# Differences in the SAS®9 and SAS® Viya™ 3.1 Platforms

SAS® Documentation

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2016. *Differences in the SAS®9 and SAS® Viya™ 3.1 Platforms*. Cary, NC: SAS Institute Inc.

**Differences in the SAS®9 and SAS® Viya™ 3.1 Platforms**

SAS Institute Inc., SAS Campus Drive, Cary, NC 27513-2414

September 2016

3.1-P1:whatsdiff

# Contents

# About This Book

## Audience

This book is intended for existing SAS 9 customers who want to learn more about SAS Viya. This document focuses on the first release, which is SAS Viya 3.1. With each successive release of the SAS Viya platform, this document will be updated.

# Chapter 1
# Introduction

## About the SAS Viya Platform: The Next Era of Analytics

SAS Viya is an open, unified, and powerful platform, ready for native Cloud deployment.

| | |
|---|---|
| Cloud Analytic Services (CAS) In-Memory Engine | This engine is the power behind SAS Viya. In SAS Viya, the SAS High-Performance Architecture and SAS LASR Analytic Server from SAS 9 combine to form a third-generation high-performance analytics engine. The CAS In-Memory Engine continues the ability to perform processing in memory and the ability to distribute processing across nodes in a cluster. The CAS In-Memory Engine adds highly efficient node-to-node communication and uses an algorithm to determine the optimal number of nodes for a given job.<br>The CAS In-Memory Engine also enables you to use distributed NFS-based storage. |
| Interfaces | You can use a variety of interfaces to interact with the CAS In-Memory Engine. These interfaces include SAS Studio, which is a browser-based interface for writing SAS code. You can also use programming interfaces for Python, Java, and Lua to access this CAS functionality. In addition, you can continue to submit SAS code in batch mode. |

For more information, see *SAS Cloud Analytic Services: Fundamentals*.

## Comparing SAS 9 and SAS Viya

| Item for Comparison | SAS 9 | SAS Viya |
|---|---|---|
| Servers | SAS 9 contains these servers:<br>• SAS Metadata Server<br>• SAS Workspace Server<br>• SAS Pooled Workspace Server<br>• SAS Stored Process Server<br>• SAS OLAP Server<br>• SAS Web Infrastructure Platform Data Server<br>• SAS/CONNECT Server<br>• SAS Content Server<br>• SAS LASR Analytic Server | SAS Viya contains these servers:<br>• SAS Cloud Analytic Services<br>• SAS Workspace Server and Object Spawner<br>• SAS/CONNECT Server and Spawner<br>• embedded web application server |
| Data Location | You must aggregate the data in some way before running the analysis or DATA step. | Your data can be distributed across various worker nodes, and the processing of that data occurs on that node. |
| Management Tools | SAS 9 uses a variety of management tools, such as SAS Management Console and SAS Deployment Manager. | SAS Viya provides cohesive analytics management. As a result, IT has centralized control over all analytic assets. |
| Installation, Configuration, and Maintenance | SAS 9 uses SAS proprietary tools, such as the SAS Deployment Wizard and SAS Deployment Manager. To apply updates and maintenance releases to your SAS 9 environment, you must take your system offline for a period of time. | SAS Viya uses industry-standard tools, such as Ansible, for installation and configuration. SAS uses RPM packages to deliver the software. At installation, these packages are either downloaded from an Internet repository or a mirror of this repository at your site. You can choose to update these RPMs independently of each other. |
| Backups | In SAS 9, backups are created using SAS Environment Manager or SAS Management Console. | In SAS Viya, backups are performed using CURL commands. In this release, the backup service does not include a scheduling feature or an email alert feature. |
| Availability | In SAS 9, a failed LASR node or job could potentially impact everyone on that server. | In SAS Viya, CAS worker nodes and CAS session processes are independent of each other. |

| Item for Comparison | SAS 9 | SAS Viya |
|---|---|---|
| Interfaces | In SAS 9, analytics procedures are accessed through SAS interfaces and REST-based services. | SAS Viya provides an open user experience. You can use your existing skill sets in SAS, Python, Lua, or Java to access SAS analytics. |

## Relationship between SAS 9.4 and SAS Viya

### *SAS Viya Does Not Replace SAS 9*

SAS 9 and SAS Viya are two unique platforms built for different use cases. However, these platforms share some functionality. For example, SAS 9 uses the SAS programming language, and SAS Viya uses the next generation of SAS programming with the new CAS programming language. The CAS language is very similar to the SAS language. Some procedures are available on both platforms, so some existing SAS code can be run on the SAS Viya platform. However, SAS Viya also contains new procedures that take advantage of the open, distributed environment. As a result, some SAS 9 procedures do not exist in the new SAS Viya environment.

Because of the advancement in technology and infrastructure since SAS 9 was released, the technology of SAS Viya is very different from SAS 9. Here are some examples:

- In SAS Viya, SAS Cloud Analytics Services is the third-generation of in-memory services from SAS. It replaces SAS LASR Analytic Server and the SAS High-Performance Analytics engine in SAS 9.
- For functionality other than CAS, there is no concept of roles or capabilities in SAS Viya. Access to functionality is managed by rules that target a service, a service endpoint, a media type, or a pseudo URL.
- Instead of the metadata-based authorization, SAS Viya maintains authorization within the CAS engine and by using host identities.
- In SAS Viya, images are a recognized data type.

### *SAS Viya Interoperates with SAS 9.4*

- You can make your SAS 9.4 data available to SAS Viya.
- You can migrate your user-defined format catalogs.
- You can remotely submit SAS 9 code to run in SAS Viya and remotely submit SAS Viya code to run in SAS 9.

SAS Viya is designed to coexist with SAS 9.4 solutions and the SAS 9 environment. SAS provides a client/server utility called SAS/CONNECT. The SAS/CONNECT toolset allows users to maintain SAS interoperability across architectures and SAS releases. SAS users can use SAS/CONNECT statements to create a virtual bridge between the two environments. The SAS/CONNECT bridge allows SAS 9.4 solution users to take advantage of SAS Viya performance while using familiar projects, interfaces, and SAS code.

For more information, see “Integrating SAS 9.4 and SAS Viya” in *Integrating SAS 9.4 and SAS Viya*.

# *Chapter 2* Programming

## UTF-8 Is the Only Session Encoding

In SAS 9, the default session encoding is different for each operating environment. Respectively, for z/OS, UNIX, and Windows, the default session encodings are OPEN_ED-1047, Latin1, and WLatin1. The session encoding can be changed when SAS starts.

In SAS Viya, UTF-8 is the only available session encoding. UTF-8 is a variable-width multi-byte encoding that encompasses all encodings.

SAS Viya requires that data set character encoding match the SAS session encoding. When the encodings do not match, the data set must be transcoded to UTF-8 before it can be used in SAS Viya. If SAS Viya can determine the data encoding, it automatically transcodes non-UTF-8 encodings to UTF-8 when it reads data. If the encoding cannot be determined, data sets must be transcoded programmatically. Caution is necessary to ensure that columns are wide enough for UTF-8 encodings. If the columns are not wide enough, data truncation might occur. You can use the CVP engine to enlarge character columns. For complete information, see *Migrating Data to UTF-8 for SAS Viya*.

### *See Also*

- *SAS Viya FAQ for Processing UTF-8 Data*
- *SAS Viya National Language Support: Reference Guide*

## SAS Programming Interfaces

SAS Studio is the SAS language code editor in SAS Viya. The SAS windowing environment and SAS Enterprise Guide are not available in SAS Viya. You can also submit SAS code by using batch mode or line mode processing. Python, Lua, and Java programmers can submit code to the CAS server using their respective programming environments.

SAS Studio provides tasks and snippets that generate code for you. A point-and-click interface makes it easy to select data sources and procedure options. After the code is generated, you can modify option values as needed.

If you are a current SAS Studio user, here are some differences between SAS Studio 3.5 (which runs on SAS 9.4) and SAS Studio 4.1 (which runs on SAS Viya):

- Only the SAS Programming perspective is available in SAS Studio 4.1. The Visual Programming perspective is not yet available.
- The list of tasks and snippets differs. SAS Studio 4.1 lists new data mining and machine learning tasks that are specific to the SAS Viya platform.
- In SAS Studio 4.1, save your data in a CAS library.

In addition, SAS Studio 4.1 does not support this functionality:

- process flows
- the query tool
- the import tool
- SAS Studio repositories
- the ability to export tables
- the ODS Graphics Designer and the ODS Graphics Editor
- FTP shortcuts
- extended attributes in the SAS Table Properties or Column Properties windows

SAS provides the ability to run CAS actions on the CAS server from Python 2.7, Lua 5.2 or Lua 5.3, Oracle JDK 8, OpenJDK 1.8.0.65, or REST APIs. SAS Wrapper for Analytics Transfer (SWAT) is the name of the package that contains the APIs. In this first release of SAS Viya, the modules that enable language-specific programming support are limited to monitoring CAS server performance and programming in Python.

### *See Also*

- *Getting Started with Programming in SAS Studio*
- *SAS Studio: User's Guide*
- *Batch and Line Mode Processing in SAS Viya*
- Programming with Actions in the SAS Viya Data Mining and Machine Learning Help Center

## Share Resources between SAS 9 and SAS Viya

If you license SAS/CONNECT, SAS/CONNECT becomes a bridge that enables you to move and share data and computing resources across environments. SAS 9.4 users can connect to and use SAS Viya, and SAS Viya users can connect to and use SAS 9.4 data. You can continue to use your SAS 9 projects and custom code while accessing SAS Viya and its new analytic algorithms.

Here are some of the tasks that you can complete when SAS 9.4 and SAS Viya are connected:

- transfer data between systems
- directly process remote data sources and receive results back locally
- use your local graphical user interfaces to process remote data sources
- run multiple independent processes asynchronously
- combine resources from multiple computers to work in parallel

For more information, see *Integrating SAS 9.4 and SAS Viya*.

### *See Also*

"SAS/CONNECT Server and Spawner: How To" in *SAS Viya Administration: SAS/CONNECT Server and Spawner*

## Client/Server Sessions

In SAS Viya, the client session with the workspace server is created when you start SAS Studio.

To use the CAS server, you must explicitly create sessions with the CAS server by using the CAS statement.

In batch mode, line mode processing, and third-party software such as Python, you must provide the server connection information and explicitly create a session with the CAS server.

Sessions are used to provide these benefits:

- user identification
- fault isolation for each session so that if a problem occurs in your session, it does not impact other sessions, the server, or the client
- efficiency:
  - using resources that are visible only to your session
  - concurrent processing of large amounts of data when needed
- resource tracking that is enabled through resource metrics

### *See Also*

- "Sessions" in *SAS Cloud Analytic Services: Fundamentals*
- "CAS Statement" in *SAS Cloud Analytic Services: Language Reference*

## Caslibs

A new concept in SAS Viya is a caslib. The CAS server operates only on tables that reside in a caslib.

A caslib is an in-memory space where tables reside. Associated with a caslib is information about who is authorized to access the tables and information about data source access. A caslib can also be a physical location when copying data from one disk location to another.

In SAS 9, librefs map to a physical location. In CAS, caslibs map to in-memory locations and physical locations.

Caslibs are available to users in one of these ways:

- configured as personal and automatically available
- predefined by an administrator
- manually added by an authorized user

When you start a CAS session, your personal caslib is available and is the active caslib. Your personal caslib is available only to you; tables cannot be shared. Authorized users can use the CASLIB statement to add, drop, or manage caslibs.

If a user has authorization, caslibs can be promoted to global scope to share tables with all users.

### *See Also*

- "Caslibs" in *SAS Cloud Analytic Services: Fundamentals*
- "CASLIB Statement" in *SAS Cloud Analytic Services: Language Reference*

## Loading Data to the CAS Server

Depending on your data source, data is loaded to CAS either serially by the CAS server controller node or in parallel by the CAS worker.

After the CAS session and caslib are in place, you can load client-side tables, SAS data sets, server-side tables, or database tables to CAS. Here are some of the methods for loading data:

- PROC CASUTIL LOADTABLE statement
- CASLIB statement using the DATASOURCE option
- DATA step SET statement
- Table actions addCaslib, addTable, loadDataSource, loadTable, and upload
- Streaming Data actions loadStream and loadSnapshot

Data in a caslib is not automatically saved. You must explicitly save tables in a caslib to its data source location.

### *See Also*

- "CASUTIL" in *SAS Cloud Analytic Services: Language Reference*
- "CASLIB Statement" in *SAS Cloud Analytic Services: Language Reference*
- "Role of the DATA Step in SAS Viya" in *SAS Cloud Analytic Services: Accessing and Manipulating Data*
- "Tables" in *SAS Cloud Analytic Services: System Programming Guide*
- "Data" in *SAS Cloud Analytic Services: Fundamentals*
- Event Stream Processing Action Set

## CAS-Based Procedures and SAS Procedures

SAS Viya contains many new statistical, data mining, and machine learning procedures to process data in memory on the CAS server. Some procedures from SAS 9 continue to run in SAS Viya.

### *Statistical Procedures*

Here are the new statistical procedures:

- The ASSESS procedure assesses and compares supervised learning models.
- The BINNING procedure performs binning, which is a common step in the data preparation stage of the model-building process.
- The CARDINALITY procedure determines a variable's cardinality or limited cardinality.
- The GENSELECT procedure provides model fitting and model building for generalized linear models.

- The KCLUS procedure performs clustering, which is a common step in data exploration.
- The LOGSELECT procedure fits binary and binomial response models.
- The MDSUMMARY procedure computes basic descriptive statistics in parallel for CAS tables.
- The NLMOD procedure uses either nonlinear least squares or maximum likelihood to fit nonlinear regression models.
- The PARTITION procedure performs simple random sampling, stratified sampling, and oversampling to produce a table that contains a subset of observations or partitioned observations.
- The PCA procedure performs principal component analysis, which is a multivariate technique for examining relationships among several quantitative variables.
- The PLSMOD procedure fits models by using any one of a number of linear predictive methods, including partial least squares (PLS).
- The QTRSELECT procedure chooses and fits quantile regression models.
- The REGSELECT procedure fits and performs model selection for ordinary linear least squares models.
- The TREESPLIT procedure builds tree-based statistical models for classification and regression.
- The VARIMPUTE procedure performs numeric variable imputation.
- The VARREDUCE procedure performs both supervised and unsupervised variable selection.

### See Also

*SAS Visual Data Mining and Machine Learning: Statistical Procedures*

## Data Mining and Machine Learning Procedures

Here are the new data mining and machine learning procedures:

- The ASTORE procedure creates score code that can use code from an analytic store as well as DS2 code.
- The BOOLRULE procedure enables you to extract Boolean rules from large-scale transactional data.
- The FACTMAC procedure implements the factorization machine model.
- The FOREST procedure creates a predictive model called a forest, which consists of several decision trees.
- The GRADBOOST procedure creates a predictive model called a gradient boosting model.
- The NETWORK procedure provides a number of network analysis algorithms that take an abstract graph or network as input, helps explain network structure, and computes important network measures.
- The SVMACHINE procedure implements the support vector machines (SVM) algorithm. A popular data mining area classification method, the SVM algorithm computes support vector machine learning classifiers for the binary pattern recognition problem. The SVM algorithm has been broadly used in fields such as image classification, handwriting recognition, financial decision, and text mining.

- The TEXTMINE procedure integrates natural language processing and statistical analysis to analyze large-scale textual data.
- The TMSCORE procedure scores textual data. In text mining, scoring is the process of applying parsing and singular value decomposition projections to new textual data.

### See Also

- *SAS Visual Data Mining and Machine Learning: Data Mining and Machine Learning Procedures*
- *SAS Visual Data Mining and Machine Learning: The NETWORK Procedure*

## PROC CASUTIL

The CASUTIL procedure manages tables. You use PROC CASUTIL to load, save, and unload CAS tables, display table metadata and tables in a caslib, and promote tables to global scope.

### See Also

"CASUTIL" in *SAS Cloud Analytic Services: Language Reference*

## PROC CAS and CASL (CAS Language)

SAS Viya adds a new analytics scripting language called CASL. You access this language by using PROC CAS, which connects to the in-memory engine. As a result, CASL has access to pooled memory.

CASL should be familiar to any experienced SAS user. It is designed to look and feel like a SAS language, such as DATA step and IML.

### See Also

*SAS Cloud Analytic Services: CAS Procedure Programming Guide and Reference*

## SAS Procedures

SAS Viya includes some of the Base SAS procedures that you are familiar with from previous releases of SAS. Some of these procedures can access and process CAS tables and table metadata in the local SAS session. Others, such as PROC DS2 and PROC FEDSQL, submit their language code to the CAS server.

- The APPEND procedure adds rows from a CAS table to the end of a SAS data set and adds rows from a SAS data set to the end of a CAS table.
- The CONTENTS procedure shows the contents of a CAS table and prints the directory of the caslib.
- The COPY procedure copies tables to and from libraries.
- The DATASETS procedure also manages CAS tables.
- The DELETE procedure deletes SAS data sets and CAS tables.
- The DS2 procedure runs DS2 code.
- The EXPORT procedure writes a SAS data set to delimited or JMP files.

- The FEDSQL procedure submits a limited set of FedSQL code.
- The FORMAT procedure creates user-defined formats.
- The HTTP procedure processes data from the web.
- The IMPORT procedure reads external data into a SAS data set.
- The JAVAINFO procedure shows information about the version of Java on your system.
- The OPTIONS procedure lists the current settings of SAS system options in the SAS log.
- The PRINT procedure prints SAS data sets and CAS tables.
- The PRINTTO procedure redirects output.
- The PRODUCT_STATUS procedure lists the SAS products that are installed on your system.
- The PWENCODE procedure encodes passwords.
- The SGGPANEL procedure creates a panel of graph cells for the values of one or more classification variables.
- The SGPLOT procedure creates statistical graphics such as histograms and regression plots, in addition to simple graphics such as scatter plots and line plots.
- The SGSCATTER procedure creates a paneled graph of scatter plots for multiple combinations of variables.
- The SORT procedure sorts SAS data sets in the local SAS session.
- The TRANSPOSE procedure creates a CAS table by restructuring the values in the CAS table and transposing selected columns into rows.

### *See Also*

*SAS Viya Data Management and Utility Procedures Guide*

## CAS Actions for SAS and Third-Party Languages

CAS actions are a new programming language. Each action performs a single task. You submit CAS actions to the CAS server either by using PROC CAS or by using third-party software Python, Lua, or Java. Actions are aggregated in action sets based on functionality.

Here are some of the types of CAS actions:

- data mining and machine learning, including image and text processing
- statistics
- DATA step processing
- event stream processing
- table management
- server management

### See Also

- *SAS Cloud Analytic Services: Data Mining and Machine Learning Programming Guide*
- *SAS Cloud Analytic Services: Statistics Programming Guide*
- *SAS Cloud Analytic Services: System Programming Guide*
- *SAS Cloud Analytic Services: Analytics Programming Guide*
- *SAS Cloud Analytic Services: CAS Procedure Programming Guide and Reference*

## DATA Step

The DATA step can run on the CAS server and on the SAS Workspace Server. If a session with the CAS server does not exist, the DATA step runs on the SAS Workspace Server. In a CAS server session, the DATA step runs in parallel across multiple computers, which is referred to as a CAS cluster. The same DATA step program is replicated across all computers in the CAS cluster, where each DATA step thread runs on a portion of the data. By default, the DATA step runs in all available threads on every computer node in the cluster. For each node, the DATA step processes only the table rows that are allocated to that node.

Not all DATA step language elements that you have used in previous releases of SAS are appropriate for distributed processing in the CAS environment. Therefore, only a select set of SAS language elements are valid for processing in CAS. For a list of these language elements, see the category table in the respective documentation:

- Data Set Options
- Formats
- Functions
- Statements: DATA Step and Global
- Hash and Hash Iterator Objects
- Java Object Methods

Character functions use character semantics and not byte semantics that are used in traditional SAS.

The DATA step has some requirements and restrictions when used on the CAS server:

- "Restrictions for DATA Step Processing in CAS" in *SAS Cloud Analytic Services: Accessing and Manipulating Data*
- "Restrictions for Using the VARCHAR Data Type in the DATA Step" in *SAS Cloud Analytic Services: Accessing and Manipulating Data*

When using BY variables in a multi-node cluster, CAS distributes table data across multiple computers, or nodes. CAS divides the rows equally among all the available nodes in the cluster and distributes the rows to each node. Sorting is not required before BY-group processing in CAS. Because ordering of rows is not guaranteed in CAS, you can use additional BY variables to impose order. In general, when the DATA step runs on a multi-node CAS server, CAS partitions and distributes the table data across the cluster and then implicitly orders the data locally on each node.

### See Also

- “DATA Step Basics” in *SAS Cloud Analytic Services: Accessing and Manipulating Data*
- “How CAS Groups Data with BY Variables” in *SAS Cloud Analytic Services: Accessing and Manipulating Data*
- “How BY Variables Affect Multi-threaded DATA Step Execution” in *SAS Cloud Analytic Services: Accessing and Manipulating Data*

## Data Types

The DATA step and the new CAS procedures support the CHAR, DOUBLE, and VARCHAR data types. Variables that are created using the VARCHAR data type are of varying width and use character semantics. The traditional SAS CHARACTER data type is of fixed-width and uses byte semantics.

The CAS language (CASL) supports the BLOB, BOOLEAN, CHAR, DATE, DATETIME, DOUBLE, INT32, INT64, ITEMS, LISTS, STRING UTF-8, TABLE, and TIME data types.

DS2 and FedSQL support the same data types in CAS as they support in SAS 9.

### See Also

“SAS Cloud Analytic Services Data Types” in *SAS Cloud Analytic Services: Language Reference*

## User-Defined Formats

SAS provides formats for controlling how variables are printed. You can use PROC FORMAT in your SAS programs, as you do in SAS 9, to create user-defined formats. PROC FORMAT supports the creation of user-defined formats in catalogs on the SAS client and supports adding formats and format libraries for a CAS session.

Format libraries are associated with a CAS session, or they can be promoted to global scope to be available to all CAS sessions. These format libraries are server-side formats that the server uses when an analysis is performed according to formatted values.

A new option in PROC FORMAT, CASFMTLIB, names the CAS format library for the active caslib where the format is stored. Without the CASFMTLIB= option, formats continue to be stored in a format catalog.

You must assign a format to a variable before the table is loaded into the server. After a table is in memory, you cannot assign a format. You can assign a format with the FORMAT statement in a DATA step or with the CASUTIL procedure.

### See Also

- “User-Defined Format Basics” in *SAS Cloud Analytic Services: Accessing and Manipulating Data*

- “Formatting Table Variables for Printing” in *SAS Cloud Analytic Services: System Programming Guide*

## Support for Macros in SAS Viya

Macros are supported in SAS Viya, but only in the SAS session. Macros can be useful in generating the code that runs in CAS, but the macro language itself is not supported in CAS.

### *See Also*

*SAS Viya Macro Language: Reference*

## Support for DS2

You can run DS2 code on the SAS Workspace Server and the CAS server. In CAS, you run DS2 code by using PROC DS2 or the runDS2 action. You cannot use FedSQL queries in DS2 code when you run on the CAS server.

### *See Also*

- *SAS Viya: DS2 Programmer’s Guide*
- *SAS Viya: DS2 Language Reference*

## Support for FedSQL

In CAS, for SQL operations, you can use PROC FEDSQL or the FedSQL.execDirect action.

You can submit SQL queries on SAS data sets and CAS tables, with some exceptions for support in CAS.

In CAS, FedSQL supports these statements:

- CREATE TABLE AS
- SELECT
- DROP TABLE

SAS data sets and CAS tables cannot be joined.

The SELECT operations have some limitations. These operations are not supported in CAS:

- SET operations
- correlated subqueries
- views

- DS2 user-defined formats
- SQL pass-through

PROC FEDSQL cannot save or promote CAS tables. You must use PROC CASUTIL to save or promote a FedSQL CAS table.

When you run PROC FEDSQL in the local SAS session, all statements that you used in SAS 9 are supported.

### *See Also*

*SAS FedSQL Language Reference for SAS Cloud Analytic Services*

## Output in SAS Viya

The data mining and machine learning procedures create in-memory output tables that contain the procedure results. Some of these tables are used as input for scoring, and some are the results of scoring. Several of these procedures save their models as an analytic store. An *analytic store* is a binary object that contains that procedure’s state after the procedure completes the training phase of data analysis. You use an analytic store with the ASTORE procedure to score data at a later time.

SAS Studio, by default, creates HTML output on the **RESULTS** tab. The **OUTPUT DATA** tab opens to the table metadata if your program creates a CAS table.

The **RESULTS** tab shows information about the procedure, the caslib, and the table columns as well as other types of information based on the code that you submit.

You can save in-memory tables to disk. By default, the suffix is .sashdat. If you specify a .csv suffix, a CSV file is saved.

# Recommended Reading

- *Integrating SAS 9.4 and SAS Viya*
- *SAS Viya Quick Start*
- *SAS Cloud Analytic Services: Fundamentals*
- *SAS Viya: Deployment Guide*
- *Getting Started with SAS Visual Data Mining and Machine Learning*

For a complete list of SAS publications, go to sas.com/store/books. If you have questions about which titles you need, please contact a SAS Representative:

SAS Books
SAS Campus Drive
Cary, NC 27513-2414
Phone: 1-800-727-0025
Fax: 1-919-677-4444
Email: sasbook@sas.com
Web address: sas.com/store/books