file (for example, less than 25% of all observations). When this occurs, the cost of processing data file pages is lower than the overhead of sequentially reading the entire data file. The smaller the subset, the larger the performance gains.
- To reduce the number of I/Os performed when you create an index, first sort the data by the key variable. Then to improve performance, maintain the data file in sorted order by the key variable. This technique reduces the I/Os by grouping like values together. That is, the more ordered the data file is with respect to the key

variable, the more efficient the use of the index. If the data file has more than one index, sort the data by the most frequently used key variable.

- An index might not be necessary to optimize a WHERE expression if the data is sorted appropriately in order to satisfy the condition. To process a WHERE expression without an index, SAS first checks for the sort indicator that is stored with the file from a previous SORT procedure. If the sort indicator is appropriate, SAS stops reading the file once there are no more values that satisfy the WHERE expression. For example, consider a file that is sorted by Age, without an index. To process the expression `where age le 25`, SAS stops reading observations after it finds an observation that is greater than 25. Note that while SAS can determine when to stop reading observations, if there is no index, there is no indication where to begin. Without an index, SAS always begins with the first observation, which can require reading a lot of observations.

Index Use Considerations

- Keep the number of indexes per data file to a minimum to reduce disk storage and to reduce update costs.
- Consider how often your applications will use an index. An index must be used often in order to make up for the resources that are used in creating and maintaining it. That is, do not rely solely on resource savings from processing a WHERE expression. Take into consideration the resources it takes to actually create the index and to maintain it every time the data file is changed.
- When you create an index to process a WHERE expression, do not try to create one index that is used to satisfy all queries. If there are several variables that appear in queries, those queries might be best satisfied with simple indexes on the most discriminating of those variables.

Key Variable Candidates

In most cases, multiple variables are used to query a data file. However, it probably would be a mistake to index all variables in a data file, as certain variables are better candidates than others:

- The variables to be indexed should be variables that are used in queries. That is, your application should require selecting small subsets from a large file, and the most common selection variables should be considered as candidate key variables.
- A variable is a good candidate for indexing when the variable can be used to precisely identify the observations that satisfy a WHERE expression. That is, the variable should be *discriminating*, which means that the index should select the fewest possible observations. For example, variables such as AGE, FRSTNAME, and GENDER are not discriminating because it is very possible for a large representation of the data to have the same age, first name, and gender. However, a variable such as LASTNAME is a good choice because it is less likely that many employees share the same last name.

For example, consider a data file with variables LASTNAME and GENDER.

- If many queries against the data file include LASTNAME, then indexing LASTNAME could prove to be beneficial because the values are usually discriminating. However, the same reasoning would not apply if you issued a large number of queries that included GENDER. The GENDER variable is not discriminating (because perhaps half the population is male and half is female).

- However, if queries against the data file most often include both LASTNAME and GENDER as shown in the following WHERE expression, then creating a composite index on LASTNAME and GENDER could improve performance.

```
where lastname='LeVoux' and gender='F';
```

Note that when you create a composite index, the first key variable should be the most discriminating.

Creating an Index

Overview of Creating Indexes

You can create one index for a data file, which can be either a simple index or a composite index, and you can create multiple indexes, which can be multiple simple indexes, multiple composite indexes, or a combination of both simple and composite.

In general, the process of creating an index is as follows:

- 1 You request to create an index for one or multiple variables using a method such as the INDEX CREATE statement in the DATASETS procedure.
- 2 SAS reads the data file one observation at a time, extracts values and RIDs for each key variable, and places them in the index file.

SAS ensures that the values that are placed in the index are successively the same or increasing. SAS determines whether the data is already sorted by the key variables in ascending order by checking the sort indicator in the data file, which is an attribute of the file that indicates how the data is sorted. The sort indicator is stored with the SAS data file descriptor information and is set from a previous SORT procedure or SORTEDBY= data set option.

If the values in the sort indicator are in ascending order, SAS does not sort the values for the index file and avoids the resource. Note that SAS always validates that the data is sorted as indicated. If not, the index is not created. For example, if the sort indicator was set from a SORTEDBY= data set option and the data is not sorted as indicated, an error occurs and a message is written to the SAS log stating that the index was not created because values are not sorted in ascending order.

If the values in the sort indicator are not in ascending order, SAS sorts the data that is included in the index file in ascending value order. To sort the data, SAS follows this procedure:

- 1 SAS first attempts to sort the data using the thread-enabled sort. By dividing the sorting into separately executable processes, the time to sort the data can be reduced. However, in order to use the thread-enabled sort, the size of the index must be sufficiently large (which is determined by SAS), the SAS system option CPUCOUNT= must be set to more than one processor, and the THREADS system option must be enabled. Adequate memory must be available for the thread-enabled sort. If not enough memory is available, SAS reduces the number of threads to one and begins the sort process again, which increases the time to create the index.
- 2 If the thread-enabled sort cannot be done, SAS uses the unthreaded sort.

Note: To display messages regarding what type of sort is used, memory and resource information, and the status of the index being created, set the SAS system option MSGLEVEL=I, that is:

```
options msglevel=i;
```

Δ

Using the DATASETS Procedure

The DATASETS procedure provides statements that enable you to create and delete indexes. In the following example, the MODIFY statement identifies the data file, the INDEX DELETE statement deletes two indexes, and the two INDEX CREATE statements specify the variables to index, with the first INDEX CREATE statement specifying the options UNIQUE and NOMISS:

```
proc datasets library=mylib;
modify employee;
    index delete salary age;
    index create empnum / unique nomiss;
    index create names=(lastname frstname);
```

Note: If you delete and create indexes in the same step, place the INDEX DELETE statement before the INDEX CREATE statement so that space occupied by deleted indexes can be reused during index creation. Δ

Using the INDEX= Data Set Option

To create indexes in a DATA step when you create the data file, use the INDEX= data set option. The INDEX= data set option also enables you to include the NOMISS and UNIQUE options. The following example creates a simple index on the variable STOCK and specifies UNIQUE:

```
data finances(index=(stock /unique));
```

The next example uses the variables SSN, CITY, and STATE to create a simple index named SSN and a composite index named CITYST:

```
data employee(index=(ssn cityst=(city state)));
```

Using the SQL Procedure

The SQL procedure supports index creation and deletion and the UNIQUE option. Note that the variable list requires that variable names be separated by commas (which is an SQL convention) instead of blanks (which is a SAS convention).

The DROP INDEX statement deletes indexes. The CREATE INDEX statement specifies the UNIQUE option, the name of the index, the target data file, and the variable or variables to be indexed. For example:

```
drop index salary from employee;
create unique index empnum on employee (empnum);
create index names on employee (lastname, frstname);
```

Using Other SAS Products

You can also create and delete indexes using other SAS utilities and products, such as SAS/CONNECT software, SAS/IML software, SAS Component Language, and SAS/Warehouse Administrator software.

Using an Index for WHERE Processing

WHERE processing conditionally selects observations for processing when you issue a WHERE expression. Using an index to process a WHERE expression improves performance and is referred to as *optimizing* the WHERE expression.

To process a WHERE expression, by default SAS decides whether to use an index or read all the observations in the data file sequentially. To make this decision, SAS does the following:

- 1 Identifies an available index or indexes.
- 2 Estimates the number of observations that would be qualified. If multiple indexes are available, SAS selects the index that returns the smallest subset of observations.
- 3 Compares resource usage to decide whether it is more efficient to satisfy the WHERE expression by using the index or by reading all the observations sequentially.

Note: SAS considers several factors when deciding whether to use an index; therefore, experimentation is the best way to determine the optimal performance. If you have a WHERE expression that is used repeatedly, compare the results using an index and without an index in order to determine which method provides the best performance. You can control index usage with the `IDXWHERE=` and `IDXNAME=` data set options. See “Controlling WHERE Processing Index Usage with Data Set Options” on page 519. Δ

Identifying Available Index or Indexes

The first step for SAS in deciding whether to use an index to process a WHERE expression is to identify if the variable or variables included in the WHERE expression are key variables (that is, have an index). Even though a WHERE expression can consist of multiple conditions that specify different variables, SAS uses only one index to process the WHERE expression. SAS selects the index that satisfies the most conditions and qualifies the fewest observations:

- Usually SAS selects one condition. The variable specified in the condition will have either a simple index or be the first key variable in a composite index.
- However, you can take advantage of multiple key variables in a composite index by constructing an appropriate WHERE expression, referred to as *compound optimization*. See “Compound Optimization” on page 516.

SAS attempts to use an index for the following types of conditions:

Table 27.6 WHERE Conditions That Can Be Optimized

Condition	Valid for Compound Optimization	Examples
comparison operators, which include the EQ operator; directional comparisons like less than or greater than; and the IN operator	yes	where empnum eq 3374; where empnum < 2000; where state in ('NC','TX');
comparison operators with NOT	yes	where empnum ^= 3374; where x not in (5,10);
comparison operators with the colon modifier	yes	where lastname gt: 'Sm';

Condition	Valid for Compound Optimization	Examples
CONTAINS operator	no	where lastname contains 'Sm';
fully bounded range conditions specifying both an upper and lower limit, which includes the BETWEEN-AND operator	yes	where 1 < x < 10; where empnum between 500 and 1000;
pattern-matching operators LIKE and NOT LIKE	no	where firstname like '%Rob_%'
IS NULL or IS MISSING operator	no	where name is null; where idnum is missing;
TRIM function	no	where trim(state)='Texas';
SUBSTR function in the form of: WHERE SUBSTR (<i>variable</i> , <i>position</i> , <i>length</i>)='string'; when the following conditions are met: <i>position</i> is equal to 1, <i>length</i> is less than or equal to the length of <i>variable</i> , and <i>length</i> is equal to the length of <i>string</i>	no	where substr (name,1,3)='Mac' and (city='Charleston' or city='Atlanta');

Note: Conditions are not optimized with an index for arithmetic operators, a variable-to-variable condition, and the sounds-like operator. Δ

The following examples illustrate optimizing a single condition:

- The following WHERE expressions could use a simple index on the variable MAJOR:

```
where major in ('Biology', 'Chemistry', 'Agriculture');
where class=90 and major in ('Biology', 'Agriculture');
```

- With a composite index on variables ZIPCODE and SCHOOLID, SAS could use the composite index to satisfy the following conditions because ZIPCODE is the first key variable in the composite index:

```
where zipcode = 78753;
```

However, the following condition cannot use the composite index because the variable SCHOOLID is not the first key variable in the composite index:

```
where schoolid gt 1000;
```

Compound Optimization

Compound optimization is the process of optimizing multiple WHERE expression conditions using a single composite index. Using a single index to optimize the conditions can greatly improve performance.

For example, suppose there is a composite index for LASTNAME and FRSTNAME. If you execute the following WHERE expression, SAS uses the concatenated values for the

first two variables, then SAS further evaluates each qualified observation for the EMPID value:

```
where lastname eq 'Smith' and firstname eq 'John' and empid=3374;
```

For compound optimization to occur, all of the following must be true.

- At least the first two key variables in the composite index must be used in valid WHERE expression conditions. For a list of conditions that are valid for compound optimization, see Table 27.6 on page 515.
- At least one condition must use the EQ or IN operator; you cannot have, for example, all range conditions.
- The conditions must be connected with the AND or the OR logical operator:
 - When conditions are connected with AND, the conditions can occur in any order. For example,

```
where lastname eq 'Smith' and firstname eq 'John';
```

- When conditions are connected with OR, the conditions must specify the same variable. For example,

```
where firstname eq 'John' and
      (lastname eq 'Smith' or lastname eq 'Jones');
```

Note: SAS transforms the OR conditions that specify the same variable into a single condition that uses the IN operator. For the above WHERE expression, SAS converts the two OR conditions into **lastname IN ('Smith', 'Jones')**, and then uses the composite index for the variables Firstname and Lastname in order to select the observations where Firstname is John and Lastname is Smith or Jones. Δ

For the following examples, assume there is a composite index for variables I, J, and CH:

- The following WHERE expression conditions are compound optimized because every condition specifies a variable that is in the composite index, and each condition uses one of the supported operators. SAS positions the composite index to the first entry that meets all three conditions and retrieves only observations that satisfy all three conditions.

```
where I = 1 and J not in (3,4) and 'abc' < CH;
```

- For the following WHERE expression, the first two conditions are compound optimized. After retrieving a subset of observations that satisfy the first two conditions, SAS examines the subset and eliminates any observations that fail to match the third condition.

```
where I in (1,4) and J = 5 and K like '%c';
```

- This WHERE expression can be compound optimized for variables I and J. After retrieving observations that satisfy the second and third conditions, SAS examines the subset and eliminates those observations that do not satisfy the first condition.

```
where X < 5 and I = 1 and J = 2;
```

- The following WHERE expression can be compound optimized on I and J:

```
where X < Z and I = 1 and J = 2;
```

- The following WHERE expression cannot be compound optimized neither J or K is the left-most variable in the composite index:

```
where J = 1 and K = 2;
```

- The following WHERE expression cannot be optimized because the comparison condition on the variable I is variable-to-variable, which is not supported for index processing:

```
where I < K and J in (3,4) and CH = 'abc';
```

Compound optimization can occur for a NOMISS composite index as long as at least one condition does not qualify missing values. That is, compound optimization cannot occur on a NOMISS index, which is an index that does not maintain missing values, if every condition could result in a missing value. The following examples illustrate compound optimization with a NOMISS composite index for variables I, J, and K.

- The following WHERE expression can be compound optimized, because the condition $K = 1$ cannot result in a missing value:

```
where I in (.,5) and J < 4 and K = 1;
```

- This WHERE expression cannot be compound optimized, because each condition could result in a missing value:

```
where I in (.,5) and J < 4 and K <= 1;
```

- The following WHERE expression cannot be compound optimized, because each condition could result in a missing value. The condition $J < 4$ qualifies observations as $J = .$, and those observations are not represented in the NOMISS composite index:

```
where I = . and J < 4 and .A < K < .D;
```

Estimating the Number of Qualified Observations

Once SAS identifies the index or indexes that can satisfy the WHERE expression, the software estimates the number of observations that will be qualified by an available index. When multiple indexes exist, SAS selects the one that appears to produce the fewest qualified observations.

SAS estimates the number of observations that will be qualified by using stored statistics called cumulative percentiles (or *centiles* for short). Centiles information represents the distribution of values in an index so that SAS does not have to assume a uniform distribution. To print centiles information for an indexed data file, include the CENTILES option in PROC CONTENTS (or in the CONTENTS statement in the DATASETS procedure).

Note that, by default, SAS does not update centiles information after every data file change. When you create an index, you can include the UPDATECENTILES option to specify when centiles information is updated. That is, you can specify that centiles information be updated every time the data file is closed, when a certain percentage of values for the key variable have been changed, or never. In addition, you can also request that centiles information is updated immediately, regardless of the value of UPDATECENTILES, by issuing the INDEX CENTILES statement in PROC DATASETS.

As a general rule, SAS uses an index if it estimates that the WHERE expression will select approximately one-third or less of the total number of observations in the data file.

Note: For performance purposes, the following can occur when SAS estimates the number of qualified observations:

- If the number of qualified observations is less than 3% of the data file (or if no observations are qualified), SAS automatically uses the index, and does not bother comparing resource usage.

- If all of the observations are qualified, by default SAS does not use the index unless the `IDXNAME=` or `IDXWHERE=` data set option is specified.

△

Comparing Resource Usage

Once SAS estimates the number of qualified observations and selects the index that qualifies the fewest observations, SAS must then decide whether it is faster (cheaper) to satisfy the `WHERE` expression by using the index or by reading all of the observations sequentially. SAS makes this determination as follows:

- If only a few observations are qualified, it is more efficient to use the index than to do a sequential search of the entire data file.
- If most or all of the observations qualify, then it is more efficient to simply sequentially search the data file than to use the index.

This decision is much like a reader deciding whether to use an index at the back of a document. A document's index is designed to enable a reader to locate a topic along with the specific page number. Using the index, the reader would go to a specific page number and read only about a specific topic. If the document covers 42 topics and the reader is interested in only a couple of topics, then the index saves time by preventing the reader from reading other topics. However, if the reader is interested in 39 topics, searching the index for each topic would take more time than simply reading the entire document.

To compare resource usage, SAS does the following:

- 1 First, SAS predicts the number of I/Os it will take to satisfy the `WHERE` expression using the index. To do so, SAS positions the index to the first entry that contains a qualified value. In a buffer management simulation that takes into account the current number of available buffers, the RIDs (identifiers) on that index page are processed, indicating how many I/Os it will take to read the observations in the data file.

If the observations are randomly distributed throughout the data file, the observations will be located on multiple data file pages, which means an I/O will be needed for each page. Therefore, the more random the data in the data file, the more I/Os it takes to use the index. If the data in the data file is ordered more like the index, which is in ascending value order, a smaller number of I/Os are needed to use the index.

- 2 Then SAS calculates the I/O cost of a sequential pass of the entire data file and compares the two resource costs.

Factors that affect the comparison include the size of the subset relative to the size of the data file, data file value order, data file page size, the number of allocated buffers, and the cost to uncompress a compressed data file for a sequential read.

Note: If comparing resource costs results in a tie, SAS chooses the index. △

Controlling WHERE Processing Index Usage with Data Set Options

You can control index usage for `WHERE` processing with the `IDXWHERE=` and `IDXNAME=` data set options.

The `IDXWHERE=` data set option overrides the software's decision regarding whether to use an index to satisfy the conditions of a `WHERE` expression as follows:

- `IDXWHERE=YES` tells SAS to decide which index is the best for optimizing a `WHERE` expression, disregarding the possibility that a sequential search of the data file might be more resource efficient.

- `IDXWHERE=NO` tells SAS to ignore all indexes and satisfy the conditions of a WHERE expression by sequentially searching the data file.
- Using an index to process a BY statement cannot be overridden with `IDXWHERE=`.

The following example tells SAS to decide which index is the best for optimizing the WHERE expression. SAS disregards the possibility that a sequential search of the data file might be more resource efficient.

```
data mydata.empnew;
  set mydata.employee (idxwhere=yes);
  where empnum < 2000;
```

For details, see the `IDXWHERE` data set option in *SAS Language Reference: Dictionary*.

The `IDXNAME=` data set option directs SAS to use a specific index in order to satisfy the conditions of a WHERE expression.

By specifying `IDXNAME=index-name`, you are specifying the name of a simple or composite index for the data file.

The following example uses the `IDXNAME=` data set option to direct SAS to use a specific index to optimize the WHERE expression. SAS disregards the possibility that a sequential search of the data file might be more resource efficient and does not attempt to determine whether the specified index is the best one. (Note that the `EMPNUM` index was not created with the `NOMISS` option.)

```
data mydata.empnew;
  set mydata.employee (idxname=empnum);
  where empnum < 2000;
```

For details, see the `IDXNAME` data set option in *SAS Language Reference: Dictionary*.

Note: `IDXWHERE=` and `IDXNAME=` are mutually exclusive. Using both results in an error. \triangle

Displaying Index Usage Information in the SAS Log

To display information in the SAS log regarding index usage, change the value of the `MSGLEVEL=` system option from its default value of `N` to `I`. When you issue `options msglevel=i;`, the following occurs:

- If an index is used, a message displays specifying the name of the index.
- If an index is not used but one exists that could optimize at least one condition in the WHERE expression, messages provide suggestions as to what you can do to influence SAS to use the index; for example, a message could suggest sorting the data file into index order or specifying more buffers.
- A message displays the `IDXWHERE=` or `IDXNAME=` data set option value if the setting can affect index processing.

Using an Index with SAS Views

You cannot create an index for a SAS view; it must be a data file. However, if a SAS view is created from an indexed data file, index usage is available. That is, if the view definition includes a WHERE expression using a key variable, then SAS attempts to use the index. Additionally, there are other ways to take advantage of a key variable when using a SAS view.

In this example, you create an SQL view named STAT from data file CRIME, which has the key variable STATE. In addition, the view definition includes a WHERE expression:

```
proc sql;
  create view stat as
  select * from crime
  where murder > 7;
quit;
```

If you issue the following PRINT procedure, which refers to the SQL view, along with a WHERE statement that specifies the key variable STATE, SAS cannot optimize the WHERE statement with the index. SQL views cannot join a WHERE expression that was defined in the view to a WHERE expression that was specified in another procedure, DATA step, or SCL:

```
proc print data=stat;
  where state > 42;
run;
```

However, if you issue PROC SQL with an SQL WHERE clause that specifies the key variable STATE, then the SQL view can join the two conditions, which enables SAS to use the index STATE:

```
proc sql;
  select * from stat where state > 42;
quit;
```

Using an Index for BY Processing

BY processing enables you to process observations in a specific order according to the values of one or more variables that are specified in a BY statement. Indexing a data file enables you to use a BY statement without sorting the data file. By creating an index based on one or more variables, you can ensure that observations are processed in *ascending numeric or character order*. Simply specify in the BY statement the variable or list of variables that are indexed.

For example, if an index exists for LASTNAME, the following BY statement would use the index to order the values by last names:

```
proc print;
  by lastname;
```

When you specify a BY statement, SAS looks for an appropriate index. If one exists, the software automatically retrieves the observations from the data file in indexed order. A BY statement uses an index in the following situations:

- The BY statement consists of one variable that is the key variable for a simple index or the first key variable in a composite index.
- The BY statement consists of two or more variables and the first variable is the key variable for a simple index or the first key variable in a composite index.

For example, if the variable MAJOR has a simple index, the following BY statements use the index to order the values by MAJOR:

```
by major;
by major state;
```

If a composite index named ZIPID exists consisting of the variables ZIPCODE and SCHOOLID, the following BY statements use the index:

```
by zipcode;
by zipcode schoolid;
by zipcode schoolid name;
```

However, the composite index ZIPID is not used for these BY statements:

```
by schoolid;
by schoolid zipcode;
```

In addition, a BY statement does not use an index in these situations:

- The BY statement includes the DESCENDING or NOTSORTED option.
- The index was created with the NOMISS option.
- The data file is physically stored in sorted order based on the variables specified in the BY statement.

Note: Using an index to process a BY statement might not always be more efficient than simply sorting the data file, particularly if the data file has a high blocking factor of observations per page. Therefore, using an index for a BY statement is generally for convenience, not performance. Δ

Using an Index for Both WHERE and BY Processing

If both a WHERE expression and a BY statement are specified, SAS looks for one index that satisfies requirements for both. If such an index is not found, the BY statement takes precedence.

With a BY statement, SAS cannot use an index to optimize a WHERE expression if the optimization would invalidate the BY order. For example, the following statements could use an index on the variable LASTNAME to optimize the WHERE expression because the order of the observations returned by the index does not conflict with the order required by the BY statement:

```
proc print;
  by lastname;
  where lastname >= 'Smith';
run;
```

However, the following statements cannot use an index on LASTNAME to optimize the WHERE expression because the BY statement requires that the observations be returned in EMPID order:

```
proc print;
  by empid;
  where lastname = 'Smith';
run;
```

Specifying an Index with the KEY= Option for SET and MODIFY Statements

The SET and MODIFY statements provide the KEY= option, which enables you to specify an index in a DATA step to retrieve particular observations in a data file.

The following MODIFY statement shows how to use the KEY= option to take advantage of the fact that the data file INVTY.STOCK has an index on the variable PARTNO. Using the KEY= option tells SAS to use the index to directly access the correct observations to modify.

```
modify invty.stock key=partno;
```

Note: A BY statement is not allowed in the same DATA step with the KEY= option, and WHERE processing is not allowed for a data file with the KEY= option. △

Taking Advantage of an Index

Applications that typically do not use indexes can be rewritten to take advantage of an index. For example:

- Consider replacing a subsetting IF statement (which never uses an index) with a WHERE statement.

CAUTION:

However, be careful because IF and WHERE statements are processed differently and might produce different results in DATA steps that use the SET, MERGE, or UPDATE statements. This is because the WHERE statement selects observations before they are brought into the Program Data Vector (PDV), whereas the subsetting IF statement selects observations after they are read into the PDV. △

- Consider using the WHERE command in the FSEDIT procedure in place of the SEARCH and FIND commands.

Maintaining Indexes

SAS provides several procedures that you can issue to maintain indexes, and there are several operations within SAS that automatically maintain indexes for you.

Displaying Data File Information

The CONTENTS procedure (or the CONTENTS statement in PROC DATASETS) reports the following types of information.

- number and names of indexes for a data file
- the names of key variables
- the options in effect for each key variable
- data file page size
- number of data file pages
- centiles information (using the CENTILES option)
- amount of disk space used by the index file

Note: The available information depends on the operating environment. △

Output 27.5 Output of PROC CONTENTS

The SAS System		1			
The CONTENTS Procedure					
Data Set Name	MYFILES.STAFF	Observations	148		
Member Type	DATA	Variables	6		
Engine	V9	Indexes	2		
Created	Monday, October 01, 2007 10:31:37 AM	Observation Length	63		
Last Modified	Monday, October 01, 2007 10:36:13 AM	Deleted Observations	0		
Protection		Compressed	NO		
Data Set Type		Sorted	NO		
Label					
Data Representation	WINDOWS_32				
Encoding	wlatin1 Western (Windows)				
Engine/Host Dependent Information					
Data Set Page Size	8192				
Number of Data Set Pages	3				
First Data Page	1				
Max Obs per Page	129				
Obs in First Data Page	104				
Index File Page Size	4096				
Number of Index File Pages	5				
Number of Data Set Repairs	0				
File Name	C:\Myfiles\staff.sas7bdat				
Release Created	9.0201B0				
Host Created	XP_PRO				
Alphabetic List of Variables and Attributes					
	#	Variable	Type	Len	
	4	city	Char	15	
	3	fname	Char	15	
	6	hphone	Char	12	
	1	idnum	Char	4	
	2	lname	Char	15	
	5	state	Char	2	
Alphabetic List of Indexes and Attributes					
	#	Index	Unique Option	# of Unique Values	Variables
	1	idnum	YES	148	
	2	name		148	fname lname

Copying an Indexed Data File

When you copy an indexed data file with the COPY procedure (or the COPY statement of the DATASETS procedure), you can specify whether the procedure also recreates the index file for the new data file with the INDEX=YES|NO option; the default is YES, which recreates the index. However, recreating the index does increase the processing time for the PROC COPY step.

If you copy from disk to disk, the index is recreated. If you copy from disk to tape, the index is not recreated on tape. However, after copying from disk to tape, if you then copy back from tape to disk, the index can be recreated. Note that if you move a data file with the MOVE option in PROC COPY, the index file is deleted from IN= library and recreated in OUT= library.

The CPORT procedure also has INDEX=YES|NO to specify whether to export indexes with indexed data files. By default, PROC CPORT exports indexes with indexed data files. The CIMPORT procedure, however, does not handle the index file at all, and the index(es) must be recreated.

Updating an Indexed Data File

Each time that values in an indexed data file are added, modified, or deleted, SAS automatically updates the index. The following activities affect an index as indicated:

Table 27.7 Maintenance Tasks and Index Results

Task	Result
delete a data set	index file is deleted
rename a data set	index file is renamed
rename key variable	simple index is renamed
delete key variable	simple index is deleted
add observation	index entries are added
delete observations	index entries are deleted and space is recovered for reuse
update observations	index entries are deleted and new ones are inserted

Note: Use SAS to perform additions, modifications, and deletions to your data sets. Using operating environment commands to perform these operations will make your files unusable. Δ

Sorting an Indexed Data File

You can sort an indexed data file only if you direct the output of the SORT procedure to a new data file so that the original data file remains unchanged. However, the new data file is not automatically indexed.

Note: If you sort an indexed data file with the FORCE option, the index file is deleted. Δ

Adding Observations to an Indexed Data File

Adding observations to an indexed data file requires additional processing. SAS automatically keeps the values in the index consistent with the values in the data file.

Multiple Occurrences of Values

An index that is created without the UNIQUE option can result in multiple occurrences of the same value, which results in multiple RIDs for one value. For large data files with many multiple occurrences, the list of RIDs for a given value might require several pages in the index file. Because the RIDs are stored in physical order, any new observation added to the data file with the given value is stored at the end of the list of RIDs. Navigating through the index to find the end of the RID list can cause many I/O operations.

SAS remembers the previous position in the index so that when inserting more occurrences of the same value, the end of the RID list is found quickly.

Appending Data to an Indexed Data File

SAS provides performance improvements when appending a data file to an indexed data file. SAS suspends index updates until all observations are added, and then updates the index with data from the newly added observations. See the APPEND statement in the DATASETS procedure in *Base SAS Procedures Guide*.

Recovering a Damaged Index

An index can become damaged for many of the same reasons that a data file or catalog can become damaged. If a data file becomes damaged, use the REPAIR statement in PROC DATASETS to repair the data file or recreate any missing indexes. For example,

```
proc datasets library=mylib;
    repair mydata;
run;
```

Indexes and Integrity Constraints

Integrity constraints can also use indexes. When an integrity constraint is created that uses an index, if a suitable index already exists, it is used; otherwise, a new index is created. When an index is created, it is marked as being “owned” by the creator, which can be either the user or an integrity constraint.

If either the user or an integrity constraint requests creation of an index that already exists and is owned by the other, the requestor is also marked as an “owner” of the index. If an index is owned by both, then a request by either to delete the index results in removing only the requestor as owner. The index is deleted only after both the integrity constraint and the user have requested the index’s deletion. A note in the log indicates when an index cannot be deleted.

Indexes and CEDA Processing

When processing a SAS data file with CEDA, indexes are not supported. For example, if you move a SAS data file with a defined index from one operating environment like Windows to a different operating environment like UNIX, CEDA translates the file for you, but the index is not available. Therefore, WHERE optimization for the file is not supported.

For information on CEDA, see Chapter 33, “Processing Data Using Cross-Environment Data Access (CEDA),” on page 567.

Compressing Data Files

Definition of Compression

Compressing a file is a process that reduces the number of bytes required to represent each observation. In a compressed file, each observation is a variable-length record, while in an uncompressed file, each observation is a fixed-length record.

Advantages of compressing a file include

- reduced storage requirements for the file
- less I/O operations necessary to read from or write to the data during processing.

However, disadvantages of compressing a file are that

- more CPU resources are required to read a compressed file because of the overhead of uncompressing each observation
- there are situations when the resulting file size can increase rather than decrease.

Requesting Compression

By default, a SAS data file is not compressed. To compress, you can use these options:

- COMPRESS= system option to compress all data files that are created during a SAS session
- COMPRESS= option in the LIBNAME statement to compress all data files for a particular SAS library
- COMPRESS= data set option to compress an individual data file.

To compress a data file, you can specify the following:

- COMPRESS=CHAR to use the RLE (Run Length Encoding) compression algorithm.
- COMPRESS=BINARY to use the RDC (Ross Data Compression) algorithm.

When you create a compressed data file, SAS writes a note to the log indicating the percentage of reduction that is obtained by compressing the file. SAS obtains the compression percentage by comparing the size of the compressed file with the size of an uncompressed file of the same page size and record count.

After a file is compressed, the setting is a permanent attribute of the file, which means that to change the setting, you must re-create the file. That is, to uncompress a file, specify COMPRESS=NO for a DATA step that copies the compressed data file.

For more information on the COMPRESS= data set option, the COMPRESS= option in the LIBNAME statement, and the COMPRESS= system option, see *SAS Language Reference: Dictionary*.

Disabling a Compression Request

Compressing a file adds a fixed-length block of data to each observation. Because of the additional block of data (12 bytes for a 32-bit host and 24 bytes for a 64-bit host per observation), some files could result in a larger file size. For example, files with extremely short record lengths could result in a larger file size if compressed.

When a request is made to compress a data set, SAS attempts to determine whether compression will increase the size of the file. SAS examines the lengths of the variables. If, due to the number and lengths of the variables, it is not possible for the compressed file to be at least 12 bytes (for a 32-bit host) or 24 bytes (for a 64-bit host) per observation smaller than an uncompressed version, compression is disabled and a message is written to the SAS log.


For example, here is a simple data set for which SAS determines that it is not possible for the compressed file to be smaller than an uncompressed one:

```
data one (compress=char);
  length x y $2;
  input x y;
  datalines;
ab cd
;
```

The following output is written to the SAS log:

Output 27.6 SAS Log Output When Compression Request Is Disabled

```
NOTE: Compression was disabled for data set WORK.ONE because compression overhead  
      would increase the size of the data set.  
NOTE: The data set WORK.ONE has 1 observations and 2 variables.
```



CHAPTER 28

SAS Views

<i>Definition of SAS Views</i>	529
<i>Benefits of Using SAS Views</i>	530
<i>When to Use SAS Views</i>	531
<i>DATA Step Views</i>	531
<i>Definition of a DATA Step View</i>	531
<i>Creating DATA Step Views</i>	531
<i>What Can You Do with a DATA Step View?</i>	532
<i>Differences between DATA Step Views and Stored Compiled DATA Step Programs</i>	532
<i>Restrictions and Requirements</i>	532
<i>Performance Considerations</i>	533
<i>Example 1: Merging Data to Produce Reports</i>	533
<i>Example 2: Producing Additional Output Files</i>	533
<i>PROC SQL Views</i>	535
<i>Comparing DATA Step and PROC SQL Views</i>	536
<i>SAS/ACCESS Views</i>	536

Definition of SAS Views

A *SAS view* is a type of SAS data set that retrieves data values from other files. A SAS view contains only descriptor information such as the data types and lengths of the variables (columns), plus information that is required for retrieving data values from other SAS data sets or from files that are stored in other software vendors' file formats. SAS views are of member type VIEW. In most cases, you can use a SAS view as if it were a SAS data file.

There are two general types of SAS views:

native view

is a SAS view that is created either with a DATA step or with PROC SQL.

interface view

is a SAS view that is created with SAS/ACCESS software. An interface view can read data from or write data to a database management system (DBMS) such as DB2 or ORACLE. Interface views are also referred to as *SAS/ACCESS views*. In order to use SAS/ACCESS views, you must have a license for SAS/ACCESS software.

Note: You can create native views that access certain DBMS data by using a SAS/ACCESS dynamic LIBNAME engine. See "SAS/ACCESS Views" on page 536, or the SAS/ACCESS documentation for your DBMS for more information. △


Benefits of Using SAS Views

SAS views provide the following benefits:

- Instead of using multiple DATA steps to merge SAS data sets by common variables, you can construct a SAS view that performs a multi-table join.
- You can save disk space by storing a SAS view definition, which stores only the instructions for where to find the data and how it is formatted, not the actual data.
- SAS views can ensure that the input data sets are always current because data is derived from SAS views at execution time.
- Since SAS views can select data from many sources, once a SAS view is created, it can provide prepackaged information to the information community without the need for additional programming.
- SAS views can reduce the impact of data design changes on users. For example, you can change a query that is stored in a SAS view without changing the characteristics of the view's result.
- With SAS/CONNECT software, a SAS view can join SAS data sets that reside on different host computers, presenting you with an integrated view of distributed company data.

The following figure shows native and interface SAS views and their relationship to SAS data files:

Figure 28.1 Native and Interface SAS Views



You can use SAS views in the following ways:

- as input to other DATA steps or PROC steps

- to migrate data to SAS data files or to database management systems that are supported by SAS
- in combination with other data sources using PROC SQL
- as pre-assembled sets of data for users of SAS/ASSIST software, enabling them to perform data management, analysis, and reporting tasks regardless of how the data is stored

When to Use SAS Views

Consider the following in order to determine whether a SAS data file or a SAS view is better for your purposes:

- Data files use additional disk space; SAS views use additional processing time.
- Data file variables can be sorted and indexed before using; SAS views must process data in its existing form during execution.

DATA Step Views

Definition of a DATA Step View

A *DATA step view* is a native view that has the broadest scope of any SAS view. It contains a stored DATA step program that can read data from a variety of sources, including:

- raw data files
- SAS data files
- PROC SQL views
- SAS/ACCESS views
- DB2, ORACLE, or other DBMS data

Creating DATA Step Views

In order to create a DATA step view, specify the `VIEW=` option after the final data set name in the DATA statement. The `VIEW=` option tells SAS to compile, but not to execute, the source program and to store the compiled code in the input DATA step view that is named in the option.

For example, the following statements create a DATA step view named DEPT.A:

```
libname dept 'SAS-library';

data dept.a / view=dept.a;
    ... more SAS statements ...
run;
```

Note that if the SAS view exists in a SAS library, and if you use the same member name to create a new view definition, then the old SAS view is overwritten.

Beginning with Version 8, DATA step views retain source statements. You can retrieve these statements using the DESCRIBE statement. The following example uses the DESCRIBE statement in a DATA step view in order to write a copy of the source code to the SAS log:

```

data view=inventory;
    describe;
run;

```

For more information on how to create SAS views and use the DESCRIBE statement, see the DATA statement in *SAS Language Reference: Dictionary*.

What Can You Do with a DATA Step View?

Using a DATA step view, you can do the following:

- directly process any file that can be read with an INPUT statement
- read other SAS data sets
- generate data without using any external data sources and without creating an intermediate SAS data file.

Because DATA step views are generated by the DATA step, they can manipulate and manage input data from a variety of sources including data from external files and data from existing SAS data sets. The scope of what you can do with a DATA step view, therefore, is much broader than that of other types of SAS views.

Differences between DATA Step Views and Stored Compiled DATA Step Programs

DATA step views and stored compiled DATA step programs differ in the following ways:

- a DATA step view is implicitly executed when it is referenced as an input data set by another DATA or PROC step. Its main purpose is to provide data, one record at a time, to the invoking procedure or DATA step.
- a stored compiled DATA step program is explicitly executed when it is specified by the PGM= option on a DATA statement. Its purpose is usually a more specific task, such as creating SAS data files, or originating a report.

For more information on stored compiled DATA step programs, see Chapter 29, “Stored Compiled DATA Step Programs,” on page 539.

Restrictions and Requirements

Global statements do not apply to a DATA step view. Global statements such as the FILENAME, FOOTNOTE, LIBNAME, OPTIONS, and TITLE statements, even if included in the DATA step that created the SAS view, have no effect on the SAS view. If you do include global statements in your source program statements, SAS stores the DATA step view but not the global statements. When the view is referenced, actual execution can differ from the intended execution.

When a view is created, the labels for the variable that it returns are also created. If a DATA step view reads a data set that contains variable labels and a label is changed after the view is created, any procedure output will show the original labels. The view must be recompiled in order for the procedure output to reflect the new variable labels.

If a view uses filerefs or librefs, the fileref or libref that is used is the one that is defined at the time that the view is compiled. This means that if you change the file that is referenced in a fileref that the view uses, the new file is ignored by the view and the file that is referred to by the fileref at the time the view was compiled continues to be used.

Performance Considerations

- DATA step code executes each time that you use a DATA step view, which might add considerable system overhead. In addition, you run the risk of having your data change between steps. However, this also means that you get the most recent data available—that is, data when the view is executed as compared to data when the view was compiled.
- Depending on how many reads or passes on the data are required, processing overhead increases.
 - When one sequential pass is requested, no data set is created. Compared to traditional methods of processing, making one pass improves performance by decreasing the number of input/output operations and elapsed time.
 - When random access or multiple passes are requested, the SAS view must build a spill file that contains all generated observations so that subsequent passes can read the same data that was read by previous passes. In some instances, the view SPILL= data set option can reduce the size of a spill file.

Example 1: Merging Data to Produce Reports

If you want to merge data from multiple files but you do not need to create a file that contains the combined data, you can create a DATA step view of the combination for use in subsequent applications.

For example, the following statements define DATA step view MYV9LIB.QTR1, which merges the sales figures in the data file V9LR.CLOTHES with the sales figures in the data file V9LR.EQUIP. The data files are merged by date, and the value of the variable **Total** is computed for each date.

```
libname myv9lib 'SAS-library';
libname v9lr 'SAS-library';

data myv9lib.qtr1 / view=myv9lib.qtr1;
  merge v9lr.clothes v9lr.equip;
  by date;
  total = cl_v9lr + eq_v9lr;
run;
```

The following PRINT procedure executes the view:

```
proc print data=myv9lib.qtr1;
run;
```

Example 2: Producing Additional Output Files

In this example, the DATA step reads an external file named STUDENT, which contains student data, and then writes observations that contain known problems to data set MYV9LIB.PROBLEMS. The DATA step also defines the DATA step view MYV9LIB.CLASS. The DATA step does *not* create a SAS data file named MYV9LIB.CLASS.

The FILENAME and the LIBNAME statements are both global statements and must exist outside of the code that defines the SAS view, because SAS views cannot contain global statements.

Here are the contents of the external file STUDENT:

```

dutterono MAT 3
lyndenall MAT
frisbee MAT 94
          SCI 95
zymeco ART 96
dimette 94
mesipho SCI 55
merlbeest ART 97
scafernia 91
gilhoolie ART 303
misqualle ART 44
xyлотone SCI 96

```

Here is the DATA step that produces the output files:

```

libname myv9lib 'SAS-library';
filename student 'external-file-specification'; ❶

data myv9lib.class(keep=name major credits)
  myv9lib.problems(keep=code date) / view=myv9lib.class; ❷
infile student;
  input name $ 1-10 major $ 12-14 credits 16-18; ❸
select;
when (name=' ' or major=' ' or credits=.)
  do code=01;
    date=datetime();
    output myv9lib.problems;
  end; ❹
when (0<credits<90)
  do code=02;
    date=datetime();
    output myv9lib.problems;
  end; ❺
otherwise
  output myv9lib.class;
end;
run; ❻

```

The following example shows how to print the files created previously. The MYV9LIB.CLASS contains the observations from STUDENT that were processed without errors. The data file MYV9LIB.PROBLEMS contains the observations that contain errors.

If the data frequently changes in the source data file STUDENT, there would be different effects on the returned values in the SAS view and the SAS data file:

- New records, if error free, that are added to the source data file STUDENT between the time you run the DATA step in the previous example and the time you execute PROC PRINT in the following example, will appear in the SAS view MYV9LIB.CLASS.
- On the other hand, if any new records, failing the error tests, were added to STUDENT, the new records would not show up in the SAS data file MYV9LIB.PROBLEM, until you run the DATA step again.

A SAS view dynamically updates from its source files each time it is used. A SAS data file, each time it is used, remains the same, unless new data is written directly to the file.

```

filename student 'external-file-specification';
libname myv9lib 'SAS--library'; ⑦

proc print data=myv9lib.class;
run; ⑧

proc print data=myv9lib.problems;
    format date datetime18.;
run; ⑨

```

- ① Reference a library called MYV9LIB. Tell SAS where a file that associated with the fileref STUDENT is stored.
- ② Create a data file called PROBLEMS and a SAS view called CLASS and specify the column names for both data sets.
- ③ Select the file that is referenced by the fileref STUDENT and select the data in character format that resides in the specified positions in the file. Assign column names.
- ④ When data in the column NAME, MAJOR, or CREDITS is blank or missing, assign a code of 01 to the observation where the missing value occurred. Also assign a SAS datetime code to the error and place the information in a file called PROBLEMS.
- ⑤ When the amount of credits is greater than zero, but less than ninety, list the observations as code 02 in the file called PROBLEMS and assign a SAS datetime code to the observation.
- ⑥ Place all other observations, which have none of the specified errors, in the SAS view called MYV9LIB.CLASS.
- ⑦ The FILENAME statement assigns the fileref STUDENT to an external file. The LIBNAME statement assigns the libref MYV9LIB to a SAS library.
- ⑧ The first PROC PRINT calls the SAS view MYV9LIB.CLASS. The SAS view extracts data on the fly from the file referenced as STUDENT.
- ⑨ This PROC PRINT prints the contents of the data file MYV9LIB.PROBLEMS.

PROC SQL Views

A PROC SQL view is a PROC SQL query-expression that is given a name and stored for later use. When you use a PROC SQL view in a SAS program, the view derives its data from the data sets (often referred to as tables) or SAS views listed in its FROM clause. The data that is accessed by the view is a subset or superset of the data in its underlying data sets or SAS views.

A PROC SQL view can read or write data from:

- DATA step views
- SAS data files
- other PROC SQL views
- SAS/ACCESS views
- DB2, ORACLE, or other DBMS data.

For complete documentation on how to create and use PROC SQL views, see *Base SAS Procedures Guide*.

For information on using PROC SQL views created in an earlier release, see Chapter 34, “SAS 9.2 Compatibility with SAS Files from Earlier Releases,” on page 577.

Comparing DATA Step and PROC SQL Views

To help you decide between a DATA step view and a PROC SQL view, consider the characteristics of each type of SAS view:

- DATA step views
 - DATA step views are versatile because they use DATA step processing, including DO loops and IF-THEN/ELSE statements.
 - DATA step views do not have update capability; that is, they cannot directly change the data that they access.
 - There is no way to qualify the data in a DATA step view before using it. Therefore, even if you need only part of the data in your SAS view, you must load into memory the entire DATA step view and discard everything that you do not need.

- PROC SQL views
 - PROC SQL views can combine data from many different file formats.
 - PROC SQL views can both read and update the data that they reference.
 - PROC SQL supports more types of WHERE clauses than are available in DATA step processing and has a CONNECT TO component that enables you to easily send SQL statements and pass data to a DBMS by using the Pass-Through Facility.
 - You can also use the SQL language to subset your data before processing it. This capability saves memory when you have a large SAS view, but need to select only a small portion of the data contained in the view.
 - PROC SQL views do not use DATA step programming.

SAS/ACCESS Views

A SAS/ACCESS view is an interface view, also called a *view descriptor*, which accesses DBMS data that is defined in a corresponding *access descriptor*.

Using SAS/ACCESS software, you can create an access descriptor and one or more view descriptors in order to define and access some or all of the data described by one DBMS table or DBMS view. You can also use view descriptors in order to update DBMS data, with certain restrictions.


In addition, some SAS/ACCESS products provide a dynamic LIBNAME engine interface. If available, it is recommended that you use SAS/ACCESS LIBNAME statement to assign a SAS libref to your DBMS data because it is more efficient and easier to use than access descriptors and view descriptors. The SAS/ACCESS dynamic LIBNAME engine enables you to treat DBMS data as if it were SAS data by assigning a SAS libref to DBMS objects. Using a SAS/ACCESS dynamic LIBNAME engine means that you can use both native DATA step views and native PROC SQL views to access DBMS data instead of view descriptors.

See Chapter 32, “About SAS/ACCESS Software,” on page 561 or the SAS/ACCESS documentation for your database for more information about SAS/ACCESS features.

For information on using SAS/ACCESS view descriptors created in an earlier release, see Chapter 34, “SAS 9.2 Compatibility with SAS Files from Earlier Releases,” on page 577.

Note: Starting in SAS 9, PROC SQL views are the preferred way to access relational DBMS data. You can convert existing SAS/ACCESS view descriptors into

PROC SQL views by using the CV2VIEW procedure, enabling you to use the LIBNAME statement to access your data. See the CV2VIEW Procedure in *SAS/ACCESS for Relational Databases: Reference*. Δ



CHAPTER 29

Stored Compiled DATA Step Programs

<i>Definition of a Stored Compiled DATA Step Program</i>	539
<i>Uses for Stored Compiled DATA Step Programs</i>	539
<i>Restrictions and Requirements for Stored Compiled DATA Step Programs</i>	540
<i>How SAS Processes Stored Compiled DATA Step Programs</i>	540
<i>Creating a Stored Compiled DATA Step Program</i>	541
<i>Syntax for Creating a Stored Compiled DATA Step Program</i>	541
<i>Process to Compile and Store a DATA Step Program</i>	541
<i>Example: Creating a Stored Compiled DATA Step Program</i>	542
<i>Executing a Stored Compiled DATA Step Program</i>	542
<i>Syntax for Executing a Stored Compiled DATA Step Program</i>	542
<i>Process to Execute a Stored Compiled DATA Step Program</i>	543
<i>Using Global Statements</i>	544
<i>Redirecting Output</i>	544
<i>Printing the Source Code of a Stored Compiled DATA Step Program</i>	544
<i>Example: Executing a Stored Compiled DATA Step Program</i>	545
<i>Differences between Stored Compiled DATA Step Programs and DATA Step Views</i>	546
<i>Examples of DATA Step Programs</i>	546
<i>Example of DATA Step Program: Quality Control Application</i>	546

Definition of a Stored Compiled DATA Step Program

A *stored compiled DATA step program* is a SAS file that contains a DATA step program that has been compiled and then stored in a SAS library. You can execute stored compiled programs as needed, without having to recompile them. Stored compiled DATA step programs are of member type PROGRAM.

Note: Stored compiled programs are available for DATA step applications only. Your stored programs can contain all SAS language elements except global statements. If you do include global statements in your source program, SAS stores the compiled program but not the global statements, and does not display a warning message in the SAS log. △

Uses for Stored Compiled DATA Step Programs

The primary use of stored compiled DATA step programs is for executing production jobs. The advantage of using these DATA step programs is that you can execute them as needed without investing the resources required for repeated compilation. The

savings are especially significant if the DATA step contains many statements. If you install a new version of SAS, you do not need to recompile your source code.

Restrictions and Requirements for Stored Compiled DATA Step Programs

The following restrictions and requirements apply for using stored compiled DATA step programs:

- Stored compiled DATA step programs are available for DATA step applications only.
- Stored compiled DATA step program cannot contain global statements. If you do include global statements such as FILENAME, FOOTNOTE, LIBNAME, OPTIONS, and TITLE in your source program, SAS stores the compiled program but not the global statements. SAS does not display a warning message in the SAS log.
- SAS does not store raw data in the compiled program.

Operating Environment Information: You cannot move a compiled program to an operating environment that has an incompatible computer architecture. You must, instead, recompile your source code and store your new compiled program.


You can, however, move your compiled program to a different host computer that has a compatible architecture. Δ

How SAS Processes Stored Compiled DATA Step Programs

You first compile the SAS source program and store the compiled code. Then you execute the compiled code, redirecting the input and output as necessary.


SAS processes the DATA step through the compilation phase and then stores an intermediate code representation of the program and associated data tables in a SAS file. SAS processes the intermediate code when it executes the stored program. The following figure shows the process for creating a stored compiled DATA step program.

Figure 29.1 Creating a Stored Compiled Program



When SAS executes the stored program, it resolves the intermediate code produced by the compiler and generates the executable machine code for that operating environment. The following figure shows the process for executing a stored compiled DATA step program.

Figure 29.2 Executing a Stored Compiled Program



To move, copy, rename, or delete stored programs, use the DATASETS procedure or the utility windows in your windowing environment.

Creating a Stored Compiled DATA Step Program

Syntax for Creating a Stored Compiled DATA Step Program

The syntax for creating a stored compiled DATA step program is as follows:

```
DATA data-set-name(s) / PGM=stored-program-name
      <(<password-option><SOURCE=source-option>>);
```

where

data-set-name

specifies a valid SAS name for the output data set created by the source program. The name can be a one-level name or a two-level name. You can specify more than one data set name in the DATA statement.

stored-program-name

specifies a valid SAS name for the SAS file containing the stored program. The name can be a one-level name, but it is usually a two-level name. Stored programs are assigned the member type PROGRAM in the SAS library.

password-option

assigns a password to a stored compiled DATA step program.

source-option

enables you to save or encrypt the source code.

For complete information about the DATA statement, see *SAS Language Reference: Dictionary*.

Process to Compile and Store a DATA Step Program

To compile and store a DATA step program, do the following:

- 1 Write, test, and debug the DATA step program you want to store.

If you are reading external raw data files or if you write raw data to an external file, use a fileref rather than the actual filename in your INFILE and FILE statements so that you can redirect your input and output when the stored program executes.

- 2 When the program runs correctly, submit it using the PGM= option in the DATA statement.

The PGM= option tells SAS to compile, but not execute, the program and to store the compiled code in the SAS file named in the option. SAS sends a message to the log when the program is stored.

Note: The default SOURCE=SAVE or SOURCE=ENCRYPT options automatically save your source code. △

Note: If you move your application to another operating environment, you need to recompile your source code and store your new compiled program. △

Example: Creating a Stored Compiled DATA Step Program

The following example uses the information in the input SAS data set IN.SAMPLE to assign a plant type based on a plant code. Note that the global LIBNAME statements are necessary to identify the storage location for your files, but are not part of STORED.SAMPLE, the DATA step that SAS stores.

```
libname in 'SAS-library';
libname stored 'SAS-library';

data out.sample / pgm=stored.sample;
  set in.sample;
  if code = 1 then
    do;
      Type='Perennial';
      number+4;
    end;
  else
  if code = 2 then
    do;
      Type='Annual';
      number+10;
    end;
  else
    do;
      Type='ERROR';
      Number=0;
    end;
  run;
```

Output 29.1 Partial SAS Log Identifying the Stored DATA Step Program

```
.
.
.
NOTE: DATA STEP program saved on file STORED.SAMPLE.
NOTE: A stored DATA STEP program cannot run under a different operating system.
NOTE: DATA statement used:
      real time      1.14 seconds
      cpu time       0.12 seconds
```

Executing a Stored Compiled DATA Step Program

Syntax for Executing a Stored Compiled DATA Step Program

The syntax for executing a stored compiled DATA step program, retrieving source code, and redirecting input or output, is as follows:

global SAS statements

```
DATA PGM=stored-program-name <(password-option)>;
  <DESCRIBE;>
  <REDIRECT INPUT | OUTPUT old-name-1 = new-name-1<. . . old-name-n =
    new-name-n>;>
  <EXECUTE;>
```

where

global SAS statements

specifies any global SAS statements that are needed by the program when it executes, such as a FILENAME or a LIBNAME statement that points to input files or routes output.

stored-program-name

specifies a valid SAS name for the SAS file containing the stored program. The name can be a one-level name or a two-level name.

password-option

specifies a password that you use to access the stored compiled DATA step program.

DESCRIBE

is a SAS statement that retrieves source code from a stored compiled DATA step program or a DATA step view.

INPUT | OUTPUT

specifies whether you are redirecting input or output data sets. When you specify INPUT, the REDIRECT statement associates the name of the input data set in the source program with the name of another SAS data set. When you specify OUTPUT, the REDIRECT statement associates the name of the output data set with the name of another SAS data set.

old-name

specifies the name of the input or output data set in the source program.

new-name

specifies the name of the input or output data set that you want SAS to process for the current execution.

EXECUTE

is a SAS statement that executes a stored compiled DATA step program.

For complete information about the DATA statement, see *SAS Language Reference: Dictionary*.

Process to Execute a Stored Compiled DATA Step Program

To execute a stored compiled DATA step program, follow these steps:

- 1 Write a DATA step for each execution of the stored program. In this DATA step, specify the name of the stored program in the PGM= option of the DATA statement and include an optional password. You can
 - submit this DATA step as a separate program
 - include it as part of a larger SAS program that can include other DATA and procedure (PROC) steps
 - point to different input and output SAS data sets each time you execute the stored program by using the REDIRECT statement.

- 2 Submit the DATA steps. Be sure to end each one with a RUN statement or other step boundary.

Using Global Statements

You can use global SAS statements such as FILENAME or LIBNAME when you store or execute a stored compiled DATA step program. However, the global statements that you use to compile and store a DATA step program are not stored with the DATA step code.

Redirecting Output

You can redirect external files using filerefs. You can use the REDIRECT statement for renaming input and output SAS data sets.

You can use the REDIRECT statement to redirect input and output data to data sets you specify. Note that the REDIRECT statement is available only for use with stored compiled DATA step programs.

Note: To redirect input and output stored in external files, include a FILENAME statement at execution time to associate the fileref in the source program with different external files. \triangle

CAUTION:

Use caution when you redirect input data sets. The number and attributes of variables in the input SAS data sets that you read with the REDIRECT statement should match those of the input data sets in the SET, MERGE, MODIFY, or UPDATE statements of the source code. If they do not match, the following occurs:

- If the variable length attributes differ, the length of the variable in the source code data set determines the length of the variable in the redirected data set.
- If extra variables are present in the redirected data sets, the stored program stops processing and an error message is sent to the SAS log.
- If the variable type attributes are different, the stored program stops processing and an error message is sent to the SAS log.

\triangle

Printing the Source Code of a Stored Compiled DATA Step Program

If you use both the DESCRIBE and the EXECUTE statements when you execute a stored compiled DATA step program, SAS writes the source code to the log. The following example executes a stored compiled DATA step program. The DESCRIBE statement in the program writes the source code to the SAS log.

```
data pgm=stored.sample;
  describe;
  execute;
run;
```

Output 29.2 Partial SAS Log Showing the Source Code Generated by the DESCRIBE Statement

```

.
.
.
26
27   data pgm=stored.sample;
28       describe;
29       execute;
30   run;
NOTE: DATA step stored program STORED.SAMPLE is defined as:

data out.sample / pgm=stored.sample;
  set in.sample;
  if code = 1 then
    do;
      Type='Perennial';
      number+4;
    end;
  else
    if code = 2 then
      do;
        Type='Annual';
        number+10;
      end;
    else
      do;
        Type='ERROR';
        Number=0;
      end;
run;

NOTE: DATA STEP program loaded from file STORED.SAMPLE.
NOTE: There were 7 observations read from the data set IN.SAMPLE.
NOTE: The data set OUT.SAMPLE has 7 observations and 4 variables.
NOTE: DATA statement used:
      real time          0.80 seconds
      cpu time           0.15 seconds

```

For more information about the DESCRIBE statement, see *SAS Language Reference: Dictionary*.

Example: Executing a Stored Compiled DATA Step Program

The following DATA step executes the stored program STORED.SAMPLE created in “Example: Creating a Stored Compiled DATA Step Program” on page 542. The REDIRECT statement specifies the source of the input data as BASE.SAMPLE. The output from this execution of the program is redirected and stored in a data set named TOTALS.SAMPLE. Output 29.3 shows part of the SAS log.

```

libname in 'SAS-library';
libname base 'SAS-library';
libname totals 'SAS-library';
libname stored 'SAS-library';

data pgm=stored.sample;
  redirect input in.sample=base.sample;
  redirect output out.sample=totals.sample;
run;

```

Output 29.3 Partial SAS Log Identifying the Redirected Output File

```

      cpu time          0.00 seconds
.
.
.
6
7  data pgm=stored.sample;
8      redirect input in.sample=base.sample;
9      redirect output out.sample=totals.sample;
10 run;
NOTE: DATA STEP program loaded from file STORED.SAMPLE.
NOTE: The data set TOTALS.SAMPLE has 7 observations and 4 variables.
NOTE: DATA statement used:
      real time          0.67 seconds

```

Differences between Stored Compiled DATA Step Programs and DATA Step Views

Stored compiled DATA step programs and DATA step views are similar in function. They both store DATA step programs that can retrieve and process data stored in other files. Both have the same restrictions and requirements (see “Restrictions and Requirements for Stored Compiled DATA Step Programs” on page 540). For information about DATA step views, see Chapter 28, “SAS Views,” on page 529.

Stored compiled DATA step programs and DATA step views differ in the following ways:

- A stored compiled DATA step program is explicitly executed when it is specified by the PGM= option on a DATA statement. The stored compiled DATA step is used primarily in production jobs.
- A DATA step view is implicitly executed when the view is referenced as an input data set by another DATA or procedure (PROC) step. Its main purpose is to provide data one record at a time to the invoking procedure or DATA step.
- You can use the REDIRECT statement when you execute a stored compiled DATA step. You cannot use this statement with DATA step views.

Examples of DATA Step Programs

Example of DATA Step Program: Quality Control Application

This example illustrates how to use a stored compiled DATA step program for a simple quality control application. This application processes several raw data files. The source program uses the fileref DAILY in the INFILE statement. Each DATA step that is used to execute the stored program can include a FILENAME statement to associate the fileref DAILY with a different external file.

The following statements compile and store the program:

```

libname stored 'SAS-library-1';

data flaws / pgm=stored.flaws;
  length Station $ 15;

```

```

infile daily;
input Station $ Shift $ Employee $ NumberOfFlaws;
TotalNumber + NumberOfFlaws;
run;

```

The following statements execute the stored compiled program, redirect the output, and print the results:

```

libname stored 'SAS-library-1';
libname testlib 'SAS-library-2';

data pgm=stored.flaws;
  redirect output flaws=testlib.daily;
run;


proc print data=testlib.daily;
  title 'Quality Control Report';
run;

```

Output 29.4 Quality Control Application Output

Quality Control Report						1
Obs	Station	Shift	Employee	Number OfFlaws	Total Number	
1	Cambridge	1	Lin	3	3	
2	Northampton	1	Kay	0	3	
3	Springfiled	2	Sam	9	12	

Note that you can use the TITLE statement when you execute a stored compiled DATA step program or when you print the results.



CHAPTER

30

DICTIONARY Tables

<i>Definition of a DICTIONARY Table</i>	549
<i>How to View DICTIONARY Tables</i>	549
<i>How to View a DICTIONARY Table</i>	550
<i>How to View a Summary of a DICTIONARY Table</i>	550
<i>How to View a Subset of a DICTIONARY Table</i>	551
<i>DICTIONARY Tables and Performance</i>	551

Definition of a DICTIONARY Table

A DICTIONARY table is a read-only SAS view that contains information about SAS libraries, SAS data sets, SAS macros, and external files that are in use or available in the current SAS session. A DICTIONARY table also contains the settings for SAS system options that are currently in effect.

When you access a DICTIONARY table, SAS determines the current state of the SAS session and returns the desired information accordingly. This process is performed each time a DICTIONARY table is accessed, so that you always have current information.

DICTIONARY tables can be accessed by a SAS program by using either of these methods:

- run a PROC SQL query against the table, using the DICTIONARY libref
- use any SAS procedure or the DATA step, referring to the PROC SQL view of the table in the SASHELP library.

For more information on DICTIONARY tables, including a list of available DICTIONARY tables and their associated SASHELP views, see *SAS 9.2 SQL Procedure User's Guide*.

How to View DICTIONARY Tables

You might want to view the contents of DICTIONARY tables in order to see information about your current SAS session, before actually using the table in a DATA step or a SAS procedure.

Some DICTIONARY tables can become quite large. In this case, you might want to view a part of a DICTIONARY table that contains only the data that you are interested in. The best way to view part of a DICTIONARY table is to subset the table using a PROC SQL WHERE clause.

How to View a DICTIONARY Table

Each DICTIONARY table has an associated PROC SQL view in the SASHELP library. You can see the entire contents of a DICTIONARY table by opening its SASHELP view with the VIEWTABLE or FSVIEW utilities. This method provides more detail than you receive in the output of the DESCRIBE TABLE statement, as shown in “How to View a Summary of a DICTIONARY Table” on page 550.

The following steps describe how to use the VIEWTABLE or FSVIEW utilities to view a DICTIONARY table in a windowing environment.

- 1 Invoke the Explorer window in your SAS session.
- 2 Select the SASHELP library. A list of members in the SASHELP library appears.
- 3 Select a SAS view with a name that starts with V, for example, VMEMBER. A VIEWTABLE window appears that contains its contents. (For z/OS, type the letter 'O' in the command field for the desired member and press ENTER. The FSVIEW window appears with the contents of the view.)

In the VIEWTABLE window the column headings are labels. To see the column names, select **View ► Column Names**.

How to View a Summary of a DICTIONARY Table

The DESCRIBE TABLE statement in PROC SQL produces a summary of the contents of a DICTIONARY table. The following example uses the DESCRIBE TABLE statement in order to generate a summary for the table DICTIONARY.INDEXES. (The SASHELP view for this table is SASHELP.VINDEX).

```
proc sql;
  describe table dictionary.indexes;
```

The result of the DESCRIBE TABLE statement appears in the SAS log:

NOTE: SQL table DICTIONARY.INDEXES was created like:

```
create table DICTIONARY.INDEXES
(
  libname char(8) label='Library Name',
  memname char(32) label='Member Name',
  memtype char(8) label='Member Type',
  name char(32) label='Column Name',
  idxusage char(9) label='Column Index Type',
  indxname char(32) label='Index Name',
  indxpos num label='Position of Column in Concatenated Key',
  nomiss char(3) label='Nomiss Option',
  unique char(3) label='Unique Option'
);
```

- The first word on each line is the column (or variable) name, the name that you need to use when writing a SAS statement that refers to the column (or variable).
- Following the column name is the specification for the type of variable and the width of the column.
- The name that follows **label=** is the column (or variable) label.

After you know how a table is defined, you can use the processing ability of the PROC SQL WHERE clause in a PROC SQL step in order to extract a portion of a SAS view.

How to View a Subset of a DICTIONARY Table

When you know that you are accessing a large DICTIONARY and you need to use only a portion of it, use a PROC SQL WHERE clause in order to extract a subset of the original. The following PROC SQL statement demonstrates the use of a PROC SQL WHERE clause in order to extract lines from DICTIONARY.INDEXES.

```
proc sql;
    title 'Subset of the DICTIONARY.INDEX View';
    title2 'Rows with Column Name equal to STATE';
    select * from dictionary.indexes
        where name = 'STATE';
quit;
```

The results are shown in the following output:

Output 30.1 Result of the PROC SQL Subsetting WHERE Statement

Subset of the DICTIONARY.INDEX View Rows with Column Name equal to STATE						
Library Name	Member Name	Member Type	Column Name	Column Index Type	Index Name	
Position of Column in Concatenated Key		Nomiss	Unique			
		Key	Option	Option		
MAPS	USAAC0	DATA	STATE	COMPOSITE	SC000000	
MAPS	USAAC8	DATA	STATE	COMPOSITE	CS000000	
MAPS	USAAS.	DATA	STATE	SIMPLE	STATE	

Note that many character values in the DICTIONARY tables are stored as all-uppercase characters; you should design your queries accordingly.

DICTIONARY Tables and Performance

When you query a DICTIONARY table, SAS gathers information that is pertinent to that table. Depending on the DICTIONARY table that is being queried, this process can include searching libraries, opening tables, and executing SAS views. Unlike other SAS procedures and the DATA step, PROC SQL can improve this process by optimizing the query before the select process is launched. Therefore, although it is possible to access DICTIONARY table information with SAS procedures or the DATA step by using the SASHELP views, it is often more efficient to use PROC SQL instead.

For example, the following programs both produce the same result, but the PROC SQL step runs much faster because the WHERE clause is processed prior to opening the tables that are referenced by the SASHELP.VCOLUMN view:


```
data mytable;
    set sashelp.vcolumn;
    where libname='WORK' and memname='SALES';
```

```
run;

proc sql;
  create table mytable as
    select * from sashelp.vcolumn
    where libname='WORK' and memname='SALES';
quit;
```

Note: SAS does not maintain DICTIONARY table information between queries. Each query of a DICTIONARY table launches a new discovery process. Δ

If you are querying the same DICTIONARY table several times in a row, you can get even faster performance by creating a temporary SAS data set (with the DATA step SET statement or PROC SQL CREATE TABLE AS statement) with the information that you desire and run your query against that data set.



CHAPTER

31

SAS Catalogs

<i>Definition of a SAS Catalog</i>	553
<i>SAS Catalog Names</i>	553
<i>Parts of a Catalog Name</i>	553
<i>Accessing Information in Catalogs</i>	554
<i>Tools for Managing SAS Catalogs</i>	554
<i>Profile Catalog</i>	555
<i>Definition of a Profile Catalog</i>	555
<i>How the Information Is Used</i>	555
<i>How Sasuser.Profile Is Created</i>	555
<i>Default Settings</i>	555
<i>How to Recover Locked or Corrupt Profile Catalogs</i>	556
<i>Catalog Concatenation</i>	557
<i>Definitions</i>	557
<i>Example 1: LIBNAME Catalog Concatenation</i>	557
<i>Example 2: CATNAME Catalog Concatenation</i>	558
<i>Rules for Catalog Concatenation</i>	560

Definition of a SAS Catalog

SAS catalogs are special SAS files that store many different kinds of information in smaller units called catalog entries. Each entry has an entry type that identifies its purpose to SAS. A single SAS catalog can contain several types of catalog entries. Some catalog entries contain system information such as key definitions. Other catalog entries contain application information such as window definitions, help windows, formats, informats, macros, or graphics output. You can list the contents of a catalog using various SAS features, such as SAS Explorer and PROC CATALOG.

SAS Catalog Names

Parts of a Catalog Name

SAS catalog entries are fully identified by a four-level name in the following form:

libref.catalog.entry-name.entry-type

You commonly specify the two-level name for an entire catalog, as follows:

libref.catalog

libref

is the logical name of the SAS library to which the catalog belongs.

catalog

is a valid SAS name for the file.

The entry name and entry type are required by some SAS procedures. If the entry type has been specified elsewhere or if it can be determined from context, you can use the entry name alone. To specify entry names and entry types, use this form:

entry-name.entry-type

entry-name

is a valid SAS name for the catalog entry.

entry-type

is assigned by SAS when the entry is created.

Accessing Information in Catalogs

In Base SAS software, SAS catalog entries are generally accessed automatically by SAS when the information stored in them is required for processing. In other SAS software products, you must specify the catalog entry in various procedures. Because the requirements differ with the SAS procedure or software product, refer to the appropriate procedure or product documentation for details.

Tools for Managing SAS Catalogs

There are several SAS features to help you manage the entries in catalogs. The CATALOG procedure and the CEXIST function are two features of Base SAS software. Another tool is SAS Explorer, which enables you to view the contents of SAS catalogs. Many interactive windowing procedures contain a catalog directory window for managing entries. The following list summarizes the tools that are available for managing catalogs:

CATALOG procedure

is similar to the DATASETS procedure. Use the CATALOG procedure to copy, delete, list, and rename entries in catalogs.

CEXIST function

enables you to verify the existence of a SAS catalog or catalog entry. See the CEXIST function in *SAS Language Reference: Dictionary* for more information.

CATALOG window

is a window that you can access at any time in an interactive windowing environment. It displays the name, type, description, and date of last update for each entry in the specified catalog. CATALOG window commands enable you to edit catalog entries. You can also view and edit catalog entries after double-clicking on a catalog file in SAS Explorer.

catalog directory windows

are available in some procedures in SAS/AF, SAS/FSP, and SAS/GRAPH software. A catalog directory window lists the same kind of information that the CATALOG window provides: entry name, type, description, and date of last update. See the description of each interactive windowing procedure for details about the catalog directory window for that procedure.

Profile Catalog

Definition of a Profile Catalog

Profile catalog (Sasuser.Profile) is a catalog that is available for customizing 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.

How the Information Is Used

The information in the Sasuser.Profile catalog is accessed automatically by SAS when you need it for processing. For example, each time you enter the KEYS window and change the settings, SAS stores the new settings with the KEYS entry type. Similarly, if you change and save the attributes for interactive window procedures, the changes are stored under the appropriate entry name and type. When you use the window or procedure, SAS then looks for information in the Profile catalog.

How Sasuser.Profile Is Created

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

At SAS startup, SAS checks for an existing uncorrupted Sasuser.Profile catalog. If this catalog is found, then SAS copies the Sasuser.Profile catalog to SASUSER.PROFBACK. The backup is used if the Sasuser.Profile catalog becomes corrupted. For more information, see “How to Recover Locked or Corrupt Profile Catalogs” on page 556.

Operating Environment Information: The Sasuser library is implemented differently in various operating environments. See the SAS documentation for your host system for more information about how the SAS user library is created. △

Default Settings

The default settings for your SAS session are stored in several catalogs in the Sashelp installation library. If you do not make any changes to key settings or other options, SAS uses the default settings. If you make changes, the new information is stored in

your Sasuser.Profile catalog. To restore the original default settings, use the CATALOG procedure or the CATALOG window to delete the appropriate entries from your Profile catalog. By default, SAS then uses the corresponding entry from the SasHELP library.

During SAS sessions, you can make customizations, such as window resizing and positioning, and save them to Sasuser.Profile.

How to Recover Locked or Corrupt Profile Catalogs

Occasionally, a Sasuser.Profile catalog becomes locked or corrupted. SAS uses SASHELP.PROFILE and SASUSER.PROFBACK to replace the locked or corrupted catalog.

If your Sasuser.Profile catalog is locked, SAS will check for SASHELP.PROFILE. If SASHELP.PROFILE exists, SAS copies it to WORK.PROFILE and then saves the customizations in Work.Profile instead of in Sasuser.Profile. The following notes will appear in the SAS log:

```
ERROR: Expecting page 1, got page -1 instead.
ERROR: Page validation error while reading SASUSER.PROFILE.CATALOG.
NOTE: Unable to open SASUSER.PROFILE. WORK.PROFILE will be opened instead.
NOTE: SASHELP.PROFILE has been copied to WORK.PROFILE.
NOTE: All profile changes will be lost at the end of the session.
```

If your Sasuser.Profile catalog is corrupted, SAS copies the corrupted catalog to SASUSER.BADPRO. SAS then checks for SASUSER.PROFBACK. If SASUSER.PROFBACK exists, SAS copies it to Sasuser.Profile. Any changes made to the Sasuser.Profile catalog during the previous session will be lost. The following notes will appear in the SAS log:

```
ERROR: Expecting page 1, got page -1 instead.
ERROR: Page validation error while reading SASUSER.PROFILE.CATALOG.
NOTE: A corrupt SASUSER.PROFILE has been detected. A PROFILE catalog can
      become corrupt when a SAS session is prematurely terminated.
NOTE: SASUSER.PROFILE.CATALOG has been renamed to SASUSER.BADPRO.CATALOG.
NOTE: SASUSER.PROFILE.CATALOG has been restored from SASUSER.PROFBAK.CATALOG.
NOTE: Changes made to SASUSER.PROFILE.CATALOG during the previous SAS session
      have been lost. The type of data stored in the PROFILE catalog is
      typically related to SAS session customizations such as key
      definitions, fonts for graphics, and window attributes.
```

If your Sasuser.Profile catalog is corrupted and there is no SASUSER.PROFBACK, SAS will check for SASHELP.PROFILE. If SASHELP.PROFILE exists, SAS copies it to WORK.PROFILE and then saves the customizations in Work.Profile instead of in Sasuser.Profile. The following notes will appear in the SAS log:

```
ERROR: Expecting page 1, got page -1 instead.
ERROR: Page validation error while reading SASUSER.PROFILE.CATALOG.
NOTE: Unable to open SASUSER.PROFILE. WORK.PROFILE will be opened instead.
NOTE: SASHELP.PROFILE has been copied to WORK.PROFILE.
NOTE: All profile changes will be lost at the end of the session.
```

Catalog Concatenation

Definitions

You can logically combine two or more SAS catalogs by concatenating them. This enables you to access the contents of several catalogs, using one catalog name. There are two ways to concatenate catalogs, using the *LIBNAME statement* and *CATNAME statement*.

LIBNAME catalog concatenation

results from a concatenation of libraries through a LIBNAME statement. When two or more libraries are logically combined through concatenation, any catalogs with the same name in each library become logically combined as well.

CATNAME catalog concatenation

is a concatenation that is specified by the global CATNAME statement in which the catalogs to be concatenated are specifically named. During CATNAME catalog concatenation, a logical catalog is set up in memory.

Example 1: LIBNAME Catalog Concatenation

This LIBNAME statement concatenates the two SAS libraries:

```
libname both ('SAS-library 1' 'SAS-library 2');
```

Members of library1	Members of library2
MYCAT.CATALOG	MYCAT.CATALOG
TABLE1.DATA	MYCAT2.CATALOG
TABLE3.DATA	TABLE1.DATA
	TABLE1.INDEX
	TABLE2.DATA
	TABLE2.INDEX

The concatenated libref BOTH would have the following:

Concatenated libref BOTH

MYCAT.CATALOG (from path 1 and 2)

MYCAT2.CATALOG (from path 2)

TABLE1.DATA (from path 1)

TABLE2.DATA (from path 2)

TABLE2.INDEX (from path 2)

TABLE3.DATA (from path 1)

Notice that TABLE1.INDEX does not appear in the concatenation but TABLE2.INDEX does appear. SAS suppresses listing the index when its associated data file is not part of the concatenation.

So what happened to the catalogs when the libraries were concatenated? A resulting catalog now exists logically in memory, with the full name BOTH.MYCAT.CATALOG. This catalog combines each of the two physical catalogs residing in 'library 1' and 'library2', called MYCAT.CATALOG.

To understand the contents of the concatenation BOTH.MYCAT, first look at the contents of both parts of the concatenation. Assume that the two original MYCAT.CATALOG files contain the following:

Contents of MYCAT.CATALOG in library1	Contents of MYCAT.CATALOG in library 2
A.FRAME	A.GRSEG
C.FRAME	B.FRAME
	C.FRAME

Then the combined catalog BOTH.MYCAT contains the following:

BOTH.MYCAT
A.GRSEG (from path 2)
A.FRAME (from path 1)
B.FRAME (from path 2)
C.FRAME (from path 1)

Example 2: CATNAME Catalog Concatenation

The syntax of the CATNAME statement is:

```
CATNAME libref.catref
      (libref-1.catalog-1 (ACCESS=READONLY)
      libref-n.catalog-n (ACCESS=READONLY));
```

To disassociate a concatenated catalog the syntax is:

```
CATNAME libref.catref | _ALL_ clear;
```

In the following example, there must be a libref that is defined and named CATDOG. The libref CATDOG establishes the scope for the CATNAME concatenation definition.

Note: If a file in CATDOG named COMBINED.CATALOG already exists, it becomes inaccessible until the CATNAME concatenation CATDOG.COMBINED is cleared.

\triangle

Members of library1	Members of library2
MYCAT.CATALOG	MYDOG.CATALOG
TABLE1.DATA	MYCAT2.CATALOG

Members of library1	Members of library2
TABLE3.DATA	TABLE1.DATA
	TABLE1.INDEX
	TABLE2.DATA
	TABLE2.INDEX

If we issue the following statement,

```
CATNAME catdog.combined
      (library1.mycat (ACCESS=READONLY)
       library2.mydog (ACCESS=READONLY));
```

then the concatenated catalog CATDOG.COMBINED combines the following catalogs:

Concatenated catalog CATALOG.COMBINED

MYCAT.CATALOG (from library 1)

MYDOG.CATALOG (from library 2)

Note: In CATNAME concatenation only the named catalogs are combined. In LIBNAME concatenation, any catalogs that have the same name in their respective libraries are concatenated when those libraries are concatenated. Δ

The previous CATNAME statement creates a catalog that exists logically in memory. This catalog, named CATDOG.COMBINED.CATALOG, combines the two physical catalogs residing in library1 and library2, called MYCAT.CATALOG and MYDOG.CATALOG respectively.

To understand the contents of the concatenation COMBINED.CATALOG, first look at the contents of both parts of the concatenation. The two original catalog files contain the following entries:

MYCAT.CATALOG	MYDOG.CATALOG
library 1	library 2
A.FRAME	A.GRSEG
C.FRAME	B.FRAME
	C.FRAME

The concatenated catalog COMBINED contains:

COMBINED.CATALOG contents

A.GRSEG (from MYDOG)

A.FRAME (from MYCAT)

COMBINED.CATALOG contents

B.FRAME (from MYDOG)

C.FRAME (from MYCAT)

Rules for Catalog Concatenation

The rules for catalog concatenation are the same, whether the catalogs are concatenated using the LIBNAME statement or the CATNAME statement.


- When a catalog entry is open for input or update, the parts are searched and the first occurrence of the specified entry is used.
- When an item is open for output, it will be created in the catalog that is listed first in the concatenation.

Note: A new catalog entry is created in the first catalog even if there is an item with the same name in another part of the concatenation. Δ

Note: If the first catalog in a concatenation that is opened for update does not exist, the item will be written to the next catalog that exists in the concatenation. Δ

- When you want to delete or rename a catalog entry, only the first occurrence of the entry is affected.
- Any time a list of catalog entries is displayed, only one occurrence of the catalog entry is shown.

Note: Even if a catalog entry occurs multiple times in the concatenation, only the first occurrence is shown. Δ



CHAPTER 32

About SAS/ACCESS Software

<i>Definition of SAS/ACCESS Software</i>	561
<i>Dynamic LIBNAME Engine</i>	561
<i>SAS/ACCESS LIBNAME Statement</i>	561
<i>Using Data Set Options with SAS/ACCESS Librefs</i>	562
<i>Embedding a SAS/ACCESS LIBNAME Statement in a PROC SQL View</i>	562
<i>SQL Procedure Pass-Through Facility</i>	563
<i>ACCESS Procedure and Interface View Engine</i>	564
<i>DBLOAD Procedure</i>	565
<i>Interface DATA Step Engine</i>	565

Definition of SAS/ACCESS Software

SAS/ACCESS software

enables you to read and write data to and from other vendors' *database management systems (DBMS)*, as well as from some *PC file formats*. Depending on your DBMS, a SAS/ACCESS product might provide one or more of the following:

- a dynamic LIBNAME engine
- the SQL Pass-Through Facility
- the ACCESS procedure and interface view engine
- the DBLOAD procedure
- an interface DATA step engine.

These interfaces are described in this section. Each SAS/ACCESS product provides one or more of these interfaces for each supported DBMS. See Chapter 36, "SAS Engines," on page 589 for more information about SAS engines.

Note: To use the SAS/ACCESS features described in this section, you must license SAS/ACCESS software. See the SAS/ACCESS documentation for your DBMS for full documentation of the features described in this section. △

Dynamic LIBNAME Engine

SAS/ACCESS LIBNAME Statement

Beginning in Version 7, you can associate a SAS libref directly with a database, schema, server, or group of tables and SAS views, depending on your DBMS. To assign

a libref to DBMS data, you must use the SAS/ACCESS LIBNAME statement, which has syntax and options that are different from the Base SAS LIBNAME statement. For example, to connect to an ORACLE database, you might use the following SAS/ACCESS LIBNAME statement:

```
libname mydblib oracle user=smith password=secret
      path='myoracleserver';
```

This LIBNAME statement connects to ORACLE by specifying the ORACLE connection options: USER=, PASSWORD=, and PATH=. In addition to the connection options, you can specify SAS/ACCESS LIBNAME options that control the type of database connection that is made. You can use additional options to control how your data is processed.

You can use a DATA step, SAS procedures, or the Explorer window to view and update the DBMS data associated with the libref, or use the DATASETS and CONTENTS procedures to view information about the DBMS objects.

See your SAS/ACCESS documentation for a full listing of the SAS/ACCESS LIBNAME options that can be used with librefs that refer to DBMS data.

Using Data Set Options with SAS/ACCESS Librefs

After you have assigned a libref to your DBMS data, you can use SAS/ACCESS data set options, and some of the Base SAS data set options, on the data. The following example associates a libref with DB2 data and uses the SQL procedure to query the data:

```
libname mydb2lib db2;

proc sql;
  select *
    from mydb2lib.employees(drop=salary)
   where dept='Accounting';
quit;
```

The LIBNAME statement connects to DB2. You can reference a DBMS object, in this case, a DB2 table, by specifying a two-level name that is comprised of the libref and the DBMS object name. The DROP= data set option causes the SALARY column of the EMPLOYEES table on DB2 to be excluded from the data that is returned by the query.

See your SAS/ACCESS documentation for a full listing of the SAS/ACCESS data set options and the Base SAS data set options that can be used on data sets that refer to DBMS data.

Embedding a SAS/ACCESS LIBNAME Statement in a PROC SQL View

You can issue a SAS/ACCESS LIBNAME statement by itself, as shown in the previous examples, or as part of a CREATE VIEW statement in PROC SQL. The USING clause of the CREATE VIEW statement enables you to store DBMS connection information in a SAS view by *embedding* a SAS/ACCESS LIBNAME statement inside the SAS view. The following example uses an embedded SAS/ACCESS LIBNAME statement:

```
libname viewlib 'SAS-library';

proc sql;
  create view viewlib.emp_view as
    select *
```

```

from mydblib.employees
using libname mydblib oracle user=smith password=secret
    path='myoraclepath';
quit;

```

When PROC SQL executes the SAS view, the SELECT statement assigns the libref and establishes the connection to the DBMS. The scope of the libref is local to the SAS view and does not conflict with identically named librefs that might exist in the SAS session. When the query finishes, the connection is terminated and the libref is deassigned.

Note: You can also embed a Base SAS LIBNAME statement in a PROC SQL view. Δ

SQL Procedure Pass-Through Facility

The SQL Procedure Pass-Through Facility is an extension of the SQL procedure that enables you to send DBMS-specific statements to a DBMS and to retrieve DBMS data. You specify DBMS SQL syntax instead of SAS SQL syntax when you use the Pass-Through Facility. You can use Pass-Through Facility statements in a PROC SQL query or store them in a PROC SQL view.

The Pass-Through Facility consists of three statements and one component:

- The CONNECT statement establishes a connection to the DBMS.
- The EXECUTE statement sends dynamic, non-query DBMS-specific SQL statements to the DBMS.
- The CONNECTION TO component in the FROM clause of a PROC SQL SELECT statement retrieves data directly from a DBMS.
- The DISCONNECT statement terminates the connection to the DBMS.

The following Pass-Through Facility example sends a query to an ORACLE database for processing:

```

proc sql;
    connect to oracle as myconn (user=smith password=secret
        path='myoracleserver');

    select *
        from connection to myconn
            (select empid, lastname, firstname, salary
                from employees
                where salary>75000);

    disconnect from myconn;
quit;

```

The example uses the Pass-Through CONNECT statement to establish a connection with an ORACLE database with the specified values for the USER=, PASSWORD=, and PATH= arguments. The CONNECTION TO component in the FROM clause of the SELECT statement enables data to be retrieved from the database. The DBMS-specific statement that is sent to ORACLE is enclosed in parentheses. The DISCONNECT statement terminates the connection to ORACLE.

To store the same query in a PROC SQL , use the CREATE VIEW statement:

```

libname viewlib 'SAS-library';

proc sql;

```

```

connect to oracle as myconn (user=smith password=secret
    path='myoracleserver');

create view viewlib.salary as
  select *
    from connection to myconn
      (select empid, lastname, firstname, salary
        from employees
        where salary>75000);

disconnect from myconn;
quit;

```

ACCESS Procedure and Interface View Engine

The ACCESS procedure enables you to create *access descriptors*, which are SAS files of member type ACCESS. They describe data that is stored in a DBMS in a format that SAS can understand. Access descriptors enable you to create SAS/ACCESS views, called *view descriptors*. View descriptors are files of member type VIEW that function in the same way as SAS views that are created with PROC SQL, as described in “Embedding a SAS/ACCESS LIBNAME Statement in a PROC SQL View” on page 562 and “SQL Procedure Pass-Through Facility” on page 563.

Note: If a dynamic LIBNAME engine is available for your DBMS, it is recommended that you use the SAS/ACCESS LIBNAME statement to access your DBMS data instead of access descriptors and view descriptors; however, descriptors continue to work in SAS software if they were available for your DBMS in Version 6. Some new SAS features, such as long variable names, are not supported when you use descriptors. Δ

The following example creates an access descriptor and a view descriptor in the same PROC step to retrieve data from a DB2 table:

```

libname adlib 'SAS-library';
libname vlib 'SAS -library';

proc access dbms=db2;
  create adlib.order.access;
  table=sasdemo.orders;
  assign=no;
  list all;

  create vlib.custord.view;
  select ordernum stocknum shipto;
  format ordernum 5.
         stocknum 4.;
run;

proc print data=vlib.custord;
run;

```

When you want to use access descriptors and view descriptors, both types of descriptors must be created before you can retrieve your DBMS data. The first step, creating the access descriptor, enables SAS to store information about the specific DBMS table that you want to query.

After you have created the access descriptor, the second step is to create one or more view descriptors to retrieve some or all of the DBMS data described by the access descriptor. In the view descriptor, you select variables and apply formats to manipulate the data for viewing, printing, or storing in SAS. You use only the view descriptors, and not the access descriptors, in your SAS programs.

The interface view engine enables you to reference your SAS view with a two-level SAS name in a DATA or PROC step, such as the PROC PRINT step in the example.

See Chapter 28, “SAS Views,” on page 529 for more information about SAS views. See the SAS/ACCESS documentation for your DBMS for more detailed information about creating and using access descriptors and SAS/ACCESS views.

DBLOAD Procedure

The DBLOAD procedure enables you to create and load data into a DBMS table from a SAS data set, data file, SAS view, or another DBMS table, or to append rows to an existing table. It also enables you to submit non-query DBMS-specific SQL statements to the DBMS from your SAS session.

Note: If a dynamic LIBNAME engine is available for your DBMS, it is recommended that you use the SAS/ACCESS LIBNAME statement to create your DBMS data instead of the DBLOAD procedure; however, DBLOAD continues to work in SAS software if it was available for your DBMS in Version 6. Some new SAS features, such as long variable names, are not supported when you use the DBLOAD procedure. Δ

The following example appends data from a previously created SAS data set named INVDATA into a table in an ORACLE database named INVOICE:

```
proc dbload dbms=oracle data=invdata append;
  user=smith;
  password=secret;
  path='myoracleserver';
  table=invoice;
  load;
run;
```

See the SAS/ACCESS documentation for your DBMS for more detailed information about the DBLOAD procedure.

Interface DATA Step Engine

Some SAS/ACCESS software products support a DATA step interface, which enables you to read data from your DBMS by using DATA step programs. Some products support both reading and writing in the DATA step interface.


The DATA step interface consists of four statements:

- The INFILE statement identifies the database or message queue to be accessed.
- The INPUT statement is used with the INFILE statement to issue a GET call to retrieve DBMS data.
- The FILE statement identifies the database or message queue to be updated, if writing to the DBMS is supported.
- The PUT statement is used with the FILE statement to issue an UPDATE call, if writing to the DBMS is supported.

The following example updates data in an IMS database by using the FILE and INFILE statements in a DATA step. The statements generate calls to the database in the IMS native language, DL/I. The DATA step reads BANK.CUSTOMER, an existing SAS data set that contains information on new customers, and then it updates the ACCOUNT database with the data in the SAS data set.

```
data _null_;
  set bank.customer;
  length ssa1 $9;
  infile accupdt dli call=func dbname=db ssa=ssa1;
  file accupdt dli;
  func = 'isrt';
  db = 'account';
  ssa1 = 'customer';
  put @1  ssnumber $char11.
      @12  custname $char40.
      @52  addr1   $char30.
      @82  addr2   $char30.
      @112 custcity $char28.
      @140 custstat $char2.
      @142 custland $char20.
      @162 custzip  $char10.
      @172 h_phone  $char12.
      @184 o_phone  $char12.;
  if _error_ = 1 then
    abort abend 888;
run;
```

In SAS/ACCESS products that provide a DATA step interface, the INFILE statement has special DBMS-specific options that enable you to specify DBMS variable values and to format calls to the DBMS appropriately. See the SAS/ACCESS documentation for your DBMS for a full listing of the DBMS-specific INFILE statement options and the Base SAS INFILE statement options that can be used with your DBMS.



CHAPTER

33

Processing Data Using Cross-Environment Data Access (CEDA)

<i>Definition of Cross-Environment Data Access (CEDA)</i>	567
<i>Advantages of CEDA</i>	568
<i>SAS File Processing with CEDA</i>	568
<i>What Types of Processing Does CEDA Support?</i>	568
<i>Behavioral Differences for Output Processing</i>	569
<i>Restrictions for CEDA</i>	569
<i>Understanding When CEDA Is Used to Process a File</i>	570
<i>Determining Whether Update Processing Is Allowed</i>	572
<i>Alternatives to Using CEDA</i>	573
<i>Creating Files in a Foreign Data Representation</i>	574
<i>Examples of Using CEDA</i>	574
<i>Example 1: Automatically Processing a Foreign File</i>	574
<i>Example 2: Creating a New File in a Foreign Environment</i>	575
<i>Example 3: Changing the Data Representation of an Existing File</i>	575

Definition of Cross-Environment Data Access (CEDA)

Cross-environment data access (CEDA) is a Base SAS feature that enables a SAS file that was created in a directory-based operating environment (for example, UNIX, Windows, or OpenVMS) to be processed in any of the following situations:

- by a SAS session that is running in another directory-based environment. For example, if you move a file from one operating environment like Windows to a different operating environment like UNIX, CEDA translates the file, which eliminates the need for you to convert the file.
- on a platform that is different from the platform on which the file was created. For example, CEDA is useful if you have upgraded to a 64-bit platform from a 32-bit platform.
- by a SAS session in which the session encoding is incompatible with the encoding for the SAS file.

With CEDA, you do not need to create a transport file, use other SAS procedures, or change your SAS program. CEDA is available for files that are created with SAS 7 and later releases.

Here are a few terms and definitions to help you understand CEDA:

data representation	the form in which data is stored in a particular operating environment. Different operating environments use different standards or conventions for storing floating-point numbers (for example, IEEE or IBM 390); for character encoding (ASCII or EBCDIC); for the ordering of bytes in memory (big Endian or little
---------------------	--

	Endian); for word alignment (4-byte boundaries or 8-byte boundaries); and for data-type length (16-bit, 32-bit, or 64-bit).
encoding	is 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. Some examples of encodings are Wlatin1 and Danish EBCDIC.
foreign	refers to a file or an environment for which the data representation contrasts with the CPU that is processing the file. For example, the data representation that is created by an IBM mainframe is considered foreign to the data representation that is created by a Windows environment.
native	refers to a file or an environment for which the data representation is compatible with the CPU that is processing the file. For example, a file that is in Windows data representation is native to a Windows environment. There are also native groupings that are compatible, such as HP_UX_64, RS_6000_AIX_64, SOLARIS_64, and HP_IA64. For more information, see Table 33.2 on page 571.

Advantages of CEDA

CEDA offers these advantages:

- You can transparently process a supported SAS file with no knowledge of the file's data representation or character encoding.
- No transport files are created. CEDA requires a single translation between native and foreign representations, rather than a translation from native representation to transport file to native representation.
- CEDA eliminates the need to perform explicit steps in order to process the file.
- CEDA does not require an explicit sign-on as is needed in SAS/CONNECT or a dedicated server as is needed in SAS/SHARE.

SAS File Processing with CEDA

What Types of Processing Does CEDA Support?

CEDA supports SAS 7 and later SAS files that are created in directory-based operating environments like UNIX, Windows, and OpenVMS. CEDA provides the following SAS file processing for these SAS engines:

BASE	default Base SAS engine for SAS 9 (V9), SAS 8 (V8), and SAS 7 (V7).
SASESOCK	TCP/IP port engine for SAS/CONNECT software.
TAPE	sequential engine for SAS 9 (V9TAPE), SAS 8 (V8TAPE), and SAS 7 (V7TAPE).

Table 33.1 SAS File Processing Provided by CEDA

SAS File Type	Engine	Supported Processing
SAS data file	BASE, TAPE, SASESOCK	input and output ¹ processing
PROC SQL view	BASE	input processing
SAS/ACCESS view for Oracle or Sybase	BASE	input processing
MDDDB file ²	BASE	input processing

1 For output processing that replaces an existing SAS data file, there are behavioral differences. For more information, see “Behavioral Differences for Output Processing” on page 569.

2 CEDA supports SAS 8 and later MDDDB files.

Behavioral Differences for Output Processing

For output processing that replaces an existing SAS data file, the BASE and TAPE engines behave differently regarding the following attributes:

encoding

- The BASE engine uses the encoding of the existing file; that is, the encoding is cloned.
- The TAPE engine uses the current SAS session encoding.
- For both the BASE and TAPE engines, the COPY procedure uses the encoding of the file from the source library (that is, the file being copied), regardless of whether the file existed in the target library.

data representation

The BASE and TAPE engines use the data representation of the native environment, except with the COPY procedure. By default, PROC COPY uses the data representation of the file that is copied. When writing out the file, if you want PROC COPY to use the data representation of the target operating system and not the data representation of the source file, you must specify the NOCLONE option.

Restrictions for CEDA

CEDA has the following restrictions:

- CEDA does not support DATA step views, SAS/ACCESS views that are not for SAS/ACCESS for Oracle or Sybase, SAS catalogs, stored compiled DATA step programs, item stores, DMDDB files, FDB files, or any SAS file that was created before SAS 7.
- Update processing is not supported.
- Integrity constraints cannot be read or updated.
- An audit trail file cannot be updated but it can be read.
- Indexes are not supported. Therefore, WHERE optimization with an index is not supported.
- On z/OS, only UNIX file system libraries fully support CEDA. SAS can create bound library members on z/OS with a character encoding that differs from the default. CEDA is used in creating such a member as well as in any subsequent attempts to read the member. However, for bound library members, the character encoding is the only aspect of the data representation that might differ from the

default. For example, it is not possible, using the OUTREP option of a LIBNAME statement, to create a member in data representation for another host besides z/OS. For more information about the various types of libraries that are supported by SAS on z/OS, see “Library Implementation Types for Base and Sequential Engines” in *SAS Companion for z/OS*.

- Because the BASE engine translates the data as the data is read, multiple procedures require SAS to read and translate the data multiple times. In this way, the translation could affect system performance.
- If a foreign data set is damaged, CEDA cannot process the file in order to repair it. CEDA does not support update processing, which is required in order to repair a damaged data set. To repair the foreign file, you must move it back to its native environment. For information on how to repair a damaged data set, see the REPAIR statement in the DATASETS procedure in *Base SAS Procedures Guide*.
- Transcoding could result in character data loss when encodings are incompatible. For information about encoding and transcoding, see the *SAS National Language Support (NLS): Reference Guide*.
- Loss of precision can occur in numeric variables when you move data between operating environments. If a numeric variable is defined with a short length, you can try increasing the length of the variable. Full-size numeric variables are less likely to encounter a loss of precision with CEDA. For more information, see “Numeric Precision in SAS Software” on page 38.
- Numeric variables have a minimum length of either 2 or 3 bytes, depending on the operating environment. In an operating environment that supports a minimum of 3 bytes (such as Windows or UNIX), CEDA cannot process a numeric variable that was created with a length of 2 bytes (for example, in z/OS). If you encounter this restriction, then use the XPORT engine or the CPORT and CIMPORT procedures instead of CEDA.

Note: If you encounter these restrictions because your files were created under a previous version of SAS, consider using the MIGRATE procedure, which is documented in the *Base SAS Procedures Guide*. PROC MIGRATE retains many features, such as integrity constraints, indexes, and audit trails. Δ

Understanding When CEDA Is Used to Process a File

Because CEDA translation is transparent, you might not be aware when CEDA is being used. Knowing when CEDA is used could be helpful because, for example, CEDA translation might require additional resources.

Starting in SAS 9, by default, SAS writes a message to the log when CEDA is used. Here is an example:

```
Note: Data file HEALTH.GRADES.DATA is in a format that is native to another
host, or the file encoding does not match the session encoding. Cross
Environment Data Access will be used, which might require additional CPU
resources and might reduce performance.
```

CEDA is used in these situations:

- when the encoding of character values for the SAS file is incompatible with the currently executing SAS session encoding.
- when the data representation of the SAS file is incompatible with the data representation of the currently executing SAS session. For example, an incompatibility can occur if you move a file from an operating environment like Windows to an operating environment like UNIX, or if you have upgraded to a 64-bit platform from a 32-bit platform.

In the following table, each row contains a group of operating environments that are compatible with each other. CEDA is used only when you create a file with a data representation in one row and process the file under a data representation of another row. (Environments are named by the operating system and the platform on which SAS is executed.)

Table 33.2 Compatibility across Environments

Data Representation Value	Environment
ALPHA_TRU64	Tru64 UNIX (table note 1)
LINUX_IA64	Linux for Itanium (table note 1)
LINUX_X86_64	Linux for x64 (table note 1)
SOLARIS_X86_64	Solaris for x64 (table note 1)
ALPHA_VMS_32	OpenVMS Alpha (table note 2)
ALPHA_VMS_64	OpenVMS Alpha (table note 2)
VMS_IA64	OpenVMS for HP Integrity servers 64-bit platform (table note 2)
HP_IA64	HP-UX for Itanium on 64-bit platform
HP_UX_64	HP-UX for PA-RISC on 64-bit platform
RS_6000_AIX_64	AIX on 64-bit platform
SOLARIS_64	Solaris on SPARC 64-bit platform
HP_UX_32	HP-UX on 32-bit platform
MIPS_ABI	ABI UNIX on 32-bit platform
RS_6000_AIX_32	AIX on 32-bit platform
SOLARIS_32	Solaris on SPARC 32-bit platform
LINUX_32	Linux for Intel architecture
INTEL_ABI	ABI UNIX for Intel on 32-bit platform
MVS_32	z/OS on 32-bit platform
OS2	OS/2 for Intel on 32-bit platform
VAX_VMS	OpenVMS VAX
WINDOWS_32	Microsoft Windows on 32-bit platform
WINDOWS_64	Microsoft Windows 64-Bit Edition (for both Itanium-based systems and x64)

TABLE NOTE 1: Although all four of the environments in this group are compatible, catalogs are an exception:

- Catalogs are compatible between Tru64 UNIX and Linux for Itanium.
- Catalogs are compatible between Linux for x64 and Solaris for x64.

TABLE NOTE 2: Although these OpenVMS environments have different data representations for some compiler types, SAS data sets that are created by the BASE engine do not store the data types that are different. Therefore, if the encoding is compatible, CEDA is not used between these environments. However, note that SAS 9 does not support SAS 8 catalogs from OpenVMS. You can migrate the catalogs with the MIGRATE procedure. For more information on the MIGRATE procedure, see the *Base SAS Procedures Guide*.

Determining Whether Update Processing Is Allowed

If a file's data representation is the same as the data representation of the processing environment, and if the encoding is compatible with the currently executing SAS session encoding, then you can manually update the file, because CEDA is not needed in order to translate the file. For example, in a Windows environment, if a file was created in a Windows environment or if the OUTREP= option was used to designate the file in Windows data representation, then you can update the file.

Otherwise, if CEDA is used to translate the file, you cannot update it. If you attempt to update the file, then you will receive an error message that says that updating is not allowed. For example:

```
ERROR: File HEALTH.OXYGEN cannot be updated because its encoding does
not match the session encoding or the file is in a format native to another
host, such as SOLARIS_32, HP_UX_32, RS_6000_AIX_32, MIPS_ABI.
```

To determine the data representation and the encoding of a file, you can use the CONTENTS procedure (or the CONTENTS statement in PROC DATASETS). For example, the data set HEALTH.OXYGEN was created in a UNIX environment in SAS 9. The file was moved to a SAS 9 Windows environment, in which the following CONTENTS output was requested:

Output 33.1 CONTENTS Output Showing Data Representation

The SAS System		1	
The CONTENTS Procedure			
Data Set Name	HEALTH.OXYGEN	Observations	31
Member Type	DATA	Variables	7
Engine	V9	Indexes	0
Created	Wednesday, January 24, 2007 10:11:39	Observation Length	56
Last Modified	Wednesday, January 24, 2007 10:11:33	Deleted Observations	0
Protection		Compressed	NO
Data Set Type		Sorted	NO
Label			
Data Representation	SOLARIS_32, HP_UX_32, RS_6000_AIX_32, MIPS_ABI		
Encoding	latin1 Western (ISO)		
Engine/Host Dependent Information			
Data Set Page Size	5120		
Number of Data Set Pages	1		
First Data Page	1		
Max Obs per Page	90		
Obs in First Data Page	31		
Number of Data Set Repairs	0		
File Name	/u/xxxxxx/myfiles/health/\oxygen.sas7bdat		
Release Created	9.0200A0		
Host Created	HP-UX		
Alphabetic List of Variables and Attributes			
	#	Variable	Type Len
	1	AGE	Num 8
	6	MAXPULSE	Num 8
	7	OXYGEN	Num 8
	4	RSTPULSE	Num 8
	5	RUNPULSE	Num 8
	3	RUNTIME	Num 8
	2	WEIGHT	Num 8

Alternatives to Using CEDA

Because of the restrictions, it might not be feasible to use CEDA. You can use the following methods in order to move files across operating environments:

XPORT engine with the DATA step or PROC COPY

In the source environment, the LIBNAME statement with the XPORT engine and either the DATA step or PROC COPY creates a transport file from a SAS data set. In the target environment, the same method translates the transport file into the target environment's native format. Note that the XPORT engine does not support SAS 7 and later features, such as long file and variable names.

XML engine with the DATA step or PROC COPY

In the source environment, the LIBNAME statement with the XML engine and either the DATA step or PROC COPY creates an XML document from a SAS data set. In the target environment, the same method translates the XML document into the target environment's native format.

CPORT and CIMPORT procedures

In the source environment, PROC CPORT writes data sets or catalogs to transport format. In the target environment, PROC CIMPORT translates the transport file into the target environment's native format.

Data transfer services in SAS/CONNECT software

Data transfer services is a bulk data transfer mechanism that transfers a disk copy of the data and performs the necessary conversion of the data from one environment's representation to another's, as well as any necessary conversion between SAS releases. You must establish a connection between the two SAS sessions by using the SIGNON command and then executing either PROC UPLOAD or PROC DOWNLOAD to move the data.

Remote library services in both SAS/CONNECT software and SAS/SHARE software

Remote library services gives you transparent access to remote data through the use of the LIBNAME statement.

Creating Files in a Foreign Data Representation

By default, when you create a new file, SAS uses the data representation of the CPU that is running SAS. You can specify the OUTREP= option to override this default.

The OUTREP= option is both a SAS data set option and a LIBNAME statement option. The data set option applies to an individual file. The LIBNAME statement option applies to an entire library. This option uses CEDA to create a file that has a foreign data representation. For example, in a UNIX environment, you can create a data set that is foreign to UNIX, but which is native to Windows.

See "Example 2: Creating a New File in a Foreign Environment" on page 575 and "Example 3: Changing the Data Representation of an Existing File" on page 575.

See the OUTREP= option for the LIBNAME statement or the OUTREP= data set option in *SAS Language Reference: Dictionary*.

Examples of Using CEDA

Example 1: Automatically Processing a Foreign File

This example shows how simple it is to move a SAS data set from one operating environment to another and to process the file in the new environment without any conversion steps.

First, the SAS data set is moved, using FTP, from an HP UNIX environment to a Windows PC.

```
C:\>ftp my.unix.node.com
FTP>binary
FTP>get unxdata.sas7bdat
FTP>quit
```

Then, using CEDA, SAS automatically recognizes the foreign data representation (which is HP UNIX) and translates it to the native data representation for the Windows environment. The log output displays a message that the file is being processed using CEDA.

```
libname unx '.';

proc print data=unx.unxdata;
run;
```

Output 33.2 Log Output from Processing a Foreign File

Note: Data file HEALTH.GRADES.DATA is in a format that is native to another host, or the file encoding does not match the session encoding. Cross Environment Data Access will be used, which might require additional CPU resources and might reduce performance.

Example 2: Creating a New File in a Foreign Environment

In this example, an administrator who works in a z/OS operating environment wants to create a SAS file in a UNIX file system directory that can be processed in an HP UNIX environment. Specifying OUTREP=HP_UX_32 as a data set option forces the data representation to match the data representation of the UNIX operating environment that will process the file. This method of creating the file can enhance system performance because the file does not require data conversion when being read by an HP UNIX computer.

```
libname foreign v9 'HFS-file-spec';


data foreign.a (outrep=HP_UX_32);
  infile file-specifications;
  input student $ test1 test2 test3 final;
  total = test1+test2+test3+final;
  grade = total/4.0;
run;
```

Example 3: Changing the Data Representation of an Existing File

To change an existing file's data representation, you can use PROC COPY with the NOCLONE option and specify OUTREP= in the LIBNAME statement. The following example copies a source library of Windows data sets to a target library of Solaris data sets. Note that if you move a data set to another platform, you must move it as a binary file.

```
libname target '~pathname-to- target-library' outrep=solaris;
libname source '~pathname-to- source-library';
proc copy in=source out=target noclone memtype=data;
run;
```

For more information, see the OUTREP= option for the LIBNAME statement or the OUTREP= data set option in *SAS Language Reference: Dictionary*.



CHAPTER

34

SAS 9.2 Compatibility with SAS Files from Earlier Releases

<i>Introduction to Version Compatibility</i>	577
<i>Comparing SAS System 9 to Earlier Releases</i>	577
<i>SAS 9 File Format</i>	577
<i>SAS 9 Filename Extensions</i>	578
<i>Using SAS Library Engines</i>	578

Introduction to Version Compatibility

SAS recognizes that SAS 9 customers often have existing data and programs. You want to seamlessly process your existing files, and possibly to simultaneously operate both SAS 9 and an earlier release of SAS. In many cases, you can use SAS 9 to process SAS files that were created in versions 8, 7, and 6 of SAS without first converting the files; however, there are some limitations.

Compatibility between versions will vary depending on the type of SAS file, the SAS release that you are running, the operating environment in which the file was created, and the type of processing you need to do. Compatibility issues are generally handled automatically by SAS. However, there are situations that require you to specify an engine name or to migrate the file.

The additional information in this topic provides general, overview information about compatibility.

For specific processing information and guidelines for migration issues, see the Migration Focus Area at support.sas.com.

The Migration Focus Area is your guide to migrating files from previous versions of SAS to SAS 9. Refer to this focus area for planning and cost analysis information, known compatibility issues and their resolutions, and step-by-step instructions. The MIGRATE procedure, which provides a simple way to migrate a library of SAS files from previous releases of SAS, is documented in *Base SAS Procedures Guide*.

Comparing SAS System 9 to Earlier Releases

SAS 9 File Format

In order to provide longer filenames and variable names, the file format used in SAS 7 and 8 is different from the file format used in SAS 6.

For SAS 9, the file format is basically the same as in SAS 7 and 8. The Base SAS engine is the same engine, except that for SAS 9, you can define and use longer format and informat names and have more than 32,767 variables in a SAS data set. You cannot use these features in SAS 7 or 8.

SAS files created with SAS 7 and 8 are compatible with SAS 9. However, a SAS file that is created with a 32-bit version of SAS has a different data representation than a SAS file from a 64-bit version of SAS. The data representation is the format in which data is represented on a computer architecture or in an operating environment. Therefore, if you have SAS files that were created on a 32-bit version of SAS and you have upgraded to a 64-bit version of SAS, you will have processing limitations due to the different data representation.

SAS 9 Filename Extensions

A filename extension reflects the engine that was used to create both the file and the SAS file member type.

Because SAS needs to distinguish among the different file types and versions, SAS automatically assigns a specific extension to each file when the file is created. For example, in order to distinguish SAS 7 and 8 files from SAS 6 files, the extensions are different.

For SAS 9, the file extensions are the same as the file extensions in SAS 7 and 8.

The following table lists the file extensions for a SAS data file (SAS data set with member type DATA) in SAS 6, 7, 8, and 9 for different operating environments:

Table 34.1 File Extensions for a SAS Data File in Different Operating Environments

Engine Name	UNIX	OpenVMS for Integrity servers	Windows	z/OS ¹
V6	.ssd01	.SASEB\$DATA	.sd2	not available
V7	.sas7bdat	.sas7bdat	.sas7bdat	.sas7bdat
V8	.sas7bdat	.sas7bdat	.sas7bdat	.sas7bdat
V9	.sas7bdat	.sas7bdat	.sas7bdat	.sas7bdat

¹ applies to SAS data sets that reside in the hierarchical file system of UNIX System Services.

Operating Environment Information: For a complete list of SAS member types and extensions, see the SAS documentation for your operating environment. Δ

Using SAS Library Engines

In order to access a SAS library, SAS needs a libref and a library engine name. For example, you assign a libref to the SAS library with the LIBNAME statement or the New Library window, but usually you do not have to specify an engine name because SAS automatically selects the appropriate engine.

If you do not specify an engine, SAS automatically assigns one based on the contents of the SAS library. For example, SAS is able to differentiate between a SAS 6 library and a SAS 9 library. Note that in SAS 9, a SAS library containing SAS 7 and 8 files is the same as a SAS 9 library, because the engine that creates a SAS file determines its format, and the file format for SAS 7, 8, and 9 is the same.

For example, in a SAS 9 session, if you issue the following LIBNAME statement to assign a libref to a SAS library containing SAS 8 SAS files, SAS automatically uses the SAS 9 engine:

```
libname mylib 'v8-SAS-library';
```

In a SAS 9 session, if you issue the following LIBNAME statement to assign a libref to a SAS library that contains only SAS 6 files, SAS automatically uses the Version 6 compatibility engine:

```
libname mylib 'v6-SAS-library';
```

SAS automatically assigns an engine based on the contents of the SAS library as shown in the following table:

Table 34.2 Default Library Engine Assignment in SAS 9


Engine Assignment	SAS Library Contents
V9	No SAS files; the library is empty
V9	Only SAS 9 SAS files
V9	Only SAS 8 SAS files
V9	Only SAS 7 SAS files
V6	Only SAS 6 SAS files
V9	Both SAS 9 SAS files and SAS files from earlier releases

Note: Even though SAS will automatically assign an engine based on the library contents, it is more efficient for you to specify the engine. For example, specifying the engine name in the following LIBNAME statement saves SAS from determining which engine to use:

```
libname mylib v6 'v6-SAS-library';
```

Δ

For more information about SAS engines, see Chapter 36, “SAS Engines,” on page 589.



CHAPTER 35

File Protection

<i>Definition of a Password</i>	581
<i>Assigning Passwords</i>	582
<i>Syntax</i>	582
<i>Assigning a Password with a DATA Step</i>	582
<i>Assigning a Password to an Existing Data Set</i>	583
<i>Assigning a Password with a Procedure</i>	583
<i>Assigning a Password with the SAS Windowing Environment</i>	584
<i>Assigning a Password Outside of SAS</i>	584
<i>Removing or Changing Passwords</i>	584
<i>Using Password-Protected SAS Files in DATA and PROC Steps</i>	584
<i>How SAS Handles Incorrect Passwords</i>	585
<i>Assigning Complete Protection with the PW= Data Set Option</i>	585
<i>Encoded Passwords</i>	586
<i>Using Passwords with Views</i>	586
<i>How the Level of Protection Differs from SAS Views</i>	586
<i>PROC SQL Views</i>	587
<i>SAS/ACCESS Views</i>	587
<i>DATA Step Views</i>	588
<i>SAS Data File Encryption</i>	588
<i>Passwords and Encryption with Generation Data Sets, Audit Trails, Indexes, and Copies</i>	588

Definition of a Password

SAS software enables you to restrict access to members of SAS libraries by assigning passwords to the members. You can assign passwords to all member types except catalogs. You can specify three levels of protection: read, write, and alter. When a password is assigned, it appears as uppercase Xs in the log.

Note: This document uses the terms *SAS data file* and *SAS view* to distinguish between the two types of SAS data sets. Passwords work differently for type VIEW than they do for type DATA. The term “SAS data set” is used when the distinction is not necessary. △

<i>read</i>	protects against reading the file.
<i>write</i>	protects against changing the data in the file. For SAS data files, write protection prevents adding, modifying, or deleting observations.
<i>alter</i>	protects against deleting or replacing the entire file. For SAS data files, alter protection also prevents modifying variable attributes and creating or deleting indexes.

Alter protection does not require a password for read or write access; write protection does not require a password for read access. For example, you can read an alter-protected or write-protected SAS data file without knowing the alter or write password. Conversely, read and write protection do not prevent any operation that requires alter protection. For example, you can delete a SAS data set that is read- or write-protected only without knowing the read or write password.

To protect a file from being read, written to, deleted, or replaced by anyone who does not have the proper authority, assign read, write, and alter protection. To allow others to read the file without knowing the password, but not change its data or delete it, assign just write and alter protection. To completely protect a file with one password, use the PW= data set option. See “Assigning Complete Protection with the PW= Data Set Option” on page 585 for details.

Note: Because of the way SAS opens files, you must specify the read password to update a SAS data set that is only read-protected. Δ

Note: The levels of protection differ somewhat for the member type VIEW. See “Using Passwords with Views” on page 586. Δ

Assigning Passwords

Syntax

To set a password, first specify a SAS data set in one of the following:

- a DATA statement
- the MODIFY statement of the DATASETS procedure
- an OUT = option in some procedures
- the CREATE VIEW statement in PROC SQL
- the ToolBox

Then assign one or more password types to the data set. The data set might already exist, or the data set might be one that you create. The following is an example of syntax:

```
password-type=password <... password-type=password>
```

where *password* is a valid eight-character SAS name and *password-type* can be one of the following SAS data set options:

```
ALTER=
PW=
READ=
WRITE=
```

CAUTION:

Keep a record of any passwords you assign! If you forget or do not know the password, you cannot get the password from SAS. Δ

Assigning a Password with a DATA Step

You can use data set options to assign passwords to unprotected members in the DATA step when you create a new SAS data file.

This example prevents deletion or modification of the data set without a password.

```

/* assign a write and an alter password to MYLIB.STUDENTS */
data mylib.students(write=yellow alter=red);
  input name $ sex $ age;
  datalines;
Amy f 25
... more data lines ...
;

```

This example prevents reading or deleting a stored program without a password and also prevents changing the source program.

```

/* assign a read and an alter password to the SAS view ROSTER */
data mylib.roster(read=green alter=red) /
  view=mylib.roster;
  set mylib.students;
run;

libname stored 'SAS-library-2';

/* assign a read and alter password to the program file SOURCE */
data mylib.schedule / pgm=stored.source(read=green alter=red);
  ... DATA step statements ...
run;

```

Note: When you replace a SAS data set that is alter-protected, the new data set inherits the alter password. To change the alter password for the new data set, use the MODIFY statement in the DATASETS procedure. △

Assigning a Password to an Existing Data Set

You can use the MODIFY statement in the DATASETS procedure to assign passwords to unprotected members if the SAS data file already exists.

```

/* assign an alter password to STUDENTS */
proc datasets library=mylib;
  modify students(alter=red);
run;

```

Assigning a Password with a Procedure

You can assign a password after an OUT= data set specification in some procedures.

```

/* assign a write and an alter password to SCORE */
proc sort data=mylib.math
  out=mylib.score(write=yellow alter=red);
  by number;
run;

```

You can assign a password in a CREATE TABLE or a CREATE VIEW statement in PROC SQL.

```

/* assign an alter password to the SAS view BDAY */
proc sql;

```

```
create view mylib.bday(alter=red) as
  query-expression;
```

Assigning a Password with the SAS Windowing Environment

You can create or change passwords for any data file using the Password Window in the SAS windowing environment. To invoke the Password Window from the ToolBox, use the global command SETPASSWORD followed by the filename. This opens the password window for the specified data file.

Assigning a Password Outside of SAS

A SAS password does not control access to a SAS file beyond the SAS system. You should use the operating system-supplied utilities and file-system security controls in order to control access to SAS files outside of SAS.

Removing or Changing Passwords

To remove or change a password, use the MODIFY statement in the DATASETS procedure. For more information, see the DATASETS procedure in *Base SAS Procedures Guide*.

Using Password-Protected SAS Files in DATA and PROC Steps

To access password-protected files, use the same data set options that you use to assign protection.

```
□ /* Assign a read and alter password
   /* to the stored program file*/ /*STORED.SOURCE */
   data mylib.schedule / pgm=stored.source
     (read=green alter=red);
     <... more DATA step statements ...>
run; /*Access password-protected file*/
proc sort data=mylib.score(write=yellow alter=red);
  by number;
run;

□ /* Print read-protected data set MYLIB.AUTOS */
   proc print data=mylib.autos(read=green); run;

□ /* Append ANIMALS to the write-protected */
   /* data set ZOO */
   proc append base=mylib.zoo(write=yellow)
     data=mylib.animals;
run;

□ /* Delete alter-protected data set MYLIB.BOTANY */
   proc datasets library=mylib;
     delete botany(alter=red);
run;
```

Passwords are hierarchical in terms of gaining access. For example, specifying the ALTER password gives you read and write access. The following example creates the data set STATES, with three different passwords, and then reads the data set to produce a plot:

```
data mylib.states(read=green write=yellow alter=red);
  input density crime name $;
  datalines;
151.4 6451.3 Colorado
... more data lines ...
;

proc plot data=mylib.states(alter=red);
  plot crime*density;
run;
```

How SAS Handles Incorrect Passwords

If you are using the SAS windowing environment and you try to access a password-protected member without specifying the correct password, you receive a dialog box that prompts you for the appropriate password. The text you enter in this window is not displayed. You can use the PWREQ= data set option to control whether a dialog box appears after a user enters a missing or incorrect password. PWREQ= is most useful in SCL applications.

If you are using batch or noninteractive mode, you receive an error message in the SAS log if you try to access a password-protected member without specifying the correct password.

If you are using interactive line mode, you are also prompted for the password if you do not specify the correct password. When you enter the password and press ENTER, processing continues. If you cannot give the correct password, you receive an error message in the SAS log.

Assigning Complete Protection with the PW= Data Set Option

The PW= data set option assigns the same password for each level of protection. This data set option is convenient for thoroughly protecting a member with just one password. If you use the PW= data set option, those who have access only need to remember one password for total access.

- To access a member whose password is assigned using the PW= data set option, use the PW= data set option or the data set option that equates to the specific level of access you need:

```
/* create a data set using PW=,
   then use READ= to print the data set */
data mylib.states(pw=orange);
  input density crime name $;
  datalines;
151.4 6451.3 Colorado
... more data lines ...
;

proc print data=mylib.states(read=orange);
```

```
run;
```

- PW= can be an alias for other password options:

```
/* Use PW= as an alias for ALTER=. */
data mylib.college(alter=red);
  input name $ 1-10 location $ 12-25;
  datalines;
Vanderbilt Nashville
Rice           Houston
Duke           Durham
Tulane         New Orleans
... more data lines ...
;

proc datasets library=mylib;
  delete college(pw=red);
run;
```

Encoded Passwords

Encoding a password enables you to write SAS programs without having to specify a password in plain text. The PWENCODE procedure uses encoding to disguise passwords. With encoding, one character set is translated to another character set through some form of table lookup. An encoded password is intended to prevent casual, non-malicious viewing of passwords. You should not depend on encoded passwords for all your data security needs; a determined and knowledgeable attacker can decode the encoded passwords.

When an encoded password is used, the syntax parser decodes the password and accesses the file. The encoded password is never written to the SAS log in plain text. SAS doesn't accept passwords longer than eight characters. If an encoded password is decoded and is longer than eight characters, SAS reads it as an incorrect password and sends an error message to the SAS log. For more information, see the PWENCODE Procedure in *Base SAS Procedures Guide*.

Using Passwords with Views

How the Level of Protection Differs from SAS Views

The levels of protection for SAS views and stored programs differ slightly from other types of SAS files. Passwords affect the actual SAS view definition or view descriptor as well as the underlying data. Unless otherwise noted, the term “view” can refer to any type of SAS view. Also, the term “underlying data” refers to the data that is accessed by the SAS view:

- read
 - protects against reading the SAS view's underlying data.
 - allows source statements to be written to the SAS log, using DESCRIBE.
 - allows replacement of the SAS view.

- | | |
|-------|--|
| write | <ul style="list-style-type: none"> □ protects the underlying data associated with a SAS view by insisting that a write password is given. □ allows source statements to be written to the SAS log using DESCRIBE □ allows replacement of the SAS view |
| alter | <ul style="list-style-type: none"> □ protects against source statements being written to the SAS log, using DESCRIBE. □ protects against replacement of the SAS view. |

An important difference between SAS views and other types of SAS files is that you need alter access to DESCRIBE an alter-protected view. For example, to use an alter-protected PROC SQL view in a DESCRIBE VIEW statement, you must specify the alter password.

In most DATA and PROC steps, the way you use password-protected views is consistent with the way you use other types of password-protected SAS files. For example, the following PROC PRINT prints a read-protected view:

```
proc print data=mylib.grade(read=green);
run;
```

Note: You might experience unexpected results when you place protection on a SAS view if some type of protection has already been placed on the underlying data set. △

PROC SQL Views

Typically, when you create a PROC SQL view from a password-protected SAS data set, you specify the password in the FROM clause in the CREATE VIEW statement using a data set option. In this way, when you use the view later, you can access the underlying data without re-specifying the password. For example, the following statements create a PROC SQL view from a read-protected SAS data set, and drop a sensitive variable:

```
proc sql;
  create view mylib.emp as
    select * from mylib.employee(pw=orange drop=salary);
quit;
```

Note: If you create a PROC SQL view from password-protected SAS data sets without specifying their passwords, when you try to use the view you are prompted for the passwords of the SAS data sets named in the FROM clause. If you are running SAS in batch or noninteractive mode, you receive an error message. △

SAS/ACCESS Views

SAS/ACCESS software enables you to edit view descriptors and, in some interfaces, the underlying data. To prevent someone from editing or reading (browsing) the view descriptor, assign alter protection to the view. To prevent someone from updating the underlying data, assign write protection to the view. For more information, see the SAS/ACCESS documentation for your DBMS.

DATA Step Views

When you create a DATA step view using a password-protected SAS data set, specify the password in the view definition. In this way, when you use the view, you can access the underlying data without respecifying the password.

The following statements create a DATA step view using a password-protected SAS data set, and drop a sensitive variable:

```
data mylib.emp / view=mylib.emp;
    set mylib.employee(pw=orange drop=salary);
run;
```

Note that you can use the SAS view without a password, but access to the underlying data requires a password. This is one way to protect a particular column of data. In the above example, **proc print data=mylib.emp;** will execute, but **proc print data=mylib.employee;** will fail without the password.

SAS Data File Encryption

SAS passwords restrict access to SAS data files within SAS, but SAS passwords cannot prevent SAS data files from being viewed at the operating environment system level or from being read by an external program. Encryption provides security of your SAS data outside of SAS by writing to disk the encrypted data that represents the SAS data. The data is decrypted as it is read from the disk.

Encryption does not affect file access. However, SAS honors all host security mechanisms that control file access. You can use encryption and host security mechanisms together.

The following example creates an encrypted SAS data set:


```
data salary(encrypt=yes read=green);
    input name $ yrsal bonuspct;
    datalines;
Muriel    34567  3.2
Bjorn     74644  2.5
Freda     38755  4.1
Benny     29855  3.5
Agnetha   70998  4.1
;
```

To print this data set, specify the read password:

```
proc print data=salary(read=green);
run;
```

Passwords and Encryption with Generation Data Sets, Audit Trails, Indexes, and Copies

SAS extends password protection and encryption to other files associated with the original protected file. This includes generation data sets, indexes, audit trails, and copies. When accessing protected or encrypted generation data sets, indexes, audit trails, and copies of the original file, the same rules, syntax, and behavior for invoking the original password protected or encrypted files apply. SAS views cannot have generation data sets, indexes, and audit trails.



CHAPTER 36

SAS Engines

<i>Definition of a SAS Engine</i>	589
<i>Specifying an Engine</i>	589
<i>How Engines Work with SAS Files</i>	590
<i>Engine Characteristics</i>	591
<i>Read / Write Activity</i>	592
<i>Access Patterns</i>	592
<i>Levels of Locking</i>	592
<i>Indexing</i>	593
<i>About Library Engines</i>	593
<i>Definition of a Library Engine</i>	593
<i>Native Library Engines</i>	594
<i>Definition of Native Library Engine</i>	594
<i>Default Base SAS Engine</i>	594
<i>Remote Engine</i>	594
<i>SASESOCK Engine</i>	594
<i>SAS Scalable Performance Data (SPD) Engine</i>	594
<i>Sequential Engines</i>	595
<i>Transport Engine</i>	595
<i>V6 Compatibility Engine</i>	595
<i>Interface Library Engines</i>	595
<i>Special-Purpose Engines</i>	596
<i>Character Variable Padding (CVP) Engine</i>	596
<i>SAS Information Maps LIBNAME Engine</i>	596
<i>SAS Metadata LIBNAME Engine</i>	597
<i>SAS XML LIBNAME Engine</i>	597

Definition of a SAS Engine

An *engine* is a component of SAS software that reads from or writes to a file. Each engine enables SAS to access files that are in a particular format. There are several types of engines.

Specifying an Engine

Usually you do not have to specify an engine. If you do not specify an engine, SAS automatically assigns one based on the contents of the SAS library.

However, even though SAS will automatically assign an engine based on the library contents, it is more efficient for you to specify the engine. In some operating

environments, in order to determine the contents of a library, SAS must perform extra processing steps by looking at all of the files in the directory until it has enough information to determine which engine to use.

For example, if you explicitly specify the engine name as in the following LIBNAME statement, SAS does not need to determine which engine to use:

```
libname mylib v9 'SAS-library';
```

In order to use some engines, you must specify the engine name. For example, in order to use engines like the XML engine or the metadata engine, you must explicitly specify the engine name and specify specific arguments and options for that engine. For example, the following LIBNAME statement specifies the XML engine in order to import or export an XML document:


```
libname myxml xml 'c:\Myfiles\XML\Myxmlfile.xml' xmltype=generic;
```

You can specify an engine name in the LIBNAME statement, the ENGINE= system option, and in the New Library window.

How Engines Work with SAS Files

The following figure shows how SAS data sets are accessed through an engine.

Figure 36.1 How SAS Data Sets Are Accessed




- Your data is stored in files for which SAS provides an engine. When you specify a SAS data set name, the engine locates the appropriate file or files.

- The engine opens the file and obtains the descriptive information that is required by SAS, for example, which variables are available and what attributes they have, whether the file has special processing characteristics such as indexes or compressed observations, and whether other engines are required for processing. The engine uses this information to organize the data in the standard logical form for SAS processing.
- This standard form is called the *SAS data file*, which consists of the descriptor information and the data values organized into columns (variables) and rows (observations).
- SAS procedures and DATA step statements access and process the data only in its logical form. During processing, the engine executes whatever instructions are necessary to open and close physical files and to read and write data in appropriate formats.

Data that is accessed by an engine is organized into the SAS data set model, and in the same way, groups of files that are accessed by an engine are organized in the correct logical form for SAS processing. Once files are accessed as a SAS library, you can use SAS utility windows and procedures to list their contents and to manage them. See Chapter 25, “SAS Libraries,” on page 447 for more information about SAS libraries. The following figure shows the relationship of engines to SAS libraries.


Figure 36.2 Relationship of Engines to SAS Libraries



Engine Characteristics

The engine that is used to access a SAS data set determines its processing characteristics. Different statements and procedures require different processing characteristics. For example, the FSEDIT procedure requires the ability to update selected data values, and the POINT= option in the SET statement requires random access to observations as well as the ability to calculate observation numbers from record identifiers within the file.

The following figure describes the types of activities that engines regulate.

Figure 36.3 Activities That Engines Regulate

Read/Write Activity

An engine can

- limit read/write activity for a SAS data set to read-only
- fully support updating, deleting, renaming, or redefining the attributes of the data set and its variables
- support only some of these functions.

For example, the engines that process BMDP, OSIRIS, or SPSS files support read-only processing. Some engines that process SAS views permit SAS procedures to modify existing observations while others do not.

Access Patterns

SAS procedures and statements can read observations in SAS data sets in one of four general patterns:

sequential access	processes observations one after the other, starting at the beginning of the file and continuing in sequence to the end of the file.
random access	processes observations according to the value of some indicator variable without processing previous observations.
BY-group access	groups and processes observations in order of the values of the variables that are specified in a BY statement.
multiple-pass	performs two or more passes on data when required by SAS statements or procedures.

If a SAS statement or procedure tries to access a SAS data set whose engine does not support the required access pattern, SAS prints an appropriate error message in the SAS log.

Levels of Locking

Some features of SAS require that data sets support different levels at which update access is used. When a SAS data set can be opened concurrently by more than one SAS session or by more than one statement or procedure within a single session, the level of locking determines how many sessions, procedures, or statements can read and write to the file at the same time. For example, with the FSEDIT procedure, you can request two windows on the same SAS data set in one session. Some engines support this capability; others do not.

The levels that are supported are record level and member (data set) level. Member-level locking enables read access to many sessions, statements, or procedures, but restricts all other access to the SAS data set when a session, statement, or procedure acquires update or output access. Record-level locking enables concurrent read access and update access to the SAS data set by more than one session, statement, or procedure, but prevents concurrent update access to the same observation. Not all engines support both levels.

By default, SAS provides the greatest possible level of concurrent access, while guaranteeing the integrity of the data. In some cases, you might want to guarantee the integrity of your data by controlling the levels of update access yourself. Use the CNTLLEV= data set option to control levels of locking. CNTLLEV= enables locking at three levels:

- library
- data set
- observation.

Here are some situations in which you should consider using the CNTLLEV= data set option:

- your application controls access to the data, such as in SAS Component Language (SCL), SAS/IML software, or DATA step programming
- you access data through an interface engine that does not provide member-level control of the data.

For more information on the CNTLLEV= data set option, see *SAS Language Reference: Dictionary*.

You can also acquire an exclusive lock on an existing SAS file by issuing the LOCK global statement. After an exclusive lock is obtained, no other SAS session can read or write to the file until the lock is released. For more information on the LOCK statement, see *SAS Language Reference: Dictionary*.

Note: SAS products, such as SAS/ACCESS and SAS/SHARE, contain engines that support enhanced session management services and file locking capabilities. Δ

Indexing

A major processing feature of SAS is the ability to access observations by the values of key variables with indexes. See “Understanding SAS Indexes” on page 505 for more information on using indexes for SAS data files. Note that not all engines support indexing.

About Library Engines

Definition of a Library Engine

A *library engine* is an engine that accesses groups of files and puts them into a logical form for processing by SAS utility procedures and windows. A library engine also determines the fundamental processing characteristics of the library and presents lists of files for the library directory. Library engines can be classified as *native* or *interface*.

Native Library Engines

Definition of Native Library Engine

A *native library engine* is an engine that accesses forms of SAS files that are created and processed only by SAS.

Operating Environment Information: Engine availability is host dependent. See the SAS documentation for your operating environment. Also, specific products provide additional engines. \triangle

Default Base SAS Engine

The default Base SAS engine writes SAS libraries in disk format. The engine processes SAS 7, 8, and 9 files. If you do not specify an engine name when you are creating a new SAS library, the Base SAS engine, which for SAS 9 is named V9, is automatically selected.

When accessing existing SAS data sets on disk, SAS assigns an engine based on the contents of the library. The Base SAS engine

- is the only engine that supports the full functionality of the SAS data set and the SAS library.
- supports view engines.
- meets all the processing characteristics required by SAS statements and procedures.
- creates, maintains, and uses indexes.
- reads and writes compressed (variable-length) observations. SAS data sets created by other engines have fixed-length observations.
- assigns a permanent page size to data sets and temporarily assigns the number of buffers to be used when processing them.
- repairs damaged SAS data sets, indexes, and catalogs.
- enforces integrity constraints, creates backup files, and creates audit trails.

Note: SAS files created in SAS 7, 8, and 9 have the same file format. \triangle

Remote Engine

The REMOTE engine is a SAS library engine for SAS/SHARE software. Using it enables a SAS session to access shared data by communicating with a SAS server. See *SAS/SHARE User's Guide* for more information.

SASESOCK Engine

The SASESOCK engine processes input to and output from TCP/IP ports instead of physical disk devices. The SASESOCK engine is required for SAS/CONNECT applications that implement MP CONNECT processing with the piping mechanisms. See *SAS/CONNECT User's Guide* for more information.

SAS Scalable Performance Data (SPD) Engine

The SAS Scalable Performance Data Engine (SPD Engine) provides parallel I/O, using multiple CPUs to read SAS data and deliver it rapidly to applications. The SPD

Engine can process very large data sets because the data can span volumes but can be referenced as a single data set. The data in these data sets is also partitioned, enabling the data to be read in multiple threads per CPU. The SPD Engine is not intended to replace the default Base SAS engine for processing data sets that do not span volumes.

See *SAS Scalable Performance Data Engine: Reference* for details about this engine's capabilities.

Sequential Engines

A sequential engine processes SAS files on storage media that do not provide random access methods (for example, tape or sequential format on disk). A sequential engine requires less overhead than the default Base SAS engine because sequential access is simpler than random access. However, a sequential engine does not support some Base SAS features such as audit trails, generation data sets, integrity constraints, and indexing.

The sequential engine supports some file types for backup and restore purposes only, such as CATALOG, VIEW, and MDDDB. ITEMSTOR is the only file type that the sequential engine does not support. DATA is the only file type that is useful for purposes other than backup and restore.

The following sequential engines are available:

V9TAPE (TAPE)	processes SAS 7, 8, and 9 files.
V6TAPE	processes SAS 6 SAS files without requiring you to convert the file to the SAS 9 format.

For more information, see “Sequential Data Libraries” on page 457.

Transport Engine

The XPORT engine processes transport files. The engine transforms a SAS file from its operating environment-specific internal representation to a transport file, which is a machine-independent format that can be used for all hosts. In order to create a transport file, explicitly specify the XPORT engine in the LIBNAME statement, then use the DATA step or COPY procedure.

For information about using the XPORT engine, see *Moving and Accessing SAS Files*.

V6 Compatibility Engine

The V6 compatibility engine can automatically support some processing of SAS 6 files in SAS 9 without requiring you to convert the file to the SAS 9 format.

For more information, see Chapter 34, “SAS 9.2 Compatibility with SAS Files from Earlier Releases,” on page 577, or see the Migration Focus Area at support.sas.com.

Interface Library Engines

An *interface library engine* is a SAS engine that accesses files formatted by other software. Interface library engines are not transparent to the user and must be explicitly specified, for example, in the LIBNAME statement.

The following are interface library engines:

SPSS

reads SPSS portable file format, which is analogous to the transport format for SAS data sets. The SPSS portable files (also called an export file) must be created by using the SPSS EXPORT command. Under z/OS, the SPSS engine also reads

SPSS Release 9 files and SPSS-X files in either compressed or uncompressed format.

OSIRIS

reads OSIRIS data and dictionary files in EBCDIC format.

BMDP

reads BMDP save files.

In addition, a *view engine* is an interface library engine that is used by SAS/ACCESS software in order to retrieve data from files formatted by another vendor's software. These engines enable you to read and write data directly to and from files formatted by a database management system (DBMS), such as DB2 and ORACLE.

View engines enable you to use SAS procedures and statements in order to process data values stored in these files without the cost of converting and storing them in files formatted by SAS. Contact your on-site SAS support personnel for a list of the SAS/ACCESS interfaces available at your site. For more information about SAS/ACCESS features, see Chapter 32, "About SAS/ACCESS Software," on page 561 and the SAS/ACCESS documentation for your DBMS.

Operating Environment Information: The capabilities and support of these engines vary depending on your operating environment. See the SAS documentation for your operating environment for more complete information. Δ

Special-Purpose Engines

Character Variable Padding (CVP) Engine

The character variable padding (CVP) engine expands character variable lengths, using a specified expansion amount, so that character data truncation does not occur when a file requires transcoding. Character data truncation can occur when the number of bytes for a character in one encoding is different from the number of bytes for the same character in another encoding, such as when a single byte character set (SBCS) is transcoded to a double byte character set (DBCS) or a multi-byte character set (MBCS).

The CVP engine is a read-only engine for SAS data files only. You can request character variable expansion by either of the following methods:

- You can explicitly specify the CVP engine, for example, with the LIBNAME statement, and using the default expansion of 1.5 times the variable lengths.
- You can implicitly specify the CVP engine with the LIBNAME statement options CVPBYTES= or CVPMULTIPLIER=. The options specify the expansion amount. In addition, you can use the CVPENGINE= option to specify the primary engine to use for processing the SAS file; the default is the default Base SAS engine.

For more information about using the CVP engine to avoid character data truncation and for details on the CVP engine options in the LIBNAME statement, see *SAS National Language Support (NLS): Reference Guide*.

SAS Information Maps LIBNAME Engine

The new SAS Information Maps LIBNAME Engine provides a read-only way to access data generated from a SAS Information Map and to bring it into a SAS session. Once you retrieve the data, you can run almost any SAS procedure against it.

To use the Information Maps engine, you must explicitly specify INFOMAPS as the engine name, along with specific arguments and options in the LIBNAME statement.

For information about how to use the Information Maps engine, see *Base SAS Guide to Information Maps*.

SAS Metadata LIBNAME Engine

The metadata engine accesses metadata that is stored on the SAS Metadata Server within a specific SAS Metadata Repository. The metadata is information about the structure and content of data, and about the applications that process and manipulate that data. The metadata contains details such as the location of the data and the SAS engine that is used to process the data.

The metadata engine works in a similar way to other SAS engines. That is, you execute a LIBNAME statement in order to assign a libref and specify an engine. You then use that libref throughout the SAS session where a libref is valid. However, instead of the libref being associated with the physical location of a SAS library, the metadata libref is associated with specific metadata objects that are stored in a specific repository on the metadata server. The metadata objects define the SAS engine and options that are necessary to process a SAS library and its members.

When you execute the LIBNAME statement for the metadata engine, the metadata engine retrieves information about the target SAS library from the metadata. The metadata engine uses this information in order to construct a LIBNAME statement for the underlying engine and assigns it with the appropriate options. Then, when the metadata engine needs to access your data, the metadata engine uses the underlying engine to process the data.

You invoke the metadata engine by explicitly specifying the engine name META, along with specific arguments and options for the metadata engine, for example, in the LIBNAME statement or in the New Library window.

For information about how to use the metadata engine, see *SAS Metadata LIBNAME Engine: User's Guide*.


SAS XML LIBNAME Engine

The SAS XML engine imports an XML document as one or more SAS data sets and exports a SAS data set as an XML document.

- The engine imports (reads from an input file) an external XML document by translating the XML markup into SAS proprietary format.
- The engine exports (writes to an output file) an XML document from a SAS data set by translating SAS proprietary format to XML markup.

To use the XML engine, you must explicitly specify XML as the engine name, along with specific arguments and options, for example, in the LIBNAME statement or in the New Library window.

For information about how to use the XML engine, see *SAS XML LIBNAME Engine: User's Guide*.



CHAPTER

37

SAS File Management

<i>Improving Performance of SAS Applications</i>	599
<i>Moving SAS Files between Operating Environments</i>	599
<i>Repairing Damaged SAS Files</i>	600
<i>Recovering SAS Data Files</i>	600
<i>Recovering Indexes</i>	601
<i>Recovering Disabled Indexes and Integrity Constraints</i>	602
<i>Recovering Catalogs</i>	602

Improving Performance of SAS Applications

SAS offers tools to control the use of memory and other computer resources. Most SAS applications will run efficiently in your operating environment without using these features. However, if you develop applications under the following circumstances, you might want to experiment with tuning performance:

- You work with large data sets.
- You create production jobs that run repeatedly.
- You are responsible for establishing performance guidelines for a data center.
- You do interactive queries on large SAS data sets using SAS/FSP software.

For information on improving performance, see Chapter 12, “Optimizing System Performance,” on page 151.

Moving SAS Files between Operating Environments

The procedures for moving SAS files from one operating environment to another vary according to your operating environment, the member type and version of the SAS files you want to move, and the methods you have available for moving the files.

For details on this subject, see *Moving and Accessing SAS Files*.

Repairing Damaged SAS Files

The Base SAS engine detects possible damage to SAS data files (including indexes, integrity constraints, and the audit file) and SAS catalogs and provides a means for repairing some of the damage. If one of the following events occurs while you are updating a SAS file, SAS can recover the file and repair some of the damage:

- A system failure occurs while the data file or catalog is being updated.
- Damage occurs to the storage device where a data file resides. In this case, you can restore the damaged data file, the index, and the audit file from a backup device.
- The disk where the data file (including the index file and audit file) or catalog is stored becomes full before the file is completely written to it.
- An input/output error occurs while writing to the data file, index file, audit file, or catalog.

When the failure occurs, the observations or records that were not written to the data file or catalog are lost and some of the information about where values are stored is inconsistent. The next time SAS reads the file, it recognizes that the file's contents are damaged and repairs it to the extent possible in accordance with the setting for the `DLDMGACTION=` data set option or system option, unless the data set is truncated. In this case, use the `REPAIR` statement to restore the data set.

Note: SAS is unable to repair or recover a SAS view (a `DATA` step view, an `SQL` view, or a `SAS/ACCESS` view) or a stored compiled `DATA` step program. If a SAS file of type `VIEW` or `PROGRAM` is damaged, you must recreate it. Δ

Note: If the audit file for a SAS data file becomes damaged, you will not be able to process the data file until you terminate the audit trail. Then, you can initiate a new audit file or process the data file without one. Δ

Recovering SAS Data Files

To determine the type of action SAS will take when it tries to open a SAS data file that is damaged, set the `DLDMGACTION=` data set option or system option. That is, when a data file is detected as damaged, SAS will automatically respond based on your specification as follows:

`DLDMGACTION=FAIL`

tells SAS to stop the step without a prompt and issue an error message to the log indicating that the requested file is damaged. This specification gives the application control over the repair decision and provides awareness that a problem occurred.

To recover the damaged data file, you can issue the `REPAIR` statement in `PROC DATASETS`, which is documented in *Base SAS Procedures Guide*.

`DLDMGACTION=ABORT`

tells SAS to terminate the step, issue an error message to the log indicating that the request file is damaged, and end the SAS session.

`DLDMGACTION=REPAIR`

tells SAS to automatically repair the file and rebuild indexes, integrity constraints, and the audit file as well. If the repair is successful, a message is issued to the log indicating that the open and repair were successful. If the repair is unsuccessful, processing stops without a prompt and an error message is issued to the log indicating the requested file is damaged.

Note: If the data file is large, the time needed to repair it can be long. Δ

DLDMGACTION=NOINDEX

tells SAS to automatically repair the data file, disable the indexes and integrity constraints, delete the index file, update the data file to reflect the disabled indexes and integrity constraints, and limit the data file to be opened only in INPUT mode. A warning is written to the SAS log instructing you to execute the PROC DATASETS REBUILD statement to correct the disabled indexes and integrity constraints and rebuild the index file. For more information, see “Recovering Disabled Indexes and Integrity Constraints” on page 602.

DLDMGACTION=PROMPT

tells SAS to provide the same behavior that exists in Version 6 for both interactive mode and batch mode. For interactive mode, SAS displays a dialog box that asks you to select the FAIL, ABORT, or REPAIR action. For batch mode, the files fail to open.

For a data file, the date and time of the last repair and a count of the total number of repairs is automatically maintained. To display the damage log, use PROC CONTENTS as shown below:

```
proc contents data=sasuser.census;
run;
```

Output 37.1 Output of CONTENTS Procedure

The CONTENTS Procedure			
Data Set Name	SASUSER.CENSUS	Observations	27
Member Type	DATA	Variables	4
Engine	V9	Indexes	0
Created	10:06 Wednesday, January 24, 2007	Observation Length	32
Last Modified	10:06 Wednesday, January 24, 2007	Deleted Observations	0
Protection		Compressed	NO
Data Set Type		Sorted	NO
Label			
Data Representation	WINDOWS		
Encoding	wlatin1 Western (Windows)		
Engine/Host Dependent Information			
Data Set Page Size	4096		
Number of Data Set Pages	1		
First Data Page	1		
Max Obs per Page	126		
Obs in First Data Page	27		
Number of Data Set Repairs	0		
File Name	c:\winnt\profiles\sasxxx\sasuser\census.sas7bdat		
Release Created	9.2000A0		
Host Created	XP_PRO		
Alphabetic List of Variables and Attributes			
#	Variable	Type	Len
2	CrimeRate	Num	8
1	Density	Num	8
4	PostalCode	Char	2

Recovering Indexes

In addition to the failures listed earlier, you can damage the indexes for SAS data files by using an operating environment command to delete, copy, or rename a SAS data

file, but not its associated index file. The index is repaired similarly to the `DLDMGACTION=` option as described for SAS data files, or you can use the `REPAIR` statement in `PROC DATASETS` to rebuild composite and simple indexes that were damaged.

You cannot use the `REPAIR` statement to recover indexes that were deleted by one of the following actions:

- copying a SAS data file by some means other than `PROC COPY` or `PROC DATASETS`, for example, using a `DATA` step
- using the `FORCE` option in the `SORT` procedure to write over the original data file.

In the above cases, the index must be rebuilt using the `PROC DATASETS INDEX CREATE` statement.

Recovering Disabled Indexes and Integrity Constraints

When the `DLDMGACTION=NOINDEX` data set or system option is used and SAS encounters a damaged data file, SAS does the following:

- automatically repairs the data file without the indexes and integrity constraints
- disables the indexes and integrity constraints
- deletes the index file
- updates the data file to reflect the disabled indexes and integrity constraints
- limits the data file to be opened only in `INPUT` mode
- writes the following warning to the SAS log:

```
WARNING: SAS data file MYLIB.MYFILE.DATA was damaged and has
        been partially repaired. To complete the repair, execute the
        DATASETS procedure REBUILD statement.
```

The data file stays in `INPUT` mode until the `PROC DATASETS REBUILD` statement is executed. You use this statement to specify whether you want to restore the indexes and integrity constraints and rebuild the index file or delete the disabled integrity constraints and indexes. For more information, see the `DATASETS Procedure REBUILD` statement in *Base SAS Procedures Guide*.

Recovering Catalogs

To determine the type of action that SAS will take when it tries to open a SAS catalog that is damaged, set the `DLDMGACTION=` data set option or system option. Then when a catalog is detected as damaged, SAS will automatically respond based on your specification.

Note: There are two types of catalog damage:

- localized damage* is caused by a disk condition, which results in some data in memory not being flushed to disk. The catalog entries that are currently open for update are marked as damaged. Each damaged entry is checked to determine whether all the records can be read without error.
- severe damage* is caused by a severe I/O error. The entire catalog is marked as damaged.

Δ

`DLDMGACTION=FAIL`

tells SAS to stop the step without a prompt and issue an error message to the log indicating that the requested file is damaged. This specification gives the

application control over the repair decision and provides awareness that a problem occurred.

To recover the damaged catalog, you can issue the REPAIR statement in PROC DATASETS, which is documented in the *SAS Procedures Guide*. Note that when you use the REPAIR statement to restore a catalog, you receive a warning for entries that have possible damage. Entries that have been restored might not include updates that were not written to disk before the damage occurred.

DLDMGACTION=ABORT

tells SAS to terminate the step, issue an error message to the log indicating that the requested file is damaged, and end the SAS session.


DLDMGACTION=REPAIR

for localized damage, tells SAS to automatically check the catalog to see which entries are damaged. If there is an error reading an entry, the entry is copied. If an error occurs during the copy process, then the entry is automatically deleted. For severe damage, the entire catalog is copied to a new catalog.

DLDMGACTION=PROMPT

for localized damage, tells SAS to provide the same behavior that exists in Version 6 for both interactive mode and batch mode. For interactive mode, SAS displays a dialog box that asks you to select the FAIL, ABORT, or REPAIR action. For batch mode, the files fail to open. For severe damage, the entire catalog is copied to a new catalog.

Unlike data files, a damage log is not maintained for a catalog.



CHAPTER 38

External Files

<i>Definition of External Files</i>	605
<i>Referencing External Files Directly</i>	606
<i>Referencing External Files Indirectly</i>	606
<i>Referencing Many External Files Efficiently</i>	608
<i>Referencing External Files with Other Access Methods</i>	608
<i>Working with External Files</i>	609
<i>Reading External Files</i>	609
<i>Writing to External Files</i>	609
<i>Processing External Files</i>	610

Definition of External Files

external files

are files that are managed and maintained by your operating system, not by SAS. They contain data or text or are files in which you want to store data or text. They can also be SAS catalogs or output devices. Every SAS job creates at least one external file, the SAS log. Most SAS jobs create external files in the form of procedure output or output created by a DATA step.

External files used in a SAS session can store input for your SAS job as:

- records of raw data that you want to use as input to a DATA step
- SAS programming statements that you want to submit to the system for execution.

External files can also store output from your SAS job as:

- a SAS log (a record of your SAS job)
- a report written by a DATA step.
- procedure output created by SAS procedures, including regular list output, and, beginning in Version 7, HTML and PostScript output from the Output Delivery System (ODS).

The PRINTTO procedure also enables you to direct procedure output to an external file. For more information, see *Base SAS Procedures Guide*. See Chapter 9, “SAS Output,” on page 115 for more information about ODS.

Note: Database management system (DBMS) files are a special category of files that can be read with SAS/ACCESS software. For more information on DBMS files, see Chapter 32, “About SAS/ACCESS Software,” on page 561 and the SAS/ACCESS documentation for your DBMS. △

Operating Environment Information: Using external files with your SAS jobs entails significant operating-environment-specific information. Refer to the SAS documentation for your operating environment for more information. Δ

Referencing External Files Directly

To reference a file directly in a SAS statement or command, specify in quotation marks its physical name, which is the name by which the operating environment recognizes it, as shown in the following table:

Table 38.1 Referencing External Files Directly

External File Task	Tool	Example
Specify the file that contains input data.	INFILE	<pre>data weight; infile 'input-file'; input idno \$ week1 week16; loss=week1-week16;</pre>
Identify the file that the PUT statement writes to.	FILE	<pre>file 'output-file'; if loss ge 5 and loss le 9 then put idno loss 'AWARD STATUS=3'; else if loss ge 10 and loss le 14 then put idno loss 'AWARD STATUS=2'; else if loss ge 15 then put idno loss 'AWARD STATUS=1'; run;</pre>
Bring statements or raw data from another file into your SAS job and execute them.	%INCLUDE	<pre>%include 'source-file';</pre>

Referencing External Files Indirectly

If you want to reference a file in only one place in a program so that you can easily change it for another job or a later run, you can reference a filename indirectly. Use a FILENAME statement, the FILENAME function, or an appropriate operating system command to assign a *fileref* or nickname, to a file.* Note that you can assign a fileref to a SAS catalog that is an external file, or to an output device, as shown in the following table.

* In some operating environments, you can also use the command '&' to assign a fileref.

Table 38.2 Referencing External Files Indirectly

External File Task	Tool	Example
Assign a fileref to a file that contains input data.	FILENAME	<code>filename mydata 'input-file';</code>
Assign a fileref to a file for output data.	FILENAME	<code>filename myreport 'output-file';</code>
Assign a fileref to a file that contains program statements.	FILENAME	<code>filename mypgm 'source-file';</code>
Assign a fileref to an output device.	FILENAME	<code>filename myprinter <device-type> <host-options>;</code>
Specify the file that contains input data.	INFILE	<code>data weight; infile mydata; input idno \$ week1 week16; loss=week1-week16;</code>
Specify the file that the PUT statement writes to.	FILE	<code>file myreport; if loss ge 5 and loss le 9 then put idno loss 'AWARD STATUS=3'; else if loss ge 10 and loss le 14 then put idno loss 'AWARD STATUS=2'; else if loss ge 15 then put idno loss 'AWARD STATUS=1'; run;</code>
Bring statements or raw data from another file into your SAS job and execute them.	%INCLUDE	<code>%include mypgm;</code>

Referencing Many External Files Efficiently

When you use many files from a single aggregate storage location, such as a directory or partitioned data set (PDS or MACLIB), you can use a single fileref, followed by a filename enclosed in parentheses, to access the individual files. This saves time by eliminating the need to type a long file storage location name repeatedly. It also makes changing the program easier later if you change the file storage location. The following table shows an example of assigning a fileref to an aggregate storage location:

Table 38.3 Referencing Many Files Efficiently

External File Task	Tool	Example
Assign a fileref to aggregate storage location.	FILENAME	<code>filename mydir 'directory-or-PDS-name';</code>
Specify the file that contains input data.	INFILE	<code>data weight; infile mydir(qrt1.data); input idno \$ week1 week16; loss=week1-week16;</code>
Specify the file that the PUT statement writes to. ¹	FILE	<code>file mydir(awards); if loss ge 5 then put idno loss 'AWARD STATUS=3'; else if loss ge 10 then put idno loss 'AWARD STATUS=2'; else if loss ge 15 then put idno loss 'AWARD STATUS=1'; run;</code>
Bring statements or raw data from another file into your SAS job and execute them.	%INCLUDE	<code>%include mydir(whole.program);</code>

¹ SAS creates a file that is named with the appropriate extension for your operating environment.

Referencing External Files with Other Access Methods

You can assign filerefs to external files that you access with the following FILENAME access methods:

- CATALOG
- FTP
- TCP/IP SOCKET
- URL
- WebDAV.

Examples of how to use each method are shown in the following table:

Table 38.4 Referencing External Files with Other Access Methods

External File Task	Tool	Example
Assign a fileref to a SAS catalog that is an aggregate storage location.	FILENAME with CATALOG specifier	<code>filename mycat catalog 'catalog' <catalog-options>;</code>
Assign a fileref to an external file accessed with FTP.	FILENAME with FTP specifier	<code>filename myfile FTP 'external-file' <ftp-options>;</code>
Assign a fileref to an external file accessed by TCP/IP SOCKET in either client or server mode.	FILENAME with SOCKET specifier	<code>filename myfile SOCKET 'hostname: portno' <tcpip-options>;</code> or <code>filename myfile SOCKET ':portno' SERVER <tcpip-options>;</code>
Assign a fileref to an external file accessed by URL.	FILENAME with URL specifier	<code>filename myfile URL 'external-file' <url-options>;</code>
Assign a fileref to an external file accessed on a WebDAV server.	FILENAME with WEBDAV specifier	<code>filename myfile WEBDAV 'external-file' <webdav-options>;</code>

See *SAS Language Reference: Dictionary* for detailed information about each of these statements.

Working with External Files

Reading External Files

The primary reason for reading an external file in a SAS job is to create a SAS data set from raw data. This topic is covered in Chapter 20, “Reading Raw Data,” on page 311.

Writing to External Files

You can write to an external file by using:

- a SAS DATA step
- the External File Interface (EFI)
- the Export Wizard.

When you use a DATA step to write a customized report, you write it to an external file. In its simplest form, a DATA step that writes a report looks like this:

```
data _null_;
  set budget;
  file 'your-file-name';
  put variables-and-text;
run;
```

For examples of writing reports with a DATA step, see Chapter 19, “DATA Step Processing,” on page 285.

If your operating environment supports a graphical user interface, you can use the EFI or the Export Wizard to write to an external file. The EFI is a point-and-click graphical interface that you can use to read and write data that is not in SAS internal format. By using the EFI, you can read data from a SAS data set and write it to an external file, and you can read data from an external file and write it to a SAS data set. See *SAS/ACCESS Interface to PC Files: Reference* for more information on the EFI.

Note: If the data file you are passing to EFI is password protected, you will be prompted multiple times for your login ID and password. Δ


The Export Wizard guides you through the steps to read data from a SAS data set and write it to an external file. As a wizard, it is a series of windows that present simple choices to guide you through the process. See *SAS/ACCESS Interface to PC Files: Reference* for more information on the wizard.

Processing External Files

When reading data from or to a file, you can also use a DATA step to:


- copy only parts of each record to another file
- copy a file and add fields to each record
- process multiple files in the same way in a single DATA step
- create a subset of a file
- update an external file in place
- write data to a file that can be read in different computer environments
- correct errors in a file at the bit level.

For examples of using a DATA step to process external files, see Chapter 20, “Reading Raw Data,” on page 311.



Industry Protocols Used in SAS

- Chapter 39*.....**The SMTP E-Mail Interface** 613
- Chapter 40*.....**Universal Unique Identifiers** 617
- Chapter 41*.....**Internet Protocol Version 6 (IPv6)** 621



CHAPTER 39

The SMTP E-Mail Interface

<i>Sending E-Mail through SMTP</i>	613
<i>System Options That Control SMTP E-Mail</i>	613
<i>Statements That Control SMTP E-mail</i>	614
<i>FILENAME STATEMENT</i>	614
<i>FILE and PUT Statements</i>	614

Sending E-Mail through SMTP

You can send electronic mail programmatically from SAS using the SMTP (Simple Mail Transfer Protocol) e-mail interface. SMTP is available for all operating environments in which SAS runs. To send SMTP e-mail with SAS e-mail support, you must have an intranet or Internet connection that supports SMTP.

Some SMTP servers require just the user identification as the login ID while others require the full e-mail address. The SAS SMTP e-mail interface authenticates the user identification in the following order.

- 1 If the FROM= option is specified on the FILENAME, EMAIL Access Method statement, the SAS SMTP e-mail interface authenticates by using the user identification—for example, *userid*. If unsuccessful, then the SAS SMTP e-mail interface authenticates by using the full e-mail address—for example, *userid@domain.com*.
- 2 If the FROM=option is *not* specified on the FILENAME, EMAIL Access Method, the SAS SMTP e-mail interface checks the EMAILID system option and authenticates by using the user identification—for example, *userid*. If unsuccessful, then the SAS SMTP e-mail interface checks the EMAILID system option and authenticates by using the full e-mail address—for example, *userid@domain.com*.
- 3 If neither the FROM=option nor the EMAILID system option is specified, the SAS SMTP e-mail interface will look up your user ID and authenticate it.

For more information about sending e-mail from SAS, refer to the SAS documentation for your operating environment.

System Options That Control SMTP E-Mail

Several SAS system options control SMTP e-mail. Depending on your operating environment and whether the SMTP e-mail interface is supported at your site, you might need to specify these options at start up or in your SAS configuration file.

Operating Environment Information: To determine the default e-mail interface for your operating environment and to determine the correct syntax for setting system options, refer to the SAS documentation for your operating environment. \triangle

The EMAILSYS system option specifies which e-mail system to use for sending electronic mail from within SAS. For more information about the EMAILSYS system option, refer to the SAS documentation for your operating environment.

The following system options are specified only when the SMTP e-mail interface is supported at your site:

EMAILAUTHPROTOCOL=

specifies the authentication protocol for SMTP E-mail. For more information, see the EMAILAUTHPROTOCOL= system option in *SAS Language Reference: Dictionary*.

EMAILFROM

specifies whether the FROM e-mail option is required when sending e-mail by using either the FILE or FILENAME statements. For more information, see EMAILFROM system option in *SAS Language Reference: Dictionary*.

EMAILHOST

specifies the SMTP server that supports e-mail access for your site. For more information, see the EMAILHOST system option in *SAS Language Reference: Dictionary*.

EMAILPORT

specifies the port to which the SMTP server is attached. For more information, see the EMAILPORT system option in *SAS Language Reference: Dictionary*.

The following system options are specified with other e-mail systems, as well as SMTP:

EMAILID=

specifies the identity of the individual sending e-mail from within SAS. For more information, see the EMAILID= system option in *SAS Language Reference: Dictionary*.

EMAILPW=

specifies your e-mail login password. For more information, see the EMAILPW= system option in *SAS Language Reference: Dictionary*.

Statements That Control SMTP E-mail


FILENAME STATEMENT

In the FILENAME statement, the EMAIL (SMTP) access method enables you to send e-mail programmatically from SAS using the SMTP e-mail interface. For more information, see the FILENAME statement, EMAIL (SMTP) Access Method in *SAS Language Reference: Dictionary*.

FILE and PUT Statements

You can specify e-mail options in the FILE statement. E-mail options that you specify in the FILE statement override any corresponding e-mail options that you specified in the FILENAME statement.

In the 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. The PUT statement directives override any other e-mail options in the FILE and FILENAME statements.



CHAPTER

40

Universal Unique Identifiers

<i>Universal Unique Identifiers and the Object Spawner</i>	617
<i>What Is a Universal Unique Identifier?</i>	617
<i>What Is the Object Spawner?</i>	617
<i>Defining the UUID Generator Daemon</i>	617
<i>Installing the UUID Generator Daemon</i>	618
<i>Using SAS Language Elements to Assign UUIDs</i>	619
<i>UUIDGEN Function</i>	619
<i>UUIDCOUNT= System Option</i>	619
<i>UUIDGENHOST System Option</i>	619

Universal Unique Identifiers and the Object Spawner

What Is a Universal Unique Identifier?

A Universal Unique Identifier (UUID) is a 128-bit identifier that consists of date and time information, and the IEEE node address of a host. UUIDs are useful when objects such as rows or other components of a SAS application must be uniquely identified. For example, if SAS is running as a server and is distributing objects to several clients concurrently, you can associate a UUID with each object to ensure that a particular client and SAS are referencing the same object.

What Is the Object Spawner?

The object spawner is a program that runs on the server and listens for requests. When a request is received, the object spawner accepts the connection and performs the action that is associated with the port or service on which the connection was made. The object spawner can be configured to be a UUID Generator Daemon (UUIDGEN), which creates UUIDs for the requesting program. Currently, SAS can generate UUIDs only in the Windows operating environment. UUIDGEN generates UUIDs for SAS sessions that execute on hosts that do not have native UUID generation support.

Defining the UUID Generator Daemon

The definition of UUIDGEN is contained in a setup configuration file that you specify when you invoke the object spawner. This configuration file identifies the port

that listens for UUID requests, and, in operating environments other than Windows, the configuration file also identifies the UUID node. If you install UUIDGEN in an operating environment other than Windows, contact SAS Technical Support (<http://support.sas.com/techsup/contact/index.htm>) to obtain a UUID node. The UUID node must be unique for each UUIDGEN installation in order for UUIDGEN to guarantee truly unique UUIDs.

Here is an example of a UUIDGEN setup configuration file for an operating environment other than Windows:

```
#
## Define our UUID Generator Daemon. Since this UUIDGEN is
## executing on a UNIX host, we contacted SAS Technical
## Support to get the specified sasUUIDNode.
#
dn: sasSpawnercn=UUIDGEN,sascomponent=sasServer,cn=SAS,o=ABC Inc,c=US
objectClass: sasSpawner
sasSpawnercn: UUIDGEN
sasDomainName: unx.abc.com
sasMachineDNSName: medium.unx.abc.com
sasOperatorPassword: myPassword
sasOperatorPort: 6340
sasUUIDNode: 0123456789ab
sasUUIDPort: 6341
description: SAS Session UUID Generator Daemon on UNIX
```

Here is an example of a UUIDGEN setup configuration file for Windows:

```
#
## Define our UUID Generator Daemon. Since this UUIDGEN is
## executing in a Windows operating environment, we do not need to specify
## the sasUUIDNode.
#
dn: sasSpawnercn=UUIDGEN,sascomponent=sasServer,cn=SAS,o=ABC Inc,
c=US
objectClass: sasSpawner
sasSpawnercn: UUIDGEN
sasDomainName: wnt.abc.com
sasMachineDNSName: little.wnt.abc.com
sasOperatorPassword: myPassword
sasOperatorPort: 6340
sasUUIDPort: 6341
description: SAS Session UUID Generator Daemon on XP
```

Installing the UUID Generator Daemon

When you have created the setup configuration file, you can install UUIDGEN by starting the object spawner program (`objspawn`), and specifying the setup configuration file with the following syntax:

```
objspawn -configFile filename
```

The `configFile` option can be abbreviated as `-cf`.

filename specifies a fully qualified path to the UUIDGEN setup configuration file. Enclose pathnames that contain embedded blanks in single or double quotation marks. On Windows, enclose pathnames that contain embedded blanks in double quotation marks. On z/OS, specify the configuration file as follows:

//dsn:myid.objspawn.log for MVS files.
 //hfs:filename.ext for OpenEdition files.

On Windows, the objspawn.exe file is installed in the core\sasext folder in your installed SAS folder.

On UNIX, the objspawn file is installed in the utilities/bin directory in your installed SAS directory.

In the VMS operating environment, the OBJSPAWN_STARTUP.COM file executes the OBJSPAWN.COM file as a detached process. The OBJSPAWN.COM file runs the object spawner. The OBJSPAWN.COM file also includes other commands that your site might need in order to run the appropriate version of the spawner, to set the display node, to define a process level logical name that points to a template DCL file (OBJSPAWN_TEMPLATE.COM), and to perform any other necessary actions before the object spawner is started. The OBJSPAWN_TEMPLATE.COM file performs setup that is needed in order for the client process to execute. The object spawner first checks to see whether the logical name SAS\$OBJSPAWN_TEMPLATE is defined. If it is, the commands in the template file are executed as part of the command sequence used when starting the client session. You do not have to define the logical name.

Using SAS Language Elements to Assign UUIDs

If your SAS application executes on a platform other than Windows and you have installed UUIDGEN, you can use the following to assign UUIDs:

- UUIDGEN function
- UUIDCOUNT= system option
- UUIDGENDHOST systems option.

UUIDGEN Function


The UUIDGEN function returns a UUID for each cell. For more information, see the UUIDGEN function in *SAS Language Reference: Dictionary*.

UUIDCOUNT= System Option

The UUIDCOUNT= system option specifies the number of UUIDs to acquire each time the UUID Generator Daemon is used. For more information, see the UUIDCOUNT= system option in *SAS Language Reference: Dictionary*.

UUIDGENDHOST System Option

The UUIDGENDHOST system option identifies the operating environment and the port of the UUID Generator Daemon. For more information, see the UUIDGENDHOST system option in *SAS Language Reference: Dictionary*.



CHAPTER

41

Internet Protocol Version 6 (IPv6)

<i>Overview of IPv6</i>	621
<i>IPv6 Address Format</i>	621
<i>Examples of IPv6 Addresses</i>	622
<i>Example of Full and Collapsed IPv6 Address</i>	622
<i>Example of an IPv6 Address That Includes a Port Number</i>	622
<i>Example of an IPv6 Address That Includes a URL</i>	622
<i>Fully Qualified Domain Names (FQDN)</i>	622

Overview of IPv6

SAS 9.2 introduces support for the next generation of Internet Protocol, IPv6, which is the successor to the current Internet Protocol, IPv4. Rather than replacing IPv4 with IPv6, SAS 9.2 will support both protocols. There will be a lengthy transition period during which the two protocols will coexist.

A primary reason for the new protocol is that the limited supply of 32-bit IPv4 address spaces was being depleted. IPv6 will use a 128-bit address scheme, which provides more IP addresses than did IPv4.

IPv6 includes these benefits over IPv4:

- larger address space (128 bits rather than 32 bits)
- simplified header format
- automatic configuration
- more efficient routing
- improved quality of service and security
- compliance with regulatory requirements
- widespread use in global markets

IPv6 Address Format

IPv6 and IPv4 use different address formats. The following table compares the features of the protocols.

Table 41.1 Comparison of Features of the IPv6 and IPv4 Address Formats

Feature	IPv6	IPv4
Address Space	128-bit	32-bit
Representation	string	integer
Length (including Field Separators)	39	15

Feature	IPv6	IPv4
Field Separator	colon (:)	period (.)
Notation	hexadecimal	decimal
Example of IP Address	db8:0:0:1	10.23.2.3

Examples of IPv6 Addresses

Example of Full and Collapsed IPv6 Address

Here is an example of a full IPv6 address:

```
FE80:0000:0000:0000:0202:B3FF:FE1E:8329
```

It shows a 128-bit address in eight 16-bit blocks in the format *global:subnet:interface*. Here is an example of a collapsed IPv6 address:

```
FE80::0202:B3FF:FE1E:8329
```

The :: (consecutive colons) notation can be used to represent four successive 16-bit blocks that contain zeros. When SAS software encounters a collapsed IP address, it reconstitutes the address to the required 128-bit address in eight 16-bit blocks.

Example of an IPv6 Address That Includes a Port Number

Here is an example of an IP address that contains a port number:

```
[2001:db8:0:1]:80
```

The brackets are necessary only if also specifying a port number. Brackets are used to separate the address from the port number. If no port number is used, the brackets can be omitted.

As an alternative, the block that contains a zero can be collapsed. Here is an example:

```
[2001:db8::1]:80
```

Example of an IPv6 Address That Includes a URL

Here is an example of an IP address that contains a URL:

```
http://[2001:db8:0:1]:80
```

The **http://** prefix specifies a URL. The brackets are necessary only if also specifying a port number. Brackets are used to separate the address from the port number. If no port number is used, the brackets can be omitted.

Fully Qualified Domain Names (FQDN)

Because IP addresses can change easily, SAS applications that contain hard-coded IP addresses are prone to maintenance problems.

To avoid such problems, use of an FQDN is preferred over an IP address. The name-resolution system that is part of the TCP/IP protocol is responsible for locating the IP address that is associated with the FQDN.

The following example restores client activity in the paused repository:

```
PROC METAOPERATE
  SERVER="d6292.us.company.com"
  PORT=2222
  USERID="myuserid"
  PASSWORD="mypassword"
  PROTOCOL=BRIDGE


  ACTION=RESUME
  OPTIONS=" "
  NOAUTOPAUSE;
```

If an IP address had been used and if the IP address that was associated with the computer node name had changed, the code would be inaccurate.

An FQDN can remain intact in the code while the underlying IP address can change without causing unpredictable results. The TCP/IP name-resolution system automatically resolves the FQDN to its associated IP address.

Here is an example of an FQDN that is specified in a SAS GUI application.

Display 41.1 Example of an FQDN in a SAS Management Console Window




The full FQDN, **d11076.na.apex.com**, is specified in the **Remote Host** field of the Connect Server Properties window in SAS Management Console.

Some SAS products impose limits on the length for computer names.

The following code is an example of an FQDN that is assigned to a SAS menu variable:


```
%let sashost=hrmach1.dorg.com;
rsubmit sashost.sasport;
```

Because the FQDN is longer than eight characters, the FQDN must be assigned to a SAS macro variable, which is used in the RSUBMIT statement.



Appendix

Appendix 1 **Recommended Reading** 627



APPENDIX

1

Recommended Reading

Recommended Reading 627

Recommended Reading

Here is the recommended reading list for this title:

- Base SAS Procedures Guide*
- SAS Language Reference: Dictionary*
- SAS Logging: Configuration and Programming Reference*
- SAS Metadata LIBNAME Engine: User's Guide*
- SAS National Language Support (NLS): Reference Guide*
- SAS Output Delivery System: User's Guide*
- SAS Scalable Performance Data Engine: Reference*
- SAS Table Server: TSPL Language Reference*
- SAS Table Server: TSSQL Language Reference*
- SAS XML LIBNAME Engine: User's Guide*

The recommended reading list from SAS Press includes the following titles:

- Learning SAS by Example: A Programmer's Guide*
- Output Delivery System: The Basics, Second Edition*
- The Little SAS Book: A Primer, Fourth Edition*

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

* For other SAS Institute business, call (919) 677-8000.

Customers outside the United States and Canada, please contact your local SAS office for assistance.

Index

A

- About SAS 9 260
- access descriptors 536, 564
- ACCESS procedure
 - interface view engine and 564
- ADD method
 - storing and retrieving data 398
- additional fonts 213, 214, 219
- address formats (IPv6) 621
- Adobe SVG Viewer 226
- alignment
 - of variable values 32
- _ALL_ name list 35
- alter protection 581, 587
- ampersands
 - in name literals 23
- AND operator 71
- appender objects 392
- appending
 - data to indexed data files 526
 - fast-append feature 481
 - files 358
 - generation groups 491
- Application Response Measurement (ARM) 163
- applications
 - CPU-bound 162
 - I/O-bound 161
 - performance of 599
- arithmetic operators 66
 - in WHERE expressions 138
- ARM 163
- array bounds 439
 - determining 440
 - HBOUND function 440
 - HBOUND function versus DIM function 440
 - in two-dimensional arrays 441
 - LBOUND function 440
 - upper and lower bounds 439
- array-name 429
- array processing 429
 - definition 430
 - examples 441
 - simple arrays 432
 - terminology 429
- Array Reference statement 432
- array references 430, 431, 435
- ARRAY statement 431
- arrays 429
 - action on all numeric variables 444
 - assigning initial values to elements 442
 - character variables in 441
 - conceptual view of 430
 - defining 431
 - defining quickly 436
 - defining the number of elements 435
 - determining the number of elements 436
 - DO loops 433
 - DO loops for selected elements 433
 - DO UNTIL expressions 436
 - DO WHILE expressions 436
 - grouping variables in 432
 - Java objects and 414
 - multidimensional 430, 437
 - one-dimensional 430
 - referencing 431, 435
 - selecting current variable 434
 - simple 432
 - syntax for defining and referencing 431
 - temporary 443
 - two-dimensional 431
 - variable lists 436
- aspect ratio 231
- assignment statement
 - creating variables 29
- _AT*_ variables 479
- _ATOPCODE_ values 479
- ATTRIB statement
 - creating variables 31
- audit trails 479
 - capturing rejected observations 485
 - CEDA processing and 483
 - considerations 481
 - data file update 485
 - definition 479
 - description 479
 - encryption with 588
 - examples 484
 - fast-append feature and 481
 - in shared environment 481
 - initiating 482, 484
 - passwords with 588
 - performance and 481
 - preservation by other operations 481
 - programming considerations 481
 - reading 482
 - resuming 482
 - status of 482
 - suspending 482
 - terminating 482

- autoexec files 9
- automatic naming convention 466
- automatic numeric-character conversion 64
- automatic variables 33
 - _ERROR_ 33
 - _IORC_ 382
 - _N_ 33

B

- base 38
- BASE engine 568
- Base SAS 4
 - concepts 10
- Base SAS engine 594
- base version 487
- batch mode 9
 - defining printers for 205
 - SAS log in 119
- batch programs
 - restarting 112
- BETWEEN-AND operator 140
- bias 39
- big endian platforms 324
- binary data 323
- binary informats 324
- bit masks 62
- bit testing 62
- blanks
 - in constants 63
 - reading data containing 314
- BMDP engine 596
- Boolean numeric expressions 72
- Boolean operators 70
 - in WHERE expressions 144
- boundaries of intervals 91
- browsers
 - printing SVG documents from 241
 - supporting PNG format 242
 - supporting SVG 226
- buffers
 - index requirements for 510
 - input 288
- BUFNO= system option
 - I/O optimization and 156
- BUFSIZE= system option
 - I/O optimization and 156
- BY-group processing 131, 329
 - data grouped by formatted values 338
 - data not in alphabetic or numeric order 338
 - in DATA step 336
 - indexing for 333, 521
 - indexing for, with WHERE processing 522
 - invoking 332
 - preprocessing 332
 - preprocessing input data for 333
 - references for 131
 - sorting observations for 333
 - syntax 330
 - terminology 329
- BY groups 330
 - DATA step identification of 334
 - engine access 592
 - processing conditionally 337
 - processing in DATA step 336
 - processing observations in 334

- with multiple BY variables 331
 - with single BY variable 331
- BY values 330
 - interleaving data sets and 362
- BY variables 329
 - BY groups with multiple 331
 - BY groups with single 331
 - interleaving data sets and 361
 - match-merge with duplicate values of 372
 - updating data sets with duplicate values of 378

C

- calculations
 - propagation of missing values in 53
- calendar intervals, retail 95
- catalog concatenation 557
 - CATNAME 557, 558
 - LIBNAME 557
 - rules for 560
 - terminology 557
- catalog directory windows 555
- CATALOG procedure 554
- CATALOG window 554
- catalogs 4, 553
 - accessing information in 554
 - concatenating 557
 - management tools 554
 - names of 553
 - Profile catalog 555
 - recovering 602
 - remote access and 14
- CATCACHE= system option
 - I/O optimization and 156
- CATNAME catalog concatenation 557, 558
- CEDA processing 567
 - advantages of 568
 - alternatives to 573
 - audit trails and 483
 - automatically processing foreign files 574
 - changing data representation in 575
 - compatibility across environments 571
 - creating files in foreign data representation 574
 - creating files in foreign environment 575
 - examples 574
 - file processing 568, 570
 - indexes and 526
 - integrity constraints and 498
 - output processing 569
 - processing types supported by 568
 - restrictions 569
 - terminology 567
 - update processing 572
 - when CEDA is used 570
- cell values
 - editing 272
- centering output 127
- CEXIST function 554
- character comparisons 69
 - IN operator in 69
- character constants 59
 - compared with character variables 59
 - in hexadecimal notation 60
 - quotation marks with 59
- character data
 - reading 314

- character missing values 322
- _CHARACTER_ name list 35
- character-to-numeric conversions
 - missing values and 53
- character values 313
- character variable padding (CVP) engine 596
- character variables 26
 - compared with character constants 59
 - converting to numeric 32, 53
 - in arrays 441
 - sort order of missing values 52
- characteristic 39
- checkpoint mode 110
 - setting up and executing 111
- CIMPORT procedure 574
- class loader 422
- CLASSPATH environment variable 409
- clearing WHERE expressions 276
- collating sequence
 - character comparisons and 69
- colon lists 465
- colon modifier (:)
 - comparing values 67
 - with IN operator 67
- colors
 - adding programmatically 172
 - adding with Registry Editor 172
 - controlling with SAS registry 171
- column-binary data 325
- column-binary data storage 325, 326
- column-binary informats 326
- column input 318
- columns
 - moving 269
 - sorting by values 271
 - temporarily changing headings 270
- combining data sets 343, 345
 - access methods 348
 - appending files 358
 - concatenating 349, 356
 - correct order 355
 - data relationships 345
 - direct access 348
 - error checking 353
 - interleaving 349, 359
 - match-merging 351, 370
 - methods for 349, 356
 - one-to-one merging 350, 366
 - one-to-one reading 350, 363
 - preparing data sets 354
 - procedures for 352
 - sequential access 348
 - statements for 352
 - testing programs 355
 - tools for 343, 352
 - troubleshooting 354
 - updating 351, 374
- command line 258
 - Help in 259
- commands
 - operating environment 460
 - window commands 260
- comment tags, unmatched 102
- comparison operators 66
 - in WHERE expressions 138
- comparisons
 - character 69
 - numeric 67
- compatibility
 - See version compatibility
- compilation phase (DATA step) 288
- component objects
 - See DATA step component objects
- composite indexes 508
 - NOMISS 518
- compound expressions 58
 - order of evaluation 73
- compound optimization 508, 516
- compound WHERE expressions 134
 - processing 145
- COMPRESS= system option
 - I/O optimization and 156
- compressing data files 526
 - definition of compression 526
 - disabling compression requests 527
 - requesting compression 527
- concatenating catalogs
 - See catalog concatenation
- concatenating data sets 349, 356
 - DATA step processing during 356
 - efficiency 359
 - SQL for 358
 - syntax 356
 - with DATA step 357
- concatenating libraries 452
- concatenation operator 72
 - in WHERE expressions 144
- configuration
 - Explorer 176
 - file shortcuts 177
 - filerefs 177
 - libraries 177
 - librefs 177
 - SAS registry 176
 - Universal Printing 176, 204
- configuration data storage
 - See SAS registry
- configuration files 9
- console log 116
- constants 58
 - bit testing 62
 - blank space in 63
 - character 59
 - date 61
 - datetime 61
 - in WHERE expressions 137
 - misinterpretation of 63
 - numeric 60
 - time 61
- constructors
 - initializing hash objects 393
- CONTAINS operator 140
- copies
 - encryption with 588
 - passwords with 588
- copying data sets 267
- CPORT procedure 574
- CPU-bound applications 162
- CPU performance 158
 - increasing memory 158
 - parallel processing and 159

- program compilation optimization and 159
 - reducing I/O 158
 - search time for executable files 158
 - storing compiled programs 158
 - variable lengths and 159
 - CPU time 152
 - cross-environment data access
 - See* CEDA processing
 - current list item 397
 - custom class loader 422
 - custom intervals 95
 - CVP engine 596
- D**
- D-floating format 41
 - damaged files 600
 - data errors 106
 - format modifiers for reporting 108
 - data files 5, 461, 477, 591
 - as DATA step output 15
 - as input data source 14
 - audit trails 479
 - compressing 526
 - creating with DATA step 299
 - encryption 588
 - generation data sets 487
 - indexes 505
 - integrity constraints 493
 - interface 477
 - native 477
 - recovering 600
 - versus views 477
 - data management
 - exporting a subset 276
 - importing data into a table 279
 - in windowing environment 263
 - subsetting data with WHERE expressions 273
 - with Explorer 264
 - with VIEWTABLE 268
 - data pairs 397
 - data relationships 345
 - many-to-many 347
 - many-to-one 346
 - one-to-many 346
 - one-to-one 346
 - data representation 567
 - foreign 574
 - output processing and 569
 - data set lists 464
 - data set names 463
 - how and when assigned 463
 - one-level 464
 - parts of 463
 - two-level 464
 - where to use 463
 - data set options
 - controlling index usage 519
 - dropping, keeping, or renaming variables with 36
 - SAS/ACCESS librefs with 562
 - data sets 5, 285, 461
 - accessing through engines 590
 - as input data source 14
 - as SAS files 4
 - assigning passwords to 583
 - automatic naming convention 466
 - BY-group processing and 131
 - calculating size for performance 160
 - combining 343, 345
 - concatenating 349, 356
 - copying 267
 - creating for I/O optimization 155
 - creating with DATA step 299
 - default 466
 - defining multiple printers 206
 - descriptor information 288, 462
 - editing 473
 - generation 487
 - interleaving 349, 359
 - logical components of 462
 - management tools 472
 - master data set 374, 381
 - match-merging 351, 370
 - missing values when reading 53
 - modifying 343
 - names for 463
 - null 465
 - one-to-one merging 350, 366
 - one-to-one reading 350, 363
 - reading 302, 343, 344
 - referencing groups of 464
 - renaming 266
 - representation of 5
 - saving hash object data in 404
 - sort indicator 466, 471
 - sorted 466
 - sorting in a library 267
 - structure and contents of 354
 - transaction data set 374
 - updating 351, 374
 - updating, with missing values 379
 - updating, with new variables 379
 - updating, with nonmatched observations 379
 - viewing 473
 - viewing contents of 268
 - viewing properties of 267
 - viewing with Explorer 264
 - writing observations to 291
 - DATA statement
 - writing to log with /NESTING option 122
 - DATA step 4
 - See also* stored compiled DATA step programs
 - assigning passwords 582
 - checking for missing values in 56
 - concatenating data sets 357
 - concepts 10
 - creating data files 299
 - creating data sets 299
 - creating HTML reports 308
 - creating views 299
 - generating data from programming statements 302
 - identifying BY groups 334
 - input data sources 299
 - language elements and 6
 - ODS and 310
 - password-protected files in 584
 - processing BY groups in 336
 - reading from data sets 302
 - reading raw data 300
 - reasons for using 285
 - SAS processing and 15
 - setting values to missing 55

- writing reports with 303
- DATA step Component Interface 392
- DATA step component objects 392
 - hash iterator object 406
 - hash object 392
- DATA step debugger 4, 114
- DATA step output 15
- DATA step processing 286
 - altering flow for a given observation 295
 - changing default sequence of execution 295
 - compilation phase 288
 - default sequence of execution 294
 - during concatenation 356
 - during interleaving 360
 - during match-merging 371
 - during one-to-one merging 366
 - during one-to-one reading 364
 - example 289
 - execution phase 288
 - flow of action 286
 - step boundaries 297
 - troubleshooting execution 298
 - with UPDATE statement 375
- DATA step programs
 - See* stored compiled DATA step programs
- DATA step views 531
 - additional output files 533
 - compared with PROC SQL views 536
 - compared with stored compiled DATA step programs 532, 546
 - creating 531
 - definition 531
 - examples 533
 - merging data for reports 533
 - passwords and 588
 - performance 533
 - restrictions and requirements 532
 - uses for 532
- data transfer services 574
- data types 312
 - Java type set 412
- data values 5, 290, 312
- database management system (DBMS) files 6
- database management systems (DBMS) 561
- DATASETS procedure
 - creating indexes 514
 - creating integrity constraints 499
 - data set lists and 464
- date constants 61
- date durations 88
- date intervals 88
 - boundaries of 91
 - by category 89
 - multi-unit 93
 - multi-week 93
 - shifted intervals 94
 - single-unit 92
 - syntax 88
- date values 77
 - as recognizable dates 86
 - calculating 87
 - formats and 80
 - in output 127
 - informats and 80
 - integrity of 80
 - reading 87
 - tools by task 81
 - writing 87
 - year 2000 and 78
 - year digits 78, 79
- datetime constants 61
- datetime intervals 88
 - boundaries of 91
 - by category 89
 - multi-unit 93
 - multi-week 93
 - shifted intervals 94
 - single-unit 92
 - syntax 88
- datetime values 78
 - as recognizable dates and times 86
 - formats and 80
 - in output 127
 - informats and 80
 - integrity of 80
 - tools by task 81
 - year 2000 and 78
 - year digits 78, 79
- DBLOAD procedure 565
- DBMS (database management systems) 561
- DBMS files 6
- debugging 4, 99
 - DATA step debugger 114
 - logic errors 114
- declaring
 - hash iterator object 406
 - hash objects 393
 - Java objects 411
- default Base SAS engine 594
- default data sets 466
- descriptor information 5, 288, 462
- device types
 - Universal Printing 209, 211
- DICTIONARY tables 549
 - performance 551
 - viewing 549
 - viewing subset of 551
 - viewing summary of 550
- DIM function
 - determining number of array elements 436
 - versus HBOUND function 440
- direct access
 - combining data sets 348
- directives
 - naming SAS logs 120
- directories, library 459
- disk space
 - index requirements for 511
 - integrity constraints and 497
- DLDMGACTION= data set option
 - recovering catalogs 602
 - recovering data files 600
- DLDMGACTION= system option 600
- DO loops 433
 - nested 438
 - processing selected array elements 433
- DO UNTIL expressions 436
- DO WHILE expressions 436
- domain names
 - fully qualified 622
- double-precision floating-point numbers 46
- drop-down menus 256

- DROP statement
 - I/O optimization and 154
 - dropping variables 36
 - examples 37
 - input or output data sets for 36
 - order of application 37
 - statements versus data set options 36
 - duration 88
 - example 91
 - dynamic LIBNAME engine 561
- E**
- e-mail
 - See SMTP e-mail interface
 - embedding SAS/ACCESS LIBNAME statement 562
 - encoded passwords 586
 - encoding 568
 - output processing and 569
 - encryption 588
 - audit trails with 588
 - copies with 588
 - generation data sets with 588
 - indexes with 588
 - engines 589
 - access patterns 592
 - characteristics of 591
 - CVP engine 596
 - data set access 590
 - dynamic LIBNAME 561
 - I/O optimization and 156
 - indexing 593
 - interface DATA step 565
 - interface view 564
 - levels of locking 592
 - libraries and 591
 - library engines 449, 593
 - library engines and compatibility 578
 - metadata LIBNAME engine 597
 - read/write activity 592
 - SAS files and 590
 - SAS Information Maps LIBNAME Engine 596
 - SAS XML LIBNAME engine 597
 - SPD 161
 - special-purpose 596
 - specifying 589
 - supported by CEDA processing 568
 - environment, foreign 568
 - changing data representation in 575
 - creating files in 575
 - environment, native 568
 - environment variables
 - CLASSPATH 409
 - _ERROR_ automatic variable 33
 - error checking
 - combining data sets 353
 - importance of 382
 - indexes and 382
 - on all statements using KEY= option 386
 - routing execution, when unexpected condition occurs 383
 - tools for 382
 - error processing 99, 108
 - checkpoint mode and restart mode 110
 - _IORC_ automatic variable and 114
 - log control options 114
 - multiple errors 109
 - other options for 114
 - return codes 114
 - syntax check mode 108
 - system options for 113
 - error reporting
 - format modifiers for 108
 - ERROR statement
 - writing to log with 122
 - error types 99
 - data 106
 - execution-time 103
 - logic 114
 - macro-related 108
 - semantic 102
 - summary of 99
 - syntax 100
 - ERRORS= system option
 - error-checking with 114
 - executable files
 - reducing search time for 158
 - execution phase (DATA step) 288
 - execution-time errors 103
 - examples 104
 - out-of-resources condition 104
 - Explorer 250
 - assigning file shortcuts 265
 - backing up SASUSER registry 170
 - configuring with SAS registry 176
 - copying data sets 267
 - data management with 264
 - renaming data sets 266
 - sorting data sets in a library 267
 - tree view 250
 - viewing data set properties 267
 - viewing libraries and data sets 264
 - exponent 38
 - exporting
 - printer definitions 208
 - registry files 176
 - subset of data 276
 - expressions 58
 - automatic numeric-character conversion 64
 - Boolean numeric 72
 - compound 58
 - constants in 58
 - DO UNTIL 436
 - DO WHILE 436
 - examples 58
 - functions in 65
 - logical (Boolean) operators and 70
 - operators in 65
 - order of evaluation 73
 - simple 58
 - variables in 64
 - WHERE expressions 58
 - external data files
 - as DATA step output 15
 - external files 5, 605
 - as raw data source 14
 - processing 610
 - reading 300, 609
 - reading, with leading blanks and semicolons 314
 - reading raw data 316
 - referencing directly 606
 - referencing indirectly 606

referencing many files efficiently 608
 referencing with FILENAME access methods 608
 writing to 609

F

fast-append feature 481
 file encryption 588
 file formats
 SAS 9 577
 file management
 applications performance 599
 moving SAS files between operating environments 599
 repairing damaged files 600
 file processing
 with CEDA 568
 file protection
 See also passwords
 assigning with PW= data set option 585
 complete protection 585
 encryption 588
 file shortcuts 265
 configuring with SAS registry 177
 FILE statement
 SMTP e-mail 614
 File Transfer Protocol (FTP) 14
 file types 448
 FILENAME access methods 608
 filename extensions
 SAS 9 578
 FILENAME statement
 SMTP e-mail 614
 filerefs 606
 configuring with SAS registry 177
 files 4
 See also external files
 See also SAS files
 appending 358
 autoexec files 9
 configuration files 9
 data files 5
 DBMS files 6
 executable 158
 external data files 15
 foreign 568, 574
 HTML files 16
 importing non-standard 282
 importing standard 279
 moving between operating environments 599
 native 568
 password-protected 584
 procedure output files 15
 repairing when damaged 600
 FIND method
 storing and retrieving data 398, 399
 FIRSTOBS= data set option
 I/O optimization and 154
 segmenting a subset 147
 FIRST.variable 330, 334
 floating-point representation 38
 double-precision versus single-precision 46
 fractions 42
 IBM mainframes 39
 IEEE standard 42
 minimum number of bytes for storage 45
 numeric comparisons 43
 OpenVMS 41
 precision versus magnitude 42
 storing numbers with less precision 43
 transferring data between operating systems 46
 troubleshooting 39
 truncating numbers 45
 fonts
 See also TrueType fonts
 additional 213, 214, 219
 footnotes 127
 foreign data representation
 creating files in 574
 foreign environment 568
 changing data representation in 575
 creating files in 575
 foreign files 568
 automatically processing 574
 foreign key integrity constraints 494
 format modifiers
 for error reporting 108
 FORMAT statement
 creating variables 30
 formats
 as variable attribute 27
 data grouped by formatted values 338
 date values and 80
 datetime values and 80
 time values and 80
 formatted input 319
 formatting characters 127
 Forms printing 183, 210
 FQDN 622
 fractions 42
 FTP 14
 FULLSTIMER system option 152
 interpreting statistics 152
 fully bounded range condition 139
 fully qualified domain names (FQDN) 622
 functions
 changing DATA step execution sequence 295
 in expressions 65
 in WHERE expressions 136

G

general integrity constraints 493
 generation data sets 487
 base version 487
 definition 487
 deleting 492
 encryption with 588
 GENMAX= data set option 488
 GENNUM= data set option 488
 historical versions 488
 invoking 488
 maintaining 488
 oldest version 488
 passwords with 588
 rolling over 488
 terminology 487
 youngest version 488
 generation groups 488
 appending 491
 copying 491
 deleting versions 492
 displaying data set information 491

- managing 491
- modifying number of versions 491
- naming 489
- passwords in 493
- processing specific versions 490
- renaming versions 493
- generation numbers 488
- GENMAX= data set option 488
- GENNUM= data set option 488
- Getting Started with SAS 259
- Ghostview previewer 208
- GIF format 184
- global statements
 - stored compiled DATA step programs and 544
- graphic symbol TrueType fonts 213
- grouped observations
 - BY-group processing 131

H

- hash iterator objects 392, 406
 - declaring and instantiating 406
 - retrieving hash object data 407
- hash objects 392
 - attributes 405
 - comparing 405
 - declaring and instantiating 393
 - defining keys and data 395
 - initializing with constructor 393
 - maintaining key summaries 399
 - non-unique key and data pairs 397
 - reasons for using 392
 - replacing and removing data in 402
 - retrieving data with hash iterator 407
 - saving data in data sets 404
 - storing and retrieving data 397
- HBOUND function 440
 - versus DIM function 440
- headings
 - temporarily changing column headings 270
- Help 259
 - in command line 259
 - in individual windows 260
 - Migration Focus Area 577
- Help menu
 - opening from toolbar 259
- hexadecimal notation
 - character constants in 60
 - numeric constants in 61
- historical versions 488
- HTML documents
 - embedding SVG documents in 237
- HTML files
 - as DATA step output 16
- HTML reports
 - creating with ODS and DATA step 308

I

- I/O
 - cost of indexes 510
 - reducing for CPU performance 158
 - threaded 161
- I/O-bound applications 161
- I/O optimization 153
 - BUFNO= system option 156
 - BUFSIZE= system option 156
 - CATCACHE= system option 156
 - COMPRESS= system option 156
 - creating data sets 155
 - DROP statement 154
 - engine efficiency 156
 - FIRSTOBS= data set option 154
 - indexes 155
 - KEEP statement 154
 - LENGTH statement 154
 - OBS= data set option 154
 - SASFILE statement 157
 - views for data access 155
 - WHERE processing 153
- IBM 370 mode 324
- IBM mainframes
 - floating-point representation 39
- IEEE standard
 - floating-point representation 42
- illegal character-to-numeric conversions 53
- illegal operations
 - missing values and 53
- images
 - in SVG documents 240
- importing
 - data into a table 279
 - non-standard files 282
 - registry files 175
 - standard files 279
- IN= data set option
 - creating variables 31
- IN operator
 - colon modifier (:) with 67
 - in character comparisons 69
 - in numeric comparisons 68
 - in WHERE expressions 139
- INDEX= data set option
 - creating indexes 514
- index files 506
- index type 28
- indexed data files
 - adding observations to 525
 - appending data to 526
 - copying 524
 - sorting 525
 - updating 525
- indexes 348, 462, 505
 - benefits of 506
 - buffer requirements 510
 - CEDA processing and 526
 - composite 508
 - compound optimization 508, 516
 - costs of 509
 - CPU cost 509
 - creating 513
 - creating with DATASETS procedure 514
 - creating with INDEX= data set option 514
 - creating with SQL procedure 514
 - data file considerations 511
 - definition 505
 - disk space requirements 511
 - displaying data file information 523
 - encryption with 588
 - engines and 593
 - error checking 382
 - for BY-group processing 333, 521

- guidelines for creating 511
- I/O cost 510
- I/O optimization and 155
- integrity constraints and 496, 526
- key variable candidates 512
- maintaining 523
- missing values 509
- multiple occurrences of values 525
- NOMISS composite 518
- passwords with 588
- recovering 601
- recovering when damaged 526
- recovering when disabled 602
- simple 507
- specifying with KEY= option 522
- taking advantage of 523
- types of 507
- unique values 508
- updating data sets 376, 382
- usage information in log 520
- use considerations 512
- WHERE and BY processing with 522
- WHERE-processing performance and 146
- WHERE processing with 515
- with MODIFY statement 376
- with views 520
- INFILE statement
 - data-reading features 320
- infix operators 65
- INFORMAT statement
 - creating variables 30
- Information Maps LIBNAME Engine 596
- informats
 - as variable attribute 28
 - binary 324
 - column-binary 326
 - date values and 80
 - datetime values and 80
 - native or IBM 370 mode 324
 - time values and 80
- initializing hash objects
 - with constructore 393
- input buffers 288
 - creating 289
- input data
 - preprocessing for BY-group processing 333
 - representing missing values in 322
- input data sets
 - dropping, keeping, or renaming variables 36
- input data sources 14, 299
- input pointer 290
- INPUT statement
 - creating variables when reading data 30
 - data-reading features 320
 - input styles 316
 - reading raw data 316
- input styles 316
 - column input 318
 - formatted input 319
 - list input 317
 - modified list input 317
 - named input 319
- instantiating
 - hash iterator object 406
 - hash objects 393
 - Java objects 411
- instream data 315
 - as raw data source 14
 - multiple input files in 301
 - reading, with leading blanks and semicolons 314
 - reading, with missing values 301
 - reading raw data from 300
 - semicolons in 316
- INTCK function
 - interval boundaries and 91
- integrity constraints 493
 - CEDA processing and 498
 - creating with DATASETS procedure 499
 - creating with SCL 500
 - creating with SQL procedure 499
 - definition 493
 - examples 499
 - foreign key constraints 494
 - general constraints 493
 - indexes and 496, 526
 - inter-libref referential 497
 - listing 498
 - locking 497
 - overlapping primary and foreign key constraints 495, 504
 - preservation of 495
 - reactivating 504
 - recovering when disabled 602
 - referential constraints 494
 - rejected observations 498
 - removing 503
 - specifying 497
- inter-libref referential integrity constraints 497
- interactive line mode 8
- interactive mode
 - SAS log in 119
- interface data files 477
- interface DATA step engine 565
- interface library engines 595
 - BMDP 596
 - OSIRIS 596
 - SPSS 595
 - view engines 596
- interface view engine 564
- interface views 529
- interleaving data sets 349, 359
 - comments and comparisons 363
 - DATA step processing during 360
 - different BY values in each data set 362
 - duplicate values of BY variable 361
 - simple interleaving 360
 - sort requirements 359
 - syntax 359
- international characters 212, 219
- Internet Protocol Version 6
 - See IPv6
- intervals 88
 - custom 95
 - date and time 88
 - ISO 8601 compliant 95
 - multi-unit 93
 - multi-week 93
 - retail calendar intervals 95
 - shifted 94
- INTNX function
 - interval boundaries and 91
- invalid data 321

- _IORC_ automatic variable
 - error checking with 382
 - error-checking with 114
- IPv6 621
 - address formats 621
 - addresses including a port number 622
 - addresses including a URL 622
 - examples of addresses 622
 - full and collapsed addresses 622
 - fully qualified domain names 622
- IS MISSING operator 141
- IS NULL operator 141
- ISO 8601 95
- iterative DO loops 433
 - processing selected array elements 433

J

- Java objects 392, 409
 - accessing object fields 411
 - accessing object methods 412
 - arrays and 414
 - calling a simple method 418
 - CLASSPATH and Java options 409
 - creating a custom class loader 422
 - creating a user interface 418
 - declaring and instantiating 411
 - examples 418
 - exceptions 417
 - passing arguments 415
 - restrictions and requirements 410
 - standard output 417
 - type issues 412
- Java type set 412
 - mapping to SAS data types 412

K

- KEEP statement
 - I/O optimization and 154
- keeping variables 36
 - examples 37
 - input or output data sets for 36
 - order of application 37
 - statements versus data set options 36
- KEY= option
 - error checking and 386
 - MODIFY statement 522
 - SET statement 522
 - specifying indexes 522
- key summaries 399
- key variables 507, 512
- keys
 - data pairs and 397
 - defining 395
 - non-unique 397
 - SAS registry and 167

L

- labels 28
 - in output 127
- language elements 6
- LAST.variable 330, 334
- Latin TrueType fonts 213
- LBOUND function 440

- leading blanks
 - reading data containing 314
- Learning SAS Programming 259
- LENGTH statement
 - creating variables 30
 - I/O optimization and 154
- LIBNAME catalog concatenation 557
- LIBNAME statement
 - See also* SAS/ACCESS LIBNAME statement
 - associating librefs 450
 - clearing librefs 450
- libraries 4, 447
 - accessing permanent files without libref 459
 - concatenation 452
 - configuring with SAS registry 177
 - defined by SAS 264
 - defining 449
 - engines and 591
 - file types 448
 - library directories 459
 - library engines 449
 - librefs 449
 - logical names 449
 - management tools 458
 - names 449
 - operating environment commands 460
 - permanent 454
 - physical names 449
 - remote access 451
 - SAS system libraries 455
 - sequential 457
 - sorting data sets in 267
 - temporary 454
 - utilities for 458
 - viewing with Explorer 264
- library concatenation 452
 - definition 452
 - library members 452
 - rules for 453
- library directories 459
- library engines 449, 593
 - definition 593
 - interface library engines 595
 - native 594
 - version compatibility and 578
- librefs 449
 - accessing permanent files without 459
 - assigning 450
 - assigning the USER libref 456
 - associating 450
 - clearing 450
 - configuring with SAS registry 177
 - fixing problems with SAS registry 178
 - reserved 451
 - SAS/ACCESS 562
 - syntax for assigning 450
- LIKE operator 141
- line mode
 - SAS log in 119
- line size 127
- links
 - SAS registry and 168
 - to SVG documents 237
- list input 317
- LIST statement
 - writing to log with 122

- listing output 126
 - viewing in Output window 253
- literals 18
 - See also* constants
 - name literals 23
- little endian platforms 324
- locking integrity constraints 497
- log 15, 115
 - altering contents of 122
 - appending to 120
 - as DATA step output 15
 - changing destination of 128
 - console log 116
 - customizing 122
 - customizing appearance of 124
 - in batch mode 119
 - in interactive mode 119
 - in line mode 119
 - in objectserver mode 119
 - index usage information in 520
 - logging facility logs 116
 - naming with directives 120
 - replacing 120
 - rolling over 120
 - specifying when to write to 120
 - structure of 117
 - suppressing parts of 122
 - writing to, in all modes 122
- log control options 114
- Log window 252
- logger objects 392
- logging facility logs 116
- logic errors 114
- logical names 449
 - See* librefs
- logical operators 70
 - combining WHERE expressions 144
 - syntax for WHERE expressions 144

M

- macro facility 6
 - definition 4
- macro-related errors 108
- magnitude versus precision 42
- mantissa 38
- many-to-many relationships 347
- many-to-one relationships 346
- Maps library 264
- master data set 374
 - updating 381
- match-merging 351, 370
 - combining observations based on a criterion 371
 - DATA step processing during 371
 - duplicate values of BY variable 372
 - nonmatched observations 373
 - syntax 370
- MAX operator 72
 - in WHERE expressions 144
- memory
 - increasing for CPU performance 158
 - optimizing usage 158
- menus
 - drop-down 256
 - Help menu 259
 - pop-up 257

- Universal Printing 185
- MERGE statement
 - data set lists and 464
 - match-merging 370
 - one-to-one merging 366
- merging
 - data for reports 533
 - match-merging 351, 370
 - one-to-one 350, 366
- metadata LIBNAME engine 597
- Migration Focus Area 577
- MIN operator 72
 - in WHERE expressions 144
- missing semicolons 102
- missing values 5
 - automatically set by SAS 52
 - character 322
 - checking for, in DATA step 56
 - creating special missing values 50, 54
 - definition 49
 - generated by SAS 53
 - illegal character-to-numeric conversions and 53
 - illegal operations and 53
 - in raw data 322
 - indexes 509
 - numeric 322
 - order of 51
 - printing 128
 - propagation of 53
 - propagation of, preventing 54
 - reading instream data lines with 301
 - representing in input data 322
 - representing in raw data 55
 - setting to missing in DATA step 55
 - special numeric 322
 - updating data sets 376, 379
- modified list input 317
- MODIFY statement
 - indexes with 376
 - primary uses 377
 - updating data sets 374, 375
 - with BY versus UPDATE statement 376
- modifying data sets 343
 - tools for 343
- monolingual Asian TrueType fonts 213
- Mozilla Firefox
 - font mapping 226
 - SVG documents and 226
- Multi Engine Architecture 449
- multi-unit intervals 93
- multi-week intervals 93
- multidimensional arrays 430, 437
 - grouping variables in 437
 - nested DO loops 438
- multilingual Unicode TrueType fonts 213
- multiple-pass access
 - engines 592

N

- _N_ automatic variable 33
- name literals 23
 - avoiding errors 24
 - examples 24
 - restrictions 23
- name prefix lists 35, 465

- name range lists 35
 - named input 319
 - names 18, 20
 - automatic naming convention 466
 - catalogs 553
 - data sets 463
 - definition 20
 - generation groups 489, 493
 - length of 20
 - library names 449
 - reserved 20, 27
 - user-supplied 20
 - variable names 22
 - native data files 477
 - native environment 568
 - native files 568
 - native library engines 594
 - default Base SAS engine 594
 - definition 594
 - REMOTE engine 594
 - SASESOCK engine 594
 - sequential engines 595
 - SPD Engine 594
 - transport engine 595
 - V6 compatibility engine 595
 - native mode 324
 - native views 529
 - nested DO loops 438
 - NOMISS composite indexes 518
 - non-standard files
 - importing 282
 - noninteractive mode 8
 - nonmatched observations
 - updating data sets with 379
 - nonstandard data 313
 - NOT operator 71
 - null data sets 465
 - numbered range lists 34, 464
 - numbers 18
 - numeric-character conversion 64
 - numeric comparisons 67
 - IN operator in 68
 - numeric constants 60
 - in hexadecimal notation 61
 - in scientific notation 61
 - in standard notation 60
 - numeric data
 - reading 313
 - numeric missing values 322
 - _NUMERIC_ name list 35
 - numeric precision 38
 - numeric precision variables 26
 - numeric values 313
 - storing 38
 - numeric variables 26
 - converting to character 32
 - sort order of missing values 51
- O**
- object fields 411
 - object methods 412
 - object spawner 617
 - UUIDs and 617
 - objects
 - See* DATA step component objects
 - objectserver mode 9
 - SAS log in 119
 - OBS= data set option
 - I/O optimization and 154
 - segmenting a subset 147
 - observations 5
 - adding to indexed data files 525
 - altering DATA step execution sequence for 295
 - BY-group processing 131
 - position of variables in 28
 - processing in BY groups 334
 - reading data sets 344
 - rejected 498
 - rejected, capturing 485
 - sorting for BY-group processing 333
 - writing to data sets 291
 - ODS destinations
 - for SVG documents 224
 - PNG images for 241
 - ODS (Output Delivery System) 4, 128
 - creating HTML reports 308
 - DATA step and 310
 - Universal Printing and 184
 - ODS PRINTER statement
 - creating PNG files 242
 - ODS styles
 - TrueType fonts and 212
 - ODS templates
 - specifying TrueType fonts 218
 - oldest version 488
 - one-dimensional arrays 430
 - one-level data set names 464
 - one-to-many relationships 346
 - one-to-one merging 350, 366
 - comments and comparisons 370
 - DATA step processing during 366
 - different values of common variables 369
 - duplicate values of common variables 368
 - equal number of observations 367
 - syntax 366
 - unequal number of observations 367
 - one-to-one reading 350, 363
 - comments and comparisons 365
 - DATA step processing during 364
 - equal number of observations 364
 - syntax 363
 - one-to-one relationships 346
 - OpenVMS
 - floating-point representation 41
 - operands 58
 - in WHERE expressions 135
 - operating environment commands
 - for libraries 460
 - operating environments
 - moving SAS files between 599
 - operators 58
 - AND 71
 - arithmetic 66, 138
 - BETWEEN-AND 140
 - Boolean 70
 - character comparisons 69
 - comparison 66, 138
 - concatenation 72, 144
 - CONTAINS 140
 - fully bounded range condition 139
 - IN 68, 69, 139

- in expressions 65
 - in WHERE expressions 138
 - infix 65
 - IS MISSING 141
 - IS NULL 141
 - LIKE 141
 - logical 70, 144
 - MAX 72, 144
 - MIN 72, 144
 - NOT 71
 - numeric comparisons 67
 - OR 71
 - prefix 65, 144
 - SAME-AND 143
 - sounds-like 142
 - OR operator 71
 - ordered observations
 - BY-group processing 131
 - OSIRIS engine 596
 - out-of-resources condition 104
 - output 115
 - See also* log
 - CEDA processing 569
 - centering 127
 - changing destination of 128
 - console log 116
 - DATA step 15
 - date and time values 127
 - default destinations 117
 - footnotes 127
 - formatting characters for 127
 - Java standard output 417
 - labels 127
 - line size 127
 - listing output 253
 - logging facility logs 116
 - page breaks 127
 - page numbering 127
 - page size 127
 - printing missing values 128
 - PROC step 16
 - program results 116
 - redirecting for stored compiled DATA step programs 544
 - reformatting values 128
 - routing 117
 - SVG 227
 - titles 127
 - traditional listing output 126
 - Universal Printing formats 184
 - viewing in Results window 255
 - output data sets
 - dropping, keeping, or renaming variables 36
 - Output Delivery System
 - See* ODS (Output Delivery System)
 - output files
 - DATA step views 533
 - Output window 253
 - overlaid transparent SVG documents 235
- P**
- packed decimal data 323
 - page breaks 127
 - page numbering 127
 - page options 196
 - page properties 200
 - page size 127
 - parallel processing
 - CPU performance and 159
 - parentheses
 - order of evaluation in WHERE expressions 145
 - password-protected files
 - in DATA and PROC steps 584
 - passwords 581
 - alter protection 581, 587
 - assigning 582
 - assigning, syntax for 582
 - assigning outside of SAS 584
 - assigning to data sets 583
 - assigning with DATA step 582
 - assigning with procedures 583
 - assigning with windowing environment 584
 - audit trails with 588
 - changing 584
 - copies with 588
 - DATA step views and 588
 - definition 581
 - encoded 586
 - generation data sets with 588
 - in generation groups 493
 - incorrect 585
 - indexes with 588
 - level of protection 581, 586
 - PROC SQL views and 587
 - read protection 581, 586
 - removing 584
 - SAS/ACCESS views and 587
 - views with 586
 - write protection 581, 587
 - pattern matching 141
 - PC file formats 561
 - PCL format 184
 - PDF format 184
 - PDV (program data vector) 288
 - input buffer and 289
 - percent sign
 - in name literals 23
 - performance
 - See also* system performance
 - ARM 163
 - audit trails 481
 - DATA step views 533
 - DICTIONARY tables 551
 - of applications 599
 - of WHERE processing 146
 - parallel processing and 161
 - sort indicator and 471
 - performance statistics 151
 - collecting and interpreting 152
 - permanent files
 - accessing without libref 459
 - permanent libraries 454
 - physical names 449
 - PNG format 184, 241
 - browsers and viewers supporting 242
 - creating with ODS PRINTER statement 242
 - for ODS destinations 241
 - Universal Printers for 241
 - pointers
 - input pointer 290
 - pop-up menus 257

- port number
 - addresses including 622
- Portable Network Graphics
 - See* PNG format
- PostScript output 184
 - previewing with Ghostview 208
- precision versus magnitude 42
- prefix operators 65
 - in WHERE expressions 144
- Preview command box 200
- previewers 197
 - creating a Ghostview definition 208
 - defining 197
 - defining with PRTDEF procedure 206
 - previewing print jobs 200
 - seeding Preview command box 200
- primary key 494
- Print dialog box
 - specifying TrueType fonts 215
- print jobs
 - previewing 200
- print previewers
 - See* previewers
- PRINT procedure
 - specifying TrueType fonts 217, 218
- printer definitions 187
 - adding, modifying, deleting 207
 - backing up 208
 - batch mode 205
 - exporting 208
 - multiple printers 206
 - multiple users 206
 - PRTDEF procedure for 206
- printers 184
 - See also* Universal Printers
 - changing default 187
 - properties for default printer 191
 - removing from selection list 187
 - setting up 186
 - specifying for current SAS session 194
- printing
 - See also* SVG documents
 - See also* Universal Printing
 - active window contents 195
 - Forms 183, 210
 - international characters 212, 219
 - missing values 128
 - page options 196
 - page properties 200
 - source code for stored compiled DATA step programs 544
 - test page 195
 - TrueType fonts 211
- PROC SQL views 535
 - compared with DATA step views 536
 - embedding SAS/ACCESS LIBNAME statement in 562
 - passwords and 587
- PROC step output 16
- PROC steps 6
 - password-protected files in 584
 - SAS processing and 16
- procedure output files
 - as DATA step output 15
- procedures 4
 - assigning passwords with 583
 - combining data sets with 352

- Profile catalog 555
 - creation of 555
 - default settings 555
 - definition 555
 - how information is used 555
 - recovering when locked or corrupt 556
- program compilation optimization 159
- program data vector (PDV) 288
 - input buffer and 289
- Program Editor window 251
- program results 116
- program statements
 - generating data from 302
 - specifying TrueType fonts 216
- propagation of missing values 53
 - preventing 54
- PRTDEF procedure
 - defining printers and previewers 206
- PS format 184
 - previewing with Ghostview 208
- punched cards 326
- PUT statement
 - SMTP e-mail 614
 - writing to log with 122
- %PUT statement
 - writing to log with 122
- PUTLOG statement
 - writing to log with 122
- PW= data set option
 - assigning complete file protection 585

Q

- quality control
 - stored compiled DATA step programs 546
- quotation marks
 - character constants with 59
 - unmatched 102

R

- radix point 39
- random access
 - engines 592
 - error checking when using indexes for 382
- raw data 311
 - as input data source 14
 - binary data 323
 - character data 314
 - column-binary data 325
 - data types 312
 - external files 316
 - instream data 315
 - instream data with semicolons 316
 - invalid data 321
 - missing values in 55, 322
 - missing values when reading 52
 - numeric data 313
 - reading 311
 - reading with DATA step 300
 - reading with INPUT statement 30, 316
 - sources of 315
 - ways to read 312
- reactivating integrity constraints 504
- read protection 581, 586

- reading data sets 343, 344
 - multiple data sets 344
 - reading and writing observations 344
 - reading and writing variables 344
 - single data set 344
 - tools for 343
 - real binary representation 38
 - real time 152
 - records
 - DATA step processing of 290, 292
 - recovering
 - catalogs 602
 - data files 600
 - disabled indexes 602
 - indexes 601
 - integrity constraints 602
 - referential integrity constraints 494
 - inter-libref 497
 - reformatting values 128
 - registering TrueType fonts
 - for Windows or UNIX 215
 - for z/OS 215
 - with SAS Registry 214
 - registry
 - See* SAS registry
 - Registry Editor 172
 - adding values or keys to registry 174
 - backing up SAS registry 170
 - backing up SASUSER registry 169
 - deleting items from registry 174
 - displaying SASUSER and SASHELP items separately 175
 - exporting registry files 176
 - finding specific data in registry 173
 - importing registry files 175
 - modifying registry values 173
 - renaming items in registry 175
 - saving registry files 176
 - starting 173
 - when to use 172
 - registry files
 - exporting 176
 - importing 175
 - in SASHELP library 166
 - in SASUSER library 166
 - saving 176
 - REGISTRY procedure
 - backing up SASUSER registry 169
 - related data 345
 - remote access
 - as input data source 14
 - REMOTE engine 594
 - remote libraries
 - SAS/CONNECT 451
 - SAS/SHARE 451
 - WebDAV servers 452
 - remote library services 574
 - renaming data sets 266
 - renaming generation groups 493
 - renaming variables 36
 - examples 37
 - input or output data sets for 36
 - order of application 37
 - statements versus data set options 36
 - repairing damaged files 600
 - reports
 - as DATA step output 15
 - creating, without creating a data set 303
 - creating customized reports 304
 - creating HTML reports 308
 - merging data for 533
 - writing with DATA step 303
 - reserved librefs 451
 - reserved names 20, 27
 - resource usage 152
 - restart mode 110
 - restarting batch programs 112
 - setting up and executing 111
 - restoring site defaults 167
 - Results window 254
 - retail calendar intervals 95
 - retrieving data
 - ADD and FIND methods for 398
 - FIND method for, and loading a data set 399
 - hash object data 407
 - hash objects and 397
 - return codes 114
 - RIDs 506
 - rolling over the log 120
 - automatically, when directives change 121
 - by log size 121
 - by SAS session 121
 - directives for naming the log 120
 - disabling rollover 122
 - names for rolled over logs 121
 - rolling over version numbers 488
 - routing output 117
 - changing destination 128
 - default destinations 117
 - rows
 - subsetting 273
- ## S
- SAME-AND operator 143
 - SAS 3
 - system-wide concepts 10
 - SAS 9
 - comparing to earlier releases 577
 - SAS/ACCESS 561
 - ACCESS procedure and interface view engine 564
 - data set options with librefs 562
 - DBLOAD procedure 565
 - dynamic LIBNAME engine 561
 - interface DATA step engine 565
 - SQL Procedure Pass-Through Facility 563
 - SAS/ACCESS LIBNAME statement 561
 - embedding in PROC SQL views 562
 - SAS/ACCESS views 529, 536
 - passwords and 587
 - SAS catalogs
 - See* catalogs
 - SAS/CONNECT
 - data transfer services 574
 - remote library access 451
 - remote library services 574
 - SAS console log 116
 - SAS constants
 - See* constants
 - SAS data files
 - See* data files

- See files
- SAS data sets
 - See data sets
- SAS engines
 - See engines
- SAS expressions
 - See expressions
- SAS files 4
 - See also files
 - concepts 10
 - engines and 590
 - moving between operating environments 599
 - repairing damaged files 600
- SAS/GRAPH
 - devices and TrueType fonts 211
 - specifying TrueType fonts 217
- SAS Help and Documentation 259
- SAS indexes
 - See indexes
- SAS Information Maps LIBNAME Engine 596
- SAS language 4
 - data sets 5
 - DBMS files 6
 - external files 5
 - language elements 6
 - macro facility 6
 - SAS files 4
- SAS libraries
 - See libraries
- SAS log
 - See log
- SAS name lists 35
- SAS names
 - See names
- SAS on the Web 259
- SAS output
 - See output
- SAS processing 13
 - DATA step 15
 - input data sources 14
 - PROC steps 16
- SAS programs
 - input data sources for 14
- SAS registry 166
 - See also Registry Editor
 - adding values or keys to 174
 - backing up SASUSER registry 169
 - backing up with Registry Editor 170
 - changing values in 173
 - color control with 171
 - configuring 176
 - configuring Explorer 176
 - configuring file shortcuts 177
 - configuring libraries 177
 - configuring Universal Printing 176
 - deleting items from 174
 - displaying 167
 - editing 166, 168
 - exporting registry files 176
 - finding data in 173
 - fixing libref problems with 178
 - importing registry files 175
 - keys and 167
 - links and 168
 - managing 168
 - recovering from failure 170
 - registering TrueType fonts 214
 - registry files 166
 - renaming items 175
 - restoring 170
 - restoring site defaults 167
 - storage location 166
 - subkeys and 167
 - terminology 167
 - values and 168
 - who should use 166
- SAS Registry file 168
- SAS Scalable Performance Data (SPD) engine 161, 594
- SAS sessions 7
 - batch mode 9
 - customizing 9
 - customizing windowing environment 10
 - default system option settings 9
 - executing statements automatically 9
 - in SAS windowing environment 7
 - interactive line mode 8
 - noninteractive mode 8
 - objectserver mode 9
 - restoring site defaults 167
 - rolling over the log by 121
 - starting 7
 - types of 7
- SAS/SHARE
 - remote library access 451
 - remote library services 574
- SAS system libraries
 - See system libraries
- SAS utilities
 - for libraries 458
- SAS variables
 - See variables
- SAS views
 - See views
- SAS windowing environment
 - See windowing environment
- SAS XML LIBNAME engine 597
- SASESOCK engine 568, 594
- SASFILE statement
 - I/O optimization and 157
- Sashelp library 264
- SASHELP library 457
 - registry files in 166
- Sasuser library 264
- SASUSER library 457
 - registry files in 166
- SASUSER registry
 - backing up 169
- Sasuser.Profile
 - See Profile catalog
- .SASXREG file 168
- Scalable Vector Graphics
 - See SVG documents
- scientific notation 38
 - numeric constants in 61
- SCL
 - creating integrity constraints 500
 - search time
 - reducing for executable files 158
 - segmenting a subset 146
 - semantic errors 102
 - semicolons
 - instream data containing 316

- missing 102
- reading data containing 314
- sequential access
 - combining data sets 348
 - engines 592
- sequential engines 595
- sequential libraries 457
- sessions
 - See* SAS sessions
- SET statement
 - concatenating data sets 356
 - data set lists and 464
 - interleaving data sets 359
 - one-to-one reading 363
- shared files
 - location of 497
- sharing disk space
 - integrity constraints and 497
- shifted intervals 94
- shortcuts 265
- sign bit 39
- simple expressions 58
- simple indexes 507
- simple WHERE expressions 134
- single-precision floating-point numbers 46
- site defaults, restoring 167
- SMP computers 161
- SMTP e-mail interface 613
 - statements for 614
 - system options for 613
- sort indicator 466
 - performance and 471
- sort order
 - missing values 51
- SORT procedure 469
- sorted data sets 466
 - validating 472
- SORTEDBY= data set option 468
- sorting
 - by column values 271
 - data sets in a library 267
 - for interleaving data sets 359
 - for UPDATE statement 376
 - hash objects and 397
 - indexed data files 525
 - observations for BY-group processing 333
- sounds-like operator 142
- source code
 - printing for stored compiled DATA step programs 544
- SPD engine 161
- SPD Engine 594
- special characters 18
- special missing values 50, 54
- special numeric missing values 322
- special SAS name lists 35
- SPSS engine 595
- SQL
 - concatenating data sets 358
- SQL procedure
 - creating indexes 514
 - creating integrity constraints 499
- SQL Procedure Pass-Through Facility 563
- standard data 313
- standard files
 - importing 279
- standard notation
 - for numeric constants 60
- statements
 - changing DATA step execution sequence 295
 - combining data sets with 352
 - controlling reading and writing 344
 - default DATA step execution of 294
 - dropping, keeping, or renaming variables with 36
 - executing automatically 9
 - for SMTP e-mail 614
 - specifying TrueType fonts 216
 - step boundaries 297
 - word placement and spacing in 19
- statistics
 - See* performance statistics
- step boundaries 297
- STIMER system option 152
 - interpreting statistics 152
- stored compiled DATA step programs 539
 - compiling and storing 541
 - CPU performance and 158
 - creating 541, 542
 - examples 546
 - executing 542, 545
 - execution process 543
 - global statements for 544
 - printing source code 544
 - processing 540
 - quality control application 546
 - redirecting output 544
 - restrictions and requirements 540
 - syntax for creating 541
 - syntax for executing 542
 - uses for 539
 - versus DATA step views 532, 546
- stored programs 4
- storing data
 - ADD and FIND methods for 398
 - hash objects and 397
- subkeys
 - SAS registry and 167
- subsets
 - segmenting 146
- subsetting data
 - exporting a subset 276
 - rows of a table 273
 - with WHERE expression 273
- subsetting IF statement
 - versus WHERE expression 148
- SVG canvas 223
- SVG documents 184, 223
 - Adobe SVG Viewer and 226
 - browser support for 226
 - creating stand-alone 231
 - creating with Universal Printers 224
 - embedding in HTML documents 237
 - environment for stand-alone 227
 - images in 240
 - linking to 237
 - Mozilla Firefox and 226
 - multi-page, in separate files 234
 - multi-page, in single file 233
 - ODS destinations for 224
 - output from Universal Printers 227
 - overlaid transparent 235
 - printing from browser 241

- reasons for creating 224
- setting Universal Printer 228
- system options for 228, 229
- tag attributes 229
- terminology 223
- titles 231
- user coordinate system 224
- user space 224
- viewBox 223
- viewBox, aspect ratio 231
- viewBox, setting 229
- viewBox, static 230
- viewport 223
- viewport, scaling to 228
- viewport coordinate system 223
- viewport space 223
- syntax check mode 108
 - enabling 109
- syntax errors 100
- SYSERR automatic macro variable
 - error-checking with 114
- SYSMSG function
 - error-checking with 114
- SYSPRINTFONT= system option
 - specifying TrueType fonts 216
- SYSRC automatic macro variable
 - error-checking with 114
- SYSRC function
 - error-checking with 114
- SYSRC macro 382
- system libraries 455
 - SASHELP 457
 - SASUSER 457
 - USER 456
 - WORK 455
- system options
 - affecting the log 125
 - altering log contents with 123
 - customizing log appearance with 124
 - default settings 9
 - for error handling 113
 - for SMTP e-mail 613
 - for threaded application processing 162
 - SVG documents and 228
 - SVG tag attributes and 229
 - Universal Printing 204
- system performance 151
 - calculating data set size 160
 - CPU performance 158
 - definition 151
 - memory usage 158
 - optimizing I/O 153
 - performance statistics 152

T

- table editor
 - See* VIEWTABLE window
- tables
 - editing cell values 272
 - importing data into 279
 - moving columns in 269
 - sorting by column values 271
 - subsetting rows 273
 - temporarily changing column headings 270
- TAPE engine 568

- TCP/IP socket 14
- temporary libraries 454
- test page 195
- testing programs 355
- threaded application processing 162
- threaded I/O 161
- threads 161
- time constants 61
- time durations 88
- time intervals 88
 - boundaries of 91
 - by category 89
 - multi-unit 93
 - shifted intervals 94
 - single-unit 92
 - syntax 88
- time values 78
 - as recognizable times 86
 - formats and 80
 - in output 127
 - informat and 80
 - tools by task 81
- titles 127
 - SVG documents 231
- tokens
 - See* words
- toolbars 257
 - opening Help menu from 259
- traditional listing output 126
- transaction data set 374
- transaction monitoring 163
- transport engine 595
- TrueType fonts 211
 - international character support 212, 219
 - locating font files 215
 - ODS styles and 212
 - portability of 212
 - registering for Windows or UNIX 215
 - registering for z/OS 215
 - registering with SAS Registry 214
 - rendering 211
 - specifying with Print dialog box 215
 - specifying with PRINT procedure 217
 - specifying with PRINT procedure and ODS template 218
 - specifying with program statements 216
 - specifying with SAS/GRAPH 217
 - specifying with SYSPRINTFONT= system option 216
 - supplied by SAS 213
- truncating numbers 45
- two-dimensional arrays 431
 - bounds in 441
- two-level data set names 464

U

- unexpected conditions 383
- Universal Printers 227
 - creating SVG documents 224
 - PNG format 241
 - setting for SVG documents 228
 - SVG output 227
- Universal Printing 183
 - See also* printer definitions
 - accessing printers 184
 - configuring 204
 - configuring with SAS registry 176

- default printer 187
 - destination 209
 - device type 209
 - host options 209
 - interface 185
 - menu 185
 - ODS and 184
 - output formats and printers 184, 208, 223, 241
 - page options 196
 - page properties 200
 - previewers 197, 208
 - previewing print jobs 200
 - printer for current SAS session 194
 - printing active window contents 195
 - printing test page 195
 - printing with 195
 - properties for default printer 191
 - removing printer from selection list 187
 - seeding Preview command box 200
 - setting up printers 186
 - system options for 204
 - TrueType fonts and 211
 - turning on and off 183
 - windows 186
 - Universal Unique Identifiers
 - See UUIDs
 - UNIX
 - registering TrueType fonts 215
 - unmatched comment tags 102
 - unmatched quotation marks 102
 - UPDATE statement
 - DATA step processing with 375
 - sort requirements 376
 - updating data sets 374, 377
 - versus MODIFY statement with BY 376
 - with duplicate values of BY variable 378
 - with nonmatched observations, missing values, and new variables 379
 - updating data sets 351, 374
 - basic updating 377
 - error checking 382
 - indexes with MODIFY statement 376
 - master data set 381
 - missing values 376
 - new variables 376
 - nonmatched observations 376
 - sort requirements for UPDATE statement 376
 - syntax 374, 375
 - UPDATE versus MODIFY with BY 376
 - with duplicate values of BY variable 378
 - with nonmatched observations, missing values, and new variables 379
 - URLs
 - addresses including 622
 - remote access and 15
 - user-defined ODS templates
 - specifying TrueType fonts 218
 - user interface
 - creating for Java 418
 - USER library 456
 - assigning USER libref 456
 - relation to WORK library 456
 - user-supplied names 20
 - length of 20
 - reserved names and 20
 - variable names 22
 - user variables 479
 - Using This Window 259
 - utilities
 - for libraries 458
 - UUID Generator Daemon 617
 - installing 618
 - UUIDCOUNT= system option 619
 - UUIDGEN function 619
 - UUIDGENHOST system option 619
 - UUIDs 617
 - assigning 619
 - object spawner and 617
- ## V
- V6 compatibility engine 595
 - values
 - SAS registry and 168
 - variable attributes 26
 - format 27
 - index type 28
 - informat 28
 - label 28
 - length 27
 - name 27
 - position in observation 28
 - type 27
 - variable labels 28
 - variable length 27
 - CPU performance and 159
 - not explicitly set 29
 - variable lists 34
 - defining arrays quickly 436
 - name prefix 35
 - name range 35
 - numbered range 34
 - special SAS name 35
 - variable names 22, 27
 - reserved 27
 - variable type 27
 - not explicitly set 29
 - variable type conversions 32
 - automatic numeric-character 64
 - illegal character-to-numeric 53
 - variable values 290
 - variables 5, 26
 - aligning values 32
 - _AT*_ 479
 - automatic 33
 - automatically set to missing 52
 - BY variables 329
 - character 26
 - creating 28
 - creating with assignment statement 29
 - creating with ATTRIB statement 31
 - creating with FORMAT statement 30
 - creating with IN= data set option 31
 - creating with INFORMAT statement 30
 - creating with INPUT statement 30
 - creating with LENGTH statement 30
 - dropping 36
 - FIRST. 330, 334
 - in expressions 64
 - in WHERE expressions 135
 - keeping 36
 - LAST. 330, 334

- maximum number of 26
- numeric 26
- numeric precision 26, 38
- position in observations 28
- reading data sets 344
- renaming 36
- setting values to missing in DATA step 55
- updating data sets with new variables 379
- user 479
- version compatibility 577
 - comparing SAS 9 to earlier releases 577
 - library engines and 578
 - SAS 9 file format 577
 - SAS 9 filename extensions 578
- view descriptors 536, 564
- view engines 596
- viewBox (SVG) 223
 - aspect ratio 231
 - setting 229
 - static 230
- viewers
 - Adobe SVG Viewer 226
 - supporting PNG format 242
- viewport (SVG) 223
 - coordinate system 223
 - scaling documents to 228
- views 5, 461, 529
 - as DATA step output 15
 - as input data source 14
 - benefits of 530
 - creating with DATA step 299
 - DATA step 531
 - I/O optimization and 155
 - indexes with 520
 - interface 529
 - level of protection 586
 - native 529
 - passwords with 586
 - PROC SQL 535
 - SAS/ACCESS 529, 536
 - versus data files 477
 - when to use 531
 - WHERE expressions and 147
- VIEWTABLE window 268
 - editing cell values 272
 - moving columns 269
 - sorting by column values 271
 - temporarily changing column headings 270
 - viewing and editing data sets 473
 - viewing data set contents 268

W

- WebDAV
 - remote library access 452
- WHERE-expression processing 133
 - comparing resource usage 519
 - compound optimization 516
 - controlling index usage 519
 - displaying index usage information 520
 - estimating qualified observations 518
 - I/O optimization and 153
 - identifying available indexes 515
 - indexes for 515
 - indexes for, with BY processing 522
 - indexes with views 520

- performance improvement 146
- WHERE expressions 58, 134
 - clearing 276
 - combining with logical operators 144
 - compound 134, 145
 - compound optimization 516
 - constants in 137
 - efficient 146
 - functions in 136
 - operands in 135
 - operators in 138
 - optimizing 515
 - order of evaluation for combining 145
 - processing views 147
 - segmenting a subset 146
 - simple 134
 - subsetting data with 273
 - syntax 135
 - variables in 135
 - versus subsetting IF statement 148
 - where to use 134
- window contents
 - printing 195
- windowing environment 4, 247
 - assigning passwords with 584
 - command line 258
 - commands 260
 - customizing 10
 - data management in 263
 - drop-down menus 256
 - Explorer 250, 264
 - Help 259
 - list of windows 260
 - Log window 252
 - main windows 248
 - navigating 255
 - Output window 253
 - pop-up menus 257
 - Program Editor window 251
 - Results window 254
 - running SAS sessions in 7
 - toolbars 257
 - VIEWTABLE window 268
- Windows
 - registering TrueType fonts 215
- words 17
 - placement and spacing in statements 19
 - types of 18
- Work library 264
- WORK library 455
 - relation to USER library 455
- wrapper classes 415
- write protection 581, 587

X

- XML engine 573
- XML LIBNAME engine 597
- XPORT engine 573

Y

- year 2000 78
- YEARCUTOFF= system option
 - year 2000 and 78
 - year digits and 78, 79

years
 four-digit 78, 79
 two-digit 78, 79
youngest version 488

Z

z/OS
 registering TrueType fonts 215
zoned decimal data 324

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



THE
POWER
TO KNOW®

