SAS® Viya® 3.4 Administration

SAS® Documentation
May 22, 2019

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2018. *SAS® Viya® 3.4 Administration*. Cary, NC: SAS Institute Inc.

SAS® Viya® 3.4 Administration

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

May 2019

3.4-P1:caladmin

Contents

About This Document

Purpose

This document is provided as a convenience to provide offline access to administrative information about SAS Viya 3.4.

Scope

This document combines the individual administrative guides for SAS Viya 3.4, presenting each individual guide as a chapter.

To reduce file size and frequency of updates, this document excludes the following guides:

- deployment guides
- *SAS Viya Administration: Multi-tenancy* (PDF)
- *SAS Viya Administration: QKB Management* (PDF)
- *SAS Viya Administration: Publishing Destinations* (PDF)
- *SAS Viya Administration: Promotion (Import and Export)* (PDF)
- *SAS Data Explorer: User's Guide* (PDF)
- *Encryption in SAS Viya: Data in Motion* (PDF)
- *SAS Cloud Analytic Services: Fundamentals* (PDF)

Access

The primary access point for the information that this document contains is the SAS Help Center.

> **TIP** For metadata about any individual guide, access the HTML version in SAS Help Center, select ? in the banner, and then select **About**. Details include the guide's title, the copyright date, and the date of the last update.

1

Orientation

SAS Viya Administration: Orientation

Documentation

This documentation supports administration of the SAS Viya 3.4 components of the SAS platform and the following products:

- SAS Visual Analytics 8.3
- SAS Visual Statistics 8.3
- SAS Visual Data Mining and Machine Learning 8.3

Other SAS Viya products supplement this documentation with their own product-specific administrative information, as needed.

Note: If this document is titled **SAS Viya 3.4 Administration: Orientation**, access the entire collection before you search.

Note: This document does not include the deployment guides for SAS Viya: *SAS Viya for Linux: Deployment Guide* and *SAS Viya for Windows: Deployment Guide*. See also the SAS Viya Deployment Guides page on the SAS support site.

Deployment Types

Some SAS Viya administrative tasks and tools vary by deployment type. Here is a brief description of each type:

Deployment Type	Description
Full	Includes all of the software to which you are entitled. This is the default type of deployment.
Programming-only	Excludes SAS Drive, most graphical user interfaces, and most services. This is the simplest and smallest type of deployment. Note: If SAS Drive is available (at your equivalent of http://*host*/SASDrive), you do not have a programming-only deployment.

For details, see “Diagrams by Deployment Type” on page 18.

Note: The programming-only deployment type relies on a subset of SAS Viya components to provide programmatic access to SAS Cloud Analytic Services for high-performance analytical processing of in-memory data. Convenient for sites that choose to install and use a minimal subset of the software, the programming-only deployment type does not correspond to a limitation in software licensing or entitlement.

Administrative Tools

Here are the main administrative tools:

	Deployment Type	
Tool	**Full**	**Programming-Only**
SAS Environment Manager	✓	
Command-Line Interface	✓	
CAS Server Monitor		✓

Routine Ongoing Tasks

Here is a summary of routine administrative tasks:

	Deployment Type		*Multi-tenant* Scope	
Task	**Full**	**Programming -Only**	**Provider-Level***	**Intra-tenant****
View logs	✓	✓	✓	
Monitor services	✓		✓	
Start and stop CAS	✓	✓	✓	
Update your software***	✓	✓	✓	
Renew your license	✓	✓	✓	
Create backups	✓	✓	✓	✓
Add caslibs	✓	✓		✓
Manage access to data	✓	✓		✓
Manage access to content	✓			✓
Manage access to functionality	✓			✓
Promote content	✓			✓
Assign users to custom groups	✓			✓
Manage mobile devices	✓			✓

Task	**Deployment Type**		***Multi-tenant* Scope**	
	Full	**Programming-Only**	**Provider-Level**[*]	**Intra-tenant**[**]
Schedule and monitor jobs	✓			✓
Start and stop microservices	✓	Not applicable[†]	✓	
Manage tenants	✓		✓	
Docker Containers	Not applicable	✓	Not applicable	Not applicable

[*] These are the most common tasks that the provider in a *multi-tenant* deployment performs.

[**] An administrator within a tenant can perform these tasks for that tenant.

[***] See your deployment guide (for example, *SAS Viya for Linux: Deployment Guide*).

[†] SAS Viya microservices are not installed on programming-only deployment types. However, you must still restart CAS and SAS Object Spawner.

See Also

- *SAS Viya Administration: Initial Tasks*

2

Initial Tasks

Initial Tasks in SAS Viya Administration

Introduction

This topic assumes that all applicable tasks in your deployment guide have been completed.

Complete only the tasks for your deployment type.

Note: If you have a multi-tenant deployment, skip these steps. Instead, see Provider Administrator: Onboard Tenants.

Initial Tasks in a Full Deployment

1 Familiarize yourself with the predefined custom groups, and add members if needed.

2 If you have users who access CAS from both programming and visual interfaces, review the host access considerations.

3 Limit access to certain system paths for CAS and for the SAS programming run-time servers.

4 Decide whether to implement CAS resource management policies.

5 Promote any supported content from previous releases.

6 Address any applicable considerations for SAS 9 and SAS Viya integration.

7 Create a backup, and set up automated backups.

8 Send programmers the following link: An Introduction to SAS Viya Programming.

Send visual interface users the following link: SAS Viya.

Initial Tasks in a Programming-Only Deployment

1 Add members to the Superuser role for your CAS server.

2 Limit access to certain system paths for CAS and for the SAS programming run-time servers.

3 Address any applicable considerations for SAS 9 and SAS Viya integration.

4 Create a backup of CAS information and configuration files.

5 Send programmers the following link for information about programming in SAS: An Introduction to SAS Viya Programming.

See Also

- "Routine Ongoing Tasks" on page 2
- *SAS Viya: Overview*

3

What’s New

What’s New in Administration of SAS Viya 3.4: Highlights

Here are the key administrative features that are introduced in SAS Viya 3.4:

- A new promotion wizard is available for importing content from SAS 9.x to SAS Viya 3.4 and SAS Viya 3.x to SAS Viya 3.4. For more information, see “Promotion: How to Import (Wizard)” in *SAS Viya Administration: Promotion (Import and Export)*.
- SAS Viya now supports additional Cloud Analytic Services (CAS) servers. For more information, see “Add a CAS Server” on page 477.
- CAS has extended support for resource management through policies for table quotas and CPU utilization. For more information, see “CAS Resource Management” on page 505.
- A new license file has been introduced in preparation for future billing enhancements.

  For more information, see “Licensing: How To” on page 182.
- In SAS Environment Manager, new and enhanced interfaces help you monitor and schedule jobs, view the properties and status of servers, manage improved dashboards, and configure user-defined formats. See SAS Viya Administration: Using SAS Environment Manager on page 663.
- SAS Viya now supports Windows Server.
- SAS Viya now supports running SAS software in a container on Docker and Kubernetes.

See Also

"What's New in Administration of SAS Viya 3.4: Details" on page 8

What's New in Administration of SAS Viya 3.4: Details

This topic provides details. A summary of highlights on page 7 is available.

Authorization

- In the authorization service, a new endpoint supports sharing of content among users. In SAS Environment Manager, share-based rules are visible on the **Rules** page. The Authorization window for content objects has been updated to reflect access that comes from sharing. See "Sharing: Details for Administrators".
- In the authorization service, the mediaType property is deprecated and replaced by the more specific media type properties acceptType, acceptItemType, and contentType. See "acceptType, acceptItemType, and contentType".
- In the Authorization window for content objects, origins information is directly included in each per-cell pop-up. See "Origins of Effective Access".
- The Authorization window for content objects is available in SAS Drive, as well as on the **Content** page in SAS Environment Manager. See "Navigation".
- The Authorization window for caslibs and tables is available in SAS Data Explorer and the Choose Data window, as well as on the **Data** page in SAS Environment Manager. See "Navigation".
- In the Authorization window, the icon for adding identities has been changed from ✚ to 👤.
- The Authorization window supports horizontal labels for permissions.
- In SAS Environment Manager, several filters have been added to the **Rules** page.
  - The **Rule Status** filter enables you to find disabled rules.
  - The **Settings** filter enables you to specify one or more settings (Grant, Conditional Grant, Prohibit, and Conditional Prohibit) as your filter criteria.
  - The **Media Type** filter helps you find rules that specify the mediaType property, which has been deprecated.
- Additional enhancements to the **Rules** page are as follows:
  - The filtering on the **Rules** page is no longer case-sensitive.
  - You can search for rules that contain specified text in the **ObjectUri**, **Description**, **Reason**, or **Condition** fields.
  - The Enabled column is now labeled and has values of `Enabled` and `Disabled`. The same text changes are present wherever rule properties are displayed in SAS Environment Manager.
- In CAS authorization, the new action, accessControl.accessPersonalCaslibs, enables administrators to see all personal caslibs and to drop promoted tables in any personal caslib. See Access to Personal Caslibs in *SAS Viya 3.4: System Programming Guide*.

Servers and Services

- In SAS Environment Manager, a new **Servers** page enables you to view and manage your SAS Viya servers.

- The following CAS server options have been added:
  - □ cas.CPUSHARES on page 514
  - □ cas.MAXCORES on page 523
- The following CAS environment variables have been added:
  - □ env.CAS_ADDITIONAL_YARN_OPTIONS on page 536
  - □ env.CAS_CONTROLLER_TEMP on page 536
  - □ env.CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT on page 538
  - □ env.SAS_RNG_METHOD on page 542
  - □ env.CAS_START_MONITOR_UI on page 540
- The CAS server option, cas.ELASTICSSL, has been deprecated.
- SAS Viya Programming Run-Time Servers now support the following system option and statement:
  - □ LOCKDOWN on page 589
  - □ LOCKDOWN statement on page 590
- The new SAS Viya service, GPU Reservation service, aids SAS processes in resource sharing and utilization of graphics processing units (GPUs) available on a system.

Command-Line Interfaces

- You can manage the SAS Quality Knowledge Bases (QKBs) on a Cloud Analytics Services (CAS) server with the new QKBs CLI. See "CLI Examples:Quality Knowledge Bases (QKBs)" on page 722.
- You can do the following with the cas CLI:
  - □ Create caslibs that are based on the following sources:

    Note: See "Manage Caslibs" on page 708.

    - DB2 database
    - file path in external file system (distributed network file system)
    - Hadoop database that is managed by Hive
    - SAP Hana
    - file path in Hadoop Distributed File System
    - Impala database
    - SAS LASR Analytic Server
    - ODBC data source
    - Oracle database
    - local file path
    - PostgreSQL database
    - Amazon Redshift
    - Teradata database
  - □ Create and manage SAS format libraries, the SAS formats that they contain, and the SAS format search order. See "CLI Examples: Formats" on page 716.
  - □ Manage the paths list that controls where caslibs can be created for a CAS server. See paths list example on page 709.

  - Manage the ability to create and delete session and global caslibs. See privileges example on page 709.
  - Create and manage CAS server loggers.
  - Manage resource management policies. See policies examples on page 709.
  - Generate sample template files for creating caslibs, user-defined formats, and resource management policies. See "Generate CAS Samples" on page 706 and "Source Files: JSON Templates " on page 692.
- You can do the following with the reports CLI:
  - Manage the themes that are used for SAS Visual Analytics reports. See themes example on page 724.
  - Manage the translations worksheets that are associated with a SAS Visual Analytics report. See translations example on page 725.
- You can do the following with the folders CLI:
  - Use additional filtering options when listing folders.
  - Reference a folder by path as well as ID. See "CLI Examples: Folders" on page 714.

Configuration Properties

The following services have new or deprecated configuration properties for SAS Viya 3.4:

- Backup service
- Report Data service
- Report Package service
- Report Renderer service

Licensing

- A new license file changes how SAS Viya licenses are renewed.

  For more information, see "Licensing: How To" on page 182.

SAS Environment Manager

Data

- The **Data** page incorporates the functions of the Data Explorer application.

Servers

The new **Servers** page enables you to do the following:

- view basic server information such as the status, host, and port.
- view information about the sessions on each server.
- view the CAS configuration properties for each server.
- view information about the nodes on each server.
- view a list of the users and groups that are members of the superuser role.
- view a list of users and groups who can create session caslibs or global caslibs on each server. Users who have the superuser role can edit and add to the entries in the list.

- view a list of the loggers and associated threshold levels for each server. Users who have the superuser role can edit the threshold levels and add loggers.

The new **Servers** page enables these users to do the following:

- Users who have assumed the superuser role can view and edit the paths that can be used for caslibs.
- Users who have assumed the SAS Administrator role can assume or relinquish the superuser role for a server.

Promotion

The **Import Wizard** enables you to do the following:

- promote internal identity groups from SAS 9 into SAS Viya (export and preparation outside of SAS Environment Manager is first required).
- promote definitions for base libraries, LASR libraries, and LASR tables from SAS 9 into SAS Viya. (Export and preparation outside of SAS Environment Manager is first required.)
- promote folders, reports, explorations, and supporting resources from SAS 9 into SAS Viya. (Export and preparation outside of SAS Environment Manager is first required.)

The **Export Wizard** and **Import Wizard** enable you to promote content from a previous version of SAS Viya into SAS Viya 3.4. You can also use the command-line interface to promote content.

User-Defined Formats

The new **User-Defined Formats** page enables you to do the following:

- display information about all of the user-defined formats and format libraries that are available for the data.
- add and import new user-defined formats.
- edit, copy, and delete existing user-defined formats.
- import formats from a SAS item store (although some steps are required outside of SAS Environment Manager).
- create, delete, and change the search order of format libraries.

Jobs

The new **Jobs** page:

- combines the functions of the Job Monitor application and the Scheduling function.
- enables you to view a table or a chart of jobs that are currently running and that have run in a specified time in the past.
- enables you to view a chart of jobs that are scheduled for the future.

My Credentials

- The **My Credentials** page (named **My Passwords** in SAS Viya 3.3) enables you to manage personal domain credentials.

Publishing Destinations

- The new **Publishing Destinations** page enables you to define, update, and delete publishing destinations that are used by SAS Decision Manager, SAS Model Manager, and Model Studio.

Encryption

- You can encrypt data at rest with additional encryption options. By default, data at rest is presumed to be behind the firewall and is not encrypted.
  - When using the SAS Environment Manager, you now enable a calib for encryption using SAS Data Explorer. When you select **Data** from the SAS Environment Manager window, you are now using SAS Data Explorer to add a caslib and then enable it for encryption. See "Manage Caslibs That Are Enabled for Encryption" on page 448.
  - In the Domains window, the icon for adding identities has been changed from ✚ to 👤.
- For data in motion, in a full deployment of SAS Viya, all external communication paths are secured by default. The following are new in this release of SAS Viya:
  - There are two ways to secure an LDAP connection, LDAPS and STARTTLS. STARTTLS upgrades a connection that is not encrypted, by wrapping it with TLS during the connection process. This allows unencrypted and encrypted connections to be handled on the same port. See "Configure the Secure LDAP Connection Using STARTTLS " in *Encryption in SAS Viya: Data in Motion*.
  - The SAS Configuration Server is secure by default. In the vars.yml file, by default, the SECURE_CONSUL setting is set to `true` and DISABLE_CONSUL_HTTP_PORT is set to `true`. Only the HTTPS port is available after the software is deployed. You can change these settings post-deployment. See "Enable or Disable TLS on the SAS Configuration Server Ports" in *Encryption in SAS Viya: Data in Motion*.

Mobile

- SAS Mobile BI is now SAS Visual Analytics App.
- SAS SDK for Android and SAS SDK for iOS are available for download. These Software Development Kits (SDKs) enable your mobile apps to include SAS Visual Analytics content. See "Mobile: Software Development Kits" on page 281.
- In the Mobile Devices window, the device status icons have changed. See Table 17.44 on page 277.

Authentication

- SAS Launcher Server supports user authentication using Kerberos. See "Configure Kerberos for SAS Launcher Server" on page 294.
- Integration between SAS Logon Manager on SAS 9.4 and SAS Logon Manager on SAS Viya is supported.
- While enabling guest access, the commands to modify the authorization rules differ, based on which of the following tasks you are doing:
  - performing a new SAS Viya 3.4 installation
  - upgrading from SAS Viya 3.3 to SAS Viya 3.4
  - upgrading from SAS Viya 3.2 and earlier to SAS Viya 3.4

  See Step 2b in "Enable Guest Access" in SAS Viya Administration: Authentication on page 327.

Tenancy

- Support for customized LDAP configurations includes configuring the LDAP server per tenant.
- A new offboarding playbook can be used to offboard tenants.

Support for Windows

- SAS Viya now supports Windows Server.
- Support is limited to a single-machine deployment.

SAS for Containers

- Create Docker container deployments of SAS Viya.
- Programming-only deployments are supported.

4

SAS Viya Overview

Elements of SAS Viya

Introduction

This section provides a concise summary for new administrators.

Here are related topics:

- To get started with SAS Viya administration, see Orientation on page 1.
- To learn about benefits of SAS Viya, see SAS Viya on the SAS website.

Key Components

Here are software components that might be of particular interest to administrators.

The analytics engine to SAS Viya	*SAS Cloud Analytic Services: Fundamentals*
A modular set of supporting services	*SAS Viya Administration: General Servers and Services*

A web application for basic administration	CAS Server Monitor on page 486
A web application for enterprise administration	*SAS Viya Administration: Using SAS Environment Manager*
A web application for writing and submitting code	*Getting Started with Programming in SAS Studio*
A web application for visual reporting, exploration, and modeling	*SAS Visual Analytics: Overview*
Multiple application programming interfaces	http://developer.sas.com

For information about other components, search the SAS Viya administration documentation.

Cumulative Functionality

Among some of the products on SAS Viya, available functionality is cumulative.

- SAS Visual Analytics provides baseline functionality, including reporting and basic analytics.
- SAS Visual Statistics provides an additional set of advanced analytic functions.
- SAS Visual Data Mining and Machine Learning provides a second additional set of advanced analytic functions.

For example, if you have SAS Visual Data Mining and Machine Learning, the objects that are available in the SAS Visual Analytics web application are as follows:

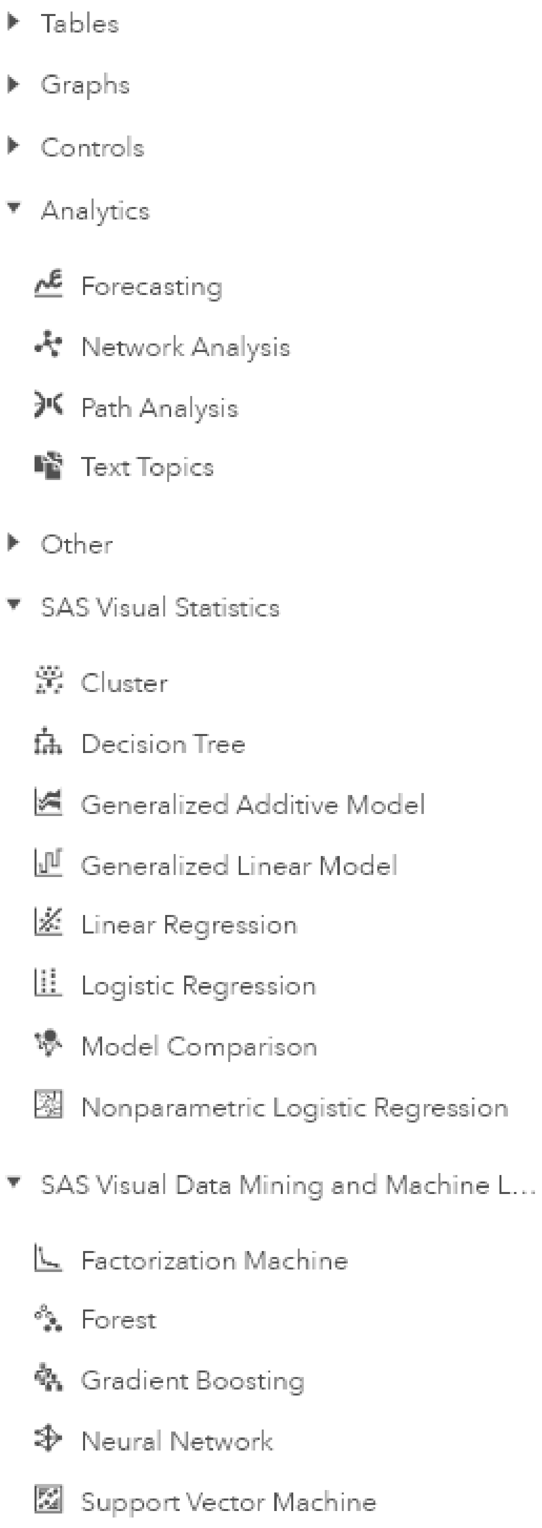

Note: All three of the preceding products offer both programming and visual interfaces.

Selective Deployment (Optional)

By default, all of your software is deployed. As a convenience for special circumstances, it is possible to deploy only a subset of components. A programming-only deployment excludes general services and visual interfaces.

For example, a programming-only deployment of SAS Visual Analytics does not include the SAS Visual Analytics web application.

Note: SAS supplies two versions of SAS Studio, version 4 and version 5. SAS Viya programming-only deployments use SAS Studio 4. For a comparison of the two SAS Studio versions, see "SAS Studio 5.1 and 4.4" in *What's New in SAS 9.4 and SAS Viya*.

Diagrams by Deployment Type

Full Deployment (Native Operating Systems)

The following diagram shows the components in a SAS Viya full deployment:

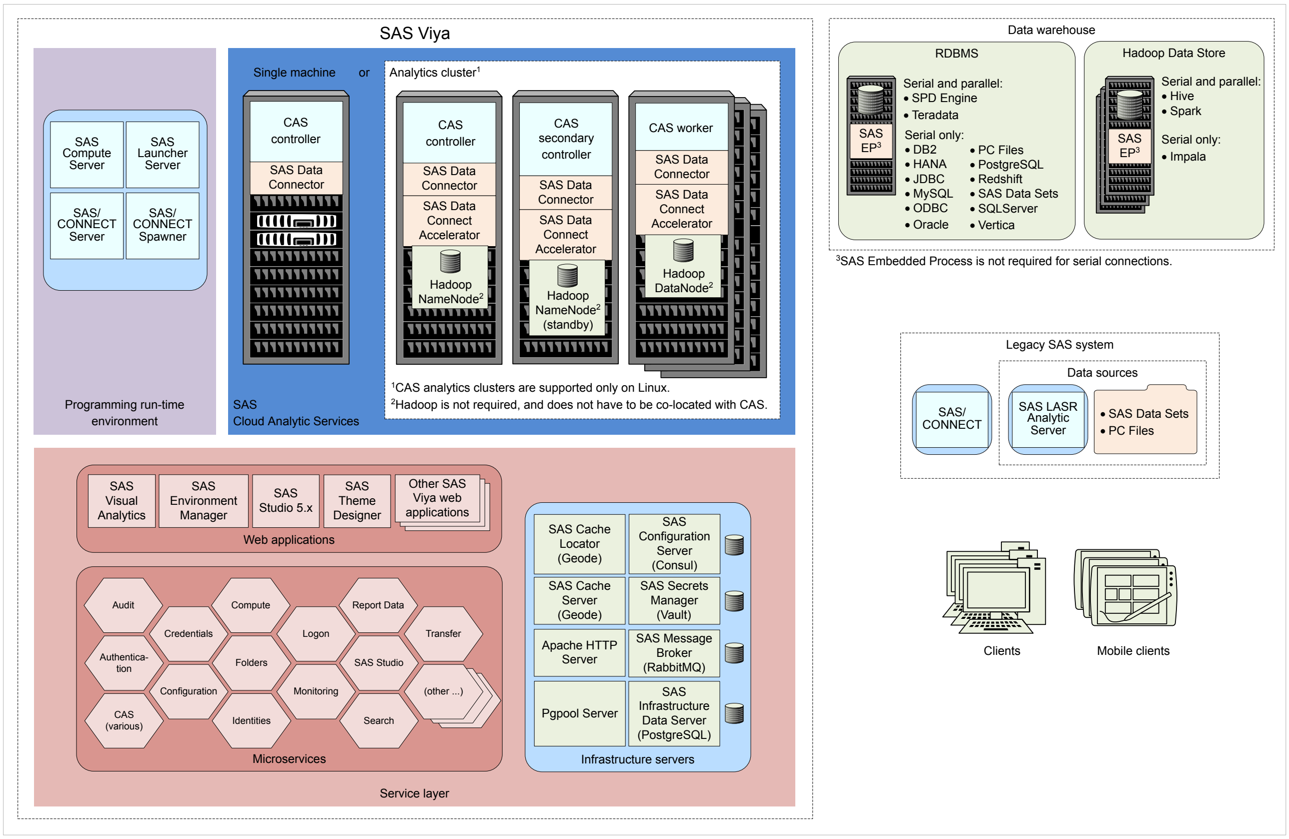

Programming-Only Deployment (Native Operating Systems)

The following diagram shows the components in a SAS Viya programming-only deployment:

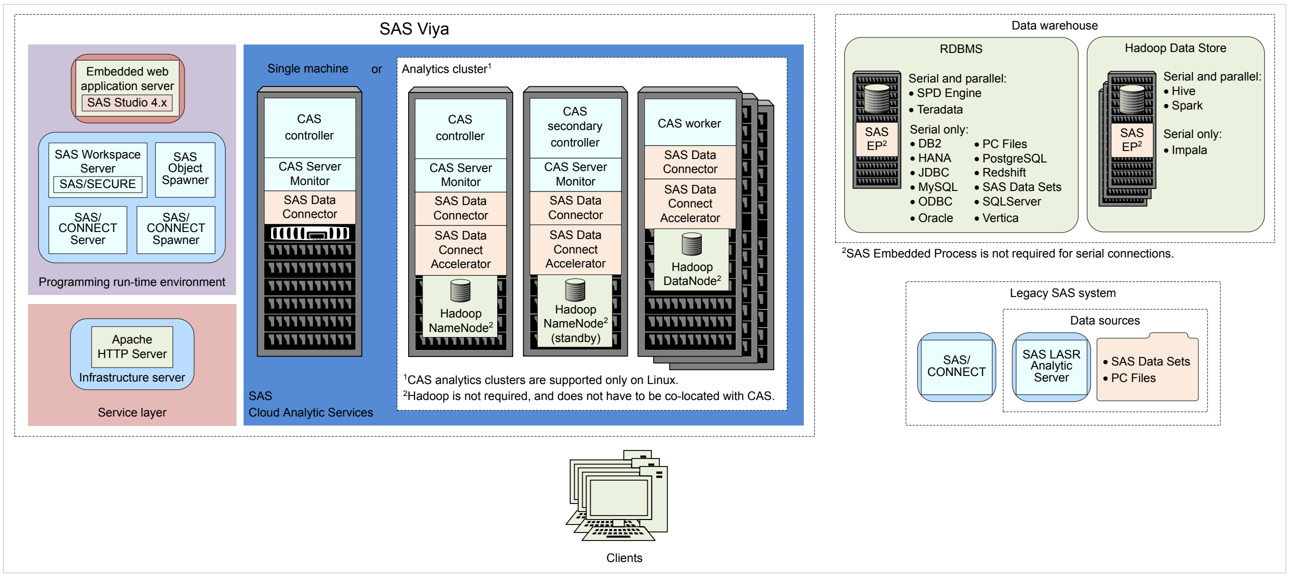

Security in SAS Viya

Authentication

Authentication is the aspect of security that verifies the identity of a user or service account.

When you sign in, one of the following authentication patterns is used:

Pattern	Description	Usage
Host authentication	Requests are sent to the appropriate host and processed by any authentication mechanism supported by that host. Programming-only deployments use this pattern exclusively. Other deployments use dual authentication for access to CAS from SAS Studio 4. **Note:** You can configure the host to use pluggable authentication modules (PAM). SAS provides starter PAM configuration files for CAS and SAS Studio 4. You can create an authinfo file for use with PAM in command-line access and batch processing programs to CAS. Credentials for the user ID that runs the program are supplied from the authinfo file.	When you sign in to SAS Studio 4 from a URL that is similar to https://*reverse-proxy-server*/SASStudio/, you are prompted for a user ID and password. The associated object spawner asks its host (which is also the host of the SAS Studio 4 web application) to validate your credentials. That validation enables the object spawner to launch a workspace server for you. When you access CAS from SAS Studio 4, you must authenticate to the host of the target CAS server. When you sign in to CAS Server Monitor, you must authenticate to the host of the target CAS server.

Pattern	Description	Usage
Direct LDAP authentication	Requests are sent to and processed by your designated direct LDAP provider, unless you configure front-end single sign-on using Kerberos, Open Authorization (OAuth), or Security Assertion Markup Language (SAML). Kerberos, OAuth, and SAML are alternate mechanisms for identity verification by the logon service, not alternate sources of user and group information for the identities service. **Note:** User and group information is always obtained from your designated direct LDAP provider.	SAS Drive enables you to access the visual interfaces, for example, SAS Visual Analytics or SAS Environment Manager. When you sign in to SAS Drive from a URL that is similar to https://*reverse-proxy-server*/SASDrive/, a user ID and password are required to authenticate to SAS Logon Manager, using this pattern. Before you can submit a command-line request to a general service (for example, the backup service or the transfer service), you must authenticate using this pattern.
Host and direct LDAP authentication	Requests are authenticated using both host authentication and direct LDAP authentication. If the servicesBaseUrl option is specified, CAS requires dual authentication. To facilitate this pattern, use one of these approaches: ■ Ensure that all requests are ultimately processed by the same authentication provider. For example, configure the SAS Studio 5 and CAS hosts to use the same LDAP provider that is designated for direct LDAP authentication requests in your deployment. ■ Ensure that each affected user has a single set of credentials that are valid for all applicable authentication providers.	In a full deployment, dual authentication occurs for access to CAS from SAS Studio 5. In a programming-only deployment, CAS Server Monitor provides a web-based interface for administration. **Note:** When you access CAS from a web application such as SAS Visual Analytics or SAS Environment Manager, your OAuth token is validated.

The following high-level conceptual drawings illustrate key points from the preceding table:

Figure A.1 *Authentication from SAS Studio or CAS Server Monitor to CAS*

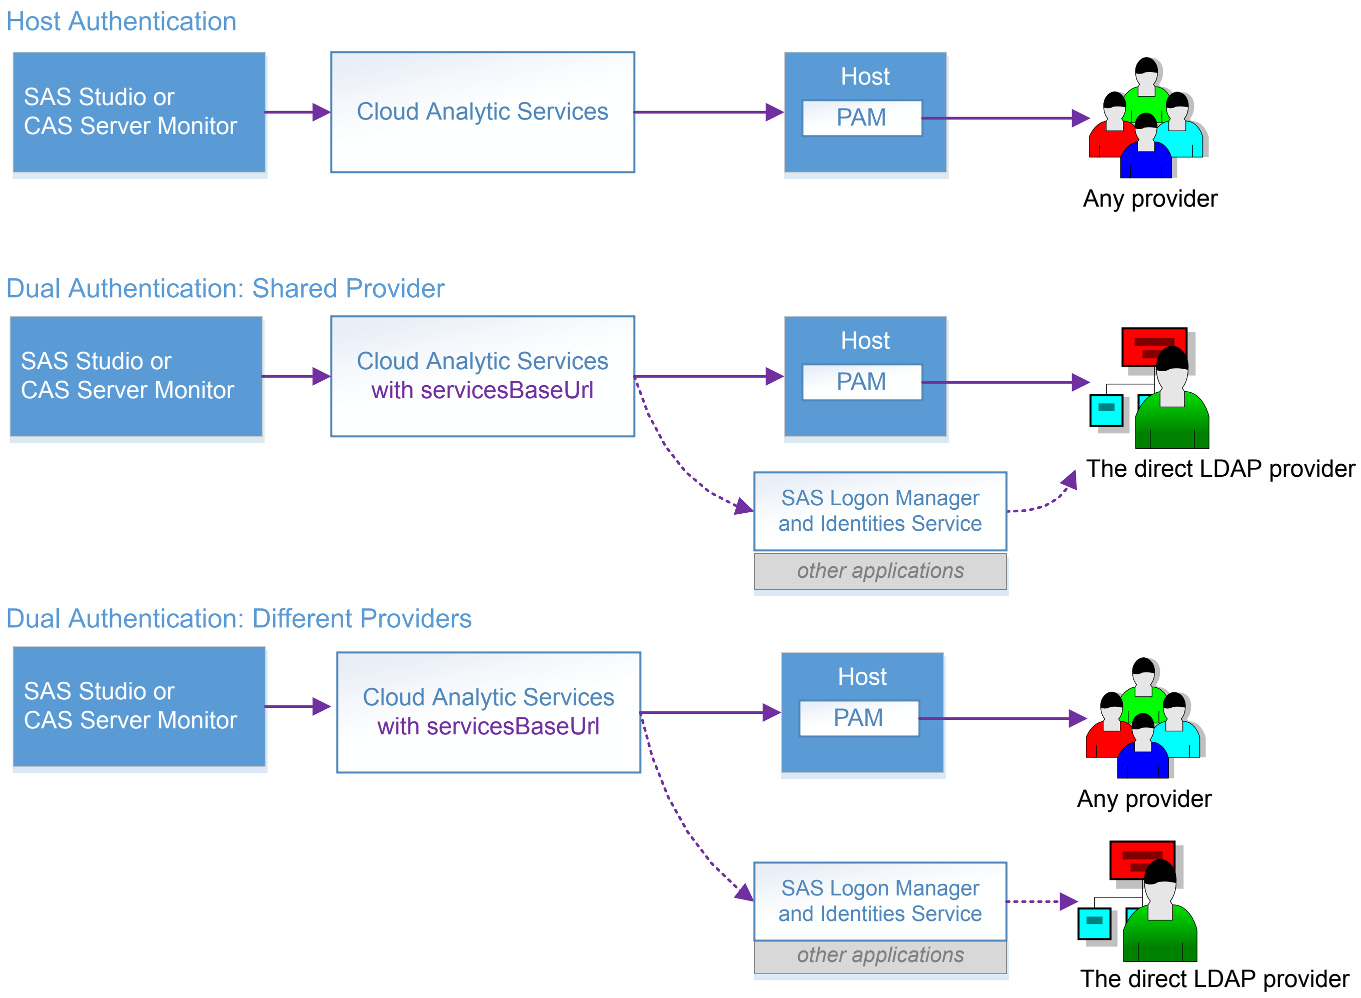
Host Authentication
SAS Studio or CAS Server Monitor
Cloud Analytic Services
Host
PAM
Any provider
Dual Authentication: Shared Provider
SAS Studio or CAS Server Monitor
Cloud Analytic Services with servicesBaseUrl
Host
PAM
The direct LDAP provider
SAS Logon Manager and Identities Service
other applications
Dual Authentication: Different Providers
SAS Studio or CAS Server Monitor
Cloud Analytic Services with servicesBaseUrl
Host
PAM
Any provider
SAS Logon Manager and Identities Service
other applications
The direct LDAP provider

Figure A.2 *Authentication from Other Applications*

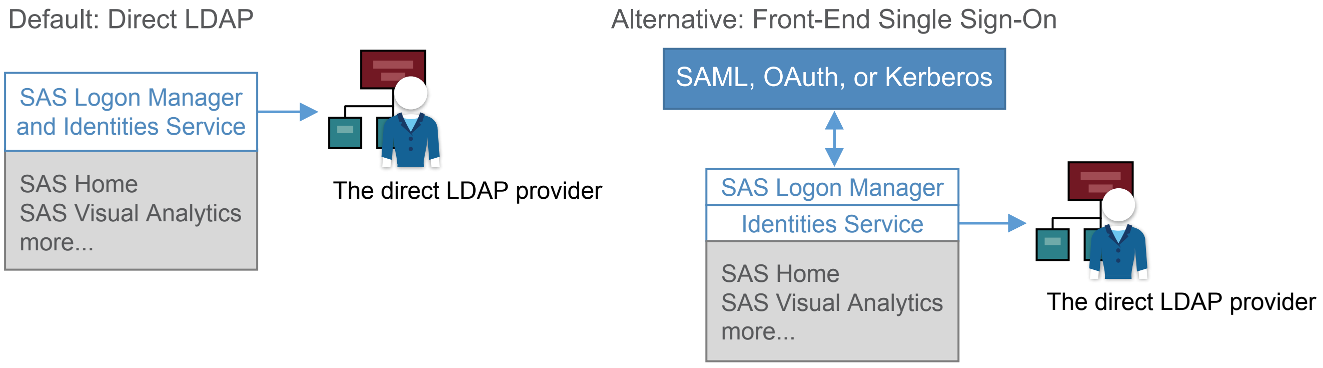

After you sign in, you have seamless access to SAS Viya and, in some contexts, to external data sources.

For more information, see the following documents:

SAS Viya Administration: Authentication

SAS Viya Administration: Identity Management

SAS Viya Administration: External Credentials

Authorization

Authorization is the aspect of security that determines which resources are available to which users. The SAS Viya authorization layer consists of two authorization systems:

- CAS authorization system
- general authorization system

Each system uses a distinct model to protect a distinct class of resources. The general authorization system is not applicable in a programming-only deployment.

Initial and default access are restrictive:

- Any access that is not granted is implicitly disallowed.
- Predefined objects are protected by predefined rules or access controls.
- Only members of special groups or roles have access to privileged administrative functionality.
- Access to objects that users add is managed by inheritance, other influencing rules, and any direct settings.
- Regular users have limited Write access. They can write to their personal folder, the shared Public folder, and the shared Public caslib.

  CAUTION! An exception is that all authenticated users initially have Read and Write access to all registered models. See “Details for Models”.

For more information, see the following documents:

SAS Viya Administration: Orientation to Authorization

SAS Viya Administration: Cloud Analytic Services Authorization

SAS Viya Administration: General Authorization

SAS Viya Administration: Identity Management

Encryption

Encryption is the aspect of security that protects data by converting it into an unintelligible form in transmission or in storage.

For data in motion, SAS Viya is deployed with Transport Layer Security (TLS) to secure network connections. It is fully compliant with SAS security standards.

- In a full deployment of SAS Viya on Linux, almost all external network connections are secured by default. You can harden the full Linux deployment by blocking external connections to port 80, by adding custom certificates on all machines in the deployment, and by upgrading the security protocol and ciphers that are enabled by default. You can also configure TLS–encrypted connections between CAS workers and take additional steps to secure the SAS Embedded Process.
- In a SAS Viya programming-only deployment on Linux, the basic framework for security is included by default, but it is not enabled by default. You can enable TLS and harden the deployment by performing post-deployment tasks.

- In a Windows deployment, the deployment provides a default level of encryption for data in motion. You can harden the deployment by blocking external connections to port 80, by adding custom certificates on Apache HTTPD, and by upgrading the security protocol and ciphers that are enabled by default. You can also upgrade to custom certificates on CAS and SAS/CONNECT.

For data at rest in a new deployment, encryption is not automatically enabled. You can configure encryption of data that is added to PATH, HDFS, and DNFS caslibs.

For more information, see the following documents:

Encryption in SAS Viya: Data in Motion

Encryption in SAS Viya: Data at Rest

Web Security

Web security is the aspect of security that deals with securing against certain types of attacks on web applications and using the security features that are available in modern web browsers.

SAS Viya provides properties that are configured, by default, to protect against the web security risks that are listed below. You can disable or change the properties, based on your environment. For example, you might have to configure Cross-Origin Resource Sharing (CORS) to allow origins in your company's domain. This allows SAS web pages to be included in other web pages inside your company’s network.

For more information about the SAS Viya configuration properties, see the following:

Property	Description	Default Settings
Cross-Origin Resource Sharing on page 127	Technique for relaxing the browser same-origin policy, allowing Javascript on a web page to consume a REST API served from a different origin.	The following cross-origin requests are configured: ■ User credentials are used ■ All HTTP headers are allowed ■ All HTTP methods are allowed ■ Same origins are allowed
Cross-Site Request Forgery (CSRF) on page 127	Prevents attacks that force a user to execute unwanted actions on a web application in which they are currently authenticated.	The following options are configured: ■ All requests that use an authenticated HTTP session, except GET and HEAD requests, must pass a CSRF token specified by the server. ■ Referrers internal to the deployment are allowed
X-Frame-Options on page 126	Avoids clickjacking attacks by making sure that your content is not embedded in other sites.	Same origin

Property	Description	Default Settings
Content-Security-Policy on page 126	Exposes and reduces the risk of data injection and cross-site scripting (XSS) attacks.	default-src 'self'; script-src 'self' 'unsafe-inline' 'unsafe-eval'; img-src 'self' *.sas.com blob: data:; style-src 'self' 'unsafe-inline'; child-src 'self' blob: data: mailto:;
X-Content-Type-Options on page 126	Prevents the browser from interpreting files as something other than what is declared by the content type in the HTTP headers (content sniffing).	nosniff
X-XSS-Protection on page 127	Stops web browser from loading pages when XSS attacks are detected.	1; mode=block

For information about these web attacks, see the following OWASP pages:

- Category:Attack
- OWASP Secure Headers Project
- Cross-Site Request Forgery (CSRF)
- Cross-Origin Resource Sharing

SAS 9 and SAS Viya

Summary

SAS 9 customers continue to benefit from their investment in SAS 9 as they begin to make use of SAS Viya functionality and features. From within familiar SAS 9 interfaces, projects, and code, customers can access the performance enhancements that SAS Viya provides.

- On most hosts, SAS 9.4M5 is tightly integrated with SAS Viya. See SAS 9.4M5 Integration with SAS Viya in *What’s New in Base SAS: Details*. (The exceptions are z/OS and 32–bit Windows.)
- All releases of SAS can use SAS/CONNECT as a bridge to SAS Viya. See the appendix Sharing Data Between SAS 9 and SAS Viya using SAS/CONNECT in *SAS/CONNECT for SAS Viya User’s Guide*.
- SAS Viya visual web applications share a single sign-on and logout with the SAS 9 environment.

Here are some of the methods for accessing SAS 9 data from SAS Viya:

- In SAS Visual Analytics, use self-service import. See SAS Data Explorer: User’s Guide.
- In SAS Environment Manager, interactively load data. See Data Administration: How to (SAS Environment Manager) in *SAS Viya Administration: Data*.
- In SAS Enterprise Guide or SAS Add-In for Microsoft Office (7.13 or later), move data from SAS 9 to CAS. See the topic "Configure Your Environment to Use the Upload to CAS Task" in the SAS Enterprise Guide or SAS Add-In for Microsoft Office chapter in *SAS Intelligence Platform: Desktop Application Administration*.
- In any programming interface, write code to load data. See Programming Interfaces in *An Introduction to SAS Viya Programming*.

- If a more seamless method is not available, use SAS/CONNECT for SAS 9 and SAS Viya to move and share data. See the appendix Sharing Data Between SAS 9 and SAS Viya using SAS/CONNECT in *SAS/CONNECT for SAS Viya User's Guide*.

Note: Not all deployments and releases include all products and support all methods.

Note: Your site must license and install SAS Viya to access SAS Viya functionality. By default, when you order SAS Viya, you receive SAS Visual Analytics. All analytical procedures are separate licenses: SAS Econometrics Procedures, SAS Optimization Procedures, SAS Forecasting Procedures, SAS Visual Data Mining and Machine Learning Procedures, SAS Statistics Procedures, and SAS Viya Procedures.

Considerations: Interacting with SAS 9 Data

Use UTF-8 Encoding

If you access SAS 9 data from SAS Viya, be aware that SAS Viya operates with UTF-8 encoded data. If your SAS 9 data is not UTF-8 encoded, you might need to re-create your data sets. See Migrating to UTF-8 for SAS Viya.

Manage User-Defined Formats

If you access SAS 9 data from SAS Viya, you must make any user-defined formats available to your CAS session. See SAS Cloud Analytic Services: User-Defined Formats.

Considerations: Accessing CAS from SAS 9.4M5

Find CAS

If a SAS 9.4M5 client session cannot find CAS, make information about the host and port of the CAS server available. For example, add the following line to your SAS Application Server sasv9_usermods.cfg or appserver_autoexec_usermods.sas file:

```
CASHOST=("primary-controller-host-name" <"backup-controller-host-name">) CASPORT=port;
```

Here is an example with a CAS backup controller:

```
CASHOST=("mysrv01" "mysrv02") CASPORT=5570;
```

Here is an example without a CAS backup controller:

```
CASHOST=("mysrv01") CASPORT=5570;
```

For more information, see CASHOST= System Option.

Authenticate to CAS

If a SAS 9.4M5 client session cannot authenticate to CAS, create an authinfo file, store CAS credentials in the SAS 9 metadata, or use a different authentication mechanism. See *SAS Viya Administration: Authentication*.

Conform to CAS Encryption Requirements

If a SAS 9.4M5 client session does not meet the encryption standards of the CAS server, make an appropriate certificate available. See "Configure SAS 9.4 Clients to Work with SAS Viya" in *Encryption in SAS Viya: Data in Motion*.

SAS Visual Analytics Administration

SAS Viya Administration documentation is applicable to SAS Visual Analytics. Links to specific SAS Visual Analytics topics that deserve special attention are included here:

- Promoting data and report content
- Granting guess access on page 327
- Managing user-defined formats on page 142
- Loading geographic polygon data as a CAS table on page 161
- Loading data for reports on page 161
- Making data available to CAS
- Using the report data service on page 108
- Using the report packages service on page 111
- Using the report renderer service on page 112
- Understanding identity management concepts on page 349
- Modifying rules that affect access to functionality on page 358

5

Accessibility Features of SAS Viya Administrative Interfaces

Accessibility Features of CAS Server Monitor for SAS Viya 3.4

Overview

CAS Server Monitor for SAS Viya 3.4 has not been tested against U.S. Section 508 standards and W3C Web Content Accessibility Guidelines (WCAG). If you have specific questions about the accessibility of SAS products, send email to accessibility@sas.com or call SAS Technical Support.

Accessibility Features of SAS Environment Manager 3.3

Overview

SAS Environment Manager 3.3 has been tested against the accessibility standards for electronic information technology that were adopted by the U.S. Government under Section 508 of the U.S. Rehabilitation Act of 1973 (2008 draft proposal initiative update). It was also tested against Web Content Accessibility Guidelines (WCAG) 2.0 levels A and AA, part of the Web Accessibility Initiative (WAI) of the Worldwide Web Consortium (W3C). For detailed information about the accessibility of this product, send email to accessibility@sas.com or call SAS Technical Support.

We recommend the following software for a better experience using our products with assistive technologies:

- On Mac OS X, use VoiceOver.
- On Microsoft Windows, use the latest version of JAWS.

Note: If you are using NVDA with Internet Explorer, some software might not load. Use the recommended software and screen reader instead.

Accessibility Settings

To customize accessibility settings, click your name in the application bar and select **Settings**. In the **Settings** menu, navigate to the **Accessibility** section.

Note: Changing these settings does not impact other users.

- Select **Enable sounds** to hear audio indicators for events that occur within the user interface.
- Select **Enable visual effects** to show animations that indicate state changes. For example, if this setting is enabled when you delete an item, a subtle animation is shown when you remove the item.
- Select **Invert application colors** to make the user interface easier to see for users with sensitivity to certain bright colors (for example, a black-on-white display). You can also use the Ctrl+` (Ctrl+back quote) keyboard shortcut to invert the application colors.

  Note: The in-field hint text might not be visible unless you increase the contrast on your browser or computer monitor.
- Select **Display tooltips when using the keyboard to navigate** to enable keyboard-only users to access tooltips. If this option is enabled, focusing on a control also displays the tooltip on the screen. By default, this option is not selected and a mouse is required to see a tooltip.
- Make the focus indicator easier to see by selecting **Customize the focus indicator settings** and adjusting the color, thickness, and opacity. If you cannot see the focus indicator after customizing it, try updating your computer’s accessibility settings to change how the focus indicator is displayed.

  Note: When using Internet Explorer 11 and JAWS to customize the focus indicator settings, focusing on any of the controls causes the entire text of the page to be read. It is recommended that you use another browser.

High Contrast Theme

You can change the appearance of web applications by using the **Theme** setting. The default theme is set by the system administrator, , though you can choose to use a high contrast theme instead. The high contrast theme presents a dark background with high-contrast foreground elements.

Note: Changing the theme does not impact other users.

1 Click your name in the application bar and select **Settings**.

2 Navigate to the **General** section.

3 Select **Choose a theme** and then select **High Contrast**.

Documentation Format

Contact accessibility@sas.com if you need this document in an alternative digital format.

User Interface Layout

The banner contains a left navigation menu, which you can use to navigate around the product. Open the left navigation menu and click the links contained within it to access different parts of the user interface. Use the first link in the left navigation menu, **Dashboard**, to access the SAS Environment Manager dashboard.

Some table columns contain pop-up menus. You can access the menu by moving focus to the column heading and pressing Shift-F10. To access the pop-up menu in iOS, first move focus to the column heading and then double-tap and hold.

Some pages contain a second search field. This search acts as a filter for the items on the page.

Keyboard Shortcuts

In the user interface, some keyboard shortcuts are displayed within parentheses in tooltips and menu item labels.

Table A.1 *Keyboard Shortcuts*

Action	Keyboard Shortcut
Open a window that lists landmarks	Ctrl+F6
Open a pop-up menu. For example, when a column heading has focus, pressing Shift+F10 opens the pop-up menu for that column. **Note:** If you are using a screen reader and Firefox, pop-up menus have some limitations. For example, the screen reader might not announce when the pop-up menu opens. In addition, when the pop-up menu opens, the focus might not be constrained to that pop-up menu.	Shift+F10
Zoom in	Ctrl+ plus (+) sign
Zoom out	Ctrl+ minus (-)sign
Reset zoom	Ctrl+zero (0)
Invert colors	Ctrl+`

Tips for Using Screen Readers

If you do not want your screen reader to announce “clickable” repeatedly throughout the user interface, update your screen reader settings. For example, in JAWS, you can create a new scheme to prevent “clickable” from being announced repeatedly.

1. In JAWS, open the **Settings Center**.
2. Expand **Speech and Sound Schemes**.
3. Select **Modify Schemes**.

4 Select the scheme that you want to modify and click **Edit Selected Scheme**.

5 Navigate to the **HTML** tab.

6 Select the **onclick** HTML attribute and set how you want JAWS to announce the onclick attribute. For example, if you do not want to hear anything when JAWS encounters the onclick attribute, select **Ignore**.

Note: When you save these changes, JAWS might create a new scheme called **Classic (Modified)**. In order to use this new scheme, in the **Settings Center** select **Speech and Sound Schemes** and make sure that the **Active Speech and Sound Scheme** is set to **Classic (Modified)**.

Exceptions to Accessibility Standards

SAS Environment Manager

Table A.2 Known Limitations for SAS Environment Manager

Accessibility Issue	Workaround
The dashboard is not supported for keyboard users.	No workaround is available.
The status of some items on the dashboard is communicated only through color, and JAWS does not announce the color.	No workaround is available.
When items on the dashboard refresh, JAWS continually announces all of the refreshed items.	No workaround is available.
The contents of fields and drop-down menus in the Logs window are not announced by JAWS	No workaround is available.
The New Launcher Context window in the **Contexts** area is not supported for keyboard use.	No workaround is available.
JAWS might stop announcing screen contents, particularly after saving or closing some dialog boxes.	Press F5 to refresh the screen.

Tables

Table A.3 Known Limitations for Tables

Accessibility Issue	Workaround
Tables containing rows that can be expanded or collapsed are not fully supported with screen readers. Limitations include the following: ■ You might not be able to expand or collapse the rows with a keyboard shortcut. ■ Extra rows might be announced after the table has ended. ■ Cell coordinates might not be announced. ■ You might not be able to update a column heading that is meant to be edited.	No workaround is available.
Tables are not fully supported for keyboard users. Columns cannot be resized using only the keyboard.	No workaround is available.
Tables are not fully supported with screen readers. Limitations include the following: ■ Screen readers cannot reliably announce whether a column is sorted. ■ NVDA incorrectly announces that rows are selected when they are not selected.	No workaround is available.

Search and Filter

Table A.4 Known Limitations for Search and Filter

Accessibility Issue	Workaround
When using an advanced filter with a screen reader, there are some limitations, including: ■ Items are not always announced as selected or deselected. ■ The reset link is not announced and you might not be able to navigate to it using a keyboard. ■ You cannot tab to the close button while the reset link is disabled. ■ The focus indicator might not be visible on the reset link.	No workaround available.

Contact Us

Email us at accessibility@sas.com if you need this document in an alternative digital format.

6

Backup and Restore

Overview

Introduction to Backup and Restore

SAS Viya can be deployed with other SAS software at your site to gain insight from analytics and business intelligence. SAS Viya consists of many software components (including SAS Cloud Analytic Services (CAS)), servers and services (including SAS Compute Server, SAS Launcher Server, Identities service, and Audit service). You should have a strong backup and restore strategy and validate your backup and restore processes regularly. This ensures that your backup process is operational and that a restore can be successfully performed. A backup strategy should include tasks for backing up all software components and data in addition to running the Backup service.

You can run the Backup service functionality through the SAS Backup Manager or using a backup command-line interface (CLI). After SAS Viya is deployed, a backup schedule is created for the Backup service. Scheduled backups are automatically performed at specified times. In addition to scheduled backups, you can perform an immediate backup. From these backups, you can successfully restore your system if necessary.

What Is Backed Up?

- Content that users have generated and saved in the SAS Infrastructure Data Server is backed up. Content can be reports (including SAS Visual Analytics reports), comments, authorization rules, attachments, audit records, user preferences, and data source definitions. Content in any additional and required SAS Infrastructure Data Server instances is backed up.

- Infrastructure configuration settings and application configuration settings that you defined based on your business requirements are backed up. These settings are saved in the SAS Configuration Server. Infrastructure configuration settings include IP addresses, hostnames, paths, ports, and certificates. Application configuration settings include user interface (UI) themes (for example, the SAS Visual Analytics UI theme), purge frequency, schedule frequency, and batch run date (for example, batch date of running SAS Visual Analytics reports).

- All SAS Message Broker exchanges, queues, bindings, users, virtual hosts, permissions, and parameters are backed up. Actual messages are not backed up.

- If CAS is configured for data management and analytics, CAS access controls and caslibs are backed up. Caslibs help with accessing data from a data source. A caslib consists of a data source definition and can include access controls. Access controls help with controlling user access to the data.

In a clustered environment, caslibs and access controls of the CAS primary node are backed up.

All of this information might reside on a single host or on different hosts.

What Is Not Backed Up?

Backup and Restore does not support the backup and restore of several items. You must take alternative steps to back up and restore these items.

- the operating environment

  In the operating environment, the operating system libraries and packages, environment variables, kernel settings, user identifiers (UIDs), group identifiers (GIDs), mount points, symbolic links, the Windows registry, Windows security and local security policy settings, Windows users and groups, file systems, and so on, are not backed up.

- SAS Viya deployment

  In the SAS Viya deployment, the deployment files that are included as part of the SAS Viya deployment (such as scripts, executable files, binaries, and so on), the SAS installation directory, and the SAS configuration directory are not backed up.

- User home directories
- Third-party applications that are used with SAS Viya
- Data sources, such as:
  - Files for predefined caslibs in `/opt/sas/viya/config/data/cas/default/`
  - Third-party databases
  - Data that is stored outside of the SAS Infrastructure Data Server
  - SAS Infrastructure Data Server metadata such as database user information, roles, and permissions that are stored in the SAS Infrastructure Data Server

    If a binary backup is run, SAS Infrastructure Data Server metadata is backed up.

  - Data that is loaded to CAS
  - Data that is stored on local file systems (such as SAS data sets)

    For example, suppose that the SAS deployment contains Model Studio. Model Studio uses a SAS data set stored on the local file system and you back up the SAS deployment. To ensure that Model Studio has access to the SAS data set, you must manually back up the SAS data set that is stored on the local file system and restore it in an appropriate location on the target environment.

  - Data that is stored on remote file systems (such as path-based data sources including PATH, DNFS, and HDFS)
  - Database data sources (such as Oracle, ODBC, and Hive)
- In a clustered CAS environment, CAS access controls and caslib information for the secondary CAS controller

  You do not need to manually back up the CAS access controls and caslib information for the secondary CAS controller because the `permstore` directory for the secondary controller (that contains the CAS access controls and caslib information) is always in sync with the `permstore` directory of the primary controller.

- SAS Message Broker messages

  Messages are not backed up.

Concepts

Default File Locations

After SAS Viya is installed, installation files are stored in a default location referred to as *SASHome*. For example, on Windows, SASHome is `C:\Program Files\SAS\Viya`.

Configuration files are stored in a default location referred to as *SAS-configuration-directory*. For example, on Windows systems, the SAS configuration directory is `C:\ProgramData\SAS\Viya`.

The following table lists the default locations of installation and configuration files.

Table A.1 Default Locations

Location	Example Path on Windows System	Example Path on Linux System
SAS-configuration-directory	`C:\ProgramData\SAS\Viya`	`/opt/sas/viya/config`
SASHome	`C:\Program Files\SAS\Viya`	`/opt/sas/viya/home`

Backup and Restore Terms

Common Terms

Here is a list of backup and restore terms that are used in a standard deployment on Windows and on Linux:

backup
the backup that can be initiated by a SAS administrator. SAS Viya can be deployed on Linux and on Windows. There are two types of backup: default and binary. For more information, see Table 6.6 on page 39.

restore
the process of restoring a backup. A SAS administrator can initiate the restore from a backup that is either marked ✓ or ✓⚠.

local vault
a local file system path located on the same host where the backup source resides. It is the location where the backup files for the data source are created. They are then moved to the shared vault. The location of the local vault is ***`SAS-configuration-directory`***`/backup` on Linux and Windows by default and cannot be changed.

It is a best practice that if a tenant administrator performs a local backup, the tenant administrator should copy the contents (local files) of the local vault to a shared vault.

pre-restore validations
validations that are done before performing a restore using a given backup. A pre-restore validation includes the following validation checks:

- Does the provided backup exist?
- Is the backup completed?
- Is the backup purged?

- If SAS Infrastructure Data Server is being restored using the default type of backup, does the list of databases in the backup match the list of databases currently present in SAS Infrastructure Data Server? You can determine the backup contents by submitting the following command:

  ```
  sas-admin backup show -i=<backup_ID>
  ```

- If you are restoring a binary type of backup, was the include-all-sources property set to TRUE?

  The *include-all-sources* property has different names depending on where you are. In the REST API, it is includeAllSourcesForBinaryBackup. For the command-line interface, it is include-all-sources-for-binary-backup. In the user interface, it is the check box labeled **Include remaining sources**. It is referred to as the include-all-sources property. For more information, see "Restore SAS Infrastructure Data Server from a Binary Backup Manually" on page 61.

- In a multi-tenant environment where the tenant list is not provided, the Backup service checks to see whether the onboarded tenants and the tenants in the backup match. If the tenant list is provided in the restore request, then the Backup service checks to see whether all of the tenants specified in the restore request are onboarded.

retention period
number of days that backups are stored before they are removed from the shared vault.

shared vault
any network location that preserves the backups from all tiers. Backup files are moved from the local vault to the shared vault. The shared vault is set by the sharedVault property in the Backup service configuration.

The shared vault directory must be accessible from all hosts in the deployment. The locations of the shared vault and local vault must be different.

On Linux, ensure that the sas user has Read, Write, and Execute permissions to the shared vault directory. On Windows, ensure that the CAS user has Read, Write, and Modify permissions to the shared vault directory.

It is a best practice to back up your shared vault.

alternate shared vault
any alternate network location that contains backups. An alternate shared vault location might be used when performing a restore to an alternate host. The permissions for an alternate shared vault are the same as the shared vault. For more information, see "shared vault" on page 37.

standard deployment
a SAS Viya deployment intended to be used only by a single tenant on Windows and on Linux.

slug
user-provided name for the backup or restore operation.

state
state of a backup or restore operation. The possible values are:

- *pending* indicates that a backup or restore job has been created, but the operation has not yet started.
- *running* indicates that a backup or restore is in progress.
- *completedWithWarning* indicates that at least one tenant backup or restore has failed, the CAS controller host is not reachable, or the recovery of SAS Configuration Server is partially successful.
- *completed* indicates that the backup or restore operation has completed successfully.
- *failed* indicates that a backup or restore for one or more data sources failed.
- *unknown* indicates that either a backup agent is not installed on the source or the backup agent is not running. If the source does not have a backup agent, contact your administrator. If the source has a backup agent, then restart the Backup service.
- *canceling* indicates that the backup is being canceled. This state is applicable only to a backup.
- *canceled* indicates that the backup is canceled. This state is applicable only to a backup.

trigger
an event generated periodically by a scheduler that signals when a new instance of a job should be executed.

Terms for Multi-Tenant Deployment

Here is a list of backup and restore terms that are used in a multi-tenant deployment on Linux:

deployment backup
a multi-tenant deployment on Linux in which only a provider administrator can initiate backup for multiple tenants. If an explicit list of tenants is not provided, all onboarded tenants are backed up.

deployment restore
a multi-tenant deployment on Linux in which only a provider administrator can initiate a restore of a deployment backup (with at least one successful tenant backup) for multiple tenants. If an explicit list of tenants is not provided, all onboarded tenants and the provider tenant are restored. If all the tenants in the list are in the backup, then the restore is triggered for each of the tenants in the list. If any of the tenants in the list are not in the backup, then the restore does not proceed and it is marked as failed.

tenant administrator
a person within the tenant organization who has administrative privileges for a tenant environment. For example, assigning users to custom groups and managing access to SAS Viya content and CAS data are tenant administrator tasks. The tenant administrator must be a member of the SAS Administrators group to perform the backup and restore functions that are specific to the tenant.

tenant environment
a collection of software and the infrastructure for use by a single tenant. Customers might have separate tenant environments to support development, staging, and production instances of SAS services.

multi-tenant deployment
a SAS Viya deployment in which multiple tenants can access the same environment in isolation without impacting the data or processes of other tenants. A multi-tenant deployment has the provider tenant by default.

provider administrator
a person within the provider organization who has administrative privileges for a provider environment. The provider administrator must be a member of the SAS Administrators group to perform the backup and restore functions for all the tenants in a multi-tenant deployment.

provider environment
a collection of software and the infrastructure to support one or more tenant environments.

provider tenant
the initial tenant (that is, tenant zero) created when a multi-tenant system is deployed. This tenant has full access to all applications in the deployment, but is intended for provider administrator access only. Users in this tenant have access to information about the entire deployment, including other tenants.

tenant
one of the customers using a shared SAS Viya deployment.

In a multi-tenant deployment, a tenant is said to be onboarded when the SAS Viya infrastructure for that tenant is created. This includes the LDAP groups, the LDAP identities, the SAS Infrastructure Data Server databases, the schemas, and the CAS instance.

alternate shared vault
any alternate network location that contains backups. In a multi-tenant deployment on Linux, the alternate shared vault contains the backups of all tenants. You might want to use the alternate shared vault to migrate a few tenants or to restore the environment. The permissions for an alternate shared vault are the same as the shared vault.

For more information, see “shared vault” on page 37.

Roles and Permitted Tasks

The following administrators can perform backups:

SAS administrator
can perform a backup in a standard deployment.

Provider administrator
can perform a backup of all onboarded tenants in a multi-tenant deployment.

Tenant administrator
can perform a backup of that tenant in a multi-tenant deployment.

Important: In a multi-tenant deployment, a tenant administrator cannot perform a binary type of backup.

The following table describes the roles and permitted tasks for each role.

Table A.2 Roles and Permitted Tasks

	Standard Deployment on Windows and in Linux	Multi-tenant Deployment	
Task	SAS Administrator	Provider Administrator	Tenant Administrator
View backup configuration	✓	✓	✓
Edit backup configuration	✓	✓	–
View and edit a default backup schedule	✓	✓	–
View and edit a binary backup schedule	✓	✓	–
Perform an immediate default backup	✓	✓	✓
Perform an immediate binary backup	✓	✓	–
Create backup job and schedule it	–	✓	–
Perform a restore	✓	✓	✓
Perform a restore from an alternate host	✓	✓	–

About Default Backup and Binary Backup

Default backup
A default backup can be used in standard and multi-tenant deployments. During a default backup, you can select tenants from a list of tenants in a multi-tenant deployment.

For detailed information about a default backup, see Table 6.7 on page 40.

Binary backup

A binary backup must be used when you want to restore, but the SAS Infrastructure Data Server is not restarting or is unresponsive.

For detailed information about a binary backup, see Table 6.7 on page 40.

TIP Whenever you deploy SAS Viya or make changes to an existing SAS Viya deployment, always perform a default backup and a binary backup of the deployment immediately.

The following table describes the default backup and binary backup in detail.

Table A.3 Detailed Information About Default Backup and Binary Backup

Category	Default Backup	Binary Backup
General information	The default type of backup.	A different type of backup that is not a replacement of the existing default backup.
	Uses the pg_dump utility.	Uses the pg_basebackup utility.
Backup-related information	Backs up the following source types: ▪ content of SAS Infrastructure Data Server ▪ SAS Configuration Server ▪ SAS Message Broker ▪ CAS The default backup does not back up the metadata of the SAS Infrastructure Data Server.	Backs up the content and metadata of the SAS Infrastructure Data Server (such as database users, roles, and permissions). If the include-all-sources property is set to TRUE while initiating the backup, the binary backup backs up remaining source types (such as SAS Configuration Server, SAS Message Broker, and CAS). For more information about the include-all-sources property, see "pre-restore validations" on page 36.
	A logical backup of each database in the SAS Infrastructure Data Server. When a default backup is restored, the Backup service re-creates each database in the SAS Infrastructure Data Server by retaining its state at the time of backup.	A binary copy of the database files as a TAR archive.
	In a multi-tenant deployment, you can select the tenants from a list of tenants to perform the default backup.	In a multi-tenant deployment, you can perform a binary backup of the entire deployment that includes all tenants.

Category	Default Backup	Binary Backup
Restore-related information	Use default backup to perform a logical backup of each database in the SAS Infrastructure Data Server.	Use binary backup to perform a binary copy of the database files in a SAS Infrastructure Data Server or in any additional required SAS Infrastructure Data Server instances. The binary copy is a TAR archive.
	A default backup can be restored using the SAS restore operation. The restore operation includes SAS Infrastructure Data Server content, configuration data, and other source types.	The SAS Infrastructure Data Server portion of a binary backup must be restored manually. Suppose that the include-all-sources property is set to TRUE and the manual restore of the SAS Infrastructure Data Server is completed. In such a case, you can then use the restore operation to restore the other portions of the binary backup. For more information about the include-all-sources property, see "pre-restore validations" on page 36. **Note:** While restoring a binary backup, the SAS Infrastructure Data Server must be manually restored before all other source types. Otherwise, restoring the other source types might fail due to an unresponsive SAS Infrastructure Data Server.
	Use default backup for restoring to an alternate host or to the same host. You can use SAS Backup Manager or a command-line interface (CLI) to restore a default backup.	Use binary backup to perform a restore on the same host. You cannot use binary backup to restore to an alternate host.
	Do not use default backup if the SAS Infrastructure Data Server is unresponsive.	Always use binary backup when you need to perform a backup and restore and the SAS Infrastructure Data Server does not start or is unresponsive.
	When you use default backup to restore in a multi-tenant deployment, only the tenants that were included in the default backup are restored.	When you use binary backup to restore in a multi-tenant deployment, the entire deployment (that includes all tenants) is restored.

Tenancy and Backup and Restore

The following table shows the tenancy and deployment type for Linux and Windows.

Operating System	Tenancy	Deployment Type	Administrator
Linux	SAS Viya can be deployed for a single tenant and for multiple tenants.	Standard deployment for single tenant.	Members of the SAS Administrators group and individual users can perform backup and restore operations.
		Multi-tenant deployment for multiple tenants.	A provider administrator can perform backup and restore operations for all tenants. A tenant administrator can perform backup and restore operations for that tenant. The provider administrator and the tenant administrator must be members of the SAS Administrators group.
Windows	SAS Viya can be deployed for a single tenant only.	Standard deployment for single tenant.	Members of the SAS Administrators group and individual users can perform backup and restore operations.

Note: The Backup service does not take the place of operating system backups or file system backups. Furthermore, you cannot use the Backup service when SAS Viya is deployed on a container-enabled infrastructure.

The Backup Directory Structure

The following diagram explains the directory structure of the shared vault and how backups are stored within this structure.

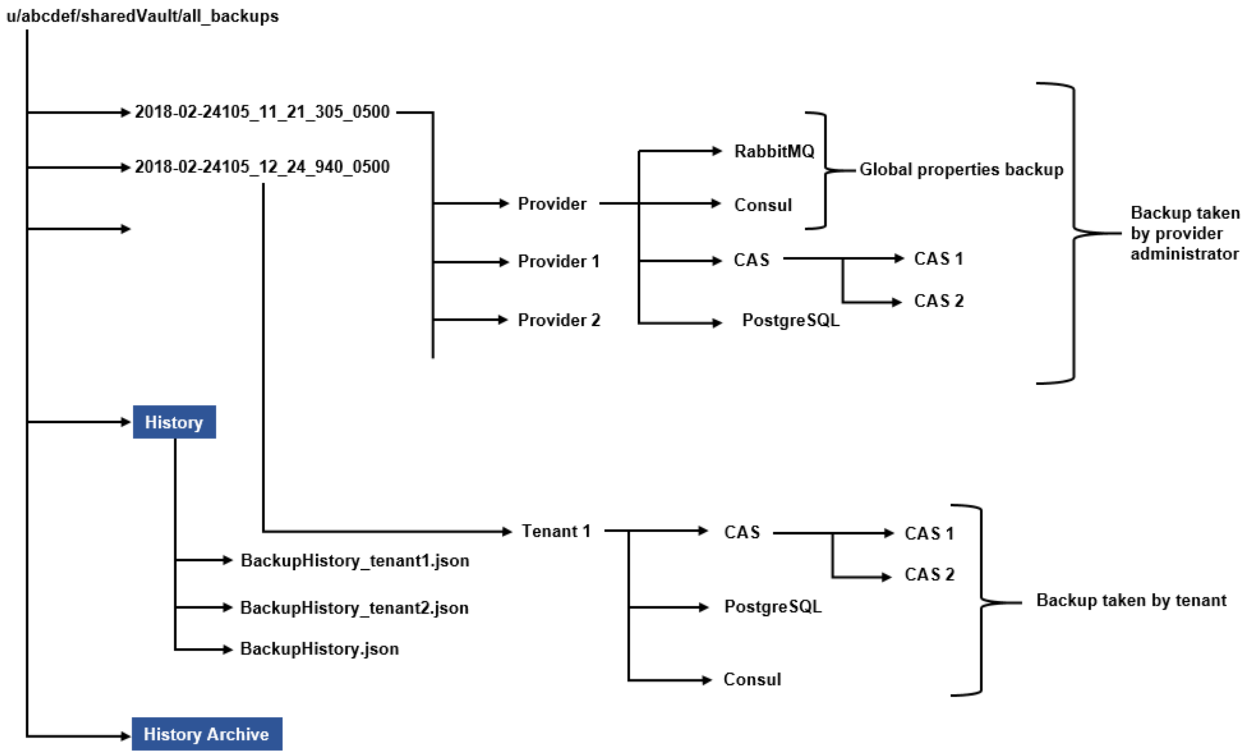

In the diagram, the path `/u/abcdef/sharedVault/all_backups` is the path of the shared vault. This directory contains the folders for backups. Folders are named using the date and time at which the backup was performed. Each backup folder contains folders for the tenants included in the backup. In a standard deployment, only one folder named __default__ exists. In a multi-tenant deployment, a *tenantID* folder is available. Within each tenant folder, there are folders for each source type to which the tenant has access. Within each source type folder, you can find the backup files for that source type.

Within the shared vault, there are folders named History and HistoryArchive. The History folder stores the history files. This includes the global history file (backuphistory.json) and tenant history files (backuphistory_<*tenantId*>.json). The HistoryArchive folder contains the backup of the History folder after each successful backup or restore operation.

Purge Backups

The backups are retained for a period that is set by an administrator. The default value for the retentionPeriod property is 30 days. The retentionPeriod property can be modified by selecting **Backup service** on the Configuration page in SAS Environment Manager. Click **New/Edit Configuration**, and then select **sas.deploymentbackup**.

The last successful backup of each type (regardless of the value for retentionPeriod) is retained. If a binary or default backup is performed for a deployment, the last successful backup for the default type is retained and the last successful backup for the binary type is retained. In a multi-tenant deployment, if a backup was explicitly performed by a tenant after a successful backup (default or binary), that backup is also retained.

Old backups are purged after the retention period. They are deleted from the file system. A previous backup for an onboarded tenant is purged in the next purge cycle after the retention period has passed. In a multi-tenant deployment, the backup of all onboarded tenants is purged after the retention period. The last successful backup of an offboarded tenant is not purged after the retention period. It permanently remains unless it is deleted manually.

SAS Infrastructure Data Server, High Availability, and Backup and Restore

By default, SAS Viya does not configure SAS Infrastructure Data Server for high availability. Before you configure it, you should perform a binary backup of the data and rely on that backup for a restore. SAS recommends that you perform a binary backup if the configuration or installation of the SAS Viya environment changes in any way. For information about restoring the SAS Infrastructure Data Server when the server is unresponsive due to SAS Infrastructure Data Server corruption or for any other reason, see “Restore SAS Infrastructure Data Server from a Binary Backup Manually” on page 61.

Suppose that the SAS Infrastructure Data Server is configured for high availability and the primary node is unresponsive. In this case, a standby node is promoted to the primary node, and the SAS Infrastructure Data Server can be started.

User Interface of SAS Backup Manager

View SAS Backup Manager

1. Log on to SAS Environment Manager as a provider administrator or a tenant administrator in a multi-tenant deployment. Or, log on as a SAS administrator in a standard deployment.
2. In the left pane, click **Backup and Restore**.

Understand SAS Backup Manager

You can perform the following common tasks in the SAS Backup Manager:

- View the history of backups and restores.
- View the backup configuration.

If you log on as a provider administrator in a multi-tenant deployment, you can perform the following tasks in addition to the common tasks:

- View the details of each backup and restore, such as the list of tenants involved in a backup or restore, start and end times, and status of each backup or restore.
- Perform an immediate (ad hoc) backup for all tenants or for selected tenants.
- View the details of a provider-specific backup and restore.
- Perform a restore.

If you log on as a SAS administrator in a standard deployment or as a tenant administrator in a multi-tenant deployment, you can perform the following tasks in addition to the common tasks:

- View the details of each backup and restore, such as the list of data sources in a backup or restore and details of each data source.
- Perform an immediate (ad hoc) backup.
- Perform a restore.

View History of Backup and Restore

To view the history of backups or restores:

- From **View**, select **Backup details** or **Restore details**.

By default, backups and restores are listed in descending order by **Local Start Time**. The list of backups and restores provides you with the following information:

- all backups or restores
- backups that have been purged due to retention period
- backups or restores currently running or that are waiting to run

Click ↻ to refresh the details. The following table describes the information that you can view for each backup and restore.

Table A.4 Description of Columns in Backup and Restore Details

Column Name	Description
Comments	Comments that were entered before performing a backup or restore.
Backup ID or **Restore ID**	Unique identifier of the backup or restore based on the date and time that it was started (for example, 2017-10-28T05_33_47_326-0400).
User ID	User ID of the user that ran the backup or restore or the identity name of the service that initiated the backup or restore.
Type	Type of backup: binary or default.

Column Name	Description
Size	Total size of the files that were backed up. (This value is not available for a provider in a multi-tenant environment.) This column is not displayed when you select **Restore details**.
Local Start Time	Date and time that the backup or restore started running.
Local End Time	Date and time that the backup or restore stopped running.
Status	Status of the backup or restore operation. One of the following icons is displayed: ◒ The backup or restore is pending; it has not started. ⟳ The backup or restore is running. ✓ The backup or restore completed without errors or warnings. ✓⚠ The backup or restore completed with warnings. ⊗ The backup or restore completed with errors. Click ⊗ to understand the reason for the failure. 🛇 The backup is purged. ● The status of the backup or restore cannot be determined. ⊘ The backup is canceled.

View Details of a Backup or Restore as a SAS Administrator or as a Tenant Administrator

View Details of a Backup or Restore

1 Log on to SAS Environment Manager as a tenant administrator in a multi-tenant deployment. Or, log on as a SAS administrator in a standard deployment.

2 From **View**, select **Backup details** to view a list of all backups or select **Restore details** to view a list of all restores.

3 Select a backup or restore, and click ☷ in the right pane.

A pane displays the following information for the selection:

- Backup ID or Restore ID.
- Status of the backup or restore.
- Total size of the files that were backed up. This information is not available for restores.
- Any comments that were entered before performing a backup or restore.
- User ID of the user that ran the backup or restore or the identity name of the service that initiated the backup or restore.
- Local start date and time for the backup or restore.
- Local end date and time for the backup or restore.
- Information about the remaining sources if they were included in the binary backup.

View Source Types of a Backup or Restore

1 Log on to SAS Environment Manager as a tenant administrator in a multi-tenant deployment. Or, log on as a SAS administrator in a standard deployment.

2 From **View**, select **Backup details** to view a list of all backups or select **Restore details** to view a list of all restores.

3 Select a backup or restore, and click 🗄 in the right pane.

The data sources for the selection are listed in the right pane. If you are viewing details for a restore, only the data sources that were restored are listed.

By default, the data sources include the following:

- SAS Message Broker (not available for a tenant in a multi-tenant deployment)
- SAS Configuration Server
- SAS Cloud Analytic Services
- SAS Infrastructure Data Server and the names of any additional SAS Infrastructure Data Server instances

4 (Optional) Click ▸ to the left of the data source to view the following information. The icon to the right of the data source displays the status of the backup or restore.

- Name of the host where the data source is deployed.
- Status of the data source’s backup or restore.
- Total size of the files that were backed up for this data source. This information is not available for restores.

5 (Optional) To collapse the data source, click ▾.

View Details of a Backup or Restore for a Tenant as a Provider Administrator

1 Log on to SAS Environment Manager as a provider administrator in a multi-tenant deployment.

2 From **View**, select **Backup details** to view a list of all backups or select **Restore details** to view a list of all restores.

3 Double-click a backup or restore. Alternatively, right-click a backup or restore, and select **Tenants**.

The following information is displayed for the selection:

- Tenant name.
- Local start date and time for the backup or restore for the tenant.
- Local end date and time for the backup or restore for the tenant.
- Status of the backup or restore for the tenant.

4 (Optional) If the provider environment is selected for backup or restore, click ▤ to view the details of the provider environment.

If the provider tenant is included in the backup or restore, the provider administrator can view its details in the Operation Details (Provider) pane.

5 (Optional) If the provider environment is selected for backup or restore, click 🗄 to view the data sources and the details of each data source that is a part of the provider environment.

If the provider tenant is included in the backup or restore, the provider administrator can view the data sources of the provider tenant in the Data Sources (Provider) pane.

You can click ☰ to view a list of all backups or restores.

Using the Command-Line Interface (CLI)

Overview

Use a command-line interface (CLI) to perform a backup or restore without using SAS Backup Manager in SAS Viya. You can enter commands on a command line and receive responses from the system. For more information, see "Command-Line Interface: Overview" on page 678.

Examples of CLI Commands for Backup

The following examples assume that you have already logged on to SAS Viya. For more information, see "Command-Line Interface: Preliminary Instructions" on page 681.

Table A.5 Examples of CLI Commands for Backup

Task	Example
List the first 50 backups	`sas-admin backup list --limit 50`
Show details of a backup with backup ID 2017-10-28T07_23_44_594-0400	`sas-admin backup show -i=2017-10-28T07_23_44_594-0400`
Obtain help for LIST command for backup	`sas-admin backup list -help`
Perform the backup	`sas-admin backup start`
Perform a binary backup (back up only SAS Infrastructure Data Server)	`sas-admin backup start -t=binary`
Perform a binary backup (including backup of SAS Message Broker, SAS Configuration Server, and SAS Cloud Analytic Services)	`sas-admin backup start -t=binary -i=true`
Cancel the backup	`sas-admin backup cancel -i=2017-10-28T07_23_44_594-0400`

Note: Commands to perform a binary backup are not available to a tenant administrator in a multi-tenant SAS Viya environment. Commands are available only to a provider administrator in a multi-tenant SAS Viya environment and to an administrator in a non-multi-tenant SAS Viya environment.

In a multi-tenant SAS Viya environment, the provider administrator has more command options.

Examples of CLI Commands for Backup (Provider Administrator Only)

The following examples assume that you have already logged on to SAS Viya. For more information, see "Command-Line Interface: Preliminary Instructions" on page 681.

Table A.6 Examples of CLI Commands for Backup (Provider Administrator Only)

Task	Example
Perform backup only for provider tenant	`sas-admin backup start -p`
List provider backups	`sas-admin backup list -p`
Show details of a backup with backup ID 2017-10-28T07_23_44_594-0400 for provider only	`sas-admin backup show -i=2017-10-28T07_23_44_594-0400 -p`
Perform backup for acme and cyberdyne tenants	`sas-admin backup start --tenants=acme,cyberdyne`

Examples of CLI Commands for Restore

The following examples assume that you have already logged on to SAS Viya. For more information, see "Command-Line Interface: Preliminary Instructions" on page 681.

Table A.7 Examples of CLI Commands for Restore

Task	Example
Show history of restores	`sas-admin restore list`
Show details of a restore with restore ID 2017-10-28T07_23_44_594-0400	`sas-admin restore show -i=2017-10-28T07_23_44_594-0400`
Obtain help for LIST command for restore	`sas-admin restore list -help`
Start a restore of a specified backup, and specify that the name of the restore is restoreA	`sas-admin restore start --<backup-name> <backup-ID> --slug restoreA`
Start a restore of a specified backup in an alternate shared vault at **`/alternate/sharedvault`**	`sas-admin restore start --backup-name=<backup_ID> --alternate-shared-vault-for-restore=/alternate/sharedvault`

Note: Commands to perform a binary backup are not available to a tenant administrator in a multi-tenant SAS Viya environment. Commands are available only to a provider administrator in a multi-tenant SAS Viya environment and to an administrator in a non-multi-tenant SAS Viya environment.

In a multi-tenant SAS Viya environment, the provider administrator has more command options.

Examples of CLI Commands for Restore (Provider Only)

The following examples assume that you have already logged on to SAS Viya. For more information, see “Command-Line Interface: Preliminary Instructions” on page 681

Table A.8 Examples of CLI Commands for Restore (Provider Only)

Task	Example
Show the history of restores for provider	`sas-admin restore list -p`

Show details of a provider restore with restore ID 2017-10-28T07_23_44_594-0400	`sas-admin restore show`
	`-i=2017-10-28T07_23_44_594-0400 -p`

Start a restore of a specified backup in an alternate shared vault at **`/alternate/sharedvault`** for acme and cyberdyne tenants	`sas-admin restore start`
	`--backup-name=<backup_ID>`
	`--alternate-shared-vault-for-restore=/`
	`alternate/sharedvault`
	`--tenants=acme,cyberdyne`

Initial Tasks

Immediately after SAS Viya is deployed, perform the following tasks:

1. Configure the backup by specifying the shared vault location.

   Check the default value of the retention period and change it if necessary.

   For more information, see “Create and Edit a Backup Configuration” on page 51.

2. In SAS Environment Manager, click ⚙ **Jobs**, and then click the **Scheduling** tab to ensure that the DEFAULT_BACKUP_SCHEDULE and BINARY_BACKUP_SCHEDULE jobs are created.

   These jobs are created automatically after SAS Viya is deployed and the Backup service starts.

   The DEFAULT_BACKUP_SCHEDULE job is set to run every Sunday at 1:00 a.m. The BINARY_BACKUP_SCHEDULE is set to run on the first Saturday of every month at 5:00 a.m.

   For more information about managing backup schedules, see “Manage Backup Schedules” on page 52.

   For more information about jobs, see Jobs on page 169.

3. (Optional) Consider your business requirements and determine the frequency of the scheduled backups. Edit the backup schedule if necessary.

   For more information about editing the backup schedule, see “Edit a Backup Schedule” on page 53.

4. If SAS Viya is deployed on Linux, check the groups for users who are running SAS Viya services and CAS controllers.

   Backup and restore uses the sas user to back up and restore data sources from each machine that is included in the backup. It creates a local vault on each machine where data sources are available such as the SAS Configuration Server, SAS Infrastructure Data Server, SAS Message Broker, and CAS.

   If a CAS controller is running with a different user (a user who is not part of the sas group), perform the following steps to set appropriate permissions on the local vault for a successful backup and restore. Perform these steps for each user running a CAS controller.

   1. Navigate to the *SAS-configuration-directory*.
   2. Set appropriate access control on the local vault directory for the user running a CAS controller. For example, if a cas user is running a CAS controller, submit the following command:
      ```
      setfacl -Rdm "u:cas:wx" backup/
      ```
   3. Set appropriate access control on the local vault directory for the sas user so that the sas user can access directories created by the cas user. For example, submit the following command:
      ```
      setfacl -Rdm "u:sas:rwx" backup/
      ```

   The setfacl permissions need to be set only on machines where CAS controllers (primary and secondary CAS controller nodes) are available.

   Note: Here are some important details to remember. During SAS Viya installation, the sas user is created with the sas group automatically. Access control lists (ACLs) should be enabled on the UNIX operating system. Whenever a new tenant is added or onboarded, set the setfacl permissions on the local vault directory on the machine where the CAS controller is available for that tenant.

5. Run the DEFAULT_BACKUP_SCHEDULE job and the BINARY_BACKUP_SCHEDULE to ensure that the backups are performed successfully.

   To run the backup schedule, perform the following steps:

   a. In SAS Environment Manager, click ⚙ **Jobs**.

   b. On the **Scheduling** tab of the Jobs page, perform the following tasks:

      - Right-click DEFAULT_BACKUP_SCHEDULE, and select **Run** to immediately run the backup.
      - Right-click BINARY_BACKUP_SCHEDULE, and select **Run** to immediately run the backup.

   c. On the **Monitoring** tab of the Jobs page, ensure that the jobs are running without any warnings and errors.

   If you do not perform these tasks, the default backup runs every Sunday at 1:00 a.m. The binary backup runs on the first Saturday of every month at 5:00 a.m.

Managing the Backup Configuration and Schedules

Manage the Backup Configuration

About the Backup Configuration

The Backup service has its own configuration. You must create the configuration initially and set the following properties:

- retentionPeriod—the number of days that backups are stored before they are deleted from the shared vault. Backups cannot be recovered after they are deleted.
- sharedVault—a shared network location for backups. For more information about a shared vault and its access rights, see “shared vault” on page 37.

View a Backup Configuration

1 In SAS Environment Manager, click **Backup and Restore**.

2 On the Backup and Restore page, click **Backup Configuration**.

Create and Edit a Backup Configuration

Create and Edit a Backup Configuration Using SAS Environment Manager

1 Log on to SAS Environment Manager as a provider administrator in a multi-tenant deployment. Or, log on as a SAS administrator in a standard deployment.

2 Click **Configuration** in the left pane.

3 Select **View** ⇨ **Basic services**.

4 Click **Backup service**.

The ✓ and ⊖ icons next to the service indicate whether information is entered in the properties. The ✓ icon indicates that the backup configuration is created. You can view it and edit it if necessary. The ⊖ icon indicates that the backup configuration is not created. You must create it.

5 If the backup configuration is not created, scroll down until you see **sas.deploymentbackup**, and click ✱ to the right of the service.

If the backup configuration is already created and you want to edit it, scroll down until you see **sas.deploymentbackup**, and click ✎ to the right of the service.

6 Enter information for the sharedVault property and other configuration properties as necessary.

For more information about a shared vault and its access rights, see “shared vault” on page 37.

Consider the guidelines and best practices for configuring backups. For more information, see “Guidelines and Best Practices for Configuring Backups” on page 52.For more information about backups, see “Backup Service” on page 94.

7 Click **Save**.

Create a Backup Configuration Using the sitedefault.yml File

You can specify backup configuration information in the sitedefault.yml file before deploying SAS Viya. For more information about the sitedefault.yml file, see “Operations” on page 82.

1 Open the sas_viya_playbook/roles/consul/files/sitedefault.yml file. Copy the following code below the config property:

```
deploymentBackup:
   sas.deploymentbackup:
          sharedVault: /opt/sas/viya/config/SharedVault
          jobTimeout: 600
          retentionPeriod: 30
          custom:
             restore.filter.sas.configuration.config.sas.deploymentbackup: "*"
          scheduledBackupAllowed: false
```

Note: Copying code can lead to extraneous characters being included in your sitedefault.yml file. Review the sitedefault.yml file carefully.

2 In the preceding code, change the values of `sharedVault`, `jobTimeout`, and `retentionPeriod` if necessary.

You must specify a value for the sharedVault property. For more information about a shared vault and its access rights, see "shared vault" on page 37.

3 Save the sitedefault.yml file, and then proceed with the deployment of SAS Viya.

Note: You cannot edit the backup configuration using the sitedefault.yml file.

Guidelines and Best Practices for Configuring Backups

- Always ensure the values of the sharedVault property and retentionPeriod property are set immediately after a new installation, an upgrade, or any modifications to your SAS deployment.

  For more information about a shared vault and its access rights, see "shared vault" on page 37.
- The shared vault location must be different from the local vault location. The local vault location is ***`SAS-configuration-directory`*****`/backup`** on Linux and ***`SAS-configuration-directory`*****`\backup`** on Windows. The local vault location is located on machines that have SAS Configuration Server, CAS (controller node), SAS Infrastructure Data Server, or SAS Message Broker.
- Set the retentionPeriod property so that you always have at least the last four backups available at any point in time. For example, if you are doing daily backups, the retention period must be four days. If you are doing weekly backups, the recommended retention period is 30 days.
- Always initiate a binary backup after a tenant is onboarded or offboarded.

Manage Backup Schedules

About Backup Schedules

After SAS Viya is deployed, the following backup schedules are created automatically whenever the backup runs for the first time:

- default backup schedule
- binary backup schedule

Default Backup Schedule

After the default backup schedule is created, the following message appears in the Backup and Restore log located at ***`SAS-configuration-directory`*****`\var\log\deploymentBackup\default`**: `Default schedule created for BackupService to run backup job every Sunday 1AM`. You can also view the backup schedule in SAS Environment Manager.

Binary Backup Schedule

After the binary backup schedule is created, the following message appears in the Backup and Restore log: `Binary schedule created for BackupService to run backup job on the first Saturday of every month at 5AM`. You can also view the backup schedule in SAS Environment Manager.

The following services must be running to run the backup schedules:

Table A.9 Names of Services on Linux and Windows

Name on Linux	Name on Windows
`sas-viya-identities-default`	SAS Identities service

Name on Linux	Name on Windows
`sas-viya-scheduler-default`	SAS Scheduling service
`sas-viya-jobdefinitions-default`	SAS Job Definition service
`sas-viya-jobexecution-default`	SAS Job Execution service
`sas-viya-restexecutionprovider-default`	SAS REST Execution Provider service

If one of the services is not running when the Backup service starts, then the Backup service retries every 5 minutes 25 times to schedule the backup. If after 25 tries the backup is still not scheduled or one of the dependent services is still not running, then the following error message is displayed: `Cannot schedule backup since maximum retry attempt is reached and one of the dependent services is still not running`. Ensure that all required services are running, and then restart the Backup service to schedule the backup.

Edit a Backup Schedule

1 In SAS Environment Manager, click ⚙ **Jobs**.

On the **Scheduling** tab, all scheduling jobs including the default backup schedule and binary backup schedule are listed.

2 Select a job, and click 🕒.

3 In the Edit Schedule dialog box, edit the required properties of the schedule.

4 (Optional) In the Edit Schedule dialog box, perform the following steps to add a trigger to the schedule:

a Click ✚ to add a trigger to the schedule.

b In the New Trigger dialog box, enter values.

You must enter values for **Name** and **Time** fields.

c Click **Save**.

5 Click **Save** again to close the Edit Schedule dialog box.

Create a Backup Job

A provider administrator can select tenants and schedule a job.

1 In the left pane of SAS Environment Manager, click **Backup and Restore**.

2 On the Backup and Restore page, click **Create Backup Job**.

3 In the Create Backup Job dialog box, enter the following information:

Table A.10 Fields in the Create Backup Job Dialog Box

Field	Description
Name	Enter a name for the backup job.
Description	Enter a description for the backup job.

Field	Description
Select tenants	Select the tenants to perform the default backup.

4 Click **Create**.

5 In SAS Environment Manager, click ⚙ **Jobs**.

6 Click the **Scheduling** tab. From the list of scheduling jobs, select the newly created schedule, and click 🕒.

7 In the Edit Schedule dialog box, edit the required properties.

8 (Optional) In the Edit Schedule dialog box, perform the following steps to add a trigger to the schedule.

   a Click **+** to add a trigger to the schedule.

   b In the New Trigger dialog box, enter values.

   You must enter values for **Name** and **Time** fields.

   c Click **Save**.

9 Click **Save** again to close the Edit Schedule dialog box.

Performing a Backup

Best Practices for Performing Backups

- Always use Backup and Restore to perform backups of the content and configuration of SAS Viya components. Backup and Restore automatically discovers the services that are deployed. It finds any new services so that they can be included in the backup. Backup and Restore also finds content and configuration data from the SAS Viya deployment. It backs up this data at the same point in time, which is required for a same point-in-time restore.
- Edit the backup configuration properties immediately after a new installation, an upgrade, or any modifications to your SAS Viya deployment. For more information, see “Create and Edit a Backup Configuration” on page 51.
- Perform a backup after any modifications to your SAS Viya deployment. Examples include but are not limited to deploying SAS Viya, installing software updates, changing the topology, modifying the SAS Viya configuration, and changing configuration properties.
- Backups are purged after the retention period. If you do not want any backups to be deleted after the retention period, you must manually archive the backups before they are purged.
- After an upgrade in place or a migration, perform an immediate backup. Do not wait for the scheduled backup to run. Use this immediate backup for a restore. Do not use previous backups of SAS Viya for a restore.
- Turn on notifications to see whether the backup failed. Use the Notifications service in SAS Environment Manager. For more information, see “Notifications Service” on page 98.
- If a tenant administrator performs a local backup, the tenant administrator should copy the contents (local files) of the local vault to a shared vault.
- It is a best practice to back up your shared vault.

Perform an Immediate Backup Using SAS Backup Manager

If necessary, you can perform an immediate backup.

1 In the left pane of SAS Environment Manager, click **Backup and Restore**.

2 On the Backup and Restore page, click **Backup**.

3 In the Backup dialog box, provide information.

Table A.11 Fields in the Backup Dialog Box

Field	Description
Comments	(Optional) Enter free-form comments describing the backup. Comments are recorded in the backup history and are displayed in the backup's **Operation Details**.
Backup type	
Binary	Select this option if you want the pg_basebackup utility to back up all binaries in SAS Infrastructure Data Server. This includes metadata (such as user information, roles, and permissions). A binary backup includes all onboarded tenants in a multi-tenant deployment by default.
Include remaining sources	Select this check box to back up remaining data sources (such as SAS Configuration Server, SAS Message Broker, and CAS). Note: This check box is available when you select the **Binary** option.
Default	Select this option if you want the pg__dump utility back up the content of the SAS Infrastructure Data Server, SAS Configuration Server, SAS Message Broker, and CAS. A default backup does not include SAS Infrastructure Data Server metadata (such as user information, roles, and permissions). Use a default backup when SAS Infrastructure Data Server metadata is not significantly changed. A default backup can be fully restored using the Backup service. The restore includes SAS Infrastructure Data Server content, configuration data, and other data sources. A default backup includes all onboarded tenants in a multi-tenant deployment by default. However, SAS Backup Manager enables you to select tenants if you do not want all of them backed up.
Tenants	If the user is a provider administrator in a multi-tenant deployment, select the tenants to be backed up from the onboarded tenants. Select **Provider** to back up the provider tenant. If the user is a tenant administrator in a multi-tenant deployment or a SAS administrator in a standard deployment, the tenant list is not displayed.

Note: Because a tenant administrator can perform only a default backup, the **Binary** backup type is not displayed.

4 Click **Backup** to start the backup.

A new row is added to the table with a running status. After the backup is completed, a message about the status of the backup is displayed. To see the status of the backup on the main page, refresh your browser.

5 (Optional) In the left pane of SAS Environment Manager, click **Logs** to view the log of a backup or restore.

Perform an Immediate Backup Using a Command-Line Interface

1. Using a command line, log on to SAS Viya.

   For more information, see "Command-Line Interface: Preliminary Instructions" on page 681.

2. Use one of the following variations of the command `sas-admin backup start`:

   - Use one of the following commands to perform a default backup with a comment:

     ```
     sas-admin backup start -c="sample comment"
     ```

     ```
     sas-admin backup start --backup-type=default -c="sample comment"
     ```

   - Use the following command to perform a default backup of multiple tenants with a comment:

     ```
     sas-admin backup start --backup-type=default -c="sample comment" --tenants=<tenant1_ID>,<tenant2_ID>
     ```

   - Use the following command to perform a binary backup with a comment:

     ```
     sas-admin backup start --backup-type=binary -c="sample comment"
     ```

   - Use the following command to perform a binary backup with all data sources:

     ```
     sas-admin backup start --backup-type=binary --includeAllSourcesForBinaryBackup=true
     ```

Canceling a Backup

The Cancel operation is new for SAS Viya 3.4 (May 2019 upgrade).

1. Using a command line, log on to SAS Viya.

   For more information, "Command-Line Interface: Preliminary Instructions" on page 681.

2. Use the following command to cancel a backup:

   ```
   sas-admin backup cancel -i=<backup_job_id>
   ```

   On Windows, you can use the Microsoft Management Console Services snap-in to stop the SAS Backup Service.

A cancel is an asynchronous operation. It returns the current state of the backup. The user can execute the SHOW command to determine whether the backup got canceled.

```
sas-admin backup show -i=<backup_job_id>
```

Ensure that the backup has a ⊘ state.

Performing a Restore

Overview of Performing a Restore

Use one of the following scenarios for a restore or recovery:

- Restore a default backup.

In this case, use SAS Backup Manager or a command-line interface (CLI) to restore the following data sources:

- SAS Infrastructure Data Server
- SAS Configuration Server
- SAS Message Broker

You must manually restore CAS. For more information, see "Restore SAS Cloud Analytic Services" on page 63.

- Restore a binary backup where the include-all-sources property was set to FALSE.

  In this case, you must manually restore SAS Infrastructure Data Server. For more information, see "Restore SAS Infrastructure Data Server from a Binary Backup Manually" on page 61. For more information about the include-all-sources property, see "pre-restore validations" on page 36.

- Restore a binary backup where the include-all-sources property was set to TRUE. For more information about the include-all-sources property, see "pre-restore validations" on page 36.

  In this case, you must first manually restore SAS Infrastructure Data Server. Then, use SAS Backup Manager or a CLI to restore the following data sources:

  - SAS Configuration Server
  - SAS Message Broker

  Then, manually restore CAS. For more information, see "Restore SAS Cloud Analytic Services" on page 63.

You cannot use SAS Backup Manager or a CLI to perform the following tasks:

- Resolve a problem where the SAS Infrastructure Data Server cannot start by redeploying SAS Infrastructure Data Server, and then attempting to restore from a backup to that redeployment.
- Restore only a specific data source from a backup.
- Restore an unresponsive SAS Infrastructure Data Server. For more information, see "Restore SAS Infrastructure Data Server from a Binary Backup Manually" on page 61.
- Perform a restore if the SAS Infrastructure Data Server does not respond for reasons such as the following:
  - If the SAS Infrastructure Data Server is configured for high availability and the primary node is unresponsive, you must first promote a standby node to the primary node before performing a restore.

    For more information, see "Restore SAS Infrastructure Data Server from a Binary Backup Manually" on page 61 or "SAS Infrastructure Data Server, High Availability, and Backup and Restore" on page 43.

  - If the SAS Infrastructure Data Server is not configured for high availability and it is unresponsive, restore it first and then restore from the latest binary backup.

    For more information, see "Restore SAS Infrastructure Data Server from a Binary Backup Manually" on page 61.

Best Practices for Performing Restores

- While performing a restore, ensure that the system is not being actively used.
- Always use SAS Backup Manager or a CLI to perform restores to ensure a same point-in-time restore of content and configuration data.
- Always choose the most recent and successful backup to perform a restore.
- The backup-agent service (called SAS Backup Agent Service on Windows) must be running on all data sources that need to be restored.

- After performing a restore, stop and restart all services that are mentioned in “General Servers and Services: Operate (Linux)” on page 458.

  Note: Stopping and restarting services is applicable only for the restore of a standard deployment and for the restore of an entire multi-tenant deployment. It is not applicable for a tenant restore.

- Turn on notifications to see whether the restore failed. Use the Notifications service in SAS Environment Manager. For more information, see “Notifications Service” on page 98.
- Suppose that you upgraded in place or migrated to the latest version of SAS Viya, and you now want to restore a backup. In this case, you cannot use a backup from the previous version of SAS Viya for the restore.

Prerequisites for a Restore

- SAS Infrastructure Data Server must be running and responding to requests.
- In a standard deployment, the status of the backup that you want to restore must be ✓ or ✓⚠. If you are a provider administrator in a multi-tenant deployment, the status of the backup that you want to restore for the tenant must be ✓ or ✓⚠. If you are a tenant administrator in a multi-tenant deployment, the status of the backup that you want to restore must be ✓ or ✓⚠.

Restore Using SAS Backup Manager

1. Identify the backup that you want to restore.
2. Select one of the following options:
   - If the backup is a default backup, log on to SAS Environment Manager as a provider administrator or a tenant administrator in a multi-tenant deployment. Or, log on as a SAS administrator in a standard deployment.
   - If the backup is a binary backup where the include-all-sources property was set to FALSE, you must restore SAS Infrastructure Data Server manually. Do not perform the remaining steps. For more information about the include-all-sources property, see “pre-restore validations” on page 36.
   - If the backup is a binary backup where the include-all-sources property was set to TRUE, you must restore SAS Infrastructure Data Server manually first. Then, log on to SAS Environment Manager as a provider administrator in a multi-tenant deployment or as a SAS administrator in a standard deployment. For more information about the include-all-sources property, see “pre-restore validations” on page 36.

   For more information, see “Restore SAS Infrastructure Data Server from a Binary Backup Manually” on page 61.
3. On the **Backup and Restore** page, select a backup, and click **Restore**.
4. In the Restore dialog box, provide information.
   - **Comments**—(Optional) Free-form comments describing the restore. Comments are recorded in the restore history and are displayed in the restore’s **Operation Details**.
   - **Force restore if some databases don't match**—Select this check box if you want to force a restore if the databases do not match.
   - **Tenants**—Select the tenants to be restored from the onboarded tenants in the backup whose backup status is ✓ or ✓⚠. Select **Provider** to restore the provider tenant.

     Note: The tenant list is displayed only for a provider administrator in a multi-tenant deployment.
5. Click **Restore**.

   After the restore is completed, a message is displayed with its status as ✓, ✓⚠, or ⊗.

6. After a successful restore, stop and restart all services. Then, manually restore the CAS server.

   For more information, see “Restore SAS Cloud Analytic Services” on page 63.

7. To restore scheduled jobs in a standard deployment on Windows and on Linux, the SAS administrator must perform the following tasks. To restore scheduled jobs for a tenant in a multi-tenant deployment on Linux, the provider administrator must perform the following tasks.

   a. Pause the Schedule service.

      Use the following command on Windows to pause the Schedule service: `curl.exe -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/state?value=PAUSED`

      Use the following command on Linux to pause the Schedule service: `curl -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/state?value=PAUSED`

   b. Restore the schedules.

      Use the following command on Windows to restore the schedules: `curl.exe -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/restore`

      Use the following command on Linux to restore the schedules: `curl -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/restore`

   c. Resume the Schedule service.

      Use the following command on Windows to resume the Schedule service: `curl.exe -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/state?value=RESUME`

      Use the following command on Linux to resume the Schedule service: `curl -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/state?value=RESUME`

   For more information about obtaining an access token, see “Obtain an Access Token Using Password Credentials” on page 306.

   Note: If the restore of the schedules fails, stop the Schedule service and delete the scheduler schema. Then, start the Schedule service, and perform a restore.

8. (Optional) In the left pane of SAS Environment Manager, click **Logs** to view the logs.

   Alternatively, the provider administrator or SAS administrator can view log paths.

   For Linux:

   ***SAS-configuration-directory*/var/log/deploymentBackup/default**

   ***SAS-configuration-directory*/var/log/backup-agent/default**

   For Windows:

   ***SAS-configuration-directory*\var\log\deploymentBackup\default**

   ***SAS-configuration-directory*\var\log\backupagent\default**

   Note: If a failure occurs in any operation initiated by a tenant administrator, the tenant administrator must contact the provider administrator.

Restore Using a Command-Line Interface

1 Identify the backup that you want to restore.

2 Using a command line, log on to SAS Viya. For more information, see “Command-Line Interface: Preliminary Instructions” on page 681.

When you log on to SAS Viya at the command line, consider the following points:

- If the backup is a default backup, log on to SAS Environment Manager as a provider administrator or a tenant administrator in a multi-tenant deployment. Or, log on as a SAS administrator in a standard deployment.
- If the backup is a binary backup where the include-all-sources property was set to FALSE, you must restore SAS Infrastructure Data Server manually. Do not perform the remaining steps. For more information about the include-all-sources property, see “pre-restore validations” on page 36.
- If the backup is a binary backup where the include-all-sources property was set to TRUE, you must restore the SAS Infrastructure Data Server manually first. Then, log on to SAS Environment Manager as a provider administrator in a multi-tenant deployment or as a SAS administrator in a standard deployment. For more information about the include-all-sources property, see “pre-restore validations” on page 36.

For more information, see “Restore SAS Infrastructure Data Server from a Binary Backup Manually” on page 61.

3 Use one of the following variations of the command `sas-admin restore start`:

- Use the following command to restore a default backup or a binary backup where the include-all-sources property was set to TRUE. For more information about the include-all-sources property, see “pre-restore validations” on page 36.

```
sas-admin restore start --backup-name=<backup_ID>
```

- Use the following command to restore a backup even if the database validation fails:

```
sas-admin restore start --backup-name=<backup_ID> force=true
```

- Use the following command to restore a backup from an alternate shared vault location:

```
sas-admin restore start --backup-name=<backup_ID>
--alternate-shared-vault-for-restore=<path_of_alternate_sharedVault>
```

- Use the following command to restore multiple tenants:

```
sas-admin restore start --backup-name=<backup_ID> --tenants=<tenant1_ID>,<tenant2_ID>
```

- Use the following command to restore multiple tenants from an alternate shared vault location:

```
sas-admin restore start --backup-name=<backup_ID>
--alternate-shared-vault-for-restore=<path_of_alternate_sharedVault> --tenants=<tenant1_ID>,<tenant2_ID>
```

When the restore is initiated by the provider administrator in a multi-tenant deployment, all onboarded tenants that are available in the backup are restored.

When the restore is initiated by a tenant administrator in a multi-tenant deployment, data for that tenant is restored.

When the restore is initiated by a SAS administrator in a standard deployment, the selected backup is restored.

4 After the restore of the backup is completed, restart all services. You must manually restore the CAS server.

For more information, see “Restore SAS Cloud Analytic Services” on page 63.

5 To restore the scheduled jobs for a tenant, the provider administrator must perform the following tasks:

a Pause the Schedule service.

Use the following command on Windows to pause the Schedule service: `curl.exe -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/state?value=PAUSED`

Use the following command on Linux to pause the Schedule service: `curl -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/state?value=PAUSED`

b Restore the schedules.

Use the following command on Windows to restore the schedules: `curl.exe -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/restore`

Use the following command on Linux to restore the schedules: `curl -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/restore`

c Resume the Schedule service.

Use the following command on Windows to resume the Schedule service: `curl.exe -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/state?value=RESUME`

Use the following command on Linux to resume the Schedule service: `curl -XPUT -H "Authorization: bearer <token>" <protocol>://<host>:<port>/scheduler/jobs/state?value=RESUME`

For more information about obtaining an access token, see "Obtain an Access Token Using Password Credentials" on page 306.

Note: If the restore of the schedules fails, stop the Schedule service and delete the scheduler schema. Then, start the Schedule service, and perform a restore.

Restore SAS Infrastructure Data Server from a Binary Backup Manually

Restore SAS Infrastructure Data Server on Linux

1 Stop all services including the SAS Infrastructure Data Server service using the following command on Linux:

```
sudo /etc/init.d/sas-viya-all-services stop
```

2 Archive or rename the existing **node0** directory using the following command on Linux:

```
cd  /opt/sas/viya/config/data/sasdatasvrc/postgres
mv node0 node0_original
```

3 Create a new **node0** directory with permissions and ownership similar to the old **node0** directory. The new **node0** directory and its extracted content should be owned by the old **node0** directory owner. Submit the following command:

```
mkdir node0
chown sas:sas node0
chmod 700 node0
cd node0
```

4 Extract the contents of the base.tar.gz file into the **node0** directory.

- In a multi-tenant deployment, use the following command:

```
tar -xvf -p ../sharedVault/<backup-Id>/provider/postgres/base.tar.gz
```

- In a standard deployment, use the following command:

```
tar -xvf ../sharedVault/<backup-Id>/__default__/postgres/base.tar.gz
```

5 Ensure that the Hot_Standby property is set to OFF in the *SAS-configuration-directory*/data/sasdatasvrc/postgres/node0/postgresql.conf file.

Use the following command to set the Hot_Standby property to OFF:

```
echo "hot_standby = off" >> postgresql.conf
```

6 Start the Consul service, and then start the SAS Infrastructure Data Server service.

Use one of the following commands to start the Consul service:

On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-consul-default
```

On Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-consul-default start
```

Use the following command to start the SAS Infrastructure Data Server service:

```
sudo /etc/init.d/sas-viya-vault-default start
sudo /etc/init.d/sas-viya-sasdatasvrc-postgres start
```

7 Ensure that the SAS Infrastructure Data Server service started successfully without any issues.

```
sudo /etc/init.d/sas-viya-sasdatasvrc-postgres  status
```

Verify that the primary data server has started without any issues and that it has a status of UP. Here is an example:

```
sudo /etc/init.d/sas-viya-sasdatasvrc-postgres status
Checking status of sas-viya-sasdatasvrc-postgres...

PGPool is running with PID=4733
```

```
PGPool is running with PID=32413
Checking Postgresql nodes status...
 node_id | hostname           | port | status | lb_weight |  role   | select_cnt | load_balance_node |
replication_delay
---------+-------------------------------------------+------+--------+-----------+---------+------------
+----------------
 0       | myhost.domain.com    5432 | up     | 1.000000  | primary | 0          | true              | 0
(1 row)
```

8 Set the Hot_Standby property in the postgresql.conf file back to ON.

```
echo "hot_standby = on" >> postgresql.conf
```

9 Restart all services.

Note: All steps should be performed for any additional required SAS Infrastructure Data Server sessions.

Restore SAS Infrastructure Data Server on Windows

1 In Microsoft Management Console, stop the SAS Services Manager service. Then, ensure that all other services are stopped.

2 Archive or rename the existing ***SAS-configuration-directory*\data\sasdatasvrc\postgres\node0** directory.

You might want to rename the directory to **node0_original**.

3 Create a new **node0** directory with permissions and ownership similar to the old **node0** directory.

4 Change the ownership of the **node0** directory to SAS Infrastructure Data Server, and provide all permissions.

Perform the following steps to change the ownership:

a Right-click the **node0** directory, and then select **Properties**.

b On the **Security** tab of the Properties dialog box, click **Advanced**.

c In the Advanced Security Settings dialog box, change the owner to SAS Infrastructure Data Server.

d In the Properties dialog box, click **Edit**.

e In the Permissions dialog box, select the SAS Infrastructure Data Server user, and provide Full control permissions to the SAS Infrastructure Data Server user.

5 Extract the contents of the base.tar file from the required backup ID of the shared vault location into the **node0** directory using any utility such as WinZip.

6 In the Microsoft Management Console, start the SAS Consul service.

7 Open the Windows PowerShell command prompt as an administrator, and then enter the following code to start the SAS Infrastructure Data Server:

```
cd "SAS-configuration-directory\etc\sasdatasvrc\postgres\node0"
.\ConfigEnvironmentVariables.ps1
cd "SASHome\libexec\sasdatasvrc\script"
.\Invoke-StartStopPostgres.ps1 -action start -batchJob
```

8 Verify that the SAS Infrastructure Data Server has started without any issues. Here is an example:

```
PS C:\Windows\system32>  cd "C:\Program Files\SAS\Viya\libexec\sasdatasvrc\script"

PS C:\Program Files\SAS\Viya\libexec\sasdatasvrc\script> .\Invoke-HealthCheck.ps1 -batchJob

Config Environment Variable file located: C:\ProgramData\SAS\Viya\etc\sasdatasvrc\postgres
\node0\ConfigEnvironmentVariables.ps1
myhost.domain.com:5432 - accepting connections
```

9 In the Microsoft Management Console, start the SAS Services Manager service.

Note: All steps should be performed for any additional required SAS Infrastructure Data Server sessions.

Restore SAS Cloud Analytic Services

1 Stop the server.

For more information, see "SAS Cloud Analytic Services: How To (Scripts)" on page 472.

2 Replace the content of ***SAS-configuration-directory*/etc/cas/default/permstore/*<hostname>*** on Linux and ***SAS-configuration-directory*\etc\cas\default\permstore\*<hostname>*** on Windows with the content of the **hostname** directory that is located in the shared vault.

The **permstore** directory is in the ***<backupID>*/*<tenantID>*/cas-shared-default** directory in the shared vault. *<tenantID>* in the path is the ID of the tenant in the multi-tenant deployment. In a standard deployment, use _Default_ instead of <*tenantID*>. The cas user needs Read, Write, and Execute access (on LInux) to the `permstore` directory.

If you specified a location in the PATH=" " option when creating a backup, then that content is what you should restore.

Note: When you are restoring CAS to an alternate host, the name of the host name directory in the source environment and in the target environment might be different.

3 Restart the server.

For more information, see "SAS Cloud Analytic Services: How To (Scripts)" on page 472.

Performing a Restore to an Alternate Host

Overview of Performing a Restore to an Alternate Host

The following terminology is used in this section:

Source machine
machine where the backups were performed. These backups need to be restored on another machine.

Target machine
alternate machine that has the same SAS Viya software installed on it as the source machine. The restore operation is executed on this alternate machine from the backup performed on the source machine.

Source shared vault
shared vault location for the source machine.

Target shared vault
network location where the content from the source shared vault was copied to.

Alternate shared vault
alternate shared vault location for the restore.

You cannot use a binary backup to perform a restore to an alternate machine. A provider administrator must perform the restore to an alternate machine in a multi-tenant deployment. A SAS administrator can perform the restore to an alternate machine in a standard deployment.

Note: **The information in this section can be used as a part of any disaster recovery strategy at a customer site.** The steps in this section restore the information that the Backup service backed up. The data sources outside of SAS Infrastructure Data Server are not backed up.

Prerequisites for Performing a Restore to an Alternate Machine

- On the source machine, do not set the source shared vault location to the path of the target shared vault. Instead, copy the content in the source shared vault to some other network location that is accessible to the target machine. This location is referred to as the target shared vault.
- The target shared vault location has the same access rights as the source shared vault location.
- The SAS Viya deployment (such as the version of SAS Viya and any hot fixes applied to it) in the source environment and in target environment must be the same.
- All the services are started on the target machine.

Perform a Restore to an Alternate Machine

Perform a Restore from a Target Shared Vault

To perform a restore from a target shared vault using CLI commands:

1. Copy the content of the source shared vault to the target shared vault.

2. Set the sharedVault property to the location of the target shared vault.

   For more information about setting the sharedVault location, see “Create and Edit a Backup Configuration” on page 51.

   After the shared vault is changed, if a backup or restore is requested or if the Backup service is restarted, the system evaluates whether to synchronize the local history with the SAS Infrastructure Data Server tables. The system checks to ensure that the last successful backup from the global history is available in the SAS Infrastructure Data Server tables. If it is not available, local history files are synchronized with the SAS Infrastructure Data Server tables.

3. Using the command line, log on to SAS Viya. For more information, see “Command-Line Interface: Preliminary Instructions” on page 681.

4. On the target machine, use the following command to view the backup list:

   ```
   sas-admin backup list
   ```

5. Select a backup to restore.

6. In a multi-tenant deployment, perform the following tasks:

   a. Restore the provider tenant using the following command:

      ```
      sas-admin restore start --backup-name=<backup_ID> --tenants=<provider>
      ```

   b. Restart all services.

   c. Restore the remaining tenants using the following command:

      ```
      sas-admin restore start --backup-name=<backup_ID> --tenants=<tenant1_ID>,<tenant2_ID>
      ```

   In a standard deployment, use the following command to restore the backup:

   ```
   sas-admin restore start --backup-name=<backup-id>
   ```

7. (Optional) Make a note of the *<jobID>* printed on the console.

8. (Optional) To view the details of a restore, submit the following command until you receive a status of completed or failed:

   ```
   sas-admin restore show -id-=<jobID>
   ```

   You might need to submit this command multiple times because when the restore is running and SAS Infrastructure Data Server is being restored, all of the services might not be able to respond. The command might not return the expected response and you might get a `403 Forbidden` error.

9. (Optional) If a restore fails, review the logs. Make any necessary changes.

   If the status of the restore is ✓⚠, review the backup-agent log. Make any necessary changes.

   The backup-agent log is located on Linux at ***SAS-configuration-directory*/var/log/backup-agent/default**.

The backup-agent log is located on Windows at ***SAS-configuration-directory*****\var\log\backupagent\default**.

10. After the restore is completed, restart all services.

11. Restore scheduled jobs.

    For more information, see Step 5 on page 60.

12. After all services are restarted, restore SAS Cloud Analytic Services.

    For more information, see “Restore SAS Cloud Analytic Services” on page 63.

13. Restore any required applications.

Perform a Restore from an Alternate Shared Vault

To perform a restore from an alternate shared vault using commands:

1. Using the command line on the target machine, log on to SAS Viya. For more information, see “Command-Line Interface: Preliminary Instructions” on page 681.

2. From the alternate shared vault, select a backup to restore.

3. In a multi-tenant deployment, perform the following tasks:

   a. Restore the provider tenant using the following command:

      ```
      sas-admin restore start --backup-name=<backup_ID>
      --alternate-shared-vault-for-restore=<path_of_alternate_sharedVault> --tenants=<provider>
      ```

   b. Restart all services.

   c. Restore the remaining tenants using the following command:

      ```
      sas-admin restore start --backup-name=<backup_ID>
      --alternate-shared-vault-for-restore=<path_of_alternate_sharedVault> --tenants=<tenant1_ID>,<tenant2_ID>
      ```

   In a standard deployment, use the following command to restore the backup:

   ```
   sas-admin restore start --backup-name=<backup_ID>
   --alternate-shared-vault-for-restore=<path_of_alternate_sharedVault>
   ```

4. (Optional) Make a note of the *<jobID>* printed on the console.

5. (Optional) To view the details of a restore, submit the following command until you receive a status of completed or failed:

   ```
   sas-admin restore show -id-=<jobID>
   ```

   You might need to submit this command multiple times because when the restore is running and SAS Infrastructure Data Server is being restored, all of the services might not be able to respond. The command might not return the expected response and it might give a 403 error.

6. (Optional) If a restore fails, review the logs. Make any necessary changes.

   If the status of the restore is ✓⚠, review the backup-agent log. Make any necessary changes.

   The backup-agent log is located on Linux at ***SAS-configuration-directory*****/var/log/backup-agent/default**.

   The backup-agent log is located on Windows at ***SAS-configuration-directory*****\var\log\backupagent\default**.

7 After the restore is completed, restart all services.

8 Restore scheduled jobs.

For more information, see Step 5 on page 60.

9 After all services are restarted, restore SAS Cloud Analytic Services from the backup that was selected on the alternate shared vault.

For more information, see “Restore SAS Cloud Analytic Services” on page 63.

10 Restore any required applications.

Backing Up and Restoring Programming-Only Deployments

Back Up CAS Access Controls and Caslib Information

Important: After global-scope caslibs and access controls are modified, you must back up each CAS server’s stored access control and caslib information.

SAS administrators can perform a backup programmatically using the createBackup and completeBackup actions as follows:

1 Run the following code in SAS Studio, replacing the PATH location with your location:

```
cas casauto host="cloud.example.com" port=5570;

proc cas;
accessControl.assumeRole /
      adminRole="SuperUser";
accessControl.createBackup /
      path="/my/backup/location";
accessControl.completeBackup;
accessControl.dropRole / adminRole="SuperUser";
quit;
```

2 Copy the backup location directory to a location where it can be saved. The cas user needs Write access to the location. If the location does not exist and the cas user has Write and Execute permission (on Linux), the location is created.

If you do not specify `path=" "`, the backup location is the directory named **`backup`**. This directory is in the PERMSTORE option location. It is under the directory with the fully qualified DNS name of the machine that runs the main controller. The cas user must have Read, Write, and Execute permission (on Linux) to both the **`permstore`** and **`backup`** directories.

Understand the Backup Configuration

The configuration information about Linux is stored in the following files for the hosts in the [sas-casserver-primary] host group in the inventory file:

`SAS-configuration-directory`**`/etc/cas/default/casconfig.lua`**

`SAS-configuration-directory`**`/etc/cas/default/cas.hosts`**

The configuration information about Windows is stored in the following file for the hosts in the [sas-casserver-primary] host group in the inventory file:

***SAS-configuration-directory*/etc/cas/default/casconfig.lua**

The configuration information is stored in the following files for the hosts in the [programming] host group in the inventory file:

```
SAS-configuration-directory/etc/sasstudio/default/init_usermods.properties
SAS-configuration-directory/etc/sasstudio/default/appserver_usermods.sh
SAS-configuration-directory/etc/spawner/default/spawner_usermods.sh
SAS-configuration-directory/etc/workspaceserver/default/autoexec_usermods.sas
SAS-configuration-directory/etc/workspaceserver/default/sasv9_usermods.cfg
SAS-configuration-directory/etc/workspaceserver/default/workspaceserver_usermods.sh
```

If your site created global folder shortcuts for SAS Studio, you must back up the directory that contains the shortcuts. By default, the shortcuts are stored in the following directory:

```
SASHome/SASFoundation/GlobalStudioSettings
```

Note: Your site might have configured a different directory for the shortcuts. For more information, see "Configuring Global Folder Shortcuts" on page 80.

Restore Programming-Only Deployment

1 Stop the CAS server.

   See "SAS Cloud Analytic Services: How To (CAS Server Monitor)" on page 486.

2 Identify the location of the permstore. Copy all content to a safe location, and then remove it from the permstore.

   The default location of the permstore is ***SAS-configuration-directory*/etc/cas/default/permstore/<*fully-qualified-domain-name*>**.

3 Copy the content of the backup directory to the permstore directory.

   Ensure that the cas account has Read and Write access to all the permstore files.

4 Start the server.

For more information about stopping and starting CAS, see "SAS Cloud Analytic Services: How To (Scripts)" on page 472.

Restore the Most Recent Permstore on Linux in the Event of a Failover

If the backup controller has taken over for the primary controller, perform the following steps to restore the most recent permstore from the backup controller:

1 Stop the CAS server.

   See "SAS Cloud Analytic Services: How To (CAS Server Monitor)" on page 486.

2 Identify the location of the primary controller's permstore. Copy all content to a safe location, and then remove it from the primary controller's permstore.

   The default location of the primary controller's permstore is ***SAS-configuration-directory*/etc/cas/default/permstore/<*fully-qualified-domain-name*>**.

3 Copy the content of the backup controller's permstore to the primary controller's permstore.

The default location of the backup controller’s permstore is *SAS-configuration-directory*/etc/cas/default/permstore/*<fully-qualified-domain-name>*.

Ensure that the cas account has Read and Write access to all the permstore files.

4 Copy the content of the backup controller’s permstore to a safe location.

5 Start the CAS server.

For more information about stopping and starting CAS, see “SAS Cloud Analytic Services: How To (Scripts)” on page 472.

Troubleshooting

Backup and Restore: Logs

Backup and Restore generates the following logs that can be used in troubleshooting:

- On the host where Backup and Restore is deployed, service logs are created in the following paths. The name of the log files is based on the time at which the backup was started.

  On Linux: ***SAS-configuration-directory*****/var/log/deploymentBackup/default**

  On Windows: ***SAS-configuration-directory*****\var\log\deploymentBackup\default**

- On each of the data sources, backup logs are created in the following paths. The name of the log files is based on the time at which the backup-agent service was started.

  On Linux: ***SAS-configuration-directory*****/var/log/backup-agent/default**

  On Windows: ***SAS-configuration-directory*****\var\log\backupagent\default**

- If a restore of the SAS Infrastructure Data Server fails, the log files for the restore are dumped to the following directory:

  <sharedVault>/*<backup_ID>*/*<tenant_ID>*/postgres/restore/*<restore_ID>*...*<restore.log>*

  This directory is on the host where the data source resides. This is the location of the local vault.

Backup and Restore: Error and Warning Messages

This service is not available which is required for scheduling default backup: "Access token denied." Cannot schedule backup since the maximum retry attempt is reached and one of the dependent services is still not running.

If a service that is starting fails to register itself with SASLogon, the service tries multiple times to register the client until the client registration is done. In this case, the default backup is not created because it cannot generate client tokens to access other services, such as Schedule or Job Execution.

If the DEFAULT_BACKUP_SCHEDULE does not exist, you must restart the sas-deploymentBackup service and check again.

This service is not available which is required for scheduling default backup: %name-of-service%. Cannot schedule backup since the maximum retry attempt is reached and one of the dependent services is still not running.

Although the Backup service is starting, other services have not yet started. Ensure that all required services are started. In this scenario, the sas-deploymentBackup service retries 25 times with an interval of 5 minutes to schedule the backups. If one or more of the services is not running after 25 retries, you get the error message.

If the DEFAULT_BACKUP_SCHEDULE does not exist, you must restart the sas-deploymentBackup service and check again.

Note: *%name-of-service%* is the name of the service that did not start.

A Backup or Restore Is Already in Progress

This message indicates that a backup or restore is in progress. You cannot initiate multiple backups and restores at the same time.

For the SAS Viya 3.4 (May 2019 upgrade), see "Canceling a Backup" on page 56.

Otherwise, here is how to stop a backup that is in progress:

1. On the command line in Linux, enter `sas-admin backup cancel -i=<backup_id>`. To determine whether the cancel sub-command is supported, enter `sas-admin backup help`. If the cancel sub-command is not supported, you must stop and then restart the Backup service.
2. Ensure that the backup has a ⊘ state.
3. Perform a backup.

   For more information, see "Performing a Backup" on page 54.

Here is how to stop a restore that is in progress:

1. On the command line in Linux, enter `sudo /etc/init.d/sas-viya-deploymentBackup-default restart`.

   On Windows, use the Microsoft Management Console Services snap-in on Windows to stop the SAS Backup Service.
2. Ensure that the restore has a ⊘ state.
3. Restart the SAS Backup Service.
4. Perform a restore.

   For more information, see "Performing a Restore" on page 56.

Database Lists in the Backup and SAS Infrastructure Data Server Database Do Not Match

The database list in the backup <*source list in backup*> and the database list in SAS Infrastructure Data Server <*source list in database*> do not match. Set the FORCE option in a restore request to TRUE to force the restore.

This indicates that the databases in SAS Infrastructure Data Server at the time of the backup and the databases in SAS Infrastructure Data Server do not match. The Backup and Restore service does not restore a database that is missing.

If you still want to restore the remaining databases, enter the following command: `sas-admin restore start --backup-name=<backup_ID> force=true`

Configuration with the ID {0} Was Not Found

This indicates that the configurationId provided in the backup request is not available or supported.

> **TIP** Currently only the DEFAULT value is supported for configurationId. Modify the backup request to set the configurationId to `default` and try again.

Error Code 403 While Retrieving Information Related to Restore Operation

While a SAS Infrastructure Data Server restore is running, it might take some time for all services to respond. In this case, the user might see a `403 Forbidden` error message.

Request Contains Invalid Values for the Start or Limit Parameters

The request contains invalid values for the start ({0}) or limit ({1}) parameter. Use positive integers as values for the start and limit parameters and resubmit the request.

Specified Backup Does Not Have a Directory in the Shared Vault or in the History File

The specified backup is not found in the shared vault or history file. This might occur because the database mode of the environment from which the backup was performed is different from the target environment. Ensure that you are using compatible environments and start another backup to initiate the restore.

Invalid Shared Vault Location

The shared vault location is invalid. The location of the shared vault and local vault must be different. Set a valid shared vault location in SAS Environment Manager. For more information, see "Create and Edit a Backup Configuration" on page 51.

List Can Only Contain Onboarded Tenants

The list of tenants should contain only tenants with a state of onboarded. Remove tenants that are not onboarded or are not valid.

backupType Value of Binary Contains a List of Tenants

A list of tenants should not be provided for the backupType value of binary. A binary backup cannot back up a subset of tenants. Remove the list of tenants or perform a default backup. In addition, a binary backup can be performed only by a provider administrator in a multi-tenant deployment.

Multi-Tenant Deployment Can Be Triggered Only by a Provider Administrator

In a multi-tenant deployment, a backup can be performed only by a provider administrator.

Shared Vault Is Not Accessible

The shared vault is not accessible. Make the shared vault accessible and restart the job. For more information about a shared vault and its access rights, see "shared vault" on page 37.

Specified Backup ID Is Incompatible with the Target System

The specified backup ID is incompatible with the target environment. The database mode of the target environment should be the same as the source environment. Use a different backup ID to restore to the target environment.

Backups or Restores Are in Pending State Even After Restarting All Services

Even after starting all services, you might see backups or restores in the pending state. This occurs because the Backup service was started before the Tenant service. Ensure that the Tenant service is running and then restart the Backup service.

After Performing UIP, the Default Backup Schedule Is Not Working

After performing an upgrade in place to SAS Viya 3.4, the default backup schedule might not work. Delete the default backup schedule and restart the Backup service. For more information about deleting the backup schedule, see “Jobs: How To” on page 170.

Cancel Operation Fails for Backup

A cancel operation failed for source {0}. Here is the error message:

```
The backup operation with the process ID "{1}" could not be canceled for the source "{0}"
```

To resolve this issue, manually terminate the OS process with ID {1}.

7

Configuration Properties

Configuration Properties: Overview

You manage configuration properties for SAS Viya servers, services, and applications using the Configuration pages in SAS Environment Manager.

Note: A programming-only deployment on page 1does not use SAS Viya services and SAS Environment Manager.

You manage configuration properties for SAS Studio 4.x by modifying its configuration file. For more information, see How To Configure SAS Studio 4 on page 80.

Configuration Properties: How to Configure Services

Introduction

These instructions explain how to view and modify service configuration properties using SAS Environment Manager.

Navigation

In the applications menu (≡), select **Administration** ⇨ **Manage Environment**. In the navigation bar, click 🔧.

The Configuration page is an advanced interface. It is available to only SAS Administrators.

Edit Configuration Instances

Note: Most SAS Viya applications and servers have a corresponding service in which you set their configuration property values.

1 Using the drop-down list, choose **All services**.

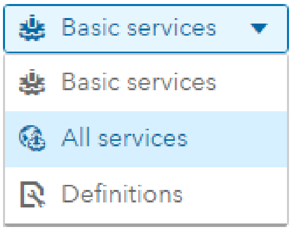

2 In the navigation pane, select a service whose configuration properties you want to change.

3 Next to the configuration instance, click ✎.

4 In the Edit Configuration dialog box, change the value in one or more of the configuration property fields.

5 When you are finished, click **Save**.

6 On a non-cloud platform, such as native Linux, some services require that you restart them when configuration changes are made. See "What Services Must Be Restarted?" on page 85.

Create Configuration Instances

In some situations, you might decide to create a configuration instance. For example, if you want to configure a logging level for a service that is not already associated with logging.level, you must create a new configuration instance of logging.level for that service.

1 Using the drop-down list, choose **All services**.

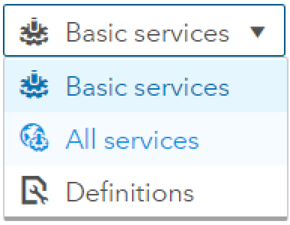

2 Select a service for which you want to create a new configuration instance.

3 At the top of the content pane, click **New Configuration**.

4 In the Select Definition dialog box, select a configuration definition from which to create a new configuration instance.

5 In **Services**, make sure that the service displayed is the one for which you want to create a new definition. If the correct service is displayed, skip to Step 6.

Otherwise, do the following:

a Next to **Services**, click ✎.

b In the Choose Services dialog box, highlight the service to which the configuration instance you are creating applies, and click ➡.

c Remove any services for which you do not want to create a configuration instance, by highlighting the service and clicking ⬅.

d When you are finished, click **OK**.

6 Continue entering values. When you are finished, click **Save**.

> **TIP** Properties with a red asterisk (*) are required to have a value.

7 On a non-cloud platform, such as native Linux, some services require that you restart them when configuration changes are made. See "What Services Must Be Restarted?" on page 85

Review Default Configuration Values

1 In the top left corner of the window, make sure that **Basic services** is selected.

2 In **Basic services** list, select a service, application, or server whose configuration instance must be created.

> **TIP** Incomplete required configuration instances are marked with a half-filled red circle.

◒ identities ◒ SASLogon

3 On the right side of the window, next to the half-filled red circle ◒, click ✱.

4 Most configuration definitions apply to only one service. In the New Configuration dialog box, if there is no edit icon (✎) next to the **Services** field, skip to Step 5.

Otherwise, do the following:

a Next to **Services**, click ✎.

b In the Choose Services dialog box, highlight the service to which the configuration instance you are creating applies, and click ➡.

c When you are finished, click **OK**.

5 Continue entering values. When you are finished, click **Save**.

> **TIP** You are required to provide a value for properties marked with a red asterisk (*).

6 Repeat steps 2 – 5 for every configuration instance that is incomplete ◒.

Set Time-out Interval for SAS Viya Web Applications

Using SAS Environment Manager, you can change the session time-out interval for one or more SAS Viya web applications. The session time-out interval is the specific period of time that a web application waits before it signs off users' inactive sessions.

Important: In SAS Studio 4.x, individual users set the session time-out interval up to a maximum that is defined by the administrator with `webdms.maxSessionTimeoutInHours`. For more information, see "SAS Studio 4.x" on page 123.

1 Using the drop-down list, choose **Definitions**.

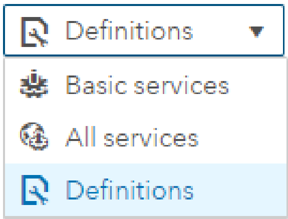

2 In the list of configuration definitions, select **server** and in top right corner of the view, click **New Configuration**.

3 In the New server Configuration dialog box, click .

4 In the Choose Services dialog box, highlight one or more SAS Viya web applications, click , and click **OK**.

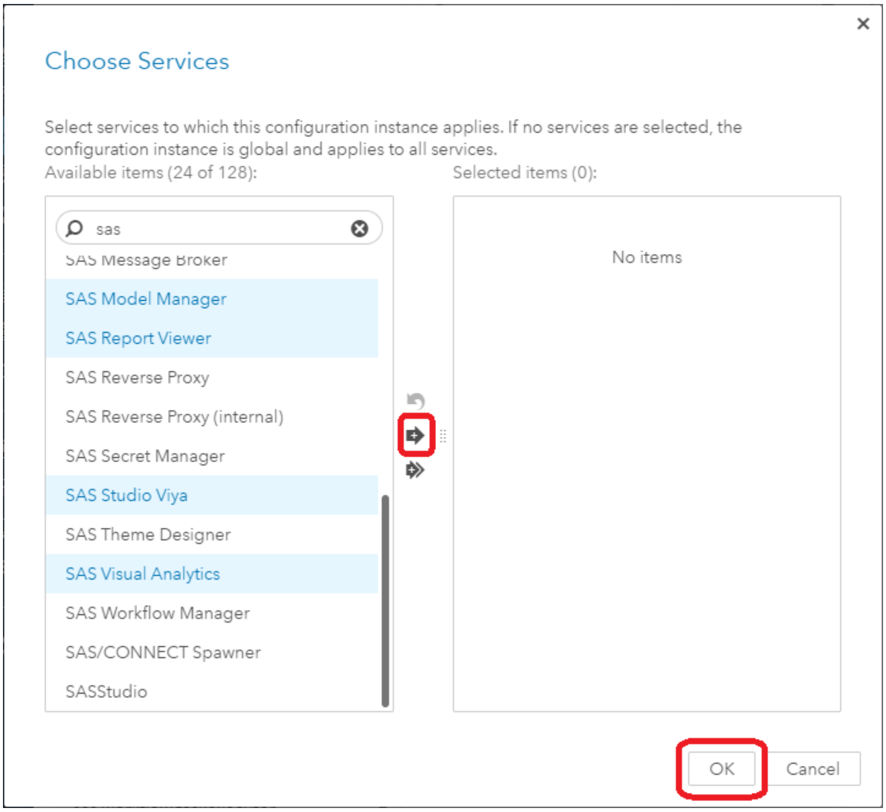

Choose Services

Select services to which this configuration instance applies. If no services are selected, the configuration instance is global and applies to all services.

Available items (24 of 128):

sas

SAS Message Broker
SAS Model Manager
SAS Report Viewer
SAS Reverse Proxy
SAS Reverse Proxy (internal)
SAS Secret Manager
SAS Studio Viya
SAS Theme Designer
SAS Visual Analytics
SAS Workflow Manager
SAS/CONNECT Spawner
SASStudio

Selected items (0):

No items

5 In the New server Configuration dialog box, click **Add property**.

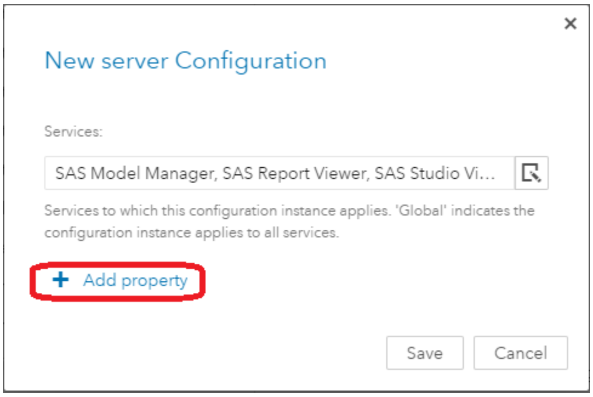

New server Configuration

Services:

SAS Model Manager, SAS Report Viewer, SAS Studio Vi...

Services to which this configuration instance applies. 'Global' indicates the configuration instance applies to all services.

6 In the Add property dialog box, in the **Name** field, enter the following Spring server property: `session.timeout`.

7 In the **Value** field, enter the number of seconds you want the SAS Viya web applications to wait before they sign off users' inactive sessions.

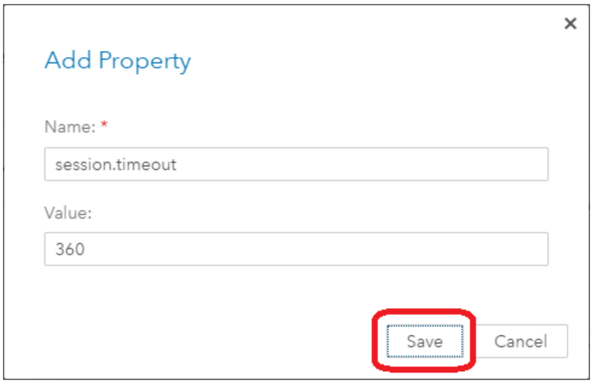

Add Property

Name: *

session.timeout

Value:

360

8 Click **Save**.

Your change takes effect for any new sign-ins to a SAS Viya web application.

Disable Opt-In Notifications

SAS Viya provides messages called *notifications* in various web applications. Users are alerted that one or more subscription-based (opt-in) notifications are available with a number circumscribed by a red circle next to the 🔔 on the right side of the web application's banner. (The number indicates the number of unread notifications.)

When you click the bell icon, a pop-up displays the notifications. The pop-up enables you to delete one or all of the notifications.

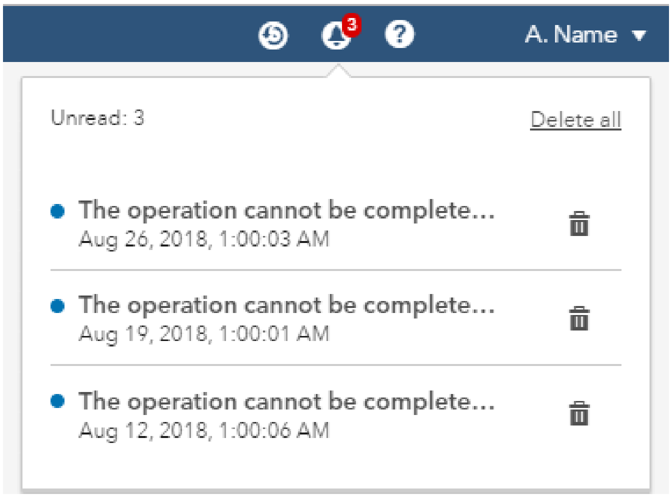

Unread: 3 Delete all

- The operation cannot be complete...
  Aug 26, 2018, 1:00:03 AM
- The operation cannot be complete...
  Aug 19, 2018, 1:00:01 AM
- The operation cannot be complete...
  Aug 12, 2018, 1:00:06 AM

The SAS Viya Notifications service processes only events that are published by participating services and applications, such as SAS Visual Analytics and the Backup service.

To disable opt-in notifications, follow these steps:

1 Using the drop-down list, choose **All services**.

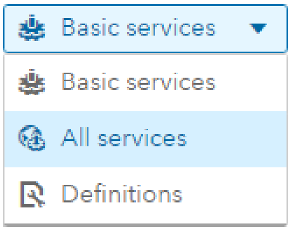

2 In the search field, enter `notifications`.

3 In list, select **Notifications service**.

> **TIP** At the top right side of the window, click ⮝ to easily locate the **Notifications service**.

4 In the content pane, locate **sas.notifications** and click ✎ on the right side of the window.

5 In the Edit sas.notifications Configuration dialog box, disable the **Enabled** slider and click **Save**.

Important: Disabling `sas.notifications.enabled`, turns off all subscription-based notification channels including subscription-based email notifications. To disable only subscription-based email notifications, disable `sas.notifications.delivery.mail.enabled`.

Configuration Properties: How to Configure SAS Studio 4.x

Update SAS Studio Configuration Properties

Note: SAS supplies two versions of SAS Studio, 4.x and 5.x. SAS Viya programming-only deployments use SAS Studio 4.x. For a comparison of the two SAS Studio versions, see "SAS Studio 5.1 and 4.4" in *What's New in SAS 9.4 and SAS Viya*.

Note: You set SAS Studio 5.x properties in the same way as the other SAS Viya applications and services. For more information, see "Configuration Properties: How to Configure Services".

To customize web application configuration properties for SAS Studio 4.x, edit init_usermods.properties, in the path appropriate for your operating system:

- Linux:

  ```
  /opt/sas/viya/config/etc/sasstudio/default/init_usermods.properties
  ```

- Windows:

  ```
  \ProgramData\SAS\Viya\etc\sasstudio\default\init_usermods.properties
  ```

Note: For sites that use Ansible: Ansible updates init_deployment.properties when it is run. Therefore, SAS Studio configuration changes that you make to init_usermods.properties are not overwritten by Ansible and are carried forward.

For a listing of configuration properties that you can update, see "SAS Studio 4.x" on page 123.

Changes take effect after you restart the web application. For more information, see instructions appropriate for your operating system on page 575.

> **TIP** Values that you specify in the init_usermods.properties file have precedence over corresponding values in other files. Unlike values in other files, values in the init_usermods.properties survive software upgrades.

Configuring Global Folder Shortcuts

In SAS Studio 4.x, you can create folder shortcuts from the **Server Files and Folders** section in the navigation pane. You might want to create global shortcuts for all the users at your site, so each user does not have to create these shortcuts manually.

1 In the init_usermods.properties file, specify a directory path for the `webdms.globalSettings` property.

  By default, this directory path is:

  - Linux:

    ```
    /opt/sas/spre/home/SASFoundation/GlobalStudioSettings
    ```

  - Windows:

    ```
    \Program Files\SAS\SPRE\SASFoundation\GlobalStudioSettings
    ```

  Important: If you choose to use this default, you must create the `GlobalStudioSettings` directory.

2 In an XML editor, create a shortcuts.xml file.

If you are trying to create a shortcut to a network location, here is the format of the shortcuts.xml file:

```
<?xml version="1.0" encoding="UTF-8"?>
<Shortcuts>
<Shortcut type="disk" name="network-location" dir="directory-path"/>
</Shortcuts>
```

3 Save the shortcuts.xml file to the global settings directory.

Set File Navigation Options

For SAS Studio 4.x, the root directories that are available in the **Server Files and Folders** section depend on whether you specify to show the system root directory with the `webdms.showSystemRoot` property. For information, see “Update SAS Studio Configuration Properties”.

- If `webdms.showSystemRoot` is set to True, the **Server Files and Folders** section displays the following information:
  - □ On Windows, the **Folder Shortcuts** folder includes a predefined shortcut to your **`My Documents`** folder.
  - □ On Linux, the **Folder Shortcuts** folder includes a predefined shortcut to your home directory. The **Files** folder is mapped to the system **`root`** directory for the server.

Figure A.1 If the webdms.showSystemRoot Property is True

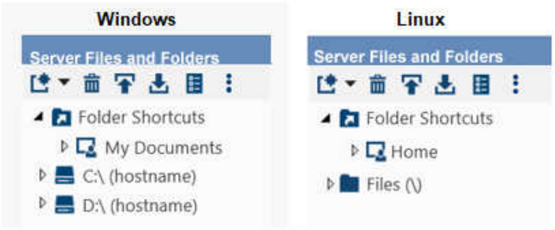

- If `webdms.showSystemRoot` is set to False, the **Server Files and Folders** section displays an empty **Folder Shortcuts** folder and a **Files** folder.
  - □ On Windows, the **Files** folder is mapped to your **`My Documents`** folder.
  - □ On Linux, the **Files** folder is mapped to your `$HOME` directory.

Figure A.2 If the webdms.showSystemRoot Property is False

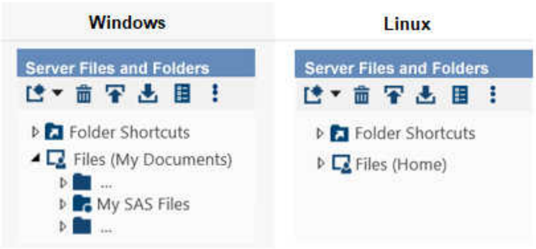

Operations

Automate Configuration Properties during Deployment (Ansible)

You can deploy SAS Viya with configuration values that are customized to your site by running your Ansible playbook with sitedefault.yml. Using sitedefault.yml, enables you to provision multiple machines in the same manner, and prevents you from having to modify configuration values with an administration interface after deployment.

Note: It is extremely important that the initial values applied with sitedefault.yml are correct. After you set a value with sitedefault.yml, you cannot re-run sitedefault.yml to change that value. You can re-run sitedefault.yml only to set properties that have not already been set. To change properties set with sitedefault.yml, you must use the sas-bootstrap-config CLI directly, or use another administration interface, such as SAS Environment Manager.

To set configuration values using sitedefault.yml, follow these steps:

1 Sign on your Ansible controller with administrator privileges, and locate the file, `/playbook/roles/consul/files/sitedefault_sample.yml`.

2 Make a copy of sitedefault_sample.yml and name the copy, sitedefault.yml.

3 Using a text editor, open sitedefault.yml and add values that are valid for your site.

   For information about the LDAP properties used in sitedefault.yml, see "sas.identities.providers.ldap" on page 101.

   For information about the all the properties that can be used in sitedefault.yml, see "Configuration Properties: Reference (Services)" on page 92.

   CAUTION! Some properties require passwords. If properties with passwords are specified in sitedefault.yml, you must secure the file appropriately. If you chose not to supply the properties in sitedefault.yml, then you can enter them using SAS Environment Manager. (Sign in to SAS Environment Manager as sasboot, and follow the instructions in "Configure the Connection to Your Identity Provider" in *SAS Viya for Linux: Deployment Guide*.)

4 When you are finished, save sitedefault.yml and make sure that it resides in the `/playbook/roles/consul/files` directory of the playbook.

5 Run your Ansible playbook using the sitedefault.yml file.

   Here is an example:

```
ansible-playbook site.yml
```

   For a complete list of playbook commands, see "Deploy the Software" in *SAS Viya for Linux: Deployment Guide*.

6 After the playbook is run, verify that the configuration values are successfully loaded into the configuration server by performing the following steps:

   a Verify that a copy of sitedefault.yml resides in `/viya/config/etc/consul.d/default/`.

   b Verify that config-kv-bulkload-sitedefault.json resides in `/viya/config/etc/consul.d/`.

   c View the configuration properties for a configuration definition such as, SAS Logon Manager, in SAS Environment Manager to verify that the specified values are present.

      For more information, follow the first five steps in "Edit Configuration Instances" on page 75.

Configuration Properties: Concepts

What Is SAS Viya Configuration?

From SAS Environment Manager, you can manage the configuration needs of the various SAS Viya services.

Configuration Components

A service’s configuration consists of the following components:

- *configuration definition*: A schema that describes a type of configuration. You create configuration instances from a configuration definition. Some examples of configuration definitions are: jvm, spring, and sas.reportdata.

  Note: Configuration definitions that apply to one or a small set of services are referred to as service configuration definitions. System configuration definitions can apply to any service.

- *configuration instance*: A collection of name-value pairs (a property) that a service uses. (These name-value pairs can sometimes be nested.)

  Note: Certain configuration instances are required for a service to be able to run. See “Review Default Configuration Values” on page 76.

How Configuration Definitions and Instances Are Displayed

The Configuration window in SAS Environment Manager contains three views: **Basic services**, **All services**, and **Definitions**.

Figure A.3 Configuration Views in SAS Environment Manager

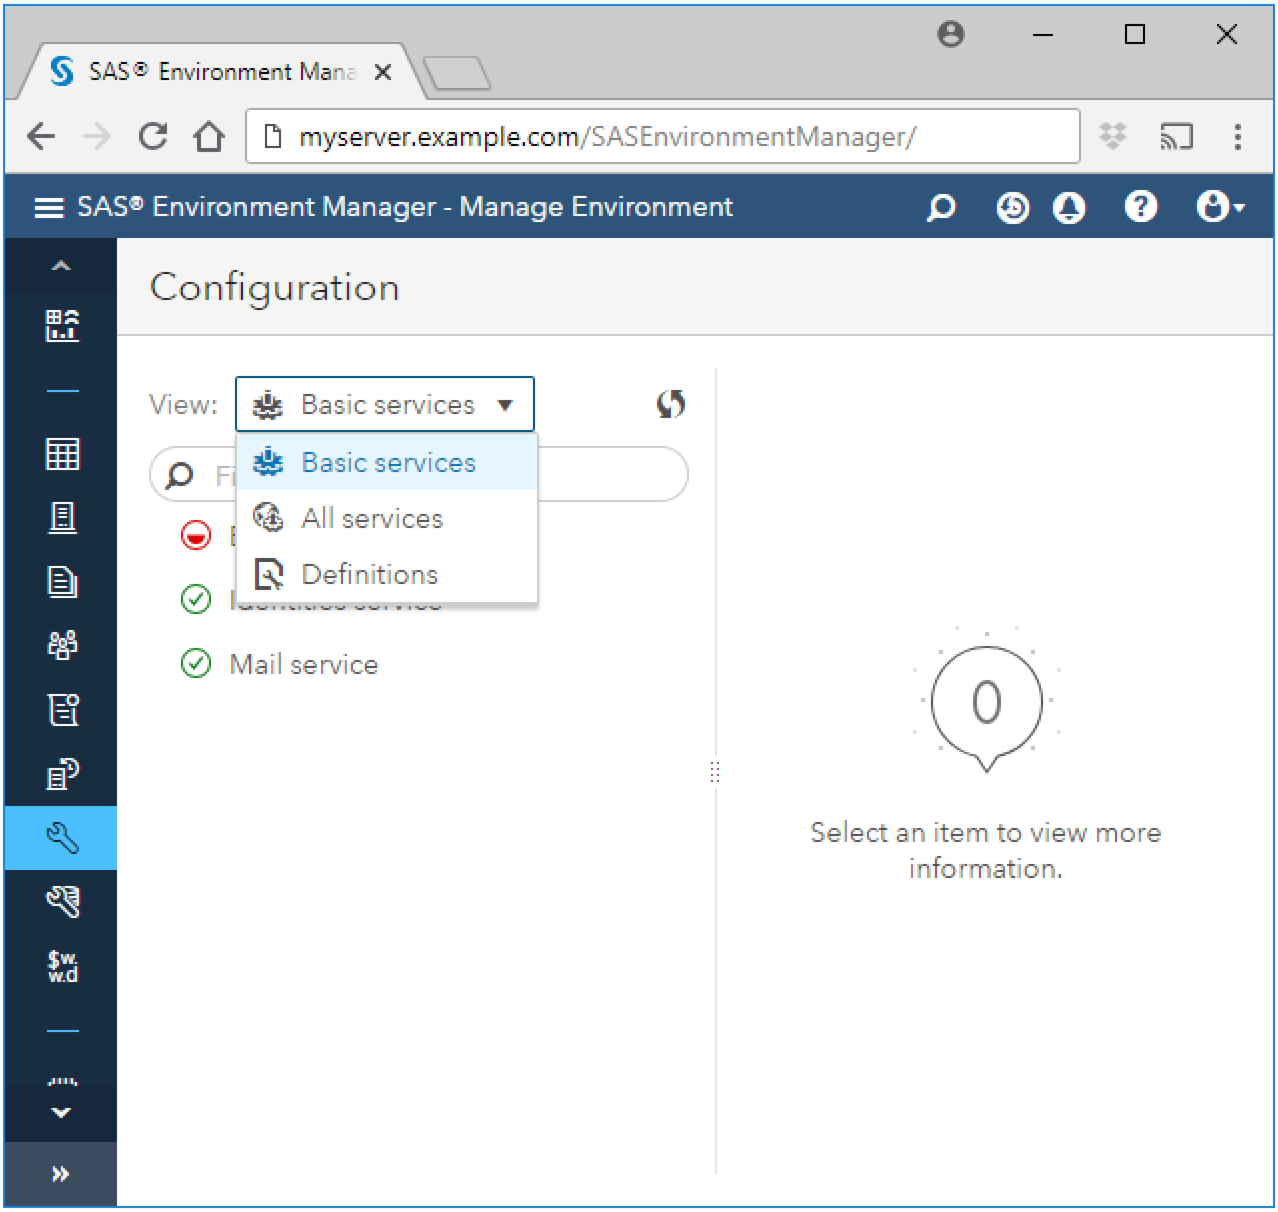
SAS® Environment Mana
myserver.example.com/SASEnvironmentManager/
SAS® Environment Manager - Manage Environment
Configuration
View:
Basic services
Basic services
All services
Definitions
Mail service
0
Select an item to view more information.

The **All services** view lists all SAS Viya services that are currently deployed and those that an administrator has not manually stopped. A SAS Viya service can be affected by one or more configuration instances. Most services have a one-to-one relationship with a configuration instance. However, some services are associated with more than one configuration instance. It is important to note that some services do not have any configuration instances, but you can set configuration properties for any of these services.

The **Definitions** view lists all the SAS Viya configuration definitions.

The **Basic services** view contains those services with configuration properties for which SAS cannot create a reasonable default (for example, the machine name for your SMTP service). As an advanced topic, these configuration properties can be set in an initial deployment using sitedefault.yml. For more information, see "Automate Configuration Properties during Deployment (Ansible)" on page 82.

For a tenant within a multi-tenant environment, the services in the **Basic services** view can appear to be incomplete configuration instances (indicated by a half-filled red circle next to the configuration instance name). For security purposes, tenants cannot see configuration values that apply to other tenants. This means that the tenant administrator sees incomplete red icons, because administrators are not allowed to see a configuration for an item that does not apply to them. For more information, see "Provider Administrator: Manage Tenants" in *SAS Viya Administration: Multi-tenancy*.

What Services Must Be Restarted?

On a non-cloud platform, such as native Linux, whenever a change is made to a Java virtual machine (JVM) configuration property (a Java option), any services that rely on that property must be restarted. For information about how to restart one or more services, see "General Servers and Services: Operate (Linux)" on page 458.

If you change configuration property values for any of the following services, you must restart the service:

- SAS Cache Locator
  - Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

    ```
    sudo systemctl restart sas-viya-cachelocator-default
    ```

  - Red Hat Enterprise Linux 6.x (or an equivalent distribution):

    ```
    sudo service sas-viya-cachelocator-default restart
    ```

- SAS Cache Server

  Important: If SAS Cache Server is restarted, then all services that are dependent on the cache server must be restarted. The list can vary depending on what SAS Viya offerings are deployed.

  > **TIP** To determine which services are dependent on SAS Cache Server, enter this command: `grep -H "Product-Name: Apache Geode" /opt/sas/viya/config/var/log/*/default/*.log`.

  - Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

    ```
    sudo systemctl restart sas-viya-cacheserver-default
    ```

  - Red Hat Enterprise Linux 6.x (or an equivalent distribution):

    ```
    sudo service sas-viya-cacheserver-default restart
    ```

- SAS Configuration Server (Consul)
  - Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

    ```
    sudo systemctl restart sas-viya-consul-default
    ```

  - Red Hat Enterprise Linux 6.x (or an equivalent distribution):

    ```
    sudo service sas-viya-consul-default restart
    ```

- SAS Message Broker (RabbitMQ)
  - Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

    ```
    sudo systemctl restart sas-viya-rabbitmq-server-default
    ```

  - Red Hat Enterprise Linux 6.x (or an equivalent distribution):

    ```
    sudo service sas-viya-rabbitmq-server-default restart
    ```

- SAS Infrastructure Data Server (PostgreSQL)

  For information, see "Operate a Cluster (Linux)" on page 603.

Note: You must be signed in to the machine where these services reside with sudo privileges to run these scripts.

See Also

- "Start and Stop a Specific Server or Service" on page 461

How SAS Viya Configuration Works

Spring-Based Microservices

For SAS Viya, Spring-based microservices, property changes are made in SAS Environment Manager and stored in SAS Configuration Server. SAS Viya triggers a refresh, and the new property values are applied to the service. In most situations, no restart of the service is required.

Figure A.4 How Configuration Properties Are Updated (Spring-Based Services)

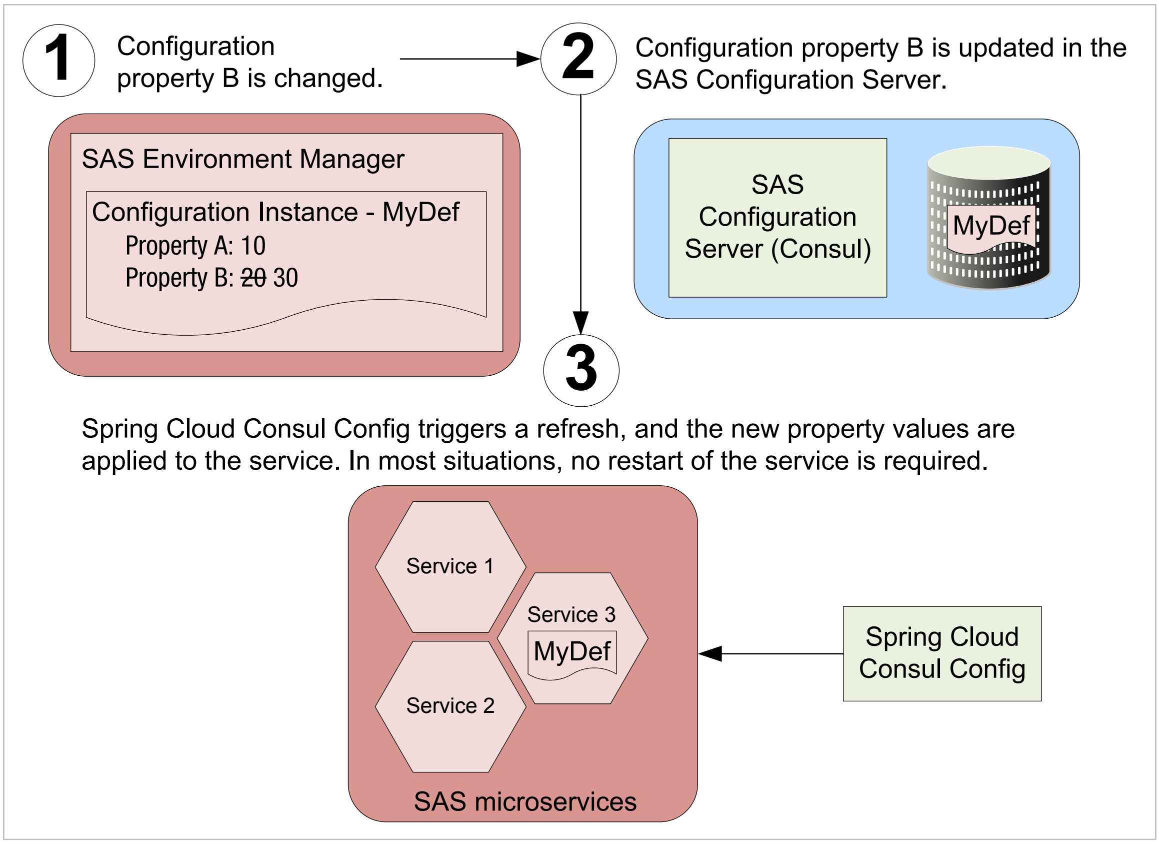
1
Configuration property B is changed.
2
Configuration property B is updated in the SAS Configuration Server.
SAS Environment Manager
Configuration Instance - MyDef
Property A: 10
Property B: ~~20~~ 30
SAS Configuration Server (Consul)
MyDef
3
Spring Cloud Consul Config triggers a refresh, and the new property values are applied to the service. In most situations, no restart of the service is required.
Service 1
Service 3
MyDef
Service 2
SAS microservices
Spring Cloud Consul Config

Non-Spring-Based Servers

For SAS Viya, non-Spring-based servers, property changes are made in SAS Environment Manager and stored in SAS Configuration Server. A tool, such as the consul-template daemon, extracts the configuration change from SAS Configuration Server and updates the appropriate service configuration file. Some servers, such as SAS Infrastructure Data Server, require you to manually restart them for their configuration changes to take effect.

Figure A.5 How Configuration Properties Are Updated (Non-Spring-Based Servers)

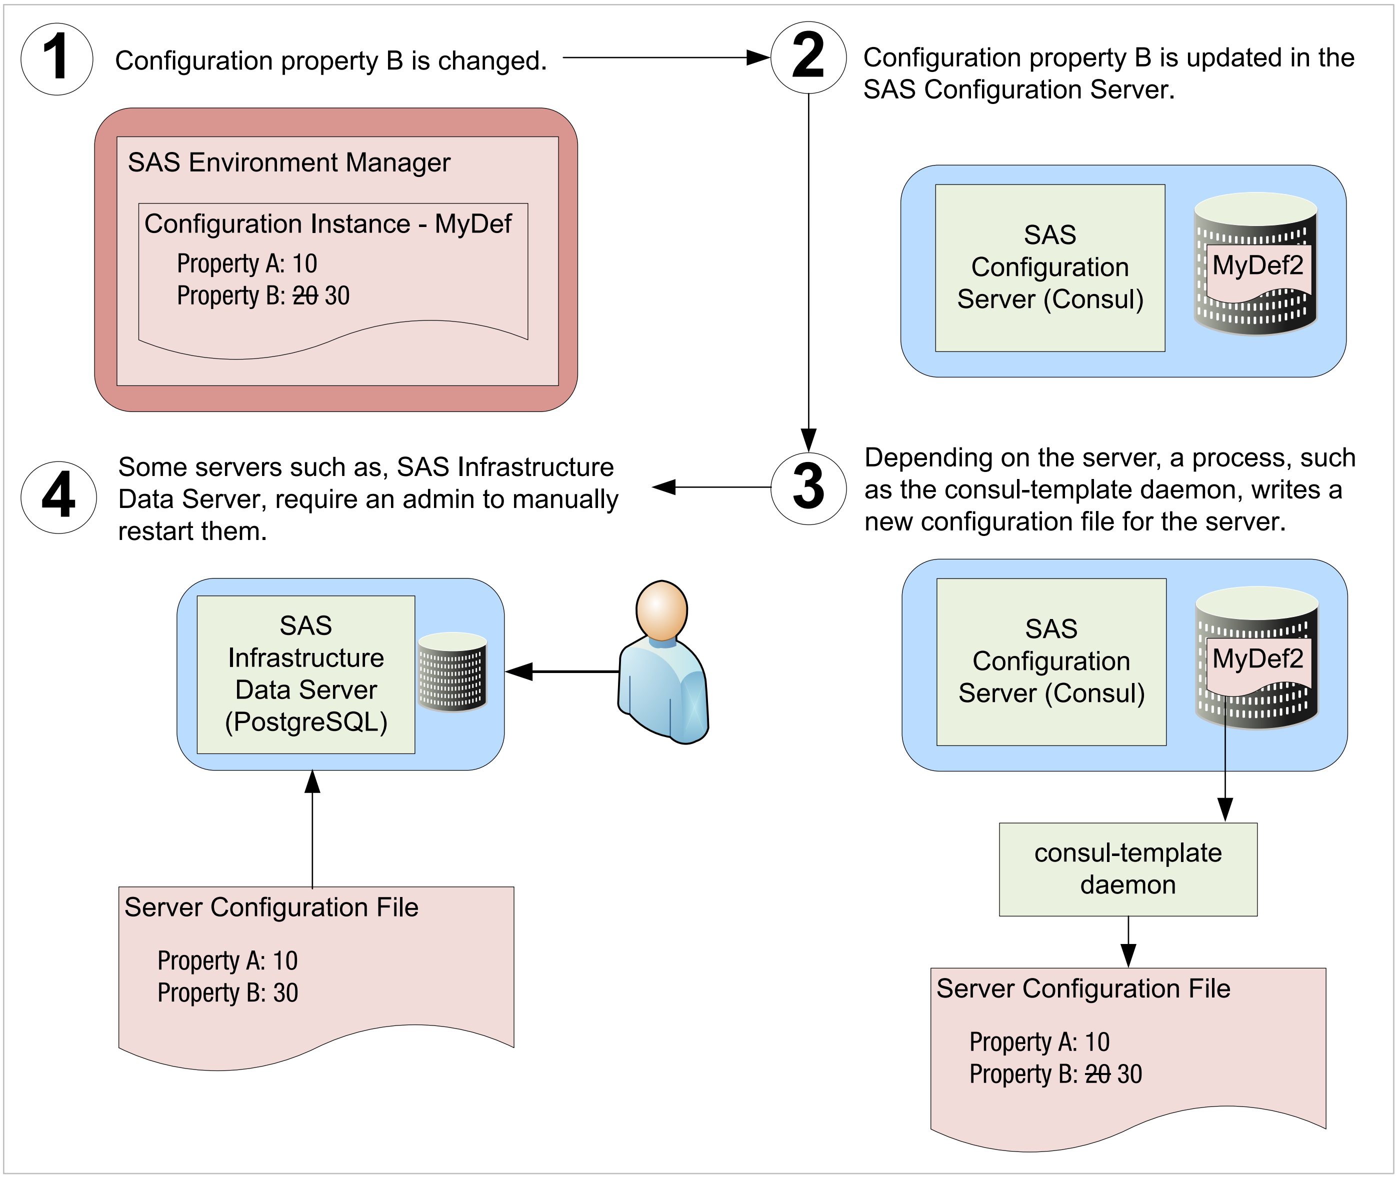
1
Configuration property B is changed.
SAS Environment Manager
Configuration Instance - MyDef
Property A: 10
Property B: ~~20~~ 30
2
Configuration property B is updated in the SAS Configuration Server.
SAS Configuration Server (Consul)
MyDef2
3
Depending on the server, a process, such as the consul-template daemon, writes a new configuration file for the server.
SAS Configuration Server (Consul)
MyDef2
consul-template daemon
Server Configuration File
Property A: 10
Property B: ~~20~~ 30
4
Some servers such as, SAS Infrastructure Data Server, require an admin to manually restart them.
SAS Infrastructure Data Server (PostgreSQL)
Server Configuration File
Property A: 10
Property B: 30

Bulk Loading of Configuration Values (sitedefault.yml)

You can deploy SAS Viya with configuration values that are customized for your site by running your Ansible playbook with sitedefault.yml. When sitedefault.yml is present in the playbook `roles/consul/files` directory, Ansible copies it to the machines that contain the SAS Configuration Server (Consul). When the configuration server starts, the watch script invokes the sas-bootstrap-config CLI to bulk load the key-value pairs that are defined in sitedefault.yml.

For more information, see "Automate Configuration Properties during Deployment (Ansible)" on page 82.

Note: The sas-bootstrap-config CLI uses a *check and set policy*. If a property currently exists in the configuration server, the CLI does not update the property. Therefore, it is extremely important that the initial values applied with sitedefault.yml are correct. After you set a value with sitedefault.yml, you cannot re-run sitedefault.yml to change that value. You can re-run sitedefault.yml only to set properties that have not already been set. To change properties set with sitedefault.yml, you must use the sas-bootstrap-config CLI directly, or use another administration interface, such as SAS Environment Manager.

Figure A.6 How Configuration Properties Are Updated (Non-Spring-Based Services)

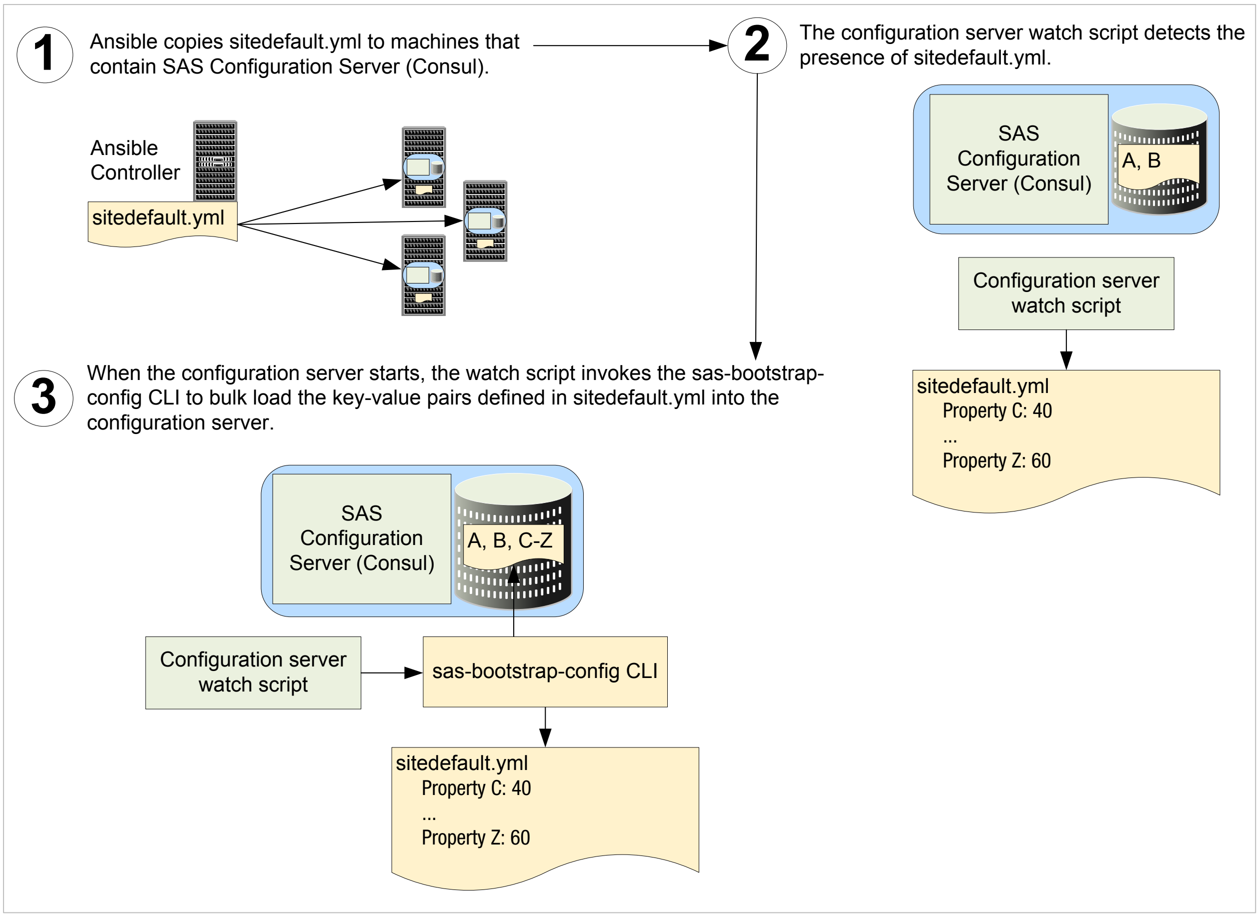
1
Ansible copies sitedefault.yml to machines that contain SAS Configuration Server (Consul).
Ansible Controller
sitedefault.yml
2
The configuration server watch script detects the presence of sitedefault.yml.
SAS Configuration Server (Consul)
A, B
Configuration server watch script
sitedefault.yml
Property C: 40
...
Property Z: 60
3
When the configuration server starts, the watch script invokes the sas-bootstrap-config CLI to bulk load the key-value pairs defined in sitedefault.yml into the configuration server.
SAS Configuration Server (Consul)
A, B, C-Z
Configuration server watch script
sas-bootstrap-config CLI
sitedefault.yml
Property C: 40
...
Property Z: 60

SAS Studio 5.x and 4.x Administration Differences

Overview

SAS Viya includes two releases of SAS Studio:

- SAS Studio 4.x

  the traditional version that has been available since the first release of SAS Viya.

- SAS Studio 5.x

  the new, microservices-based version, with a different interface. Studio 5.x integrates with other SAS Viya components (such as SAS Drive, Launcher Server, and Compute Server).

This section provides information about administration differences between the two versions of SAS Studio. For an overview of SAS Studio 5.x functionality, see “What’s New in SAS Studio 5.1” in *SAS Studio: User’s Guide*.

Deployment Environments

SAS Studio 4.x is used primarily in programming-only deployments. SAS Studio 4.x is an embedded web application that does not rely on external services. Therefore, SAS Studio 4.x has relatively few dependencies, such as a SAS Object Spawner and SAS Workspace Server.

SAS Studio 5.x is itself a microservice and provides functionality by relying on SAS Configuration Server, SAS Message Broker, SAS Infrastructure Data Server, and other microservices. For this reason, SAS Studio 5.x is available only in full deployments of SAS Viya.

Authentication Differences

SAS Studio 4.x uses host authentication such as Pluggable Authentication Modules (PAM) and Integrated Windows Authentication (IWA) to authenticate users. For more information, see "Authentication for Programming Interfaces" on page 310.

SAS Studio 5.x authentication fully integrates with the features provided by SAS Logon Manager, and provides more authentication possibilities, such as the following:

- LDAP provider

  standard user name and password form.

- Kerberos or Integrated Windows Authentication (IWA)

  single sign-on from the client host to the visual interfaces.

- Security Assertion Markup Language (SAML) provider

  single sign-on from third-party provider.

- OAuth and OpenID connect provider

  single sign-on from third-party provider.

- Pluggable Authentication Modules (PAM)

  multi-factor authentication via third-party tools.

- SAS 9.4

  single sign-on from SAS 9.4.

For more information, see "Authentication for Visual Interfaces" on page 308.

Accessing SAS Content

Microservices are stateless and do not directly consume resources exposed by the underlying operating system, including the file system. They are designed to be running in a cloud or a containerized environment that might not even have a file system. SAS Studio 5.x follows the same pattern: when you open or save a file, you do not have any access to the back-end file system. So how do you access your programs? There are at least two ways.

- Drag and drop:

  If your program file is available in your client, you can simply drag and drop it in the SAS Studio window. SAS Studio automatically opens the file.

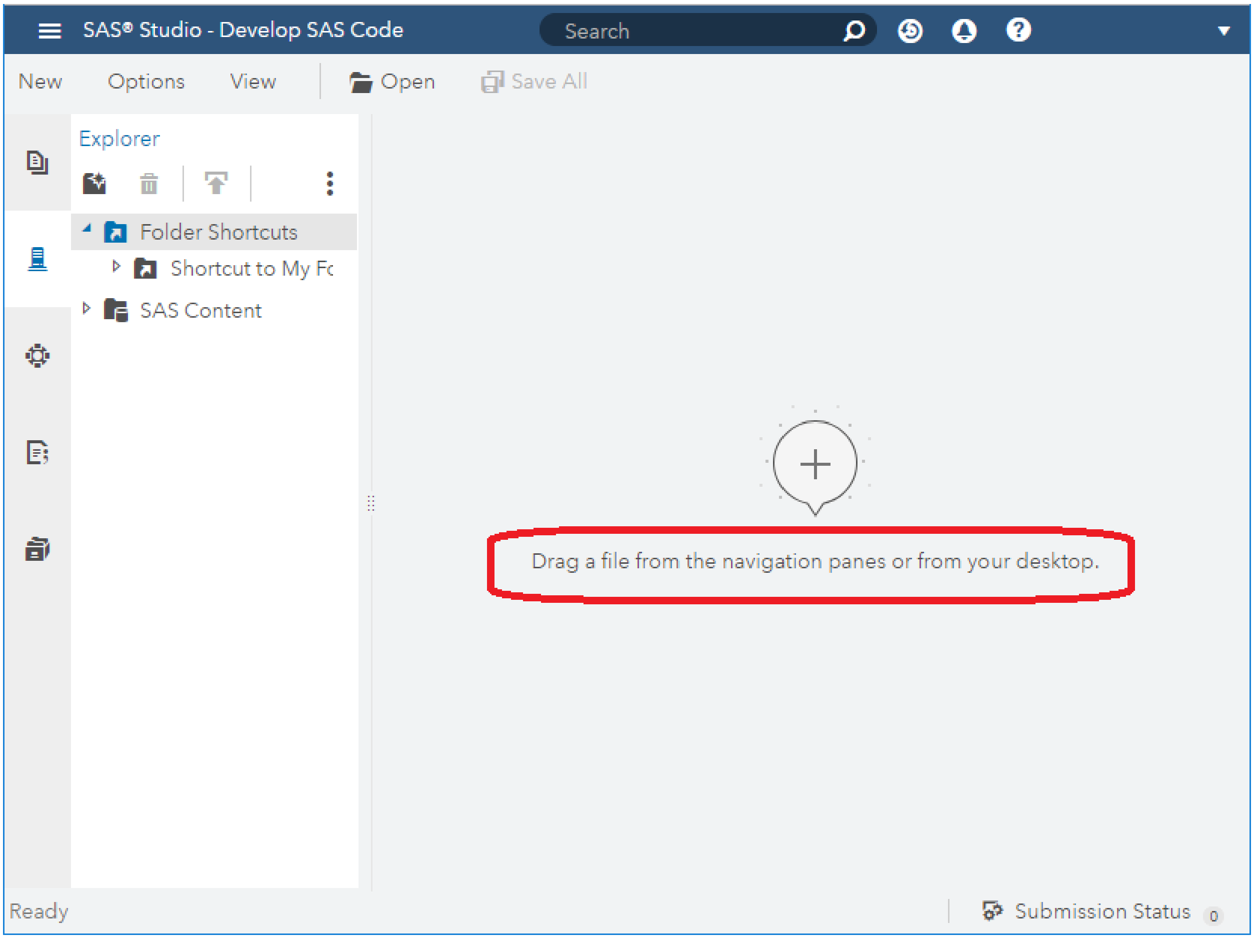

Drag a file from the navigation panes or from your desktop.

- **SAS Content** folder:

  All files, including program and data, can be read from and written to folders inside **SAS Content**. **SAS Content** is available from the **Explorer** pane.

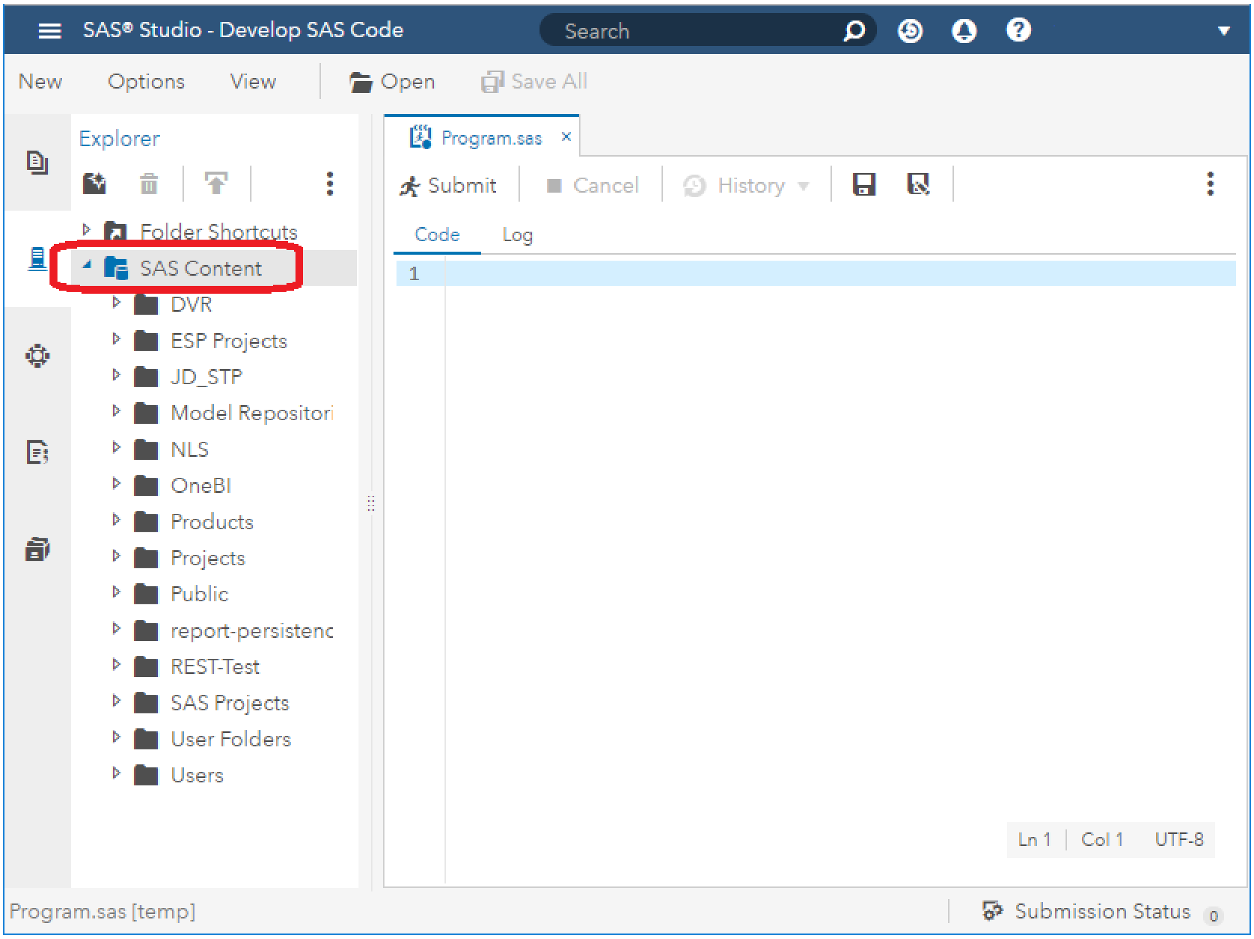

SAS Content is shared across all SAS Viya applications and managed by the SAS Drive web application.

Important:

Although any file can be written to or uploaded to folders in **SAS Content**, you cannot access data (.sas7bdat files) from SAS Content using libraries (because libraries cannot be assigned to **SAS Content** folders). Only data that can be read using a fileref can be read from **SAS Content**: CSV, TXT, XLSX, SAS, CTM (tasks), and CTK (task templates). For more information, see "Working with Data" in *SAS Studio: User's Guide*.

Improved Recovery

Previous releases of SAS Studio can be clustered for scalability and high availability, providing multiple instances of the web application. However, before SAS Studio 5.x, end-user sessions were bound to a specific instance known as a "sticky session."

To understand the disadvantages of a sticky session, consider the following example. Suppose that you have two SAS Studio 4.x instances running on server1 and server2. When you sign in to SAS Studio 4.x, the front-end load balancer directs you to server2. Until you sign out, your session remains on server2. If server2 abnormally shuts down—even though server1 is unaffected—your session is lost and you must sign out, possibly losing your work. When you sign in again, the load balancer redirects you to the available server, in this case, server1.

There is also another, more subtle issue. Each instance of SAS Studio 4.x can access only an object spawner and workspace servers running locally on the same machine. Suppose that all SAS Studio 4.x instances are running fine, but the object spawner on server2 goes down. Even if your SAS Studio 4.x session is unaffected, you are unable to perform any work. If you decide to sign out and back in, there is a chance that you will again be routed to server2, because the load balancer continues to see SAS Studio 4.x running on server2.

SAS Studio 5.x solves all these issues. Because it is a stateless microservice, Studio 5.x end-user sessions are not bound to any specific instance. If one instance stops abnormally, SAS Studio 5.x continues working using another microservice instance. Also, to run SAS code, SAS Studio 5.x does not use an object spawner and workspace server. Instead, Studio 5.x submits code to a compute server started by a launcher server. The compute and launcher server infrastructure ensures that there are no dependencies to a specific server machine. Therefore, SAS Studio 5.x can use a launcher and compute server located on any machine.

Architectural Comparison

To understand the architectures of SAS Studio 4.x and 5.x compare these diagrams:

- full deployment (SAS Studio 5.x)
- programming-only deployment (SAS Studio 4.x)

Configuration Properties: Troubleshooting

Service fails to start

Explanation:

A service might fail to start for many reasons. One error that can cause start failure is if you recently modified a configuration instance on which the service depends, and either the property name is misspelled or its value is incorrect.

Resolution:

Check any recently modified configuration instances, looking for a property misspelling or incorrect value. To do this, from the Configuration window in SAS Environment Manager, choose **Definitions**. Select the name of the configuration definition that you modified and look for any problems.

Configuration Properties: Reference (Services)

Application Registry Service

The Application Registry service registers applications to enable integration with SAS Drive and with the Application Switcher (side menu).

sas.appregistry
The set of configuration properties for the Application Registry service.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

Audit Service

The Audit service provides a framework for reporting on audit events.

sas.audit.archive
The set of properties that are used to archive audit records.

enabled
Enables archiving of audit records.

batchSize
The number of audit records to process in a single batch during an archive request.

scanSchedule
The schedule that determines when an archive request is initiated.

localRetention
The retention period for persisting audit records within the service.

storageType
The external storage mechanism to use for archiving audit records. Must be set to 'none' or 'local'.

storage.local.destination
The file location to use when 'storageType' is set to 'local'.

sas.audit.record
The set of properties that are used to control how audit events are recorded.

type
Customizable properties to control which type of audit events are recorded.

application
Customizable properties to control which application-specific audit events are recorded.

Authorization Service

The Authorization service provides the general authorization system.

sas.authorization
The set of configuration properties for the Authorization service.

maxAncestryCacheSize
Specifies the maximum number of ancestors to cache per object. The default value is `1000`. The cache enhances performance in container-based inheritance.

reshareEnabled
Specifies whether resharing is enabled. When set to `Off`, prevents downstream sharing by users who lack Secure access to the object that is being shared. The default value is `On`. See "Sharing: Details for Administrators" on page 437.

Note: If sharingEnabled is `Off`, setting this property has no effect.

Note: You cannot enable or disable downstream sharing on a per-object or per-identity basis. The ability to perform downstream sharing is controlled only by this deployment-wide configuration property.

sharingEnabled
Specifies whether sharing is enabled. When set to `Off`, prevents all sharing. The default value is `On`. See "Sharing: Details for Administrators" on page 437.

rules.executorThreads (a supplemental property.)
Specifies the number of threads that are available for bulk processing of authorization rules. The default value is `20`. Modify this value only if you are directed to do so by SAS Technical Support.

remote (a supplemental property.)
Disables enforcement in the general authorization system, if set to `false`. The default value is `true`. An administrator might temporarily disable authorization if rules that inadvertently prevent access are introduced. Do not disable authorization while the system is available to other users.

warnOnCycles (a supplemental property.)
Prevents cyclic warnings from being written to the authorization service log, if set to **`false`**. The default value is **`true`**. An administrator might temporarily set this property to **`false`** if cyclic warnings are causing the log to grow rapidly. After cyclic rules are corrected, set this property back to **`true`**.

Backup Service

The Backup service manages the backup and recovery of configuration information and user-created content in a SAS deployment.

sas.deploymentbackup
The set of configuration properties for the Backup service.

agentType
Important: This property is deprecated. Do not use this property.

The type of communication (messaging or SSH) that is used between the Backup service and the Backup agent.

jobTimeout
The number of minutes a backup job or a restore job is allowed to run before the job is marked 'Failed.'

retentionPeriod
The number of days that backups are stored before they are removed from the backup vault.

scheduledBackupAllowed
Important: This property is deprecated. Do not use this property.

Allows scheduled backups to run. In this release, the default value is false and cannot be changed.

sharedVault
A network location where all the backups are stored. This location should be accessible to the user identity installing deploymentBackup and backup-agent services.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

vaultLocation
Important: This property is deprecated. Do not use this property.

The location where all backups are stored. In a multi-machine deployment, the install user must have Read and Write permissions on every machine.

Cache Services

The Cache services provide other microservices the ability to distribute cached data across instances. The Cache services consist of both a Cache Locator and a Cache Server.

sas.cache.default

Important: The configuration definition, `sas.cache.default`, has been deprecated. Use instead, `sas.cache.config`.

sas.cache.config

The set of properties that provide customization for caching.

ackSevereAlertThreshold
: The number of seconds the distributed system waits after the ack-wait-threshold for an acknowledgment from a system member before sending a severe alert. A value `0` (zero) disables this feature.

ackWaitThreshold
: The number of seconds a distributed message waits for an acknowledgment from a system member before sending an alert.

conserveSockets
: Allows sockets to be shared by a system member's threads.

criticalHeapPercentage
: The percentage of Java old generation heap. When reached, prevents Write operations of cached data from occurring.

deployWorkingDir
: The working directory used when deploying JAR application files to distributed system members. This directory can be local and unique to the member, or it can be a shared resource.

disableAutoReconnect
: Disables the ability of a system member to reconnect and re-initialize after it has been forced out of the distributed system.

diskStoreDir
: The directory used when region data is overflowed to disk. This directory can be local and unique to the member or can be a shared resource.

diskStoreOpLogSize
: The maximum size for a region overflow file when region data is overflowed to disk. The value is specified in Megabytes.

diskStoreSize
: The maximum amount of data on disk when region data is overflowed. The value is specified in megabytes.

distributedCache
: Specifies that the cache should be distributed. Valid values are **`true`** or **`false`**. SAS Cache Server requires the value **`true`**'.

enableNetworkPartitionDetection
: Enables the distributed system to detect and handle splits in the distributed system. Splits are typically caused by a partitioning of the network (split brain) where the distributed system is running.

evictionHeapPercentage
: The percentage of Java old generation heap when reached, causes cached data to overflow to disk.

groups
: The list of groups that this system member belongs to. Use a comma to separate group names.

locatorDiscoveryAttempts
: The number of service discovery attempts allowed before a registered cache locator is found. A value of `0` (zero) allows for an unlimited number of attempts.

locatorWaitTime
: The number of seconds that a system member waits for a locator to join the distributed system.

logLevel
: Indicates the lowest diagnostic log level (TRACE, DEBUG, INFO, WARN, ERROR, and FATAL) that is processed. Log events whose levels are below the specified value are ignored.

maxConnections
: The maximum number of connections to pool when connected to a cache server.

membershipPortRange
: The port range used when selecting ephemeral ports for members of the distributed system. Values are `32768` to `61000`.

memberTimeout
: The number of milliseconds the distributed system waits before it determines that a system member has timed out.

mode
: The mode of operation to use when connecting to the cache servers. Valid values are **client** or **local**. SAS Cache Server requires **local**.

overflowEnabled
: Allows cached data to overflow to disk in low memory situations.

persistentEnabled
: Allows cached data to persist to disk.

pingInterval
: The ping interval for the cache client to check the availability of servers in milliseconds.

retryAttempts
: The number of retry attempts for operations if a time-out or exception occurs.

subscriptionEnabled
: Configures the client to register with the cache server for subscription events.

tcpPort
: The TCP port a member of the distributed system listens on for cache communications.

Cache Locator Service

The Cache Locator service provides discovery information to SAS Viya microservices for the purpose of forming a distributed data cache. SAS Cache Locator is based on the open-source Apache Geode project.

sas.cache.locator
: The set of properties that provide customization for the Cache Locator service.

host
The host where the service is deployed.

hostnameForClients
The external host name of the cache locator if different from the local bind address or host name.

port
The port registered for the cachelocator-listener.

retryCount
The number of attempts the service makes to register the cachelocator-listener.

retryPeriod
The amount of time between registration attempts for the cachelocator-listener.

timeout
The amount of time this service waits to start the locator process.

timeoutInterval
The amount of time between attempts checking for the start of the locator process.

Cache Server Service

The Cache Server service hosts long-lived data regions (a cache) and serves the contents to SAS Viya microservices. Like SAS Cache Locator, SAS Cache Server is based on the open-source Apache Geode project.

sas.cache.server

The set of properties that provide customization for the cache server.

autoStartup
: Specifies whether the cache server should be started automatically on start-up.

host
: The host where the service is deployed.

hostnameForClients
: The external host name of the cache server if different from the local bind address or host name.

maxTimeBetweenPings
: The maximum time in milliseconds between messages or a ping from a cache client.

port
: The port registered for the cache server.

CAS Management Service

The CAS Management service provides access to shared data for users and applications. The service also provides information about the SAS Viya system for operations such as monitoring and auditing.

sas.casmanagement

The set of properties that are used to configure private settings for the CAS Management service.

supplementalProperties
: The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

sas.casmanagement.global

The set of properties that are used to configure global settings for the CAS Management service.

The sas.casmanagement.global configuration instance applies to all SAS Viya servers and services (global).

- The set of properties used by applications to access shared data and analytics, such as map data.

  application.casServer
  : The name of the CAS server used for application work.

  application.caslib
  : The name of the caslib used for application data.

- The set of properties used to identify the default CAS server for users.

  default.casServer
  : The name of the default CAS server for users.

  default.caslib
  : The name of the default caslib for users.

- The set of properties used by the system for data produced during normal operation, such as audit records and monitoring data.

  system.casServer
  : The name of the system CAS server.

  system.caslib
  : The name of the system caslib.

Compute Service

The Compute service enables clients to submit SAS programs and stored procedures in the form of jobs for processing. The SAS Compute Server implements the Compute service. See “SAS Compute Server and Compute Service” on page 562.

sas.compute

The set of properties used to configure the compute and related servers.

domain.default
The default authentication domain to use for looking up host credentials.

kerberos.enabled
Authenticate to compute servers using Kerberos.

serviceAccount.default:
The default service account that should be used to run jobs on the host.

Notifications Service

The Notifications service stores and retrieves system and application event notifications. There are two types of notifications, subscription-based (also known as opt-in notifications) and directed (explicit notifications).

You can manage subscription-based notifications using the Notifications service. Directed notifications are managed by the SAS applications that use them (such as SAS Workflow Manager).

sas.notifications

The set of properties used to control subscription-based (opt-in) notifications.

enabled
Enables opt-in notifications on all channels.

Important: Turning off `sas.notifications.enabled`, turns off all subscription-based notification channels, such as notifications in SAS Viya web applications’ banner and email notifications.

sas.notifications.delivery

The set of properties used to control subscription-based (opt-in) notifications on specific delivery channels.

mail.enabled
Enables opt-in notification delivery to the email channel.

Important: To enable subscription-based notifications to email, `sas.notifications.enabled` must be on (TRUE).

CAS Proxy Service

The set of configuration properties for the CAS Proxy service.

jobExecutionProvider

Configurable values for the CAS language (CASL) job execution provider job.

caslJESExpiresAfter
The amount of time after job completion before the job execution service deletes the job. Specify time in W3C XML duration format (for example, PT5M = 5 minutes). A null value indicates that job is not deleted.

caslJESHeartbeatInterval
The heartbeat value (in seconds) to use with a CASL job execution provider job. A zero or negative value indicates that the heartbeat is not checked.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

Collections Service

The Collections service enables access to personal and shared collections.

sas.collections
The set of configuration properties for the Collections service.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

Configuration Service

The Configuration service manages changes to the configuration of services. See “Configuration Properties: Concepts” on page 83.

sas.configuration
The set of configuration properties for the Configuration service.

forceWrite.enabled
Enables writing to the persistence layer for every operation even when that operation made no changes.

locking.enabled
Enables locking between multiple instances of the SAS Configuration Service. Locking must be enabled when more than one SAS Configuration Service instance is present in the deployment.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

Cross Domain Proxy Service

The Cross Domain Proxy service provides access to external web resources over HTTP.

sas.crossdomainproxy

The set of configuration properties for the Cross Domain Proxy service.

allowedDomains
The list of domains (a whitelist) that the cross-domain proxy is allowed to access. The value is a Java regular expression. Use the Or character (|) to separate multiple domains (for example, `https?://*\.sas\.com(:\d+)?/|https?://*\.foo\.bar\.org/`).

> **TIP** SAS recommends that you escape dot (.) characters in regular expressions with a slash (\) (for example, *\.sas\.com). Also, add a trailing forward slash (/) with each domain (for example, *\\.sas\.com/).

allowSystemDomains
Enables the list of trusted domains (if any) required by SAS. If the cross-domain proxy is denied access to one or more of these domains, certain SAS features are disabled. (This list of trusted domains is displayed in the property description in SAS Environment Manager.)

sas.crossdomainproxy.system

The set of system configuration properties for the Cross Domain Proxy service.

excludeRequestHeaders
The list of header names (a blacklist) which the cross-domain proxy excludes from requests sent to a destination URL. The value is a Java regular expression. Use the Or character (|) to separate multiple header names (for example, cookie|Authorization).

maxPooledConnectionsPerRoute
The maximum number of pooled connections per route. (This value must be a positive integer.)

maxPooledConnections
The maximum number of total pooled connections. (This value must be a positive integer.)

connectionTimeoutInMinutes
The number of minutes allowed before the connection to the HTTP client times out. A value of zero (0) specifies no time-out.

Device Management Service

The Device Management service provides the means to maintain the server’s device blacklist and whitelist tables, including controlling which security model is in place. See *SAS Viya Administration: Mobile*.

sas.devicemanagement
The set of configuration properties for the Device Management service.

offlineLimitDays
The number of days before the mobile application goes off-line.

passcodeAttempts
The number of passcode attempts before the user is locked out of the mobile application.

passcodeTimeoutMinutes
The number of minutes before the passcode expires on the mobile application.

whitelistSupportEnabled
Enables whitelist support for mobile device security on the server.

Identities Service

The Identities service retrieves information about identities (users or groups) from your identity provider. It also enables the creation and management of custom groups. For detailed information about this functionality, see *SAS Viya Administration: Identity Management*. Here are the configuration properties for the Identities service:

sas.identities

The set of properties that are used to configure global settings for the Identities service.

cache.enabled
Enables identities information to be cached. Caching is enabled by default.

cache.providerPageLimit
The number of identities to process in a given request when loading the cache. The default value is `1000`.

cache.cacheRefreshInterval
The refresh interval for the identities cache.

Note: Do not set **cache.cacheRefreshInterval** below 20 minutes. Doing so might have a significant impact on your overall system, especially on the LDAP and SAS Infrastructure Data (PostgreSQL) servers.

Use the following conventions to specify the unit of time for the refresh interval:

- d - refers to days (for example, 6d).
- h - refers to hours (for example, 6h).
- m - refers to minutes (for example, 6m).
- s - refers to seconds (for example, 6s).
- ms - refers to milliseconds (for example, 6ms).

defaultProvider
The default provider. The default value is `local`. (For this release, SAS recommends that you do not change this value.)

sas.identities.providers.ldap

The set of properties that are used to configure your LDAP provider.

Important: The set of properties for `sas.identities.providers.ldap` are global settings that apply to all LDAP configurations. For multi-tenancy, these global settings apply to the provider and all tenants.

membershipCacheRefreshInterval
Specifies the interval that is used to refresh the membership cache for the LDAP provider. The default value is `6h`.

pagedResults
Enables the LDAP server to use pagination when processing requests. Pagination is enabled by default.

pageSize
The number of identity requests per page to be processed by the LDAP server (if pagination is enabled). The default value is `500`.

primaryGroupMembershipEnabled
Enables processing of primary group memberships (used only for posixGroup schema).

sas.identities.providers.ldap.connection

The set of properties that are used to configure your LDAP provider.

anonymousBind
Defines whether Read-Only operations are performed using an anonymous (unauthenticated) context.

customEnvironmentProperties
The set of user-added, advanced properties for configuring the LDAP client environment.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

host
The host name of the LDAP server to connect to.

password
: The password for logging on to the LDAP server. If **anonymousBind** is enabled, specify a value of `none`.

pool.enabled
: Enables pooling of LDAP connections. Pooling is enabled by default.

pool.evictionTimePeriodMillis
: The number of milliseconds that the idle-object evictor thread sleeps between runs. If the value is non-positive, the idle object evictor thread does not run. The default value is `240000`.

pool.idleTimeMillis
: The minimum amount of time in milliseconds that objects can sit idle in the pool before becoming eligible for eviction by the idle-object evictor, if present. The default value is `480000`.

pool.maxActive
: The maximum number of active connections of a given type (either Read-Only or Read-Write) that can be allocated from the pool at the same time. The default value is `8`.

pool.maxIdle
: The maximum number of active connections of a given type (either Read-Only or Read-Write) that can remain idle in the pool without extra connections being released. For no limit, specify a non-positive value. The default value is `8`.

pool.maxSize
: The maximum number of active connections of all types that can be allocated from the pool at the same time. For no limit, specify a non-positive value. The default value is `-1`.

pool.maxWait
: The maximum amount of time in milliseconds that the pool waits for a connection to be returned before throwing an exception. For no limit, specify a non-positive value.

pool.minIdle
: The minimum number of active connections of a given type (either Read-Only or Read-Write) that can remain idle in the pool without extra connections being created. To create no extra connections, specify zero. The default value is `0`.

pool.testOnBorrow
: Enables validation of objects before they are borrowed from the pool. If an object fails validation, it is dropped from the pool and an attempt is made to borrow another object. This option is enabled by default.

pool.testOnReturn
: Enables validation of objects before they are returned to the pool. This option is disabled by default.

pool.testWhileIdle
: Enables validation of objects by the idle-object evictor, if present. If an object fails validation, it is dropped from the pool. This option is enabled by default.

pool.whenExhaustedAction
: An integer that indicates the behavior when the pool is exhausted. Valid values are: `0` (fail), `1` (block), or `2` (grow).

port
: The port for connecting to LDAP.

  Note: When connecting via LDAP, port values are set to 389. When connecting via Lightweight Directory Access Protocol over TLS (LDAPS), port values are set to 636.

startTLS.mode
: When using StartTLS, which mode to enable. The possible values are none (default) or simple.

url
: The URL for connecting to LDAP.

  The default is: `url: ldap://${sas.identities.providers.ldap.connection.host}:${sas.identities.providers.ldap.connection.port}`

When the host and port properties have been specified, the `url` must be changed if you are connecting via the LDAPS protocol.

userDN

The distinguished name (DN) of the user account for logging on to the LDAP server (for example, `cn=AdminUser,dn=example,dn=com`). If **anonymousBind** is enabled, specify a value of **`none`**.

sas.identities.providers.ldap.group (Field Mappings)

The set of properties that are used to configure the mapping of group-level fields in your LDAP provider to group-level fields in SAS. For each of the following SAS fields, you specify the corresponding field in your LDAP provider. The default values are valid for most implementations of Microsoft Active Directory. For other LDAP providers, you must provide different values for some fields.

Property	Description	Default (valid for most implementations of Microsoft Active Directory)
accountId	The field in the LDAP provider that is used to populate the group's ID.	sAMAccountName
createdDate	The field in the LDAP provider that is used to populate the group's account created date.	whenCreated
description	The field in the LDAP provider that is used to populate the group's description.	description
member	The field in the LDAP provider that is used to populate the members of the group.	member
memberOf	The field in the LDAP provider that is used to populate the groups that this group is a member of. Set `memberOf` to **`none`** if the LDAP group attribute is not a fully qualified DN value.	memberOf
modifiedDate	The field in the LDAP provider that is used to populate the date on which the group's account was last modified.	whenChanged
name	The field in the LDAP provider that is used to populate the group's name.	displayName
objectClass	The object class value to use when loading groups.	group

sas.identities.providers.ldap.group (Additional Properties)

The set of properties that are used to configure information for retrieving group information from your LDAP provider.

Note: The Identities service does not process referrals.

baseDN

The point from which the LDAP server searches for groups (for example, `ou=groups,dc=example,dn=com`).

distinguishedName

The field in the LDAP provider that is used to populate the group's distinguished name value.

Note: If your LDAP server does not support an explicit distinguished name attribute (for example, OpenLDAP), you must set this property to `none`.

objectFilter
: The filter for customizing results that are returned when groups are queried [for example, (objectClass=group)].

You can create a custom filter to exclude identities whose accounts are disabled or expired, or to exclude objects that represent computer resources rather than actual groups. If you have a large number of groups, using a custom filter can improve performance and reduce memory requirements. In addition, user management tasks can be performed more efficiently if only relevant identities are listed in SAS Environment Manager.

searchFilter
: The filter that is used to find a group account. The default filter is `${sas.identities.providers.ldap.group.accountId}={0}`.

sas.identities.providers.ldap.user (Field Mappings)

The following properties specify the mapping of user-level fields in your LDAP provider to user-level fields in SAS. For each of the following SAS fields, you specify the corresponding field in your LDAP provider. The default values are valid for most implementations of Microsoft Active Directory. For other LDAP providers, you must provide different values for some fields.

Property	Description	Default (valid for most implementations of Microsoft Active Directory)
accountId	The field in the LDAP provider that is used to populate the user’s ID.	sAMAccountName
address.country	The field in the LDAP provider that is used to populate the user’s country.	co
address.locality	The field in the LDAP provider that is used to populate the user’s city.	l
address.postalCode	The field in the LDAP provider that is used to populate the user’s postal code.	postalCode
address.region	The field in the LDAP provider that is used to populate the user’s region or state.	region
address.street	The field in the LDAP provider that is used to populate the user’s street address.	street
createdDate	The field in the LDAP provider that is used to populate the user’s account created date.	whenCreated
description	The field in the LDAP provider that is used to populate the user’s description.	description
emailAddress.other	The field in the LDAP provider that is used to populate the user’s alternate email address.	otherMailbox
emailAddress.work	The field in the LDAP provider that is used to populate the user’s work email address.	mail

Property	Description	Default (valid for most implementations of Microsoft Active Directory)
emailAddress.sms	The field in the LDAP provider that is used to populate the user's SMS email address.	
memberOf	The field in the LDAP provider that is used to populate the groups that this user is a member of.	memberOf
modifiedDate	The field in the LDAP provider that is used to populate the date on which the user's account was last modified.	whenChanged
name	The field in the LDAP provider that is used to populate the user's name.	displayName
objectClass	The type of user objects that are being searched for.	organizationalPerson
phone.business	The field in the LDAP provider that is used to populate the user's work phone number.	telephoneNumber
phone.businessFax	The field in the LDAP provider that is used to populate the user's work fax number.	facsimileTelephoneNumber
phone.home	The field in the LDAP provider that is used to populate the user's home phone number.	homePhone
phone.mobile	The field in the LDAP provider that is used to populate the user's mobile phone number.	mobile
phone.pager	The field in the LDAP provider that is used to populate the user's pager number.	pager
title	The field in the LDAP provider that is used to populate the user's title.	title

sas.identities.providers.ldap.tenancy

The set of specific, multi-tenancy properties that can be used when implementing an LDAP provider.

Important: The purpose of `sas.identities.providers.ldap.tenancy` is for a default tenant configuration for a single, shared LDAP server. These properties are not used when customizing the structures for each tenant.

groupRdn
: The relative distinguished name group (RDN) value.

tenantKey
: The default value (OU) for tenantKey.

userRdn
: The relative distinguished name user (RDN) value.

sas.identities.providers.ldap.user (Other Properties)

The set of properties that are used to configure additional information for retrieving user information from your LDAP provider.

Note: The Identities service does not process referrals.

baseDN
The point from which the LDAP server searches for users.

distinguishedName
The field in the LDAP provider that is used to populate the user's distinguished name value.

Note: If your LDAP server does not support an explicit distinguished name attribute (for example, OpenLDAP), you must set this property to `none`.

objectFilter
The filter for customizing results that are returned when querying users.

You can create a custom filter to exclude identities whose accounts are disabled or expired, or to exclude objects that represent computer resources rather than actual users. If you have a large number of users, using a custom filter can improve performance and reduce memory requirements. In addition, user management tasks can be performed more efficiently if only relevant identities are listed in SAS Environment Manager.

Here is an example of a filter that excludes identities that represent computers and identities that are inactive. This filter is compatible with Microsoft Active Directory.

```
(&(objectCategory=person)(objectClass=user)(!
(userAccountControl:1.2.840.113556.1.4.803:=2)))
```

For OpenLDAP, the filter `(objectclass=person)` excludes identities that represent resources other than users.

searchFilter
The filter used for locating a user account in the LDAP provider so that the user can make a connection using an ID and password.

The default filter is `${sas.identities.providers.ldap.user.accountId}={0}`.

Mail Service

The Mail service provides a client the ability to send email to a configured SMTP server using a REST API.

For more information, see "Configure the Connection to the Mail Service" in *SAS Viya for Linux: Deployment Guide*.

sas.mail
The set of configuration properties for the Mail service.

allowAllSenders
Provides the ability to override restriction on the 'from' mail address allowed to send mail.

fromAddress
Default 'from' mail address to use when mail is sent directly from a service. (The default is `noreplies@company.com`.)

fromPersonalName
Default personal name to use when mail is sent directly from a service. (The default is `Service`.)

host
The mail server host (machine).

password
The optional password for connecting to the mail server.

port
The mail server port. (The default is 25.)

properties
Optional properties set on the remote mail server.

sizeLimit
The maximum size of mail sent to the configured mail server (in megabytes).

username
The optional user name for connecting to the mail server.

Maps Service

The Maps service returns polygon information for selected identifiers from a given table.

SAS Viya supports several third-party map services. Whether the map service uses HTTP or secured HTTP (HTTPS), depends on the following:

- Open Street Map

  Controlled by the `defaultOSMCommunicationProtocol` configuration property.

- ArcGIS Online

  Controlled by SAS Visual Analytics, which always uses secure HTTP (HTTPS).

- local Esri server

  Controlled by the URL that is entered for the server *protocol*://*host-name*:*port*/*path*. The protocol that is entered is used.

  Important: SAS Viya currently supports only token-based authentication for Esri. For example, an Esri server configured for Integrated Windows Authentication (IWA) is incompatible with SAS Viya.

sas.maps
The set of configuration properties for the Maps service.

defaultOSMCommunicationProtocol
The protocol (HTTP, HTTPS) that is used for the default Open Street Map servers.

localEsriServicesRequiresAuthentication
Indicates that the local Esri map services URL requires an authentication token for access.

localEsriServicesUrl
The URL to the local Esri map services. The URL consists of a protocol, host, port, and path (for example, http://myserver:6080/arcgis/rest/services/).

Note: If your on-premises Esri servers use a different network domain than your SAS Viya system, then you must add the necessary map URLs to the whitelist of domains that the cross-domain proxy is allowed to access. For more information, see “allowedDomains” on page 99.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

useArcGISOnlineMaps
Enable access to background maps from the Esri ArcGIS Online catalog.

The set of custom configuration properties for Open Street Map server settings.

customOSM.maxResolution
The maximum resolution in meters per pixel of each map tile. The value must be a decimal number.

customOSM.numResolutions
The number of tile levels configured on the tile servers. The value must be a positive integer.

customOSM.servers
A comma-separated list of servers, with paths to tiles (for example, http://myhost1.myorg.com/tiles/, http://myhost2.myorg.com/tiles/).

Monitoring Service

The monitoring service provides information about the machines and services in your environment. See *SAS Viya Administration: Monitoring*.

sas.monitoring
The set of configuration properties for the Monitoring service.

Report Data Service

The Report Data service retrieves data from reports.

sas.reportdata.system
The set of system configuration properties for the Report Data service.

Note: In a multi-tenant configuration, sas.reportdata.system properties apply to all tenants.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

executorExpirationIntervalMinutes
The number of minutes before inactive data executor sessions are closed if they have no active queries.

executorForceExpirationIntervalMinutes
The number of minutes before inactive data executor sessions are forced closed even if they have active queries.

packageResultFileTimeToIdleSeconds
The number of seconds allowed for a client to retrieve package result data files before the files are removed from the cache.

resultCacheErrorExpirationSeconds
The number of seconds before the error cases for a report result are removed from the cache.

resultCacheTimeToLiveSeconds
The number of seconds before a report result is removed from the cache.

tempCacheTimeToIdleSeconds
The number of seconds allowed for a client to retrieve temporary result data files before the files are removed from the cache.

xmlParserPoolSize
The number of XML parsers to be instantiated during application start-up.

sas.reportdata.properties
The set of configuration properties for the Report Data service.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

comparisonEpsilon
The number in E notation that is the variability allowed when comparing floating point numbers for equality.

decisionTreePredictorCardinalityLimit
The maximum cardinality of independent variables allowed to run a decision tree.

decisionTreeResponseCardinalityLimit
The maximum cardinality of a dependent variable allowed to run a decision tree.

defaultInteractiveDrillDepth
The number of interactive drill levels included in the offline data for report viewers.

defaultMaxCellsProduced
The maximum number of data cells delivered for each query result to report viewers.

enableResultCache
Enable report result caching.

exportExcelRowLimit
The maximum number of rows allowed for export files formatted for Excel.

exportExcelColumnLimit
The maximum number of columns allowed for export files formatted for Excel.

exportTSVandCSVRowLimit
The maximum number of rows allowed for tab- and comma-separated export files.

exportTSVandCSVColumnLimit
The maximum number of columns allowed for tab- and comma-separated export files.

ignoreMissingValuesInCountDistinct
Ignore missing values in count distinct.

maxTiesToIncludeOnRank
The maximum number of ties allowed for a rank.

modelingClassCardinalityLimit
The maximum number of class values allowed to run on fit models.

modelingGroupByCardinalityLimit
The maximum number of group by values allowed to run on fit models.

socketTimeoutLiveCancellableMillis
The number of milliseconds allowed for executing live, cancelable data queries.

socketTimeoutLiveNonCancellableMillis
The number of milliseconds allowed for executing live, non-cancellable data queries.

socketTimeoutSubscribeMillis
The number of milliseconds allowed for executing subscribe data queries.

A map of the maximum result rows values for the supported visual types.

maxRowsLookup.bubble
The maximum result rows for a bubble visual.

maxRowsLookup.buttonBar
The maximum result rows for a button bar visual.

maxRowsLookup.crossTab
The maximum result rows for a multidimensional table visual.

maxRowsLookup.customContent
The maximum result rows for a custom content.

maxRowsLookup.dropdown
The maximum result rows for a drop-down control.

maxRowsLookup.dualAxisTimeSeries
The maximum result rows for a dual axis time series visual.

maxRowsLookup.geoBubble
The maximum result rows for a geo bubble visual.

maxRowsLookup.geoContour
The maximum number of result rows for a geo contour visual.

maxRowsLookup.geoHeatmap
The maximum result rows for a geo heat map visual.

maxRowsLookup.geoRegion
The maximum result rows for a geo region visual.

maxRowsLookup.geoScatter
The maximum result rows for a geo scatter visual.

maxRowsLookup.graphDefault
The maximum result rows for a default graph visual.

maxRowsLookup.heatbox
The maximum result rows for a heat box visual.

maxRowsLookup.heatmap
The maximum result rows for a heat map visual.

maxRowsLookup.kpi
The maximum result rows for a kpi visual.

maxRowsLookup.list
The maximum result rows for a list visual.

maxRowsLookup.listTable
The maximum result rows for a list table visual.

maxRowsLookup.scatter
The maximum result rows for a scatter visual.

maxRowsLookup.textInput
The maximum result rows for a text input control.

maxRowsLookup.timeSeries
The maximum result rows for a time series visual.

maxRowsLookup.treeMap
The maximum result rows for a treemap visual.

maxRowsLookup.wordCloud
The maximum result rows for a word cloud visual.

sas.reportdata.debug
The set of debug configuration properties for the Report Data service.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

Report Packages Service

The Report Packages service executes reports to generate corresponding *report packages*. A report package includes the report.xml, CSS style sheets, images, CSV data files, and so on, that are required to render the report.

sas.reportpackages.system
The set of system configuration properties for the Report Packages service.

Note: In a multi-tenant configuration, sas.reportpackages.system properties apply to all tenants.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

backgroundThreadMonitorSecs
The frequency in seconds at which the background thread monitor runs. A value of zero indicates that the monitor is disabled.

packageExpirationTime
The amount of time in seconds after which the report package expires from the cache.

useProxyServiceForExternalImages
Enable the Cross Domain Proxy service to retrieve the external images in the report.

xmlParserPoolSize
The number of XML parsers to be instantiated during application start-up.

sas.reportpackages.properties
The set of configuration properties for the Report Packages service.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

highContrastTheme
The name of the theme to be used when reports are displayed in high contrast.

imageDefaultMaxBytes
The maximum number of bytes for an image. Larger images are scaled down, unless 'noscale' is specified in the report.

subscribeConcurrentRequestLimit
The maximum number of report packages that can be generated concurrently per user.

subscribeConcurrentRequestLimitGuest
The maximum number of report packages that can be generated concurrently for the Guest user.

sas.reportpackages.debug
The set of debug configuration properties for the Report Packages service.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

Report Renderer Service

The Report Renderer service creates PDF documents from report packages.

sas.reportrenderer.system
The set of system configuration properties for the Report Renderer service.

Note: In a multi-tenant configuration, sas.reportrenderer.system properties apply to all tenants.

cacheDuration
The number of seconds allowed before rendered reports are deleted from the cache.

workingDirectory
Override the working directory used for building rendered reports.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

sas.reportrenderer.properties
The set of configuration properties for the Report Renderer service.

timeoutMillis
The number of milliseconds allowed before the rendering process times out.

footerContentFormatted
The HTML formatted footer to be included on each PDF rendered page.

webContentRendererLink
The URL to a service that is set up to provide images for printing of web content.

Note: In order for web content to be included when you print a report, you must set up an application to provide images for the web content. A third-party tool such as Rendertron (https://github.com/GoogleChrome/rendertron) is one option that can be installed and configured to generate such images dynamically using Headless Chrome.

The Report Renderer service supports URLs with the following substitution tokens: `__WebContentURL__`, `__ImageWidth__`, and `__ImageHeight__`.

`_WebContentURL_` is a token that the Report Renderer service substitutes with the web content URL found in the report. The Report Renderer service substitutes `__ImageWidth__`, and `__ImageHeight__` with the requested image size, as allocated by the report layout.

Note: `webContentRendererLink` supports only standard web content (that is, content that does not require data or authentication). For this reason, data-driven content and SAS stored processes are not supported and cannot be rendered when printing reports in this release.

You must add this domain to the cross-domain proxy whitelist using the `sas.crossdomainproxy.allowedDomains` property. For more information, see "sas.crossdomainproxy".

Here is an example:

```
http://my-server.example.com:3000/screenshot/__WebContentURL__?
width=__ImageWidth_&height=__ImageHeight__
```

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

sas.reportrenderer.debug
The set of debug configuration properties for the Report Renderer service.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

Secret Manager Service

The Secret Manager service manages certificates generated by and stored by HashiCorp Vault. Vault provides a secure interface to secrets. These connections are secured by a Public Key Infrastructure (PKI) based on HashiCorp Vault, which is configured by SAS. The certificates are all signed by a Vault-generated root CA and intermediate certificate.

sas.vault

The set of configuration properties used to configure SAS Secret Manager (Vault).

Important: When modifying `sas.vault.*_ttl` properties, you must not violate precedence rules or SAS Secret Manager does not start. For more information, see "TTL Precedence Rules" on page 602.

certificate_role
Role of the certificates.

certificate_role_allow_any_name
The common name represents the name protected by the TLS certificate. Any name is allowed for the common name. The certificate is valid only if the requested host name matches the certificate common name. It consists of a single host name (for a single-name certificate) or a wildcard name (for a wildcard certificate, *.example.com).

certificate_role_key_bits
The key-length (in bits) of the certificate generated from Vault.

certificate_role_key_type
The certificate encryption algorithm. RSA or Elliptic-Curve (EC) can be specified.

certificate_role_key_usage
A comma-separated list that defines the allowed uses of a generated certificates. You can specify one or all of the following:

```
DigitalSignature,KeyAgreement,KeyEncipherment
```

certificate_role_max_ttl
The maximum length of time in hours that a Vault-issued certificate lasts before expiring. Must be greater than the **`certificate_role_ttl`**.

certificate_role_ttl
The default length of time that a Vault-issued certificate lasts before expiring.

intermediate_ca
An intermediate certificate is a subordinate certificate issued by the trusted root specifically to issue end-entity server certificates.

intermediate_ca_common_name
The common name (CN) for the Vault-issued intermediate certificate authority (CA) certificate. For SAS Viya, the name is SAS Viya intermediate CA. The CN identifies the host name associated with the certificate.

intermediate_ca_desc
The name for the Vault-issued intermediate CA certificates and SAS Viya Intermediate CA.

intermediate_ca_key_bits
The key-length (in bits) of the intermediate certificate generated from Vault.

intermediate_ca_max_ttl
The maximum length of time in hours that a Vault-issued intermediate certificate lasts before expiring. Must be greater than the **`intermediate_ca_ttl`**.

intermediate_ca_ttl
The default length of time that a Vault-issued intermediate certificate lasts before expiring.

root_ca
A public key certificate that identifies a root CA. A root certificate is the top-most certificate of the tree. The private key is used to sign other certificates.

root_ca_common_name
The common name for the Vault-issued root CA certificate and SAS Viya Root CA. The CN identifies the host name associated with the certificate.

root_ca_desc
The name for the Vault-issued Root CA and SAS Viya Root CA.

root_ca_key_bits
The key-length (in bits) of the root certificate generated from Vault.

root_ca_max_ttl
The maximum length of time in hours that a Vault-issued root CA certificate lasts before expiring. Must be greater than the **`root_ca_ttl`** value.

root_ca_ttl
The default length of time that a Vault-issued root CA certificate lasts before expiring.

systems
The time in hours that secrets and tokens are managed.

system_max_lease_ttl
The maximum amount of time (in hours) that Vault-issued secrets and tokens are valid. This value must be larger than the **`vault_token_default_lease_ttl`** value for the token configuration instance.

tokens
The time in hours that secrets and tokens are valid.

vault_token_default_lease_ttl
The default length of time (in hours) that Vault-issued tokens are valid.

Note: Changes to this value take effect after running the Ansible renewal playbook.

Configuration Properties: Reference (Applications)

SAS Data Explorer

SAS Data Explorer enables you to discover data and copy it to a CAS server.

sas.dataexplorer
The set of configuration properties for SAS Data Explorer.

casSessionImportNumNodes
The number of nodes on which to start CAS sessions when running import jobs. A value of `0` means all nodes.

casSessionInteractiveNumNodes
The number of nodes on which to start CAS sessions used for browsing data sources. A value of `0` means all nodes.

filterAvailableTab
Turn the toggle switch on to automatically populate the **Available** and **Data Sources** tabs with tables only from the default caslib for your site. The default is `off`.

By default, the **Available** and **Data Sources** tabs display all tables that have been loaded to memory, from any CAS server to which you have access. If some nodes are slow to respond to queries from the **Available** and **Data Sources** tabs, your browser might freeze while it is waiting for a response. If slow performance persists, an administrator can set an option so that the **Available** and **Data Sources** tabs are automatically populated with tables only from the default caslib for your site.

jobExecutionProvider
Configurable values for the SAS Data Explorer job execution provider job:

- **heartbeatInterval**. Default is `300`. The maximum number of items that can be submitted in a single import.
- **scheduledJobExpiresAfter**. Not set by default. The amount of time after a scheduled job completion before the job execution service deletes the job. Specify time in W3C XML duration format (for example, PT5M = 5 minutes). A null value indicates that job is not deleted.

maxImportQueueSize
The maximum number of items that can be imported with the **Import All** option on the **Import** tab. The default is `100`F.

availableTabEntryTimeoutMS
Specifies how long (in seconds) Data Explorer waits for a given library to report back its list of available tables. The default value is 10 seconds. If the request does not come back within the allotted time, the library is not included in the available tables list. This setting affects the list of available tables in the **Available** tab and the `Data Sources` tab.

supplementalProperties
The set of user-added, advanced properties.

Note: Do not enter a property name-value pair unless you are directed to do so by SAS Technical Support.

SAS Data Studio

SAS Data Studio provides a way for you to prepare data, including data transformations.

sas.datastudio
The set of configuration properties for SAS Data Studio.

casSessionNumNodes
The number of nodes on which to start a SAS Data Studio CAS session (0 means all nodes).

casSessionTimeout
The number of seconds to set as the CAS session time-out for sessions that SAS Data Studio creates.

interactiveJobExpiresAfter
The amount of time after an interactive data plan job completes before the job execution service deletes the job. Specify time in W3C XML duration format (for example, PT5M = 5 minutes). A null value indicates that job is not deleted.

saveTableJobExpiresAfter
The amount of time after a data plan job, which saves a table, completes before the job execution service deletes the job. Specify time in W3C XML duration format (for example, PT5M = 5 minutes). A null value indicates that job is not deleted.

SAS Infrastructure Data Server

SAS Infrastructure Data Server is based on PostgreSQL version 9 and is configured specifically to support SAS software. See “Overview” on page 602. Here is the list of SAS Infrastructure Data Server configuration definitions that consist of third-party PostgreSQL and pgpool-II configuration properties.

sas.dataserver.common

The set of properties that are common to a cluster (that is, both to pgpool-II and to PostgreSQL nodes).

For a list of the valid PostgreSQL property names (configuration parameters) and descriptions, see https://www.postgresql.org/docs/9.1/static/runtime-config.html.

sas.dataserver.conf

The set of properties that are used to set up the SAS Infrastructure Data Server node database configuration file, postgresql.conf.

For a list of the valid PostgreSQL property names (configuration parameters) and descriptions, see https://www.postgresql.org/docs/9.1/static/runtime-config.html.

sas.dataserver.hba

The set of properties that are used to set up the SAS Infrastructure Data Server node host-based authentication file, pg_hba.conf.

For a list of the valid PostgreSQL property names (authorization records) and descriptions, see https://www.postgresql.org/docs/9.1/static/auth-pg-hba-conf.html.

sas.dataserver.pool

The set of properties that are used to set up the pgpool-II node configuration file, pgpool.conf.

For a list of the valid pgpool-II property names (key-value pairs) and descriptions, see http://www.pgpool.net/docs/pgpool-II-3.5.4/doc/pgpool-en.html#pgpool_conf.

sas.dataserver.pool.hba

The set of properties that are used to set up the pgpool-II node host-based authentication file, pool_hba.conf.

For a list of the valid pgpool-II property names (key-value pairs) and descriptions, see http://www.pgpool.net/docs/pgpool-II-3.5.4/doc/pgpool-en.html#hba.

SAS Logon Manager

SAS Logon Manager provides OAuth2 and OpenID Connect services for authentication, and a user interface for sign-in, sign out, and other related functions. See “Authentication: Overview” on page 289. Here are the configuration properties for SAS Logon Manager:

sas.logon.callback

The set of properties that are used to configure the whitelist of URIs for trusted applications.

allowed.uris
: The comma-delimited list of URIs that users can be redirected to after signing in following a time-out or logoff.

sas.logon.custom

The set of properties that are used to provide custom content that is included on the Sign In to SAS page.

login
: The URI of the custom content included on the Sign In to SAS page.

logout
: The URI of the custom content included on the Sign In to SAS page when users sign out of the system.

timedout
: The URI of the custom content included on the time-out page.

sas.logon.groups

The set of properties that are used to customize lookup of group authorities.

assumable
: Specifies groups that have an elevated level of access that the user must approve at sign-in in order to assume those groups in a session.

approvalExpirySeconds
: When approving or denying access to a third-party application, specifies the number of seconds that the approval or denial should be remembered.

groupLookupRequired
: Requires groups to be determined for authentication to succeed.

recursion.enabled
: Allows recursive lookups of authorities for groups assigned to users.

requiresRecursion
: The comma-separated list of groups that require a recursive lookup to determine externally assigned authorities.

sas.logon.initial

The set of properties that are used to initially configure the system.

Note: Modifying one of these property values requires you to restart one or more SAS Viya services. For more information, see "General Servers and Services: Operate (Linux)" on page 458.

reset.enabled
: Displays a password reset link for the initial user account at start-up.

user
: The user name for the initial user account.

passwordResetLifetime
: The number of milliseconds for which that the password reset code is valid after restart.

redirectUri
: The URI to which the initial user should be redirected after resetting the password.

sas.logon.jasig.cas

The set of properties that are used to enable sign-ins using a Java in Administration Special Interest Group (Jasig) central authentication server (CAS) provider.

enabled
: Enable sign-ins using CAS.

single.signOn.enabled
Enable single sign-on.

single.signOut.enabled
Local sign-out should sign user out of the CAS server also.

single.signOut.logoutParameterName
The parameter name that identities a back channel single logoff request from the CAS server.

keepalive.enabled
Keep the CAS server session active by obtaining proxy tickets.

keepalive.proxyTicketUrl
The URL to obtain proxy tickets for keepalive purposes.

sasLogonUrl
The URL of the SAS Logon Manager (for example: https://sas.example.com/SASLogon).

casServerUrl
The URL for the CAS server.

service.authenticateAllArtifacts
Process all tickets, including proxy tickets.

service.sendRenew
Force the server to authenticate the user again, even if the user has previously authenticated.

service.artifactParameter
The request parameter to look for when attempting to see whether a ticket was sent from the server.

service.serviceParameter
The request parameter to send the server for this service.

validate.renew
Determine whether the ticket validation request should include a renew.

validate.encoding
Encoding of the ticket validation response from the server.

linkText
The hyperlink to display on the sign-in page.

showLinkText
Show the link text on the sign-in page.

autoLink
Automatically open the link to this provider when the login page is displayed.

sas.logon.jwt

The set of properties that are used to configure how JSON web tokens are issued.

signingKey
Either a Base64-encoded RSA private key that is used to digitally sign tokens, or a simple passphrase for HMACs. Enter a value only if you want to override the system-generated RSA private key.

issuer.uri
The URI of the application, for the issuer claim in tokens (for example, https://example.com/SASLogon).

claims.exclude
The comma-separated list of claims that should be excluded from the JSON web token.

policy.accessTokenValiditySeconds
The default number of seconds that access tokens are valid for after being issued in the default zone.

policy.refreshTokenValiditySeconds
The default number of seconds that refresh tokens are valid for after being issued in the default zone.

policy.global.accessTokenValiditySeconds
The default number of seconds that access tokens are valid for after being issued in all zones.

policy.global.refreshTokenValiditySeconds
The default number of seconds that refresh tokens are valid for after being issued in all zones.

refresh.restrictGrant
Grant refresh tokens only to clients with a scope of refresh_token for offline access.

sas.logon.kerberos

The set of properties that are used to enable sign-ins using Integrated Windows Authentication (IWA).

Note: Modifying one of these property values requires you to restart one or more SAS Viya services. For more information, see "General Servers and Services: Operate (Linux)" on page 458.

servicePrincipal
The name of the service principal in the keytab.

spn
The HTTP service principal name (SPN), if different than the principal name in the keytab.

keyTabLocation
The URL of the keytab file (for example, file:////opt/sas/viya/conf/etc/my_keytab).

stripRealmForGss
Removes the @... from the end of the user name.

holdOnToGSSContext
Enables Kerberos delegation to SAS Cloud Analytic Services.

debug
Enables the debug mode of the JAAS Kerberos login module.

disableDelegationWarning
Disables the warning message displayed to users if they are unable to perform Kerberos credential delegation from their browser to the SAS Logon Manager.

sas.logon.oauth.providers.external_oauth

The set of OAUTH provider properties that are used to enable sign-ins using an external provider. Modifying one of these property values requires you to restart SAS Logon Manager. For more information, see SAS Viya Administration in SAS Help Center.

authUrl
The URL to the authorization endpoint.

tokenUrl
The URL to the token endpoint.

tokenKey
The HMAC key or RSA public key used to sign tokens.

tokenKeyUrl
The URL to obtain the token key.

emailDomain
The comma-delimited list of possible email address domains of users that can sign on with this provider.

issuer
The principal that issued the token, as a case-sensitive string or URI.

linkText
The text that should be displayed on the sign-in page for this provider.

relyingPartyId
The client ID registered in the provider.

relyingPartySecret
The secret registered in the provider for the client ID.

scopes
The comma-delimited list of scopes for the authorization request.

addShadowUserOnLogin
Adds a local shadow user upon successful authentication.

showLinkText
Shows the link text on the sign-in page.

type
Either 'oidc1.0' or 'oauth2.0'.

attributeMappings.user_name
The attribute claim to use as the user name.

sas.logon.pam

The set of properties that are used to enable sign-ins using PAM.

enabled
Enables sign-in using PAM.

serviceName
The service name in the PAM configuration.

sas.logon.provider.guest

The set of properties that are used to configure guest access to the system.

Apply configuration only to this tenant (provider)
When this property is set to `off`, the configuration applies to all tenants, including the provider. Each tenant can override the configuration from within its own environment.

Note: **Apply configuration only to this tenant (provider)** is available only to provider tenants in a multi-tenant environment.

enabled
Enable anonymous guest access to web applications.

sas.logon.saml

The set of Security Assertion Markup Language (SAML) properties that are used to enable sign-ins using an external identity provider.

Note: Modifying one of these property values requires you to restart one or more SAS Viya services. For more information, see "General Servers and Services: Operate (Linux)" on page 458.

entityBaseURL
The URL of the application where SAML assertions are accepted, (for example: https://example.com/SASLogon).

setProxyParams
Allows the base URL to reside behind an HTTP proxy.

CAUTION! Do not modify setProxyParams. The value should remain off (false).

entityID
The entity ID of the service provider.

serviceProviderKey
The PEM-encoded, RSA private key that is used by the service provider.

serviceProviderKeyPassword
The password for the private key.

serviceProviderCertificate
The PEM-encoded, X.509 certificate that is used by the service provider.

wantAssertionSigned
Specifies that the assertions must be signed.

signatureAlgorithm
The algorithm for SAML signatures. The accepted values are SHA1, SHA256, and SHA512.

signMetaData
Specifies that the local service provider should sign metadata.

signRequest
Specifies that the local service provider should sign SAML requests.

socket.connectionManagerTimeout
The number of milliseconds before the connection pooling times out for HTTP requests for SAML metadata.

socket.soTimeout
The number of milliseconds before the read times out for HTTP requests for SAML metadata.

sas.logon.saml.providers.external_saml

The set of Security Assertion Markup Language (SAML) identity provider properties that are used to enable sign-ins using an external provider.

idpMetadata
The identity provider metadata or the URL to the metadata.

metadataTrustCheck
Specifies that the identity provider certificate must be trusted.

assertionConsumerIndex
The index of the assertion consumer service to use from identity provider metadata. The value must be a positive integer.

nameID
The default name ID format.

linkText
The hyperlink to display on the sign-in page.

addShadowUserOnLogin
Adds a local shadow user upon successful authentication. If set to `false`, users must be pre-created in the database to log on.

skipSslValidation
Skips Transport Layer Security (TLS) validation of the certificate.

showSamlLoginLink
Displays a link to the identity provider on the sign-in page.

sas.logon.sas9

The set of properties that are used to enable sign-ins using SAS 9.4 and later.

autoLink
: Automatically open the link to SAS 9 when the login page is displayed.

enabled
: Enable sign-ins using SAS 9 credentials.

linkText
: The hyperlink to display on the sign-in page.

sas9LogonUrl
: The URL of the SAS 9 Logon Manager (for example, https://sas9.sas.example.com/SASLogon).

showLinkText
: Show the link text on the sign-in page.

single.signOn.enabled
: Redirect to SAS 9 for single sign-on.

single.signOut.enabled
: Local sign-out should sign user out of SAS 9 also.

viyaLogonUrl
: The URL of the SAS Viya Logon Manager (for example, https://viya.sas.example.com/SASLogon).

sas.logon.sessions

The set of properties that are used to configure how concurrent sessions are handled.

maxConcurrentSessions
: The maximum number of allowed concurrent sessions. A value of -1 allows an unlimited number of sessions.

rejectNewSessionsIfMaxExceeded
: Rejects new sessions if the maximum number of sessions is exceeded. If false, when the maximum number of sessions is reached, an existing session is invalidated to allow a new one to be created.

sas.logon.tenancy

The set of properties that are used to configure multi-tenancy.

bootstrap.enabled
: Automatically configure identity zones and LDAP when tenants are onboarded or access policy is changed.

autoUpdateLdapConfiguration
: Automatically update all identity zones' LDAP configurations when the provider LDAP configuration is changed.

SAS Studio 5.x

SAS Studio is a development application for SAS that you access through your web browser. SAS Studio 5.x relies on the SAS Viya service layer.

Note: SAS Studio 5.x is installed with a *full deployment*.

sas.studiov
: The set of configuration properties for SAS Studio 5.x.

allowDownload
Allow users to download data.

allowUpload
Allow users to upload data.

longPollingHoldTimeSeconds
The maximum number of seconds to wait for a message from the client.

maxUploadSize
The maximum file size (bytes) allowed for upload.

SAS Studio 4.x

SAS Studio is a development application for SAS that you access through your web browser. SAS Studio 4.x relies on an embedded web application that is part of SAS Viya.

Note: SAS Studio 4.x is installed with a *programming-only deployment*.

For more information, see "Update SAS Studio Configuration Properties" on page 80.

Table A.1 SAS Studio: Configuration Properties

Property	Default Value	Description
sasstudio.appserver.https.keystorefile	(blank)	Specifies the keystore file to use for HTTPS.
sasstudio.appserver.https.keystorepass	(blank)	Specifies the keystore password to use for HTTPS.
sasstudio.appserver.https.port	38443	Specifies the port to use for HTTPS.
sasstudio.appserver.port	38080	Specifies the port to use for HTTP.
webdms.allowBackgroundSubmit	true	Specifies whether the **Background Submit** option is available when you right-click a .sas file in the navigation tree in the SAS Studio workspace.
webdms.allowFolderShortcuts	true	Specifies whether you can create folder shortcuts in the user interface.
webdms.batchSubmissionResultsRetentionPeriod	24	Specifies the number of hours to keep the output files from a background submission.
webdms.customPathRoot	(blank)	Specifies a path that determines the root node in the Folders tree. **Note:** You can use `<userid>` substitution for a directory path (for example, `/home/<userid>`).
webdms.defaultEncoding	UTF-8	Specifies the default SAS encoding method.
webdms.defaultVVN	ANY	Specifies the default value for the VALIDVARNAME option.

Property	Default Value	Description
webdms.globalSettings	`/opt/sas/spre/home/SASFoundation/GlobalStudioSettings`	Specifies the directory location for global XML files.
webdms.longPollingHoldTimeSeconds	30	Specifies the maximum number of seconds to wait for a message from the client.
webdms.maxNumActiveBatchSubmissions	3	Specifies the maximum number of active background jobs for the current SAS Studio user.
webdms.maxNumActiveBatchSubmissionsSystem	24	Specifies the maximum number of background jobs that can be submitted for a given instance of SAS Studio across all users.
webdms.maxSessionTimeoutInHours	1	Specifies the maximum number of hours a user can specify for the session time-out value in preferences.
webdms.maxUploadSize	10485760 (10MB)	Specifies the number of bytes allowed for file upload. **Note:** Large files can take a long time to load. If you have a large amount of content to upload, divide your content into smaller files if possible.
webdms.showSystemRoot	true	Specifies that the system root location be displayed in the Folders tree. **Note:** Set the value to false when the LOCKDOWN statement or the LOCKDOWN system option is used. For more information, see “References ” on page 589.
webdms.studioDataParentDirectory	(blank)	Specifies the location of SAS Studio preferences, snippets, my tasks, and more. This preference is specific to the local computer. The default value is blank. An administrator must mount a shared location to access data from any workspace server session. **Note:** You can use `<userid>` substitution for a directory path (for example, `/home/<userid>`).

Property	Default Value	Description
webdms.workspaceServer.allowGetRecordCount	true	Specifies whether to retrieve all of the rows for database tables. If you set this property to `false`, performance improves, but you might not see all rows of the table. For example, for large tables, total rows and filtered rows appear as Unavailable in the user interface. If the table has fewer than 100 rows or you scroll to the last page of the table, the values for the total rows and filtered rows are shown.
webdms.workspaceServer.cacheTableRow	true	Specifies whether to cache the rows from database tables to improve performance. If you use caching, the row count could be wrong if you modify the table. You must click **Refresh** to remove the value from the cache and to force a re-query of the database. If correct row count is more important than performance improvement, set this property to `false` to disable caching.
webdms.workspaceServer.hostName	localhost	Specifies the host to use to connect to the workspace server.
webdms.workspaceServer.largeTableRows	50,000	Specifies the maximum number of rows to display in the table viewer. If the number of rows in the table is unknown or greater than the value specified for the `webdms.workspaceServer.largeTableRows` property, the following behavior occurs: ■ SAS Studio displays a warning that sorting could take a long time. ■ SAS Studio does not generate a list of distinct values to select from when SAS Studio filters the data.
webdms.workspaceServer.port	8591	Specifies the port to use to connect to the workspace server.

Configuration Properties: Reference (System)

Commons REST Client

The following are properties to configure the commons REST client library, a library that all SAS Viya microservices incorporate.

sas.commons.rest.client

The set of configuration properties for the commons REST client library.

Bypass HTTP proxy
: Enables requests to be routed directly to the service rather than through the HTTP proxy. Changing **Bypass HTTP proxy** requires you to restart all SAS Viya services.

Java Virtual Machine (JVM)

The set of properties (Java options) that are used to configure the Java Virtual Machine when it is launched. Each JVM property defined in SAS Environment Manager corresponds to a single Java option.

To define service or global options for the JVM, follow the steps listed in "Create Configuration Instances" on page 75.

Note: Creating or modifying one of these property values requires you to restart one or more SAS Viya services. For more information, see "General Servers and Services: Operate (Linux)" on page 458.

When adding each JVM property, remember these guidelines:

- For the list of the valid Java options and descriptions, see http://docs.oracle.com/javase/6/docs/technotes/tools/windows/java.html.
- The property name for each Java option that you add must start with the string, java_option_ (for example, java_option_xmx).
- The property value is a single Java command-line option (for example, -Xmx512m).
- When the property names match, Java options specified at the service level override global Java options.
- Matching a Java option's property name (with a value consisting of a zero-length string) is the only way to disable Java option values.
- There is no control over the order that the JVM processes Java options.

Note: If you are using SAS Studio version 4.x, you must modify the JVM options in these configuration files:

- **`/opt/sas/viya/home/SASStudio/bin/appserver.sh`**
- `/opt/sas/viya/home/SASStudio/bin/sas.sasstudio.host`

Security

The following are properties to configure web security.

sas.commons.web.security

The set of properties that are used to configure web security.

content-security-policy
: The string used for the Content-Security-Policy HTTP header.

content-security-policy-enabled
: Sends the Content-Security-Policy header in HTTP responses to prevent injection attacks.

x-content-type-options
: The string used for the X-Content-Type-Options header for unsecured endpoints.

x-content-type-options-enabled
: Sends the X-Content-Type-Options header in HTTP responses for unsecured endpoints.

x-frame-options
: The string used for the X-Frame-Options HTTP header. A restart is required to pick up changes to this property.

x-frame-options-enabled

Sends the X-Frame-Options header in HTTP responses. A restart is required to pick up changes to this property.

x-xss-protection

The string used for the X-XSS-Protection header for unsecured endpoints.

x-xss-protection-enabled

Sends the X-XSS-Protection header in HTTP responses for unsecured endpoints.

sas.commons.web.security.cors

The set of properties that are used to configure Cross-Origin Resource Sharing (CORS) security. By default, CORS is enabled. For more information about CORS, see CORS support in Spring Framework.

allowCredentials

Allows credentials to be used in cross-origin requests. By default, this property is set to **On**.

allowedHeaders

The comma-separated list of HTTP headers that are allowed, by default, in cross-origin requests. Specify an asterisk ('*') to match any header.

allowedOrigins

The comma-separated list of origins that are allowed by default. The list can contain regular expressions. Specify an asterisk ('*') to match any origin.

allowedMethods

The comma-separated list of HTTP methods that are allowed, by default, in cross-origin requests. Specify an asterisk ('*') to match any method.

sas.commons.web.security.csrf

The set of properties that are used to configure Cross-Site Request Forgery (CSRF) security. By default, CSRF is enabled. To disable it, create a new configuration for the security definition. Specify the property name as **`enable-csrf`** and the value as **`false`**. For more information, see “Create Configuration Instances” on page 75.

SAS Viya protects against CSRF using the following:

- Synchronizer Tokens: Randomly generated tokens that are associated with the user’s current session. CSRF is checked only on requests with authenticated sessions, and is always skipped on GET, HEAD, TRACE, and OPTIONS requests. For more information, see Cross-Site Request Forgery (CSRF) Prevention Cheat Sheet.
- Header Checking: A filter that checks that the HTTP Referer header has the host and port of the requested URI or matches an optional whitelist of URIs that is configured as a comma-separated list in the *sas.commons.web.security.csrf.allowedUris* property.

For more information about CSRF, see Common application properties.

allowedReferers

This property is currently not supported and should be left blank.

allowedUris

The comma-separated list of referer URIs that are allowed by default. The list must contain regular expressions.

failIfNoHeaders

Blocks requests if both the Origin and Referer headers are absent.

sas.security

The set of configuration properties that are used to configure security for SAS Viya servers and services. The sas.security configuration instance applies to all SAS Viya servers and services (global).

network.databaseTraffic.enabled
    Toggle security for database traffic.

network.sasData.enabled
    Toggle security for other SAS information.

network.serverControl.enabled
    Toggle security for serverControl.

network.web.enabled
    Toggle security for web-based traffic.

Spring Boot Services

Here is the list of third-party, Spring Boot services that you can configure. For a list of the valid property names and descriptions, see Common application properties.

CAUTION! When adding a property, be extremely careful. Entering the wrong property name or an invalid data type can cause SAS Viya to become inoperable.

Endpoints
    The set of properties that are used to configure Spring Actuator endpoints.

Flyway
    The set of properties that are used to configure Spring Flyway integration.

Liquibase
    The set of properties that are used to configure Spring Liquibase integration.

Logging
    The set of properties that are used to configure logging.

Logging.Level
    The set of properties that are used to configure logging levels.

Management
    The set of properties that are used to configure Spring application management.

Multipart
    The set of properties that are used to configure Spring multipart handling.

Security
    The set of properties that are used to configure Spring security.

Server
    The set of properties that are used to configure the embedded Spring server.

Shell
    The set of properties that are used to configure the Spring remote shell.

Spring
    The set of properties that are used to configure other Spring features.

zones

The set of properties that are used to configure zone information for multi-tenancy. Modifying one of these property values requires you to restart the service.

internal.hostnames
The comma-separated list of internal host names that are used to access the provider zone, or that are used in a subdomain to access other zones.

> **TIP** Be sure to specify the base host name without any tenant prefixes.

Configuration Properties: Interfaces

There are several interfaces that you can use to manage configuration properties for SAS Viya servers, services, and applications. The following table lists these interfaces and the shading indicates the relative amount of SAS Viya configuration that each covers:

Table A.2 Interfaces to CAS Administration

◔	Ansible	A software orchestration tool that provides a straightforward approach to deploying and provisioning SAS Viya. You can set configuration property defaults at installation.
◕	SAS Environment Manager	A graphical enterprise web application that enables you to modify and view SAS Viya configuration properties.
●	Command-line interface	A command-line interface that enables you to manage SAS Viya configuration properties.

8

Content Management

Content Management: Overview

The Content page contains objects (such as SAS content and reports) that you save and that are organized into folders. When you open the Content window, you have access to your own data in the folder named **My Folder** unless you are an administrator.

When you log on to SAS Environment Manager for the first time, a folder named **My Folder** is automatically created for you. This folder contains items that you do not want to share with other users. The contents of this folder are visible only to you and administrative users.

When you identify an item such as a report as a favorite, the reference to that item is stored in the **My Favorites** folder. You can select the entries in this folder to quickly access frequently used reports and data.

As you work with reports and data, a list of the items that you access is kept in the **My History** folder. You can select entries in this folder to quickly return to items that you have recently worked with. The folder stores the 40 most recent items that you access. These folders exist at the same level as **My Folder**.

As you work, data that is used by various applications is stored in the **Application Data** folder.

Other folders are predefined because they are used for special purposes. See “Predefined Folders” on page 134 for more information.

Content Management: How To

Navigate the Folders

To open the Content page, select 🗎 **Content** from the navigation menu. The left side of the page displays a list of the folders to which you have access. Click › to the right of a folder name to open the folder and to view the content and subfolders.

Click ⤴ to move up one level in the folder hierarchy.

The name of the folder that you are currently accessing is displayed in the menu at the top of the page. To move to a different level in the folder hierarchy, click the folder name and select the folder that you want to access from the menu.

Click ⇅ to sort the folders and objects by name or date of last use.

View and Edit Authorization for a Folder and Folder Objects

By default, the authorization for a top-level folder allows the owner full access to the folder. See “Inheritance” on page 421 for information about folder authorization.

To view the authorization for a folder or an object inside a folder, select the folder or object in the list, select 🔒 and then select **View Authorization**, or select **View Authorization** from the pop-up menu. The View Authorization window appears, where you can view the permission values for the selected folder.

To edit the authorization for a folder or object, select 🔒 and then select **Edit Authorization**, or select **Edit authorization** from the pop-up menu. You can also select **Edit** on the View Authorization window. The Edit Authorization window appears, where you can change the permission values for the selected folder or object. See “General Authorization: How to (Authorization Window)” on page 412 for information about using the Authorization window to specify permissions for a folder or object.

View Content Properties

The right side of the Content page displays the properties for the currently selected object or folder. Expand the **Basic Properties** or **Advanced** sections to view the properties for the selected object or folder.

Click ✎ to change the name and description for a folder.

Create a New Folder

Click the **New Folder** icon ✱ to create a new folder. The folder appears in the current location in the hierarchy, with the default name of **New Folder**. If you do not have permission to create a folder at the current level, the icon is not selectable. By default, only a SAS administrator can create a top-level folder.

Search Folders

Click 🔍 to search the folders. After the results are displayed in the Search Results dialog box, you can filter the results by object type, the user that modified the object, and the date on which the object was modified. Click ⓘ next to an object in the **Results** list to view the object’s type, location, date created, and date modified.

Add a Shortcut

You can create a shortcut to a folder or an object and save the shortcut in a folder that you choose. Shortcuts enable you to quickly access folders or objects rather than having to navigate to them each time.

1 Right-click on the folder or object for which you want to create a shortcut and select **Add as shortcut** from the pop-up menu.

2 In the Add as shortcut window, select the location to which you want to save the shortcut. Click if you need to create a new folder.

3 Click **OK** to save the shortcut. The shortcut is named **Objectname-Shortcut** in your selected location.

Note: You cannot specify authorization values for a shortcut. Authorization values are specified for the child member that is associated with the shortcut.

Export Content

You can export the contents of a folder or a single item in a folder as a JSON package. Conveyed permissions on a report are not exported along with the report. Only permissions that are directly specified on a report are exported.

1 Select the folder or item that you want to export.

2 Click or select **Export** from the pop-up menu.

3 The Export dialog box appears, and displays the location and type of the item that you selected.

4 In the **Export file** field, specify the filename, without a file extension, of the exported package file. The export function adds the `.json` extension to the filename that you specify. If you do not specify a filename, the file is exported using the default name **`Package.json`**. If a file with the name **`Package.json`** already exists, the file is exported using the convention **`Package(1).json`**, **`Package(2).json`**, and so on.

   Note: If you attempt to export the contents of a folder that does not contain any reports, the export process creates a file named `undefined.json`, which does not contain any content.

   Note: If you export and then import a single theme, the theme's logo image is not included. However, the logo image is included if you export and then import the folder that contains the theme.

5 Click **Export** to create the package file. The file is saved to the default location where files are saved for your web browser.

   Note: If you are exporting a folder that contains many reports, the export process might be lengthy.

Import a Package to a Folder

You can import a JSON package file that contains reports to a folder. All the reports in the package file are imported. You cannot selectively import one or more reports from a file.

For the basic steps to import a package:

Note: For complete information about importing, see "Promotion: How to Import (Wizard)" in *SAS Viya Administration: Promotion (Import and Export)* .

1 Click .

2 The **Import Wizard** appears. Select the source file (either a .json or a .spk file) to import. You can view the properties and dependencies of the selected file.

3 To go to the second step of the wizard, click **Mappings**. Specify the mapping by selecting target servers, caslibs, and tables.

4 To go to the third step of the wizard, click **Import**. Verify the mappings that you selected, and click **Import** to import the source content.

Move an Object or Folder

If you have the appropriate permissions, you can move an object or subfolder to a folder that you choose. You must have these permissions:

Source folder
Read, Remove, Read (convey), Update (convey)

Object in the source folder that you are moving
Read, Update

Target folder
Read, Add

See “Inheritance” on page 421 for more information about permissions.

1 Right-click on the object that you want to move and select **Move to folder** from the pop-up menu.

2 In the Move to folder window, select the location to which you want to move the object. Click 📁 if you need to create a new folder.

3 Click **OK** to move the object.

Delete Content

Select a folder or other content in the hierarchy and click 🗑 to delete the folder or content.

If you are deleting content in a folder, you must have Remove permission for the folder and Delete permission for the object that you are deleting.

Content Management: Concepts

Folders

Information that you or other users save is stored and organized in folders. A folder is a virtual container rather than a representation of a physical file system. A folder contains members, which are URIs for other folders, SAS resources, or resources that are external to SAS software.

A member in a folder can be either a child or a reference.

Predefined Folders

Predefined folders are used for special purposes:

Users
contains folders for individual users. Each user’s folder contains their personal folders such as My Folder, My Favorites, and Application Data. See “Content Management: Overview” on page 131 for more information about these folders.The folders in the Users folder are created automatically when users first log on. Administrators should not create these folders.

Public
contains content that will be shared with anyone who has authenticated access to the system.

Products
contains items, such as samples and static content, that are generated by SAS applications. All authenticated users should have only Read access to these folders.

Your deployment might contain other predefined folders, depending on your organization’s needs and configuration.

Child Members

The URI of a resource is a child member to only one folder (its parent). Because a child member can have only one parent, you cannot copy or duplicate the child member to another folder. However, if you have the proper authorization, you can move a child member to another folder, which then becomes its new parent.

An example of a folder that contains child members is a department’s folder. It contains the reports that everyone in the department uses. This folder is shared, so everyone in the department can access the folder and its reports. Examples of child members include the following:

- Subfolders
- Reports
- Data preparation plans

Reference Members

A reference member is a pointer to a resource that exists as a child in another folder. Reference members in multiple folders can point to the same object

An example of a reference member is an entry in a history folder. If you access a report in your department’s shared folder, an entry is saved in your history folder that contains the URI for the report. Many other users can access the same report, so each of those users also have an entry in their history folder with a reference to the same report. Reference members include the following:

- History (reports that have been recently accessed)
- Favorites
- Shortcuts to objects (child members) or folders

Access to a reference member does not affect access to the child member that is associated with the reference member.

9
Themes

Overview

The Themes service is accessed through SAS Theme Designer 3.2. SAS Theme Designer enables a user with the appropriate authorization to create and manage SAS themes. A theme consists of the following components:

- style data in JSON format
- a display name
- a content type (application or report)
- a type (custom, legacy, retired, sample, or system)
- the ID of the base theme
- the last modified time stamp
- the source modified time stamp
- the published time stamp
- the published status (true or false)
- default theme status (true or false)
- the theme data version
- the root URL
- the UI framework version that the theme was created with
- the UI framework version that the theme was updated with

In addition to these components, report themes have the following components:

- the report style sheet path and name
- the HTML style sheet path and name

The base theme is the theme against which a theme is built. There are two types of base themes that you can use in order to create new themes:

Application theme
: an application theme applies a look and feel to the application user interface.

Report theme
: a report theme is a theme that applies a look and feel to report content. It typically includes styles for graphs and user interface controls that can be included in reports.

Users can save the theme data as a new theme resource and at that time a unique ID is assigned by the service. The administrator can then publish the theme. The theme is applied to all web applications or reports in the same deployment.

Each time a theme is edited and saved, it must be published again in order for users to see the latest changes. If the theme administrator no longer wants the theme to be visible by all users, the theme can be unpublished or deleted. All web applications supplied by SAS consume custom themes and a default application theme by means of the Themes service. SAS Visual Analytics also consumes report themes and a default report theme by means of the Themes service.

Note: The Themes service is not used in a programming-only deployment.

10

Preferences

Overview

The Preferences component enables clients to create, update, delete, and retrieve preference data for the currently authenticated user. This component is maintained and retrieved by applications in order to provide a custom experience for the end user. Some common uses include providing the settings window for web applications or providing an implicit user-specific state when there is no user interface.

Concepts

Preference data consists of a string identifier and a string value that can be used by an application to customize or enhance the user experience. String values can be simple strings, or more complex constructs like JSON documents.

Guidelines

The Preferences components provide access to preferences for a single person in a given request. A typical user sees only the preferences that belong to them. Preference identifiers are unique on a per-user basis. Applications are responsible for ensuring that the identifiers that they use do not collide with identifiers used by other applications. It is recommended that applications have a name that is applied to every identifier, followed by a dot, then the preference name. For example: *app1.preference*.

The identifier can be any valid printable ASCII character, except for the following characters, which will be rejected: %, &, /, !, #, ', (,), *, +, ,, :, ;, =, ?, @, [,], {, }, |, $.

11

Data Administration

Data Administration: Overview

This document assumes that you are familiar with the data and caslib concepts that are explained in SAS Cloud Analytic Services: Fundamentals.

Use the interface that best meets your needs. Here are suggestions:

- To manage caslibs and CAS tables interactively, use the Data area in SAS Environment Manager. See Making Data Available to CAS and Working with Data in CAS.
- If SAS Environment Manager is not deployed, use CAS Server Monitor to manage caslibs interactively. See "Data Administration: How to (CAS Server Monitor)".
- To programmatically manage CAS data, use the Tables Action Set. To get started, see SAS Viya Quick Start.

This document provides instructions for the following administrative tasks:

- Managing User-Defined Formats in SAS Viya
- CAS Table State Management
- Cross-Loading Data Tables
- Loading Geographic Polygon Data as a CAS Table

Managing User-Defined Formats in SAS Viya

Overview

In SAS Viya 3.4, when working with CAS tables and the columns that access user-defined formats, you must make the formats visible to the CAS server. User-defined formats are not directly available to a CAS server. The SAS catalogs that store user-defined formats must be converted to a CAS format library and stored in a file in a caslib. And, for all users to access a CAS format library, it must be promoted in a global caslib.

SAS Environment Manager 3.3 enables you to manage format libraries and user-defined formats. The SAS Viya 3.4 configuration for user-defined formats creates an initial caslib and a CAS formats library. These are the initial libraries that are created:

- the `Formats` global path-based CAS library. This library is located at `/opt/sas/viya/config/data/cas/default/formats/`.

  If you are accessing SAS Viya 3.4 from a Docker container, the path is: `/cas/data/formats/`.
- the `sassuppliedformats.sashdat` CAS format library.

Note: For more information about the initial configuration for format libraries, see "Additional User-Defined Format Tasks for SAS Viya 3.4" on page 150.

Here are the different file types from which you can import formats:

- .sashdat files in a caslib.
- a SAS item store that was created with the FMTC2ITM procedure. This file is accessible from the CAS controller.

  Note: For information about the FMTC2ITM procedure, see "Converting a SAS Catalog to an Item Store with the FMTC2ITM Procedure" on page 148.
- SAS7BCAT files (Base SAS catalog). This file must be accessible from the machine that runs the compute service. Also, the file must be UTF-8 encoded in order to import. See "SAS7BCAT File Exceptions" on page 150.

  Important: When you import formats from a SAS7BCAT file, the execution of that file occurs on the compute server. This requires that you have a host account defined on the compute server. In addition, the compute server must be active during file execution.

  Note: For more information see SAS Workspace Server and SAS Compute Server.

  Note: If you are working on a Windows machine, a SAS7BCAT file must be formatted for Window's use.

Here are additional requirements that pertain to caslibs:

- If user-defined format libraries are saved to path-based caslibs, then the directory must be a network location that is available on the controller and the backup controller.
- For server-based caslibs that use database drivers, the drivers must available on the controller and the backup controller.

You can access the User-Defined Formats area from the left navigation bar of SAS Environment Manager. To open the User-Defined Formats area:

1. Select the applications menu ≡.
2. Select **Administration** ⇨ **Manage Environment**.
3. In the navigation bar, select $w w.d.

   Important: When logging in to SAS you must assume the SAS administrator role in order to see the User-Defined Formats area in the navigation bar.

The User-Defined Formats area lists available CAS servers, promoted format libraries, and global user-defined formats in each library. In this area you can add new format libraries, add new user-defined formats, and import formats from the different file types. You can then manage the format libraries and user-defined formats. You must be a SAS Administrator to manage user-defined formats.

In the User-Defined Formats area, the **Format Filter** pane displays the following items:

- **Server** - You can select the CAS server that you want to work with from the **Server** list.

  Note: Your CAS server must be active and running in order to see available format libraries and formats.
- **Format Library** - You can select specific format libraries.
- **Format Name** - You can select specific formats. You can also search for specific formats.

Items that you select on the **Format Filter** pane are displayed on the **Format name** table. You can also search for specific formats in the **Format name** table.

You can refresh the current view of either the **Format Filter** pane or the **Format name** table by selecting **Refresh** ↻. **Refresh** updates the display of format libraries and formats that are available in your current user session. Changes made by other users are not dynamically updated.

Note: Session formats are not shown in the **User-defined formats** area.

Assume the Superuser Role

Some library functions in the User-Defined Formats area require superuser access. Select **Assume Superuser Role** to obtain superuser access. Functions that require superuser access are now available. Once you have assumed the superuser role, the **Relinquish** option becomes available. After you have completed a needed task, select **Relinquish** and click the displayed server. You no longer have superuser access.

Note: For more information about the superuser role see CAS Server Roles.

Manage Format Libraries

You can create and manage format libraries in the User-Defined Formats area of SAS Environment Manager.

Create a New Format Library

This function enables you to create a CAS format library and also import user-defined formats from a supported file in one process. This is advantageous when you have a large format library that contains a large number of formats or formats that contain large amounts of data. You can also create an empty CAS format library if needed.

To create a new format library complete the following steps:

1. From the **Server** list, select the CAS server that you want to add a library to.
2. Select **Assume Superuser Role** .
3. Select **New format library** . The New Format Library window appears.
4. Enter the following values:
   - **Format library name**
   - **Source path**- this is the file that is imported into the new library. The full path and file name are included. This can be for one of the following file types:
     - .sashdat files.
     - a SAS item store that was created with the FMTC2ITM procedure.
     - SAS7BCAT files (Base SAS catalog).

       Important: See “SAS7BCAT File Exceptions” on page 150.

   The source path must be in a location that the CAS server can access. It must also have the needed permissions to open and read the file.
5. Select the **Search order** for the library. You can select one of the following options from the menu:
   - `Append`- The new format library is appended to the existing format search order.
   - `Prepend`- The new format library is prepended to the existing format search order.
   - `Replace`- The new format library replaces the libraries that are listed in the existing format search order.

   CAUTION! The Replace option overwrites the existing format library search order and replaces it with this newly created format library.

   - `None` - The new format library is not added to the existing format search order.
6. Select **Save** to create the new library.

If needed, you can create an empty format library. And then, at a later time, import or add formats. While creating a new library, select **Create empty format library** on the New Format Library window. The option **Source path** is removed. The newly created library is empty.

Note: You cannot add subdirectories to a format library in SAS Environment Manager.

Note: After you have created the new format library, you can relinquish the Superuser privilege. Select **Relinquish** and click the displayed CAS server. You no longer have Superuser privilege.

Change the Format Library Search Order

The format library search order defines the sequence in which format libraries are searched. If there are duplicate named formats in different format libraries, the format engine uses the first format that it finds in the search order list of format libraries. To change the sort order of format libraries on a CAS server, complete the following steps:

1. From the **Server** list, select the CAS server that you want to change the format library search order for.
2. Select **Assume Superuser Role** 👤☆.
3. Select **Format library search order** ☰. The Format Library Search Order window appears.
4. On the Format Library Search Order window, you can move libraries between the columns: **Libraries not in search order** and **Search Order**. If you want a library included in the search for user-defined formats, assign the library to the **Search Order** column. You can then move the library up or down in the list of format libraries. Libraries are searched in order of their position in the **Search Order** column.

Note: After you have changed the search order of format libraries for the selected CAS server, you can relinquish the Superuser privilege. Select **Relinquish** and click the displayed CAS server. You no longer have Superuser privilege.

Note: If you delete a format library using the command line interface (CLI), but do not remove it from the search order path, the deleted format library will still be listed in the Format Library Search Order window. However, it will not appear in the **Format Library** list that is displayed on the left side of the User-Defined Formats window. For more information about the command line interface (CLI) see: "Managing Format Libraries and the Command Line Interface " on page 153.

Delete a Format Library

The **Delete format library** function deletes the format library from the search order and removes the source file for the deleted format library. To delete a format library complete the following steps:

1. From the **Server** list, select the CAS server that you want to delete a format library from.
2. Select **Assume Superuser Role** 👤☆.
3. Select **Delete format library** 🗑. The Delete Format Library window appears.
4. On the Delete Format Library window, select the check box for the format library that you want to delete.
5. Select **Delete**. Confirm the delete action and select **Delete** again on the Delete window. The format library is deleted from the CAS server.

Note: After you have deleted the format library, you can relinquish the Superuser privilege. Select **Relinquish** and click the displayed CAS server. You no longer have Superuser privilege.

Manage Formats

The User-Defined Format area enables you to import, add, and manage user-defined formats.

Import User-Defined Formats

The Import function is best used with moderate sized data sets. It is useful when you need to view the formats that you import and compare them to any existing formats.

You can import user-defined formats from the following SAS data sources:

- .sashdat files.
- a SAS item store that was created with the FMTC2ITM procedure.
- SAS7BCAT files (Base SAS catalog).

  Important: See “SAS7BCAT File Exceptions” on page 150.

To import user-defined formats from one of these data sources, follow these steps:

1 Click ↗ to open the Import Formats window.

2 Enter a directory path for the format source on the **Source path** field. The path must include the format source name and extension. For example, the following path points to a SASHDAT file:

```
c:\users\myformats\importformats1.sashdat
```

3 Select a format library from the `Target format library` drop-down list.

  Important: In a new installation of SAS Environment Manager 3.3, the SASSUPPLIEDFORMATS library is the only provided library. Use the **New format library** function to create a new format library.

4 Click **Compare**. The Importing message pane opens. The source file and the format library that you selected are displayed. Select **OK**. Formats that are found in the data source are populated in the **Comparison Results** list.

  If needed, you can select **Cancel** to cancel the compare and exit the Importing message pane.

  The number of formats that are found is displayed. The status of the formats that are located is also indicated. The following information values are updated:

  - **Information**ⓘ - This value shows the number of formats that are found in the data source that can be loaded without conflict.
  - **Warning**⚠ - This value indicates that there is a possible conflict with one or more of the formats that are listed in the **Comparison Results** list. For example, a duplicate format has been found with an identical range as an existing format.
  - **Error** ⊗ - This value indicates that there is a conflict with one or more of the formats that are listed in the **Comparison Results** list. For example, a duplicate format of an existing format has been found. However, the range of value is different between the two formats.

5 In the **Comparison Results** list, select the formats that you want to import.

  For any formats that have a warning or error conflict, you can analyze the conflict on the **Format Properties** panel. Select a format that has a conflict by clicking on the format name. The format values are displayed in the **New format** and **Existing format** columns.

  For formats that are in conflict with existing formats, you can assign the new format a different name than the existing format. In the **New format** field, enter a name for the new format. Click **Apply**. When you select **Import**, the new format is imported with the new name.

  TIP The following characters cannot be used for a format name: ~!@#%^&*()|+-?;:’,. or any DBCS.

6 You can also click **Select all** to select all of the found formats. Select **Import**. The new formats are now available.

Add a User-Defined Format

To add a new user-defined format:

1. Select + to open the New Format window.
2. Select one of the available format libraries from the **Format Library** drop-down list. The new format will be assigned to this library.
3. Select either `Character` or `Numeric` from the **Type** drop-down list.
4. Enter a name for the format in the **Name** field.

   > **TIP** The following characters cannot be used for a format name: ~!@#%^&*()|+-?;:’,. or any DBCS. If you enter one of these characters in the **Name** field, the **Save** button is disabled.

5. Check the **Save without locale** option if you do not want to include a default locale. The value **Current locale prefix** displays the currently set locale. By default a format is created with a locale.

   If you create a format with the locale included, the format name is prepended with the name of the locale. If you check **Save without locale**, the format name will not include the locale.

   > **TIP** To change the locale prefix for a format select the **Application options** menu in the upper right corner of SAS Environment Manager. Select **Settings** ⇨ **Global** ⇨ **Region and Language** ⇨ **Locale for regional formats and sorting**. Select a locale from the menu.

6. Add rows to the **Range** table. Select +. Enter values in the **Name** or **Value** columns for each row.

   Note: For **Range** values, the decimal separator is dot (.), regardless of the set locale. A comma is used to separate discrete values.

   You can delete a row by selecting 🗑.

   You can also select a range of multiple rows to delete.

7. After you have added any needed rows, click **Save**. The new format is listed in the **Format name** table.

Edit a User-Defined Format

You can edit range values for a user-defined format. To edit a user-defined format:

1. Select a format and click ✎.
2. On the Edit Format window, make any needed changes to the **Range** table of rows. You can edit values in the **Name** or **Value** columns for existing rows.

   Note: For **Range** values, the decimal separator is dot (.), regardless of the set locale. A comma is used to separate discrete values.

   Click + to add rows and 🗑 to delete rows.

3. Click **Save** when finished.

Copy a User-Defined Format

You can copy a format from one library to another, or within the current library. To copy a user-defined format:

1. Select a format and click ⧉. The Copy Format window appears.

2. On the **Format Library** list, select the library that you want to add the new copy to. You can add the copy to the current library or you can add it to a different library.

3. When you copy a format, the extension `_copy` is automatically added to the name of the copy. The **Name** field is automatically populated. For example, if you copy a format that is named *$gender*, the new copy is automatically named *$gender_copy*.

   If needed, you can modify the format name. Click on the **Name** field and enter a new, unique name for the copied format.

   > **TIP** The following characters cannot be used for a format name: ~!@#%^&*()|+-?;:',. or any DBCS. If you enter one of these characters in the **Name** field, the **Save** button is disabled.

4. Make any needed changes to the **Range** table of rows. You can edit values in the **Name** or **Value** columns for existing rows.

   Note: For **Range** values, the decimal separator is dot (.), regardless of the set locale. A comma is used to separate discrete values.

   Click ✚ to add rows and 🗑 to delete rows.

5. Click **Save** when finished. The new format is listed in the **Format name** table.

Delete a User-Defined Format

To delete a user-defined format:

1. Select a format and click 🗑. The Delete pane opens.

2. Select **Delete**.

View Properties for a User-Defined Format

To view properties for a user-defined format:

1. Select a format and click ☰. The Format window appears.

2. You can now view properties and row values for the format.

Here are some additional navigation features:

- You can customize which columns are displayed by selecting ☰ and selecting **Manage Columns**. The Manage Columns window contains all available columns for the currently selected format. From here you can choose to display or hide columns in the **Format name** table. On this window you assign table columns to the **Hidden columns** list or to the **Displayed columns** list.

  These preferences persist until the end of your SAS Environment Manager login session.

- You can reorder columns by selecting and dragging a column to the left or right.

- You can refresh the current view by selecting **Refresh**⟳. **Refresh** updates the display of your current user session. Changes made by other users are not dynamically updated.

Converting a SAS Catalog to an Item Store with the FMTC2ITM Procedure

One way to move formats stored in a SAS catalog to a CAS server is to use the FMTC2ITM procedure. The FMTC2ITM procedure is used to convert one or more format catalogs into a single file. This procedure copies the formats in the catalog to a physical file, known as an *item store*. SAS Environment Manager can then import formats from an item store.

The FMTC2ITM procedure uses a SAS format catalog as the input source. It then produces an item store that is read in CAS. The physical file must be accessible from the CAS server controller and the user must have permission to read the file. The FMTC2ITM procedure is available in SAS 9.4M3 and later.

Important: SAS catalogs are operating-system-type specific. Catalogs must be converted to the operating system in use. For example, a Windows catalog cannot be used to import on a Linux operating system and vice versa.

Syntax

FMTC2ITM <options>;

SELECT<member-list>;

RUN;

Options

CATALOG=memname | libname.memname | (list)
specifies catalog(s) that will be converted to an item store.

DEBUG
debugs information about the records that are being written.

ENCODING=encoding-name
specifies an encoding for a catalog or for all of the catalogs in a list.

ITEMSTORE= fileref | 'filename'
specifies the item store file that is being created.

LOCALE
specifies locale-sensitive prefixes that are added to item store memnames.

PAGESIZE=n
specifies the page size.

PRINT
displays information about each member that is being written.

XMLFILE= fileref
specifies an XML file that is created to accompany the item store file.

FMTC2ITM Procedure Example

The following example converts the catalogs *formats.sales* and *formats.finance* to the item store *format1*.

```
proc fmtc2itm catalog=(formats.sales, formats.finance)
print locale itemstore="/users/dsmith/formats/format1";
run ;
```

About the FMTC2ITM Procedure

Note that the item store is always written as new, so if the **`ITEMSTORE=`** option refers to an existing item store, it will be completely overwritten. For a file-based item store, the **`XMLFILE=`** option can also be provided. It is populated with a small XML stream that accompanies the item store file.

If the **CATALOG=** option is not given, then the default value is **WORK.FORMATS**. If the **CATALOG=** option is given, it can be a single-level name, which is interpreted as a catalog name in WORK. It can also be a two-level name, which is interpreted as a libname.memname for a catalog. It can also be a list of catalog names that is enclosed in parentheses.

For an item store that is a file, either a fileref or a quoted string pathname is given. For a list of catalog names that is enclosed in parentheses, each catalog is opened in order and the members of the catalogs are written to the item store. Only the first occurrence of the member is written out.

The SELECT statement is optional. If specified, it lists the formats that are selected to be placed in the item store. If a SELECT statement is not given, all formats are written to the item store.

Note: For more information about the FMTC2ITM procedure see FMTC2ITM Procedure.

SAS7BCAT File Exceptions

A SAS7BCAT file must be UTF-8 encoded in order to import. Currently, you cannot import native encoded SAS7BCAT files. As an alternative, you can convert a SAS7BCAT file to a UTF-8 encoded SAS item store. You can use the **ENCODING=** option with the FMTC2ITM procedure.

Additional User-Defined Format Tasks for SAS Viya 3.4

For SAS Viya 3.4, the CAS server uses start-up scripts to populate an initial caslib and a CAS formats library. The start-up scripts run before the CAS server accepts any client connections. The two start-up scripts that configure user-defined formats are the **Casstartup.lua** file and optionally, the **Casstartup_usermods.lua** file.

Using the Casstartup.lua File

The **casstartup.lua** file is processed as a LUA client session into the CAS server. It is used to perform static, default deployment tasks as the CAS server starts. The **casstartup.lua** file adds format libraries, promotes the format libraries, and creates the new FmtSearch list. It is pre-configured with default settings as part of the deployment process. The **casstartup.lua** file is included in the start-up processing hierarchy and is used during CAS server start-up. The command line default start-up value is**-startup casstartup.lua**.

During server start-up and in a default SAS Viya installation, the following files are created:

- the **Formats** global path-based CAS library. This library is located at **/opt/sas/viya/config/data/cas/default/formats/**.

  If you are accessing SAS Viya 3.4 from a Docker container, the path is: **/cas/data/formats/**.

- the **sassuppliedformats.sashdat** CAS format library.

The **casstartup.lua** file contains the default **addFmtLib** action. It also contains the **setServOpt** action that is used to establish the format search list that each session starts with. CAS processes the **casstartup.lua** file before any of the other start-up files that reside in the start.d/ directory. The **casstartup.lua** file searches for any format libraries that are specified in the **casstartup_usermods.lua** file.

Below is an example **casstartup.lua** file:

```
----------------------------------------------
Default Lua Startup file for Cloud Analytic Services
----------------------------------------------------

[...]

unique id for the deployment of CAS
current_dir = debug.getinfo(1).source:match("@?(.*/)")
```

```
deployment_instance = string.gsub(string.sub(current_dir,
1, -2), "(.*/)(.*)", "%2")
config_loc = current_dir

--------------------------------------------------
Default SAS Supplied Formats
--------------------------------------------------
s:sessionProp_addFmtLib{caslib="Formats",fmtLibName="SASSuppliedFormats",
name="sassuppliedformats.sashdat",promote=true}
newFmtSearch = "SASSuppliedFormats"
newFmtSearch = ((cas.fmtsearch or "") .. " " ..newFmtSearch)
s:configuration_setservopt{fmtsearch="SASSuppliedFormats"}

--------------------------------------------------
Process a user specific file if available
--------------------------------------------------
pass,err = pcall(dofile, config_loc ..
'casstartup_usermods.lua')

[...]
```

The recommended deployment best practice is to target a configuration path for all default filenames. If present, the default names will be found. For example, the following **`cfgpath`** option

```
-cfgpath /my/config/path
```

will find

```
/my/config/path/casstartup.lua
```

Using the Casstartup_usermods.lua File

After the initial deployment tasks have been completed, the **`casstartup.lua`** file invokes the **`casstartup_usermods.lua`** file. This file can contain modifications, such as adding global-scope caslibs and loading global-scope tables.

Site-specific format libraries associated with the **`addFmtLib`** action should be placed in **`casstartup_usermods.lua`**. The **`casstartup_usermods.lua`** file is automatically applied when present. The local system administrator can add any site-specific start-up processes (such as table loading) to **`casstartup_usermods.lua`**.

The **`casstartup_usermods.lua`** file is not replaced when software is updated. Therefore, using **`casstartup_usermods.lua`** ensures that custom modifications are not overwritten when the CAS deployment is upgraded.

Follow these steps to add a format caslib to the **`casstartup_usermods.lua`** file:

1 If you choose to add a caslib or use an existing caslib, set the following permissions for Authenticated Users: **`Read`**, **`Select`**, **`Promote`**, and **`CreateTable`**.

2 Log on to the CAS controller machine as the SAS install user (sas) or with pseudo privileges.

3 Make sure that the format library is saved to the caslib.

4 Use a text editor to open and edit the **`casstartup_usermods.lua`** file in the path that is appropriate for your operating system:

- Linux:

  `/opt/sas/viya/config/etc/cas/default/casstartup_usermods.lua`

- Windows:

\ProgramData\SAS\Viya\etc\cas\default\casstartup_usermods.lua

For more information, see Working with User-Defined Formats in the *SAS Viya: System Programming Guide*

Here is a sample `casstartup_usermods.lua` file:

```
-----------------------------------------------------------
 USERMODS Lua Startup file for Cloud Analytic Services

-----------------------------------------------------------------

[...]

-- Define the permanent and reloadable format libraries
s:sessionProp_addFmtLib{caslib="Formats",fmtLibName="SalesFmt",name="SalesFmt.sashdat",promote=true}
s:sessionProp_addFmtLib{caslib="Formats",fmtLibName="FinancesFmt",name="FinancesFmt.sashdat",promote=true}
s:sessionProp_addFmtLib{caslib="Formats",fmtLibName="HumanResourcesFmt",name="HumanResourcesFmt.sashdat",
promote=true}

-- Create the new custom FMTSearch list
newFmtSearch = "SalesFmt FinancesFmt HumanResourcesFmt"

-- Create the new global FMTSearch list
newGolbalFmtSearch = (newFmtSearch .. " " .. (cas.fmtsearch or "")) -- or (after new list)

-- Set the new global FMTSearch list
s:configuration_setServOpt {fmtsearch= newGlobalFmtSearch}

[...]
```

Persisting User-Defined Formats across Server Restarts

During session start-up, the format search order is used to automatically establish which format libraries to search, when a table format is needed. Another use of the `casstartup_usermods.lua` file is to pre-load tables to the CAS server to ensure that they are immediately available for use. In order to maintain the search order for user-defined formats that you have added or updated in SAS Environment Manager 3.3, you can make the following changes to the `casstartup_usermods.lua` file:

- load the CAS format library into memory in the global FORMATS caslib.
- add the CAS format library to the CAS server search order.

The `casstartup_usermods.lua` file runs every time that the CAS server starts. Below is an example `casstartup_usermods.lua` file with these changes:

```
-----------------------------------------------------------
 USERMODS Lua Startup file for Cloud Analytic Services

-----------------------------------------------------------------

[...]

-- Additional User Formats
s:sessionProp_addFmtLib{caslib="Formats",fmtLibName="SalesFmt",
name="SalesFmt.sashdat",promote=true}

-- Create the new custom FMTSearch list
customFmtSearch = "SalesFmt"
```

```
-- Create the new global FMTSearch list
newGlobalFmtSearch = (customFmtSearch .. " " .. (cas.fmtsearch or ""))

-- Set the new global FMTSearch list
s:configuration_setServOpt {fmtsearch= newGlobalFmtSearch}

[...]
```

Update-In-Place Exceptions

If you are performing an update-in-place of legacy format libraries to your SAS Viya 3.4 environment, legacy files are addressed during update:

- On an update-in-place system, when casFormats executes, it checks for legacy format libraries that are no longer used. If empty libraries are found, they are removed from the search path and will not be available in SAS Environment Manager. Legacy format libraries that do contain formats will be made available in SAS Environment Manager.
- If there is an existing **casstartup.lua** file in the configuration directory during an update-in-place, that **casstartup.lua** file will be renamed to **casstartup.lua_n**, where *n* is an integer string that represents an elapsed time count. Any existing entries in the "Default User formats" section of the **casstartup.lua_n** file should be copied to the newly installed **casstartup.lua** file under the "Default SAS Supplied Formats" heading.

Managing Format Libraries and the Command Line Interface

In SAS Viya 3.4 the command line interface (CLI) enables an administrator to create format libraries, update the format search order, and import formats. For more information about the command line interface (CLI) see SAS® Viya 3.4 Administration: Using the Command-Line Interfaces .

Note: As with SAS Environment Manager, the CLI also does not persist user-defined formats across server restarts. You must modify the **casstartup_usermods.lua** file. See "Persisting User-Defined Formats across Server Restarts" on page 152 .

CAS Table State Management

Overview of CAS Table State Management

CAS table state management enables you to manage the import, load, and unload of source files in CAS. CAS table state management is performed through the use of jobs that are created from sample jobs that are provided by SAS.

For SAS Viya 3.4, you can import batch data that is directly accessible using a caslib, but might not be in the desired caslib or format. This type of import is used to take source data and make a copy in SASHDAT format.

For example, business processes might produce new data each night in CSV or SAS7BDAT format. It is possible to access the data directly using a global caslib that points to the source of the data. However, for performance reasons, it might be desirable to make a copy of the data in SASHDAT format.

Sample jobs should be used as a starting point. These sample jobs can be used as is for the Public caslib that is associated with the cas-shared-default CAS server. The jobs can be copied, and the copies can be edited or deleted. A job includes the specific options required by the job. In the context of CAS table state management, a job performs an import, load, or unload operation on input files, tables, or loaded tables. Jobs can be listed, copied, updated, and deleted on the SAS Environment Manager Jobs page. The Jobs page contains a **Monitoring** tab and a **Scheduling** tab. For CAS table state management sample jobs, the **Scheduling** tab enables you to make copies that are used to import, load, and unload batch data. Each job can be submitted manually, or scheduled for later execution.

Note: For SAS Viya 3.4, only CSV, SAS7BDAT, or EXCEL files can be imported to a SASHDAT format file.

Sample Jobs in SAS Environment Manager

For SAS Viya 3.4, there are three sample jobs that are provided by SAS for managing table state. These jobs are available on the Jobs page, **Scheduling** tab of SAS Environment Manager. Below are the sample jobs:

`Sample: Import cas-shared-default Public data`
This job demonstrates settings that import all CSV, SAS7BDAT, and EXCEL files in the Public caslib to SASHDAT files in the same caslib.

`Sample: Load cas-shared-default Public data`
This job demonstrates how to load all SASHDAT files found in the Public caslib.

`Sample: Unload cas-shared-default Public data`
This job demonstrates how to unload all loaded CAS tables in the Public caslib that have not been accessed within the past 7 days.

On the Jobs page of SAS Environment Manager, you can access the sample jobs by selecting the **Scheduling** tab and then viewing the **Jobs** pane. The sample jobs operate on a CAS server named **`cas-shared-default`**. You cannot edit or delete the sample jobs. However, you can copy the sample jobs to create unique jobs that you can further customize. Copied jobs contain the options that you can define and update as needed.

To create a new job:

1. On the Jobs page, select the **Scheduling** tab.
2. On the **Jobs** pane, select a sample job and select **Copy** ⧉. A new job is created.
3. You can now customize the options for the new job by selecting ☰.

For further information about the Jobs page in SAS Environment Manager, see Jobs Overview.

Viewing and Editing Properties for a Job

On the Jobs page, you can view and edit properties for copied jobs by selecting the **Scheduling** tab. On the **Jobs** pane, select the copied job and then select ☰. The Job Properties window appears.

On the Job Properties window, you can select the following tabs:

- **General**- displays general information about the job including **Name**, **Description**, and **ID**.

  Note: The **ID** value is used with the **`successJobId`** option. See "Chaining Jobs Together" on page 159. .

- **Properties**- displays the **Name** and **Value** for job properties.
- **Arguments** - displays options, filters, and settings for a job. For a copied job, the ✎ option is available. On a copied job, click ✎. The Edit Argument window appears, where you can now make changes.
- **Job Definition**- displays properties and parameters for a job. These settings are read-only.

Job Options

The following options are common to the import, load, and unload jobs. They can be viewed on the **Arguments** tab and edited in the Edit Argument window.

Server name
The name of the CAS server on which the operation will be performed.

Input Caslib
The caslib name that is used as input for the job. For import and load jobs, this is the caslib that contains source files or tables. For an unload job, this is the caslib that contains potential tables to unload.

Output Caslib
The caslib name that is used for output of the job. For an import job, this is the caslib where output files are written. For a load job, this is the caslib where CAS tables are loaded. The output caslib is not applicable for an unload job.

Filter
The filter is used to subset the list of items from the input caslib upon which job operations are performed.

See "Job Filter Syntax" on page 155 for more details and example filters.

Job Filter Syntax

Job options can also contain filters. In its simplest form, a filter selects an item based on whether a condition passes. For example, to select an in-memory table whose name is exactly MYDATA, the following example filter could be used:

```
eq(name,'MYDATA')
```

In the next example, the filter is used to select a source table name ending in lowercase .sashdat:

```
endsWith(sourceTableName,'.sashdat')
```

There are several operators that can be used in a filter. The following table contains these operators:

Table A.1 Filter Operators

Operator	Description	Example	Example Result
contains	True if the value of the first parameter contains the value of the second parameter.	`contains(name,'SPECIAL')`	Only tables whose name contains SPECIAL are selected. For example: MYSPECIALDATA, SPECIALDATA, THISSPECIALDATA.
endsWith	True if the value of the first parameter ends with the value of the second parameter.	`endsWith(sourceTableName,'.sashdat')`	Only source files or table names ending in lowercase .SASHDAT are selected.
eq	True if the parameters specified are equal.	`eq(name,'MYDATA')`	Only the table named MYDATA is selected.
in	True if the value of the first parameter contains any following values.	`in(name,'TABLE1','TABLE2','TABLE3')`	Only tables TABLE1, TABLE2, or TABLE3 are selected.
startsWith	True if the value of the first parameter begins with the value of the second parameter.	`startsWith(sourceTableName,'DEPTA_')`	Only source tables beginning with 'DEPTA_' are selected. For example: DEPTA_CUSTOMERS.csv, DEPTA_ADDRESSES.sas7bdat.

The following table contains filter fields that can be used in expressions:

Table A.2 Filter Fields

Field Name	Content
name	This field represents the CAS table name (whether loaded or unloaded).
sourceTableName	This field represents the name of the source file in the input caslib.
tableReference.sourceTableName	This field is an alias for the sourceTableName field and can be used in place of it.

The following table contains filter examples:

Table A.3 Filter Examples

Example	Filter
by file extension (.sashdat, .csv, .sas7bdat)	`or(endsWith(sourceTableName,'.sashdat'), endsWith(sourceTableName,'.csv'), endsWith(sourceTableName,'.sas7bdat'))`
by exact match	`eq(name,'MAILORDER')`
by list of inputs	in(name,’AIRLINE’,’CUSTOMERS’,’WORLDBANK’)
by substring (contains some string)	`contains(name,'DATA')`
using multiple conditions where either are true	`or( eq(name,’MYDATA’), endsWith(name,’YOURDATA’) )`
using multiple conditions where both are true	`and( contains(tableReference.sourceTableName,’DEPTA_’), endsWith(tableReference.sourceTableName,’.sashdat’) )`

Importing Data

The **`Sample: Import cas-shared-default Public data`** job imports CSV, SAS7BDAT, and EXCEL files to SASHDAT files. It imports to the **`Public`** library on an example CAS server named **`cas-shared-default`**.

This job enables an import from the caslib source defined for the CAS server. By default, the import job imports CSV, SAS7BDAT, and XLS, XLSX (EXCEL) files. It imports those files to the target caslib’s source location as SASHDAT files of the same name. Source files can therefore be placed in a path-based caslib (PATH, and DNFS for example) that is accessible by the CAS server controller. The default path for imported files is **`/opt/sas/viya/config/data/cas/default/public/`**.

If you are accessing SAS Viya 3.4 from a Docker container, the path is: **`/cas/data/public/`**.

Note: For situations where the SASHDAT copy is not required, the load job can be used to load the SASHDAT file directly into memory as a CAS table.

On the **Jobs** page of SAS Environment Manager, select the **Scheduling** tab. The sample job **`Sample: Import cas-shared-default Public data`** is available from the **Jobs** pane. You can copy the job and edit the new copied job. Here are the possible settings for the **`Sample: Import cas-shared-default Public data`** job:

Table A.4 Import Job Settings

Setting	Value Type	Default Value	Sample Values
allowTruncation	Boolean	true	true, false
charMultiplier	decimal	2	1,2,2.5,3,4
delimiter	character	,	,
encoding	string	utf-8	utf-8
getNames	Boolean	true	true, false
guessRows	integer	200	20,50,500
refresh	Boolean	true	true, false
refreshMode	String	newer	always, newer
stripBlanks	Boolean	false	true, false
successJobId	String	There is no default value. You must enter the ID number of the next job that you execute.	7ff6124c-de57-4ee3-a0ed-01fefd7f883d
varChars	Boolean	false	true, false

Loading Data

The **`Sample: Load cas-shared-default Public data`** job performs a load operation on managed files or tables in the target caslib. It then creates an in-memory CAS table of the same name in the target caslib. For SAS Viya 3.4, only SASHDAT format files can be loaded. For CSV, SAS7BDAT, or EXCEL files, you must first import the files to SASHDAT format files.

Note: CAS table names are all uppercase.

This job enables you to preload tables for which there is a high user demand. Or, for scenarios where the amount of time needed to load the table is too long due to data size.

On the **Jobs** page of SAS Environment Manager, select the **Scheduling** tab. The sample job **`Sample: Load cas-shared-default Public data`** is available from the **Jobs** pane. You can copy the job and edit the new copied job. Here are the possible settings for the **`Sample: Load cas-shared-default Public data`** job:

Table A.5 Load Job Settings

Setting	Value Type	Default Value	Sample Values
allowTruncation	Boolean	true	true, false
charMultiplier	decimal	2	1,2,2.5,3,4
delimiter	character	,	,
encoding	string	utf-8	utf-8
getNames	Boolean	true	true, false
guessRows	integer	200	20,50,500
refresh	Boolean	false	true, false
refreshAccessThreshold	Number	0	300 (seconds)
refreshMode	String	newer	always, newer
scope	String	global	global, scope
stripBlanks	Boolean	false	true, false
successJobId	String	There is no default value. You must enter the ID number of the next job that you execute.	7ff6124c-de57-4ee3-a0ed-01fefd7f883d
varChars	Boolean	false	true, false

Note: When the refresh option is set to `true`, each table selected by the filter is unloaded first. If the table is not sourced from the input caslib, it is not reloaded. Therefore, it is important to ensure that the filter is properly set to select only the tables for which you want a refresh. Tables are refreshed only if they are sourced from the caslib that is specified with the inputCaslib setting.

Unloading Data

The `Sample: Unload cas-shared-default Public data` job unloads tables in the target caslib either immediately, or based on recent access. This enables you to schedule forced unloads of tables on a routine basis. Or you can schedule an unload request that is based on how often a table is used. The sample job unloads infrequently accessed data in the `Public` table on the `cas-shared-default` server.

On the **Jobs** page of SAS Environment Manager, select the **Scheduling** tab. The sample job `Sample: Unload cas-shared-default Public data` is available from the **Jobs** pane. You can copy the job and edit the new copied job. Here are the possible settings for the `Sample: Unload cas-shared-default Public data` job:

Table A.6 Unload Job Settings

Setting	Value Type	Default Value	Sample Values
successJobId	String	There is no default value. You must enter the ID number of the next job that you execute.	7ff6124c-de57-4ee3-a0ed-01fefd7f883d
unloadAccessThreshold	String	P7D	P0D, P5M, PT4H, PT5M

The setting *unloadAccessThreshold* is available in the settings for this job. If *unloadAccessThreshold* is set to a specific time period, those tables that are not accessed within the set time period are unloaded. The default value for the *unloadAccessThreshold* setting is P7D (7 days). The following example time threshold values are possible:

Table A.7 unloadAccessThreshold Values

Value	Description
P0D	zero days. This setting results in an immediate unload. There is no threshold.
P7D	period of 7 days
P5M	period of 5 months
PT4H	period of time of 4 hours
PT5M	period of time of 5 minutes
PT45S	period of time of 45 seconds

Execution and Monitoring of Jobs

On the Jobs page of SAS Environment Manager, the **Scheduling** tab enables you to schedule and execute the jobs that you define and customize. You can choose to run jobs as the SAS Administrator or as a different user. You can also schedule or unschedule a job.

In the **Jobs** pane of the **Scheduling** tab, right-click on a job. If you select **Run** for a job, you can check the execution of that job by accessing the **Monitoring** tab on the Jobs page. From the **Monitoring** tab, you can view the different jobs that have been executed and download the log file that contains details of the execution.

Note: The log file is updated as progress is made. So downloading the log file while the job is running shows progress until that point only. To see later progress, you must download the log again for those jobs.

For further information about the Monitoring tab in SAS Environment Manager, see Monitor Jobs.

Chaining Jobs Together

When scheduling jobs, you might need to trigger the execution of one job upon the successful completion of another job. You can trigger one job after another with the `successJobId` setting. This setting can be used with import, load, and unload jobs. The `successJobId` setting specifies the job request ID of the next job to submit, if the initial job succeeds. The job ID value can be located from the **Jobs** page, on the **Scheduling** tab. To implement this setting, identify the jobs that you want to chain together and follow these steps:

1. On the **Jobs** pane, right-click on the job that you want to trigger, after running another job. Select **Properties**.
2. On the Job Properties window, select the **General** tab, copy the value from the **ID** field.
3. On the **Jobs** pane, right-click on the job that you want to run first. Select **Properties**. On the Job Properties window, select the **Arguments** tab and click ✎. The Edit Argument window appears.
4. On the **Settings** table, scroll to the `successJobId` setting. Enter the copied job ID value (of the job that you want to trigger) in the **Value** column. Save the setting changes.

   When you run the initial job, the second job will trigger, once the first job completes successfully.

Refreshing Tables

For SAS Viya 3.4, you can refresh import and load jobs with one of the refresh settings that are available from the Jobs page, **Scheduling** tab in SAS Environment Manager. On the **Scheduling** tab, the **Jobs** pane lists available jobs. Right-click on the job that you want to edit. Select **Properties**. On the Job Properties window, select the **Arguments** tab and click ✎. The Edit Argument window appears. The following refresh settings are available on the **Settings** table:

`refresh`

Can be set to `true` or `false`. When this setting is set to `true`, it indicates that existing targets should be refreshed according to the `refreshMode` setting.

`refreshMode`

When the `refresh` setting is set to `true`, this value determines the manner in which a refresh will occur. Below are possible values:

- `always` - The target will be refreshed every time the job runs.
- `newer` - The target will be refreshed only if the source is newer than the existing target. This is the default value.

The load job also contains the following additional setting:

`refreshAccessThreshold`

This option enables you to stop refreshing tables that have been accessed within the number of seconds specified. When the `refresh` and `refreshMode` settings normally result in refreshing a loaded CAS table, the setting `refreshAccessThreshold` provides an additional gate to prohibit a refresh from occurring. Possible values are in seconds. For example, setting this value to 300 would skip refreshing a table (that would otherwise have been refreshed) if the table was accessed within the last 300 seconds (5 minutes).

Upgrade-in-Place

When you perform an upgrade-in-place from SAS Viya 3.2, CAS Table State Management jobs are added to your environment.

When you perform an upgrade-in-place from SAS Viya 3.3, CAS Table State Management jobs remain the same. You can manually edit copies of these jobs and add the new options that are described in “Chaining Jobs Together” on page 159 and “Refreshing Tables ” on page 160.

Cross-Loading Data Tables

When working with a SAS CAS library (caslib) it is possible to load data from different data sources. A caslib uses two separate data storage areas:

- a temporary, in-memory space to hold CAS tables
- a permanent data source to contain and backup those CAS tables.

Data moves back and forth (loading and saving) between a temporary, in-memory location and the permanent data source location. When a CAS server stops, the in-memory data disappears but the data source copy remains. The data source can be any number of physical data servers, including connector sources such as SQL Server and Oracle. It can also be a platform source such as Linux PATH files, DNFS, and HDFS.

If you want to load data from a data source other than the caslib's defined data source, you can cross-load data from unrelated data sources. For example, if you need to load data from multiple, disparate data sources and combine the data into a single target caslib.

You can cross-load data with the *CASUTIL* procedure, *LOAD* statement. The *LOAD* statement parameters *INCASLIB* and *OUTCASLIB* enable you to specify different caslibs. These parameters are used to facilitate reading from one caslib data source into a different caslib in-memory space. For more information see the CASUTIL Procedure, LOAD Statement.

Loading Geographic Polygon Data as a CAS Table

Overview of Loading Geographic Polygon Data

Some SAS Viya applications such as SAS Visual Analytics can display geographic maps with colored map regions. By default, countries and their first-level subdivisions can be displayed as a region map. To display other types of map regions, such as postal codes or sales regions, you must define a custom polygon provider that contains the polygons (geographic region shapes).

You can load two types of polygon data into CAS for use in a polygon provider: Esri shapefiles, and SAS map data sets.

After you have loaded the polygon data into CAS, you must define a polygon provider that specifies the parameters for the polygon data. For details about defining a polygon provider in SAS Visual Analytics, see “Create a Geography Data Item By Using Custom Polygonal Shapes” in *SAS Visual Analytics: Working with Report Data*.

Note: By default, SAS Visual Analytics can retrieve up to 250,000 polygon vertices at a time. If you encounter an error message in a geo map object about the number of polygon vertices, then you might need to reduce the density of your polygon data or filter the data query for your geo map object. In some cases, a very wide ID column in your polygon data can further limit the number of polygon vertices that are retrieved. Check the width of your ID column in SAS Data Explorer if you encounter this message.

Loading Polygon Data from Esri Shapefiles

Overview

To load Esri shapefile data into CAS, you must first convert the shapefile into a SAS data set.

SAS provides two autocall macros to help you inspect and load Esri shapefiles:

%SHPCNTNT
: displays the contents of the specified shapefile.

%SHPIMPRT
: converts a shapefile into a SAS data set and loads it into CAS.

> **TIP** Where possible, use shapefiles with unprojected latitude and longitude values. Configuring a polygon provider for projected data can be difficult for users who are inexperienced with map data.

%SHPCNTNT Autocall Macro

The %SHPCNTNT macro displays the contents of the specified shapefile. You can use the %SHPCNTNT macro to identify which variable in the shapefile should be used as an ID variable.

The syntax for the %SHPCNTNT macro is as follows:

%SHPCNTNT(SHAPEFILEPATH=*path-to-shapefile*)

SHAPEFILEPATH=*path-to-shapefile*
: specifies the full path to the shapefile with the .SHP extension. Do not enclose the file path in quotation marks.

%SHPIMPRT Autocall Macro

The %SHPIMPRT macro converts the shapefile into a SAS data set and then loads it into CAS.

Note: To load tables into CAS, you must configure an authentication file. See *Client Authentication Using an Authinfo File*,

The syntax for the %SHPIMPRT macro is as follows:

%SHPIMPRT(*options*)

SHAPEFILEPATH=*path-to-file*
: specifies the full path to the shapefile with the .SHP extension.

ID=*id-column*
: specifies a field in the shapefile that identifies the polygons in the map.

Requirement The ID column must contain character data, and cannot contain special characters or double-byte characters.

OUTTABLE=*table-name*
: specifies the name of the output table that is loaded into CAS.

CASHOST=*machine-name*
: specifies the machine name of the CAS server.

CASPORT=*port-number*
: specifies the port for the CAS server.

CASLIB=*library-name*

specifies the library on the CAS server where the output table is loaded.

REDUCE=0|1

(Optional) specifies whether to reduce the density of the polygon data. A value of 1 specifies that the data density is reduced, and a value of 0 specifies that the data density is not reduced.

Reducing the density of your polygon data can improve performance and might enable a greater number of map regions to be displayed at one time.

Default	0
Requirement	A license for SAS/GRAPH software is required to reduce the density.

The following example loads a shapefile without reducing the polygon density:

```
%shpimprt(shapefilepath=/tmp/myfile.shp, id=GEOID, outtable=mytable, cashost=cloud.example.com,
   casport=5570, caslib='casuser');
```

The following example loads a shapefile and reduces the polygon density:

```
%shpimprt(shapefilepath=/tmp/myfile.shp, id=GEOID, outtable=mytable, cashost=cloud.example.com,
   casport=5570, caslib='casuser', reduce=1);
```

Loading Polygon Data from SAS Map Data Sets

To use a SAS map data set as a polygon provider, you must perform the following steps:

1 Create a sequence variable to enable the polygon segments to be read in the correct order. In a SAS DATA step, you can use the _n_ automatic variable to store the observation number as a sequence variable. For example, the following DATA step creates a sequence variable for the MYMAP data set:

```
data mymap;
  set mymap;
  sequence = _n_;
run;
```

2 (Optional) Subset your polygon data to decrease the level of detail and improve performance. Reducing the level of detail might also enable you to display a greater number of map regions at one time.

If you have a license for SAS/GRAPH, then you can use the GREDUCE procedure to create a DENSITY variable that enables you to reduce the density of your polygon data. Depending on the source of your map data sets, a DENSITY variable might already be present. For more information about the DENSITY variable and the GREDUCE procedure, see SAS/GRAPH and Base SAS 9.4: Mapping Reference.

You can use the DENSITY variable in a WHERE statement in a DATA step to reduce the detail in your polygon data. For example, the following DATA step reduces the MYMAP data set to exclude segments that are density level 4 or greater:

```
data mymap;
  set mymap;
  where(density<4);
run;
```

3 Load the data set in your SAS Cloud Analytic Services environment.

Data Administration: How to (CAS Server Monitor)

Introduction

CAS Server Monitor enables you to monitor and administer your CAS server. Within CAS Server Monitor, the **System State** view contains various CAS server properties and settings, including the **Global Caslibs** table. This table displays the global caslibs for your environment. From here you can add and delete global caslibs and modify access controls for users and groups.

These instructions explain how to manage global caslibs using CAS Server Monitor.

Add a Global Caslib

1 On the **System State** page, select **Global Caslibs**.

2 Click **Add**.

> **TIP** If the **Add** button is disabled, you are not authorized to add a global caslib. For details see Caslib Management Privileges.

3 On the Add Global Caslib pane, specify general settings as follows:

Table A.8 Global Caslibs

Setting	Description
Caslib	Enter a caslib name.
Description	Enter a description for the caslib.
Path	Enter data source-specific information.
Subdirectories	For a path-based caslib, specifies whether tables and files in subdirectories of the specified path are accessible from the caslib.
Create directory	For a path-based caslib, creates the host directory that you specify in the **Path** field, if that directory does not already exist.
Permission	For a path-based caslib, sets host-layer permissions on the directory. See Using CAS to Modify Host Access .
Active on add	Specifies whether the new caslib becomes the active caslib in your current session.
Hidden	Makes the caslib and its tables unlisted in certain contexts. See Reduced Visibility: Hidden Caslibs.
Transient	Specifies that the caslib is scoped to the current session only.
Data source	Specifies the type of source data for the caslib.

Setting	Description
Data encryption password	Specifies the encryption password for the caslib.
Encryption domain	Specifies the encryption domain for the caslib.

4 Specify additional settings as needed. For information about caslib properties, see addCaslib Action.

5 Make sure your settings are as intended. In CAS Server Monitor, caslib properties are not editable.

6 Click **OK**.

Delete a Global Caslib

CAUTION! When you delete a caslib, all associated in-memory tables are immediately dropped.

Note: Deleting a caslib does not affect persisted files in the corresponding data source.

1 On the **System State** page, click **Global Caslibs**.

2 At the end of the row for the caslib, click ⋮ , and select **Drop Caslib**. On the Drop Global Caslib pane click **OK**.

Manage Access to a Global Caslib

1 On the **System State** page, click **Global Caslibs**.

2 At the end of the row for the caslib, click ⋮ , and select **Edit Access Controls**. The Edit Access Controls window appears. From here you can grant or deny permission settings to different users.

See SAS Viya Administration: Cloud Analytic Services Authorization.

Data Administration: Reference

Data Administration: Interfaces

Interfaces

All CAS data management requirements and constraints are always fully enforced. Not all interfaces enable you to see and interact with all CAS data management features.

In the following table, the shaded part of each circle is an approximation of the amount of CAS data management functionality that a particular interface exposes.

Table A.9 Interfaces to Data Administration

Value	Interface	Description
●	Tables Action Set	A programmatic interface for CASL (the CAS procedure), Python, and Lua.

Value	Interface	Description
◕	SAS Environment Manager	The enterprise graphical web application for administration.
◒	CAS Server Monitor	A graphical web application that is embedded in the CAS server. Supports adding and deleting global caslibs.
◔	CASLIB statement	A programmatic interface for adding caslibs. See CASLIB statement.

Predefined Caslibs

The following caslibs are automatically created during deployment. Each caslib has a default assignment and specifications.

Table A.10 Predefined Caslibs

Caslib	Default Assignment
AppData*	**`/opt/sas/viya/config/data/cas/default/appData/`** Stores data that specific applications use for internal purposes.
Formats	**`/opt/sas/viya/config/data/cas/default/formats/`** If you are accessing SAS Viya 3.4 from a Docker container, the path is: **`/cas/data/formats/`**. A shared location for user-defined formats. All users can read. Administrators can read and write.
Models*	**`/opt/sas/viya/config/data/cas/default/models/`** Stores models created by SAS Visual Analytics for use in SAS Studio.
Public	**`/opt/sas/viya/config/data/cas/default/public/`** If you are accessing SAS Viya 3.4 from a Docker container, the path is: **`/cas/data/public/`**. A shared location for data. All users can read and write. See Access to Files in the Public Caslib .
ReferenceData*	**`/opt/sas/viya/config/data/cas/default/referenceData/`** Stores per-server data that specific applications use for internal purposes.
Samples	**`/opt/sas/viya/config/data/cas/default/samples/`** If you are accessing SAS Viya 3.4 from a Docker container, the path is: **`/cas/data/samples/`**. Stores sample data, supplied by SAS.
SystemData*	**`/opt/sas/viya/config/data/cas/default/sysData/`** Stores application-generated data that is used for general reporting.

Caslib	Default Assignment
VAModels	`/opt/sas/viya/config/data/cas/default/vamodels/` If you are accessing SAS Viya 3.4 from a Docker container, the path is: `/cas/data/vamodels/`. This is a library for ASTORE objects that are used within a SAS Visual Analytics report.
ProductData	`/opt/sas/viya/home/share/productData/` Stores product data supplied by SAS.

* Not included in a programming-only deployment.

Note: Some predefined caslibs are hidden or have limited access. For more information about hidden caslibs, see Reduced Visibility: Hidden Caslibs.

Disabling the CASUSERHDFS Caslib

By default, a pre-assigned caslib named CASUSERHDFS is created for each user that points to an HDFS location of `/user/<username>` in HDFS. If these directories for each user do not exist in HDFS, users that select this caslib when loading data in SAS Visual Analytics will receive an error. There are two options to resolve this problem:

- Set `cas.HDFSUSERLOC` to another location in HDFS where all users have Write access.
- Disable the CASUSERHDFS caslib by commenting out `cas.HDFSUSERLOC`.

The `cas.HDFSUSERLOC` option is found in the `casconfig_usermods.lua` file. Changes to this file require a restart of CAS. For more information, see SAS Cloud Analytic Services: Reference.

Access to SAS 9.4 Data

If you are moving data from SAS 9.4 to SAS Viya, you will need to consider some preliminary information:

- You can move and you can share data between SAS 9 and SAS Viya environments using SAS/CONNECT.
- SAS Viya operates with UTF-8 encoded data. If your SAS 9 installation is not UTF-8 compliant, you might need to re-create your data sets.

See the following topics for more information:

- Comparing SAS 9 and SAS Viya
- SAS 9 and SAS Viya
- Sharing Data Between SAS 9 and SAS Viya using SAS/CONNECT
- Migrating Data to UTF-8 for SAS Viya 3.4

12

Jobs

Jobs: Overview

The Jobs page enables you to monitor and schedule jobs from a variety of sources in SAS Viya.

The **Monitoring** tab enables you to view a table or a chart of jobs that are currently running and that have run in a specified time in the past. You can filter the jobs to narrow the number of jobs displayed and change the time period for displaying jobs. You can also rerun jobs and delete jobs from the list.

The **Scheduling** tab enables you to schedule jobs to run at a particular time or in response to a specific trigger. You can run a job immediately, or you can specify a time interval (from daily to yearly) to control when the job runs. You can also unschedule, delete, and view the properties of jobs.

Jobs that are available for scheduling are from these sources:

SAS Data Explorer
Creates jobs that you can schedule using SAS Environment Manager.

SAS Data Studio
Creates jobs that you can schedule using SAS Environment Manager.

SAS Visual Analytics
Creates jobs that are scheduled in SAS Visual Analytics. You can view and modify the schedules in SAS Environment Manager.

CAS table state management
Three jobs are provided to manage CAS tables.

- Import cas-shared-default Public data
- Load cas-shared-default Public data
- Unload cas-shared-default Public data

You can schedule these jobs, but you cannot delete them, and you can modify the job options only on copies of the jobs. If you schedule one of these jobs and then make a copy of the job, only the job is copied, not any triggers that are associated with the job. For more information about these jobs, see “CAS Table State Management ” on page 153.

To access the Jobs page, click ⚙ **Jobs** in the SAS Environment Manager navigation menu.

Jobs: How To

Monitor Jobs

View a Table of Job Executions

By default, when you open the **Monitoring** tab, the **Monitor** table displays a list of all jobs that have executed in the previous 24 hours. The table displays the job name, the start and end date and time, the run time, the job status, and the user that submitted the job. If the job has completed, the table also includes a link to download the log for the job, if one was created. You can also list the environment in which the job ran, although this column is not displayed by default. For information about changing the columns that are displayed, see "Work with Information Displayed in Tables" on page 673 .

Note: If you change to a different time zone, the new time zone is not automatically reflected in the **Monitor** table. Close and reopen your browser to use the new time zone in the **Monitor** table.

If a job did not complete successfully, the message **Failed** appears in the **Status** column. Click the message to view information (if available) about the reason for the failure.

View a Chart of Job Executions

From the **Monitoring** tab, click ☰ to display a chart of the jobs that have executed in the selected time period (the default is the previous 24 hours). The sliders below the graph enable you to zoom into a specific time window within the selected time period.

Jobs that ran successfully are displayed in green. Jobs that failed are displayed in red.

Place your cursor over a bar in the chart to display the name, start time, and status of the job.

Filter Details about Jobs

You can specify filters to narrow the jobs that are displayed in the **Monitoring** tab. For example, you can specify that only jobs that failed or only jobs that were created by a specific user are displayed.

To filter by job status, select one or more check boxes in the **Status** list that you want to display.

To filter by creator, select one or more check boxes in the **Created By** list. You can enter text in the **Filter** text box to find an existing creator or to specify a creator. You can filter by creator only if you opted in to the SAS Administrators group when you signed in to SAS Environment Manager.

After you have selected all the filters that you want to use, click **Apply**. The filters affect the jobs that are displayed in both the table of jobs and the jobs bar chart.

To remove a filter, deselect its check box and click **Apply**. To remove all filters in either the **Status** or **Created By** list, click **Reset** next to the list. To remove all filters, click **Reset all**.

Rerun a Job from the Monitoring Tab

To rerun a job, right-click the job in the **Jobs** table and select **Run Now** from the pop-up menu. A copy of the job is created and is displayed in the list.

Delete a Job from the Monitoring Tab

Details about jobs remain in the list on the **Monitoring** tab unless you delete the entry. To delete the entry for a job execution, right-click the entry in the **Jobs** table and select **Delete** from the pop-up menu.

View the Job Log

If the job execution component for a job generated a log, you can download the log file for further analysis. Not all jobs create a log. Click **Download** in the **Log** column to save or open a local copy of the log file. The specific behavior depends on your browser.

Schedule Jobs

Schedule a Job

1 On the **Scheduling** tab, click ▦ to display the **Jobs** table. The table is displayed by default when you open the **Scheduling** tab. By default, the table displays the job name, scheduled status, description, and the date on which the job was created. You can also choose to display the ID of the user that created the job, the date on which the job was last modified, the ID of the user who last modified the job, the job ID, the scheduled job ID, and the job type. These columns are not displayed by default. For information about changing the columns that are displayed,see "Work with Information Displayed in Tables" on page 673.

2 Select a job in the **Jobs** table.

3 Click 🕒 in the toolbar or select **Schedule** from the pop-up menu.

4 (Optional) To run the job under credentials other than your own, in the Schedule Job window, specify the user ID under which the job should be run in the **Run as** field. Click to select from specified identities. The user that you select must have previously signed in to SAS since it was installed.

5 Activate the **Enabled** control for one or more triggers in the **Available triggers** table. A trigger controls when the job runs. See "Create a Time Trigger" on page 171 to define a new trigger. You can use a trigger only with the job for which it was created.

Note: Currently, **Time Event** is the only supported trigger type.

6 Click **Save**.

7 Verify that the listing for the job in the **Jobs** table contains 🕒 in the **Scheduled** column.

Create a Time Trigger

1 In the Schedule Job window, click + above the **Available triggers** table.

2 In the New Trigger window, assign a name to the new trigger. The name is specified as **New trigger** by default.

3 Use the **Frequency** field to specify how often the trigger should be repeated (such as a specified number of minutes, hours, or days).

4 Depending on your choice for the frequency interval, different fields appear in the window to enable you to completely specify a frequency for the trigger. For example, if you select **Yearly** in the **Frequency** field, you can specify a day of a month (such as the first of January), the last day of a month, or a specific weekday in a month (such as the third Thursday in February). If you specify **Minutes** in the **Frequency** field, you can specify that the job runs every 5, 10, 15, 20, or 30 minutes. Use these fields to specify the criteria for the trigger interval.

Note: If you select **Date List** in the **Frequency** field, you cannot select a date more than once.

5. In the **Start time** field, specify when the job schedule should start. Click the entry in the **Start time** field to select a time. Times are specified in 24-hour format.

   For example, if you use the **Frequency** fields to specify that the job runs every hour, and you specify **10:15** in the **Start time** field, the job runs at 10:15, 11:15, 12:15, and so on. If you use the **Frequency** fields to specify that the job runs every 20 minutes, and you specify **09** in the **Start time** field, the job runs at 9:00, 9:20, 9:40, and so on.

6. Specify the time zone to use when evaluating the time for the trigger, and the date on which the trigger starts.

   Note: If you choose **Date List** in the **Frequency** field, you must select the same value in the **Time zone** field for every scheduled date.

7. Specify when the trigger ends. You can specify that the trigger never ends, that it ends after a certain number of times, or that it ends on a specific date.

8. Click **Save**.

9. Repeat these steps to create other triggers for the job.

Edit a Scheduled Job

After a job is scheduled, you can edit the schedule for the job. Follow these steps:

1. Select a scheduled job in the **Jobs** table on the **Scheduling** tab. Scheduled jobs with at least one enabled trigger contain ◷ in the **Scheduled** column. Scheduled jobs with disabled triggers contain a disabled icon ❚❚ in the **Scheduled** column.

2. To modify the schedule for the job, click ◷ or select **Edit Schedule** from the pop-up menu. In the Edit Schedule window, you can add, edit, and remove triggers for the job. Click **Save** when you have finished modifying the schedule

View a Graph of Scheduled Jobs

From the **Scheduling** tab, click ☰ to display a chart of the jobs that are scheduled over a selected time period. The default time period is one year, and all scheduled jobs are shown. The sliders below the graph enable you to zoom into a specific time window within the selected time period.

Each scheduled job is listed on a separate line in the graph. Bars in the chart represent each scheduled execution of a job.

Place your cursor over a bar to display the job name, status, and the date and time of the scheduled execution.

Disable the Schedule for a Job

To prevent a job from running its specified schedule, you can either unschedule the job or disable the triggers. Unscheduling the job prevents the job from running on the defined schedule and also removes the triggers that are specified for the job. Disabling the triggers prevents the job from running the schedule but preserves the defined triggers.

To unschedule a job, select a scheduled job in the **Jobs** table in the **Scheduling** tab. Click ⊘ from the toolbar or select **Unschedule** from the pop-up menu.

CAUTION! When you unschedule a job, any enabled triggers that are associated with the job are deleted. To unschedule a job and keep the triggers, instead of selecting **Unschedule**, edit the schedule and manually disable the triggers.

To disable the triggers, select a scheduled job in the **Jobs** table in the **Scheduling** tab. Click 🕒 from the toolbar or select **Edit Schedule** from the pop-up menu. In the Edit Schedule dialog box, disable all slider controls in the **Enabled** column of the **Available triggers** table. If you disable all triggers for a job, the disabled icon ⏸ appears in the **Scheduled** column of the **Jobs** table. A 🕒 appears in the column if any of the triggers for the job are enabled.

Run a Job

1 To run a job from the **Scheduling** tab, right-click a job in the **Jobs** table in the **Scheduling** tab.

2 To run the job under your own credentials, click ▶ in the toolbar or select **Run** from the pop-up menu.

3 To run the job under credentials other than your own, select **Run As** from the pop-up menu.

The Select Identities window appears, and you can select the user ID under which the job should run.

Note: The user ID that you select must have previously signed in to SAS.

You can run a job regardless of whether it has been scheduled.

View Execution History for a Job

You can view information about previous runs of a job that is available for scheduling.

1 In the **Scheduling** tab, right-click a job in the **Jobs** table.

2 Select **Execution history** in the pop-up menu or select 📋 from the toolbar.

3 Information about previous runs of the selected job is displayed in the **Monitoring** tab.

View Job Properties

To view properties for a job, select a job in the **Jobs** table and click ▤ in the toolbar or select **Properties** from the pop-up menu. The information in the Job properties window is read-only.

Delete a Job from the Schedule Tab

Jobs remain in the list on the **Scheduling** tab unless you delete them. To delete a scheduled job, follow these steps.

1 Select a job in the **Jobs** table.

2 Click 🗑 in the toolbar or select **Delete** from the pop-up menu.

Note: You cannot delete any of the provided CAS table state management jobs (Import cas-shared-default Public data, Load cas-shared-default Public data, and Unload cas-shared-default Public data).

Scheduling Command Line Interface

Scheduling: How to (Command Line Interface)

Run a Job

In order to run a job using the command-line interface, follow these steps:

1 Create a template file for the job definition. This file contains the fields that are needed for a job definition.

```
sas-admin job definitions generate-template --template-filename
```

Here is an example of the job definition template:

```
template:
{
    "name": "Replace with name of the Job Definition",
    "type": "Replace with type of the Job Definition",
    "code": "Replace with code of the Job Definition"
}
```

2 Modify the job definition template file to supply information for the job that you want to run.

3 Use the job definition file that you created in the previous step to create the job definition.

```
sas-admin job definitions create --definition-filename
```

This command returns a URI for the job definition.

4 Create a template file for the job request. This file contains the fields that are needed for a job request.

```
sas-admin job requests generate-template --template-filename
```

Here is an example of the job request template file:

```
template:
{
    "version": 0,
    "name": "Replace with name of the Job Request",
    "description": "Replace with description of the Job Request",
    "jobDefinitionUri": "(Mutually exclusive with Definition) Replace with uri to the Job Definition",
    "arguments": null,
    "properties": null
}
```

5 Modify the generated job request template file to supply information for the job that you want to run.

6 Use the job request file that you created in the previous step to create the job request.

```
sas-admin job requests create --request-filename
```

The command returns an ID for the job request.

7 Execute the request. The job runs immediately.

```
sas-admin job requests execute --request-ID
```

Schedule a Job

In order to schedule a single job using the command-line interface, follow these steps:

1 Create a template file for the job definition. This file contains the fields needed for a job definition.

```
sas-admin job definitions generate-template --template-filename
```

Here is an example of the job definition template:

```
template:
{
    "name": "Replace with name of the Job Definition",
    "type": "Replace with type of the Job Definition",
    "code": "Replace with code of the Job Definition"
}
```

2. Modify the job definition template file to supply information for the job that you want to run.

3. Use the job definition file that you created in the previous step to create the job definition.

   ```
   sas-admin job definitions create --definition-filename
   ```

   This command returns a URI for the job definition.

4. Create a template file for the job request. This file contains the fields needed for a job request.

   ```
   sas-admin job requests generate-template --template_filename
   ```

   Here is an example of the job request template:

   ```
   template:
   {
       "version": 0,
       "name": "Replace with name of the Job Request",
       "description": "Replace with description of the Job Request",
       "jobDefinitionUri": "(Mutually exclusive with Definition) Replace with uri to the Job Definition",
       "arguments": null,
       "properties": null
   }
   ```

5. Modify the generated job request template file to supply information for the job that you want to run.

6. Use the job request file that you created in the previous step to create the job request.

   ```
   sas-admin job requests create --request-filename
   ```

   The command returns an ID of the job request.

7. Create a JSON file and include the time triggers for the job. See “Time-Based Triggers” on page 175 for information about specifying the triggers.

8. Schedule the job, and specify the file that contains the time triggers.

   ```
   sas-admin job requests schedule --triggers-file
   ```

Scheduling: Command-Line Interface Reference

Time-Based Triggers

Use the following syntax when specifying a time-based trigger to schedule a flow or a job.

Here is the general form of the syntax:

```
"triggers":[
               {"type" : "timeevent",
               "active":true,
               "event": {"recurrence":{"type":"recurrence-type"}, options,
                          "hours":hours,
                          "minutes":minutes,
                          "duration":duration,
                          "timeZone":zone,
                          "maxOccurrence":occurrences,}
               }
               ],
```

This list identifies the options that are used for each type of `recurrence` interval.

Minutes

Type

`"type":"minutely"`

Options

- `startDate` (specifies when to start the recurrence)
- `endDate` (specifies when to stop the recurrence)
- `skipCount` (specifies how many minutes pass between executions). For example, `"skipCount":"15"` specifies that the job runs every 15 minutes. Valid values are 1, 2, 3, 4, 5, 6, 10, 12, 15, 20, and 30.
- `minutes` (specifies the offset from the beginning of the hour for when the executions start). The maximum value is `skipCount`-1. For example, `"minutes":"10"` specifies that the timing for `skipCount` starts at 10 minutes past the hour.

Here is an example that specifies a trigger that starts at six minutes past the hour and then causes a job to run every 15 minutes (06, 21, 36, 51, and so on).

```
"triggers":[
               {"type" : "timeevent",
               "active":true,
               "event": {"recurrence":{"type":"minutely"}, "skipCount":"15"}
               "minutes":"6", }
               ],
```

Hours

Type

`"type":"hourly"`

Options

- `startDate` (specifies when to start the recurrence)
- `endDate` (specifies when to stop the recurrence)
- `skipCount` (specifies how many hours pass between executions). For example, `"skipCount": "3"` specifies that the job runs every 3 hours. Valid values are 1, 2, 3, 4, 6, 8, and 12.

Here is an example that specifies a trigger that starts on June 13, 2018 and causes a job to run every 4 hours:

```
"triggers":[
               {"type" : "timeevent",
               "active":true,
               "event": {"recurrence":{"type":"hourly","startDate":"2018'-'06'-'13",
                    "skipCount":"4"}}
               }
               ],
```

Days

Type

`"type":"daily"`

Options

- `startDate` (specifies when to start the recurrence)
- `endDate` (specifies when to stop the recurrence)
- `skipCount` (specifies how many days pass between executions). For example, a value of 3 specifies that the job runs every 3 days.

- `daysOfWeek` (specifies the days on which the job runs). For example, `"daysOfWeek":"thursday"` specifies that the job runs every Thursday. Valid values are names of days: monday–sunday.

Here is an example of a trigger:

```
"triggers":[
                {"type" : "timeevent",
                "active":true,
                "event": {"recurrence":{"type":"daily"}, "hours":"12","minutes":"0"}
                }
                ],
```

Weeks

Type

`"type":"weekly"`

Options

- `startDate` (specifies when to start the recurrence)
- `endDate` (specifies when to stop the recurrence)
- `skipCount` (specifies how many weeks pass between executions). For example, `"skipCount": "3"` specifies that the job runs every 3 weeks.
- `daysOfWeek` (specifies the days on which the job runs). For example, `"daysOfWeek":"thursday"` specifies that the job runs every Thursday. Valid values are names of days (monday–sunday).

Example:

```
"triggers":[
                {"type" : "timeevent",
                "active":true,
                "event": {"recurrence":{"type":"weekly"}, "startDate":"2017'-'08'-'15",
"skipCount":"4", "daysOfWeek":"tuesday"}
                }
                ],
```

Months

Type

`"type":"monthly"`

Options

- `startDate` (specifies when to start the recurrence)
- `endDate` (specifies when to stop the recurrence)
- `skipCount` (specifies how many months pass between executions). For example, `"skipCount": "3"` specifies that the job runs every 3 months.
- `daysOfWeek` (specifies the days on which the job runs). For example, `"daysOfWeek":"thursday"` specifies that the job runs on Thursday. Valid values are names of days (monday-sunday). `daysOfWeek`and `dayOfMonth` are mutually exclusive. If `daysOfWeek` is specified, then `occurrence` is required.
- `occurrence` (used with `daysOfWeek` to specify the days on which the job runs) Valid values are first, second, third, fourth, and last.
- `dayOfMonth` (specifies the day of the month on which to run. Valid values are 1–31. A value of 32 specifies the last day of the month. `dayOfMonth` and `daysOfWeek` are mutually exclusive.

Example 1:

```
"triggers":[
                {"type" : "timeevent",
                "active":true,
                "event": {"recurrence":{"type":"monthly"}, "startDate":"2017'-'05'-'14",
"daysOfWeek":"sunday", "occurrence":"second"}
                }
                ],
```

Example 2:

```
"triggers":[
                   {"type" : "timeevent",
                   "active":true,
                   "event": {"recurrence":{"type":"monthly"}, "dayOfMonth":"15"}
                   }
                   ],
```

Years

Type

`"type":"yearly"`

Options

- `startDate` (specifies when to start the recurrence)
- `endDate` (specifies when to stop the recurrence)
- `skipCount` (specifies how many years pass between executions)
- `daysOfWeek` (specifies the days on which the job runs). For example, `"daysOfWeek":"thursday"` specifies that the job runs on Thursday. Valid values are names of days (monday-sunday). `daysOfWeek` and `dayOfMonth` are mutually exclusive. If `daysOfWeek` is specified, then `occurrence` is required.
- `occurrence` (used with `daysOfWeek` to specify the days on which the job runs). Valid values are first, second, third, fourth, and last.
- `dayOfMonth` (specifies the day of the month on which to run). Valid values are 1–31. A value of 32 specifies the last day of the month. `dayOfMonth` and `daysOfWeek` are mutually exclusive.
- `monthOfYear` (specifies the month in which the job run)Valid values are january–december.

Example 1:

```
"triggers":[
                   {"type" : "timeevent",
                   "active":true,
                   "event": {"recurrence":{"type":"yearly"}, "daysOfWeek":"friday",
                      "occurrence":"last", "monthOfYear":"june"}
                   }
                   ],
```

Example 2:

```
"triggers":[
                   {"type" : "timeevent",
                   "active":true,
                   "event": {"recurrence":{"type":"yearly"}, "dayOfMonth":"32", "monthOfYear":"may"}
                   }
                   ],
```

Specified dates

Type

`"type":"dateList"`

Options

- `startDate` (specifies when to start the recurrence)

- `endDate` (specifies when to stop the recurrence)
- array of dates in the form `yyyy '-' mm '-' dd`

Example:

```
"triggers":[
                {"type" : "timeevent",
                "active":true,
                "event": {"recurrence":{"type":"dateList"}, "2017 '-' 06 '-' 13",
                   "2017 '-' 08 '-' 02", "2017 '-' 10 '-' 05"}
                }
                ],
```

For the `"hours":hours` parameter, specify a set of hours when the job runs. You can specify a list of hours separated by commas (1,2,3), a range of hours (2–4), a combination of a range and a list (1–3,5,7), or an asterisk to specify all hours. If you specify a recurrence of hourly, only the first value is used, and it must be equal to or less than the value of skipCount. If you specify a recurrence of minutely, the hours parameter is ignored.

For the `"minutes":minutes` parameter, specify a set of minutes when the job runs. You can specify a list of minutes separated by commas (0,10,30), a range of minutes (20–25), a combination of a range and a list (0,10–15), or an asterisk to specify every minute. If you specify a recurrence of minutely, only the first value is used, and it must be equal to or less than the value of skipCount.

For the `"duration":duration` parameter, specify the number of minutes the event is to remain true.

For the `"timeZone":zone` parameter, specify the time zone to use when evaluating the time trigger. Specify the value using the Olson time zone database, in the form `region/city`. For example, `America/New_York`.

For the `"maxOccurrence":occurrences` parameter, specify the maximum number of times the job can execute.

13

SAS Licensing

Licensing: Overview

SAS Viya provides a license in two file formats: a traditional text file and a new JSON web token file. The JSON file includes the same information as the text file, and enables additional licensing information to be provided. Some SAS Viya products use the text file format. Other products use the JSON file format. Both SAS Cloud Analytic Services (CAS) and the SAS programming run-time environment use the same license.

During installation, a license is applied to both the CAS in-memory compute engine and the SAS programming run-time environment. You apply a new license to enable new products or to extend expiration dates on existing products.

The following diagram identifies where the license file resides in SAS Viya.

For more information, see How To on page 182.

Figure A.1 Where the SAS License File Resides

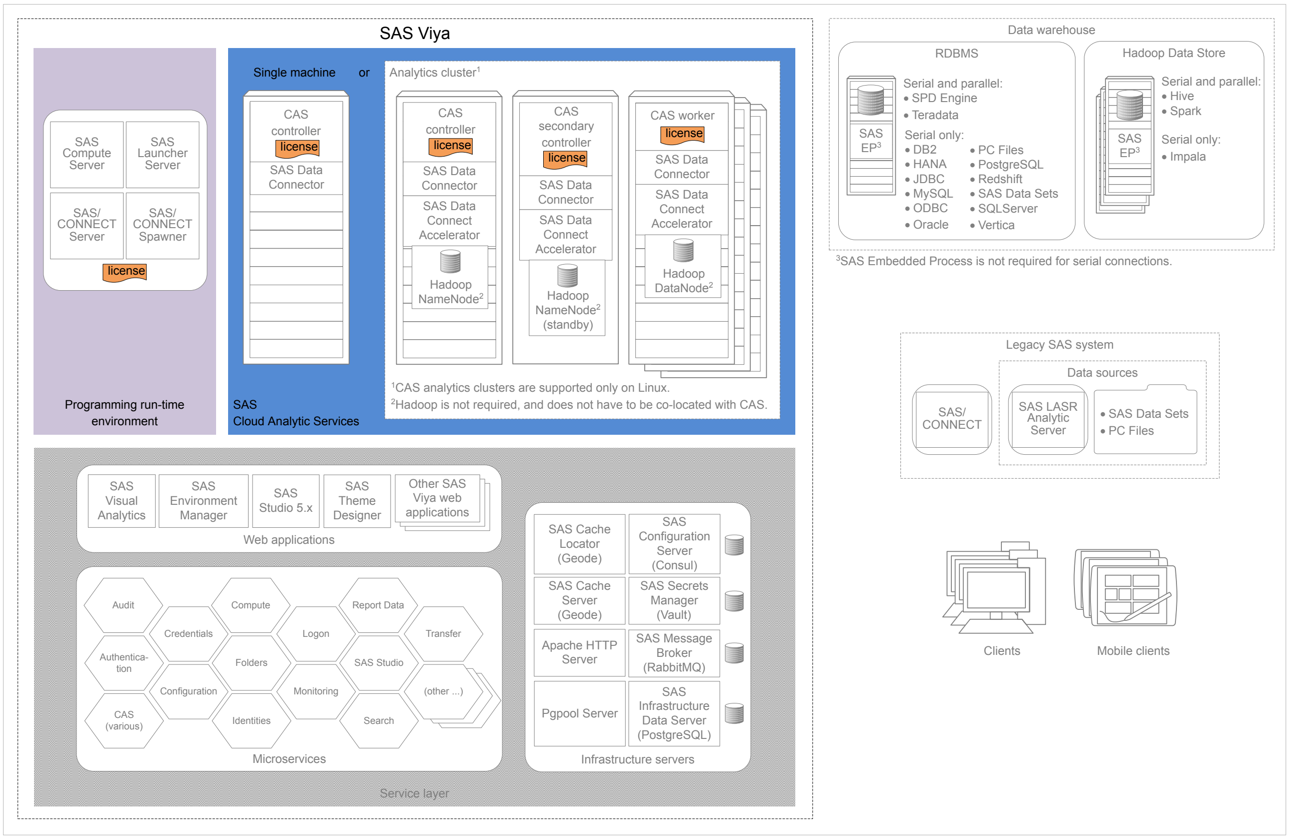

Licensing: How To

Apply New Licenses (Linux)

Apply New Licenses Using Ansible

You apply a new SAS license when your current license has expired, or when you are adding new SAS products to your deployment. If your deployment was performed using Ansible, you can use Ansible to apply a new license. Ansible applies your new license to the CAS controllers—primary and backup—and also to the SAS programming run-time environment.

Note: To add a new license without using Ansible, see "Apply New Licenses Manually" on page 183.

1 Log on to your Ansible controller machine with a user that meets the requirements in "Set Up the User Account that Deploys the Software" in *SAS Viya for Linux: Deployment Guide*.

2 Move the current license files into a backup location.

   Copies of your current license files should reside in your Ansible playbook directory (`sas_viya_playbook/`, by default).

   The license files are named, SASViyaV0300_*order-number_site-number*_Linux_x86-64.jwt and SASViyaV0300_*order-number*_Linux_x86-64.txt.

3. SAS distributes renewal licenses to customers as file attachments in a renewal order email (ROE). Make sure that your new license files (a .txt file and a .jwt file) reside in your Ansible playbook directory.

   Note: Some SAS Viya products use the text file (.txt). Other products use the JSON web token file (.jwt). Both SAS Cloud Analytic Services (CAS) and the SAS programming run-time environment use the same license.

4. Modify your Ansible playbook to point to the new license files.

   Open **sas_viya_playbook/vars.yml**, locate `LICENSE_FILENAME` and `LICENSE_COMPOSITE_FILENAME`, and replace the current license filename with the corresponding new license filename.

   Note: The JSON web token license file (.jwt) is also referred to as a composite license.

   Here is an example:

   ```
   # The name of the license file on the Ansible machine.
    LICENSE_FILENAME: "SASViyaV0300_09MMMV_Linux_x86-64.txt"

   # The name of the composite license file on the Ansible machine.
   # If both files are present, the playbook will use the
   # composite license file.
    LICENSE_COMPOSITE_FILENAME: "SASViyaV0300_09MMMV_70180938_Linux_x86-64.jwt"
   ```

5. Run the following Ansible command for the default inventory file:

   ```
   ansible-playbook apply-license.yml
   ```

   Important: If you deployed additional CAS servers, run the `ansible-playbook` command with the `-i` option using the appropriate inventory file.

   CAS sessions created after you apply the new license automatically update with license information from the new license.

6. Verify that your SAS Cloud Analytic Services license has been renewed by following the steps in "View SAS Cloud Analytic Services License Information" on page 188.

7. Verify that your SAS programming run-time license has been renewed by following the steps in "View SAS Programming Run-Time License Information" on page 188.

8. If you deployed additional CAS servers, then perform Step 5 – Step 6 on your additional controller and your additional backup controller machines.

Apply New Licenses Manually

You apply a new SAS license when your current license is about to expire, or when you are adding new SAS products to your deployment. You must apply your license to all CAS controllers—primary and secondary—and also to the SAS programming run-time environment.

Note: To add a new license using Ansible, see "Apply New Licenses (Linux)" on page 182.

1. On the machine where the SAS programming run-time environment is deployed, log on as a user that meets the requirements in "Set Up the User Account that Deploys the Software" in *SAS Viya for Linux: Deployment Guide*.

2. Move the current license files into a backup location.

   Your current license files reside in **/opt/sas/spre/home/SASFoundation/**.

   The license file is named license.

3 SAS distributes renewal licenses to customers as file attachments in a renewal order email (ROE). Make sure that your new license files (a .txt file and a .jwt file) reside in location that is accessible from your SAS programming run-time machine.

Note: Some SAS Viya products use the text file (.txt). Other products use the JSON web token file (.jwt). Both SAS Cloud Analytic Services (CAS) and the SAS programming run-time environment use the same license.

4 Run the following command to apply the license to your SAS programming run-time environment:

```
sudo su -s "/bin/sh" -c
"/opt/sas/spre/home/SASFoundation/utilities/bin/apply_license
/path/SASViyaVrelease-number_order-number_site-number_Linux_x86-64.jwt" sas
```

where *path* is the location where the new license file resides.

Here is an example:

```
sudo su -s "/bin/sh" -c
"/opt/sas/spre/home/SASFoundation/utilities/bin/apply_license
/opt/sas/installfiles/SASViyaV0300_09MMMV_70180938_Linux_x86-64.jwt" sas
```

You receive a message that your license has been applied.

5 Verify that your SAS programming run-time license has been renewed by following the steps in "View SAS Programming Run-Time License Information" on page 188.

6 On the machine where the CAS controller is deployed, log on as a user that meets the requirements in "Set Up the User Account that Deploys the Software" in *SAS Viya for Linux: Deployment Guide*.

7 Make sure that your new license files (a .txt file and a .jwt file) reside on your CAS controller machine in the following directory: **`/opt/sas/viya/config/etc/cas/default/`**.

Important: On machines that contain additional CAS servers, the path for the license file is **`/opt/sas/viya/config/etc/cas/`*****`cas-instance-name`***.

8 Update the symbolic link for sas_license.txt to point to the new CAS license file. (This should be the .jwt file.)

Here is an example:

```
cd /opt/sas/viya/config/etc/cas/default
ln -sf SASViyaV0300_09MMMV_70180938_Linux_x86-64.jwt sas_license.txt
```

Here is an example for an additional CAS server:

```
cd /opt/sas/viya/config/etc/cas/casserver2
ln -sf SASViyaV0300_09MMMV_70180938_Linux_x86-64.jwt sas_license.txt
```

CAS sessions created after you apply the new license automatically update with information from the new license file.

9 If you are running a multi-tenant deployment and have multiple CAS servers on a single machine, repeat Step 7 – Step 8 for each tenant instance of the CAS server (for example, **`/opt/sas/tenant-1/config/etc/cas/default`**, **`/opt/sas/tenant-2/config/etc/cas/default`**, and so on).

10 If you have a distributed CAS server, repeat Step 6 – Step 8 for each CAS worker node machine.

11 Verify that your SAS Cloud Analytic Services license has been renewed by following the steps in "View SAS Cloud Analytic Services License Information" on page 188.

12 If you deployed a CAS backup controller (also referred to as a secondary controller), then perform Step 7 – Step 11 on your backup controller machine.

13 If you deployed additional CAS servers, then perform Step 7 – Step 11 on your additional controller and your additional backup controller machines.

14 Check the administration documentation for your SAS Viya product in case there are additional steps required for applying a license.

Apply New Licenses (Windows)

You apply a new SAS license when your current license is about to expire, or when you are adding new SAS products to your deployment. The license file is used by both SAS Foundation and SAS Cloud Analytic Services (CAS).

1 Log on to the SAS Viya machine as a user that is a member of the Windows Administrators group.

2 Move the current license files into a backup location.

Copies of your current license files should reside in the directory where your SAS Viya deployment scripts were created.

> **TIP** If you followed the recommendation in the *SAS Viya for Windows: Deployment Guide*, your deployment scripts and license files reside underneath **`\sas\install`**.

The license files are named, SASViyaV0300_*order-number*_Win_x64_Wrkstn_Srv.jwt and SASViyaV0300_*order-number*_Win_x64_Wrkstn_Srv.txt.

3 SAS distributes renewal licenses to customers as file attachments in a renewal order email (ROE). Make sure that your new license files (a TXT file and a JWT file) reside in the same directory as your SAS Viya deployment scripts.

Note: Some SAS Viya products use the text file (TXT). Other products use the JSON web token file (JWT). Both SAS Cloud Analytic Services (CAS) and SAS Foundation use the same license.

4 Change to the deployment scripts directory, open the file, vars.psd1, and modify the license file names to match the new license file name.

Here is an example:

```
COMPOSITE_LICENSE_FILENAME = "SASViyaV0300_09ML1N_70196364_Win_x64_Wrkstn_Srv.jwt"
LICENSE_FILENAME = "SASViyaV0300_09ML1N_Win_x64_Wrkstn_Srv.txt"
```

5 Run the following command to apply your new SAS license:

```
setup.bat -apply-license
```

6 Verify that your SAS Foundation license has been renewed by following the steps in "View SAS Programming Run-Time License Information" on page 188.

7 Verify that your SAS Cloud Analytic Services license has been renewed by following the steps in "View SAS Cloud Analytic Services License Information" on page 188.

Set Up Metered Billing (Linux)

Overview

Metered billing is a pricing model where the fees that you pay SAS are based on your usage.

Important: If you have contracted with SAS for one or more metered products, your contract requires that you set up metered billing. If you have questions, contact your SAS Sales representative.

If metered billing is part of your SAS contract, then the **License** section of your Software Order Email (SOE) indicates that **Your order has a metered offering**. In this release of SAS Viya, metered billing is available only on full deployments running on Linux.

Here is an overview of the steps that must be performed to set up metered billing:

1 Open the HTTPS port for the SAS Viya Metered Billing agent service to connect to edge-metering.sas.com.

2 Set the cas.MAXCORES option on the primary CAS controller for all CAS servers in the deployment.

Set Up Metered Billing

1 On the machine on which the SAS Viya Metered Billing agent runs, make sure that the agent can connect externally to edge-metering.sas.com over the HTTPS port (TCP port 443, by default).

> **TIP** You can identify the machine that hosts the Metered Billing agent by examining inventory.ini in your Ansible playbook (**`/sas/install/sas_viya_playbook/`**) and locating the machine name mapped to the `[Operations]` host group.

From a Linux prompt, enter the following command:

```
curl -I https://edge-metering.sas.com/ --insecure
```

You should see a response similar to the following:

```
HTTP/1.1 302 Found
Location: https://support.sas.com/
Date: Fri, 01 Feb 2019 17:05:21 GMT
Content-Type: text/plain; charset=utf-8
```

Note: The HTTPS port must be opened for outgoing traffic. The Metered Billing agent does not listen to the port.

2 If you want to ensure that your site stays within your SAS license threshold, then go to Step 3. Otherwise, you are finished setting up metered billing.

3 Configure your CAS server with the limit for the total number of physical cores.

Log on to your Ansible controller machine as an administrator, and open the vars.yml file. Under `CAS_CONFIGURATION: cfg:`, add the following line, and then re-run your playbook:

```
maxcores='number-of-cores'
```

where *number-of-cores* specifies the limit for the total number of physical cores that are available to a CAS server.

Here is an example:

```
CAS_CONFIGURATION:
   env:
     CAS_DISK_CACHE: /mydisk/mydiskcache
     CAS_VIRTUAL_HOST: 'loadbalancer.example.com'
     CAS_VIRTUAL_PROTO: 'https'
     CAS_VIRTUAL_PORT: 443
   cfg:
     gcport: 0
     httpport: 8777
     port: 5570
     colocation: 'none'
     SERVICESBASEURL: 'https://loadbalancer.company.com'
     maxcores: '36'
```

Important: Pay attention to indentions when you add content to vars.yml. For information about the vars.yml file, see “Modify the vars.yml File” in *SAS Viya for Linux: Deployment Guide*.

The core count limit is server-wide, and for distributed CAS servers the value should be at least the same as the total number of machines. The total number of machines includes the primary controller and workers. (The backup controller is not included in this total.) For example, if a distributed CAS server has one controller and one worker, and `maxcores: '4'`, the maximum number of cores that the worker can use is two. If you set `maxcores` too low, CAS writes a licensing error.

In this example, we want to ensure that exactly 128 hyperthreads per worker are run. (Hyperthreads equal two times the number of cores.)

- For a single-machine CAS server, you enter `maxcores: '64'`.
- For a distributed CAS server, use the formula, `(Number of workers + 1) * 64`. For example, to ensure that 128 hyperthreads per worker are run for a distributed CAS server that has a controller plus eight workers, you enter `maxcores: '576'`.

4 If you have more than one CAS server, repeat Step 3 for each primary CAS controller in your SAS Viya deployment.

Licensing: How to (SAS Environment Manager)

Introduction

These instructions explain how to view product license information using SAS Environment Manager.

Navigation

In the applications menu (≡), under **Administration**, select **Manage Environment**. In the navigation bar, click 📜.

The Licensed Products page is an advanced interface. It is available to only SAS Administrators.

Licensed Products Page

Use the Licensed Products page to view and filter license status for one or more SAS products.

For each product, the following icons depict the effective license status:

✓	The SAS license is current.
✓⚠	The SAS license is due for renewal (grace period). The grace period is a predetermined range of days immediately after the license expiration date. For example, if the expiration date is 30 June, the grace period might extend 45 days: from 1 July - 14 August.
⚠	The SAS license is about to expire (warning period). The warning period is a predetermined range of days that follows the grace period. For example, if the expiration date is 30 June, the warning period might extend 56 days: from 15 August - 09 October.

⊗	The SAS license has expired. License expiration occurs immediately after the warning period ends. An expired license means that SAS does not run. For example, if the warning period ends on 09 October, SAS stops running at 12:00 a.m. on 10 October.

Licensing: How To (SAS Studio)

View SAS Programming Run-Time License Information

1 Open a web browser and sign in to SAS Studio with administrator privileges.

Here is an example:

`https://mysasserver.example.com/SASStudioV`

Note: If your site uses a programming-only deployment, or your site uses SAS Studio 4.x, then in the preceding example, replace **`SASStudioV`** with **`SASStudio`**.

2 In the **Code** tab, enter the following command:

```
proc setinit; run;
```

3 Click 🏃.

You should see output similar to the following:

```
 56         proc setinit;
 57
 58
 59         OPTIONS NONOTES NOSTIMER NOSOURCE NOSYNTAXCHECK;
 Original site validation data
 Current version: V.03.03M0P040416
 Site name:    'smp statistics, ml, data connectors pkg chg 3.30'.
 Site number:  70068118.
 Expiration:   22MAY2018.
 Grace Period:  45 days (ending 06JUL2018).
 Warning Period: 56 days (ending 31AUG2018).
 System birthday:   24MAR2016.
 Operating System:   LIN X64 .
 Product expiration dates:
 ---Base SAS Software            22MAY2018
 ---SAS/CONNECT                  22MAY2018
.
.
.
```

View SAS Cloud Analytic Services License Information

1 Open a web browser and sign in to SAS Studio with administrator privileges.

Here is an example:

`https://mysasserver.example.com/SASStudioV`

Note: If your site uses a programming-only deployment, or your site uses SAS Studio 4.x, then in the preceding example, replace **`SASStudioV`** with **`SASStudio`**.

> **TIP** To obtain license information without running SAS Studio, run the getLicenseInfo Action from any CAS programming language client. For more information, see getLicenseInfo Action in the *SAS Viya: System Programming Guide*.

2 In the **Code** tab, enter the following commands:

```
cas casauto;
```

```
cas casauto listabout;
```

3 Click 🏃.

You should see output similar to the following:

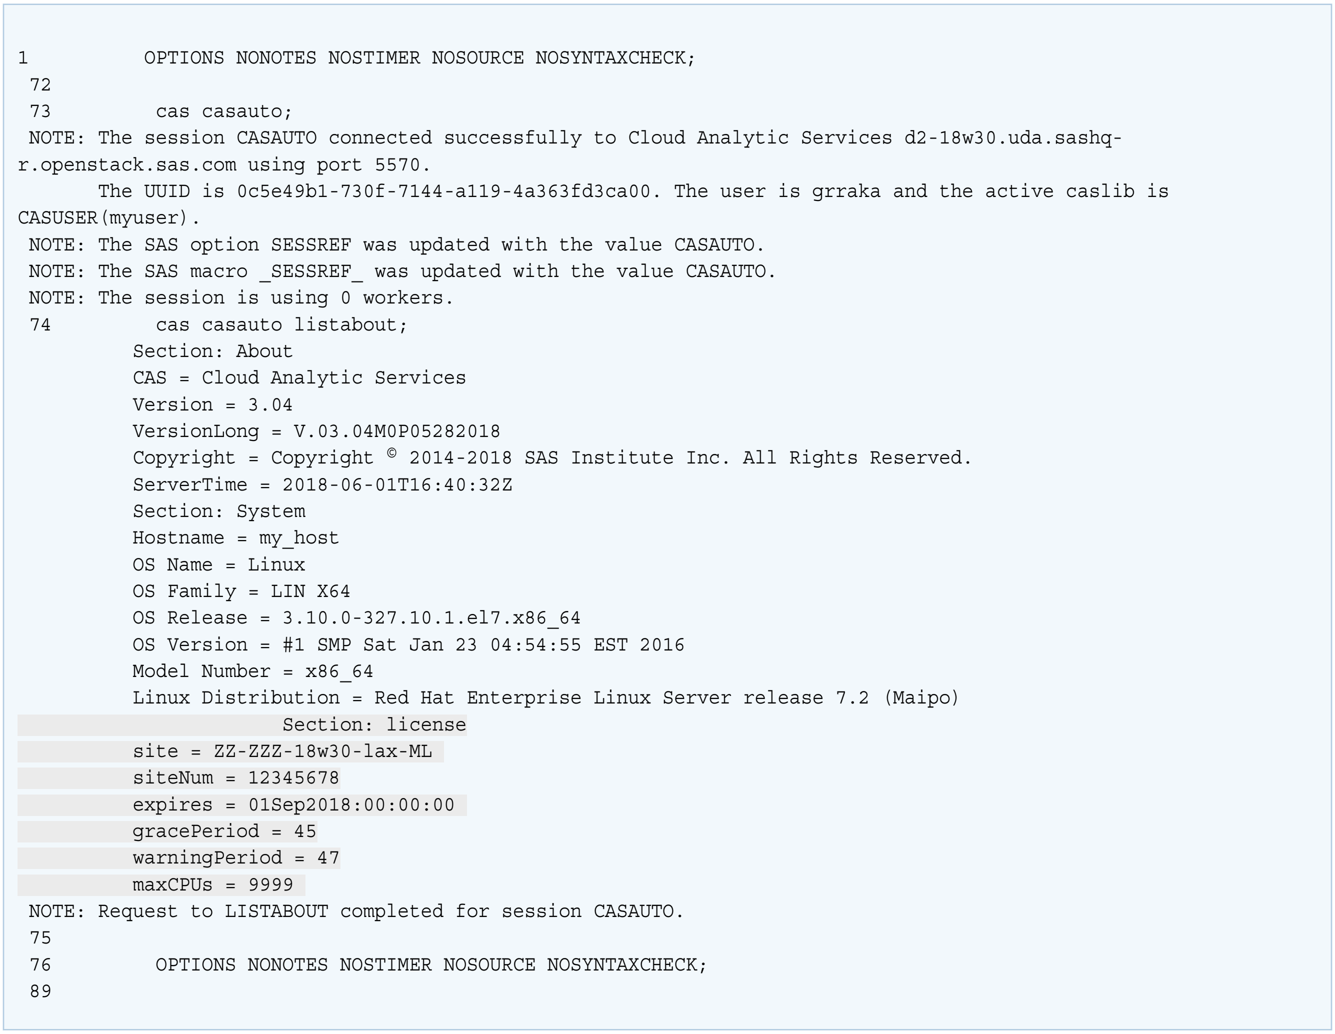

```
1          OPTIONS NONOTES NOSTIMER NOSOURCE NOSYNTAXCHECK;
 72
 73          cas casauto;
 NOTE: The session CASAUTO connected successfully to Cloud Analytic Services d2-18w30.uda.sashq-
r.openstack.sas.com using port 5570.
       The UUID is 0c5e49b1-730f-7144-a119-4a363fd3ca00. The user is grraka and the active caslib is
CASUSER(myuser).
 NOTE: The SAS option SESSREF was updated with the value CASAUTO.
 NOTE: The SAS macro _SESSREF_ was updated with the value CASAUTO.
 NOTE: The session is using 0 workers.
 74          cas casauto listabout;
          Section: About
          CAS = Cloud Analytic Services
          Version = 3.04
          VersionLong = V.03.04M0P05282018
          Copyright = Copyright © 2014-2018 SAS Institute Inc. All Rights Reserved.
          ServerTime = 2018-06-01T16:40:32Z
          Section: System
          Hostname = my_host
          OS Name = Linux
          OS Family = LIN X64
          OS Release = 3.10.0-327.10.1.el7.x86_64
          OS Version = #1 SMP Sat Jan 23 04:54:55 EST 2016
          Model Number = x86_64
          Linux Distribution = Red Hat Enterprise Linux Server release 7.2 (Maipo)
                       Section: license
          site = ZZ-ZZZ-18w30-lax-ML
          siteNum = 12345678
          expires = 01Sep2018:00:00:00
          gracePeriod = 45
          warningPeriod = 47
          maxCPUs = 9999
 NOTE: Request to LISTABOUT completed for session CASAUTO.
 75
 76          OPTIONS NONOTES NOSTIMER NOSOURCE NOSYNTAXCHECK;
 89
```

Licensing: Troubleshooting

Licensed Products page cannot be viewed.

Explanation:

Users without SAS administrator privileges and intra-tenant administrators do not have access to the Licensed Products page.

Resolution:

Contact your SAS administrator.

Licensing: Interfaces

There are several interfaces that you can use to manage and to view SAS license information. The following table lists these interfaces, and the shading indicates the relative amount of SAS license administration that each covers:

Table A.1 Interfaces to SAS Viya Licensing

◒	Ansible	A software orchestration tool that provides the only interface for renewing a license.
◒	Command-line interface	(Read-Only) A command-line interface that enables you to query SAS license information.
◒	SAS Environment Manager	(Read-Only) A graphical enterprise web application used to view SAS license information.
◔	CAS Server Monitor	(Read-Only) A graphical web application that is embedded in the CAS server. Used to view CAS license information.

14

Logging

Logging: Overview

In SAS Viya, logs are produced not just by applications and servers, but also by each SAS Viya service. In order to manage the large number of logs and to enable you to locate messages of interest, the operations infrastructure provides components to collect and store log messages.

The sas-watch log command continuously collects and sends log messages to the RabbitMQ exchange. The sas-stream command then pulls the messages from RabbitMQ and writes them to disk as a tab-separated value (TSV) file. Every five minutes, the etl_driver.sas job extracts the log messages from the TSV file and loads them into the VIYALOGS CAS search index. SAS Environment Manager uses the information in the VIYALOGS table and the VIYALOGS_SOURCES table to display log messages and graphs that contain the frequency and trends of messages. By default, log messages that are more than three days old are removed from the VIYALOGS table. Messages are removed from the table once a day.

The SAS Environment Manager Dashboard displays a graph of the number of ERROR-level and FATAL-level messages from the current top five sources of log messages. The graph displays messages from the past 30 minutes. The graph displays a separate line for each SAS Viya application so that you can see at a glance the applications that might be having trouble.

The Logs page in SAS Environment Manager displays detailed information about the logged messages. In addition to a chart of log messages over the past 30 minutes, the page displays the content of each message. You can display a graph of messages that are grouped by level or source, or that display a time series graph of all ERROR-level and FATAL-level messages.

Note: In a SAS Viya Docker image, log messages are sent to stdout and stderr, rather than being processed as described in this document.

Logging: How To

View Log Activity and Messages

In SAS Environment Manager, select 🗒 **Logs** from the left navigation menu to display the Logs page.

The **Messages** table displays log messages from SAS Viya components, subject to the specified filters and time constraints. By default, messages from the past 30 minutes are displayed. The table displays the first three lines of each log message. To view the full message, select the message and click ☰.

The chart on the Logs page displays a graph of the log messages over the selected time range. The default time range is 30 minutes. You can choose from these graph types:

By Level
Number of messages grouped by logging level. Place your pointer over a bar to view the logging level and count. The information that is displayed in this chart is changed only by filtering based on a time range or message text. The information that is displayed in this chart changes based on any filters that you select.

By Source
Number of messages grouped by source (the component or service that generated the message). Place your pointer over a bar to view the logging level, source and message count. The information that is displayed in this chart changes based on any filters that you select.

Time Series
Number of ERROR-level or FATAL-level messages, if any, for the current time period (the default is 30 minutes). Counts are displayed for the top five sources of these messages. If there were no ERROR-level or FATAL-level messages during the selected time period, the chart is replaced with a message stating that no error messages are present. The information that is displayed in this chart can be changed by filtering based on a time range, message text, or sources.

The left side of the Logs page contains options that enable you to find specific log messages. You can filter the messages using these conditions:

- display messages from a specified time period
- display messages that contain specific messages or text in the message
- display messages that are at specified levels
- display messages from specified sources

The filters or searches that you apply affect the messages that are displayed in the **Messages** table.

Filter Log Messages

1 In SAS Environment Manager, select 🗒 **Logs** from the left navigation menu to display the Logs page.

By default, the graph table display information and the Messages table display information from the past 30 minutes of log activity.

2 To specify a different relative time period to use, click the **Recent log entries** radio button and then select a value (such as **Last hour** or **Last day**) from the menu. Log entries are periodically removed from the CAS table that provides information for this page. The page header displays the date of the oldest available log entry. Selecting a value of **Last week**, **All**, or a custom time range does not display messages that are older than the oldest available log data. By default, log entries are removed from the CAS table every three days. For information about changing the length of time for which log messages are retained, see "Modify Property Values" on page 239. Log files that are older than the retention period are compressed in a ZIP file and stored in the directory `/var/log/sas/viya`.

3 To specify a different time period, click the **Custom time range** radio button and select a start and end date and time. Specifying a time range that is longer than the retention period (the default is three days) displays messages only from the past three days.

4 To display messages that contain specified text anywhere in the message, specify the text in the **Message** field.

5 To display messages only of a certain level, select the levels that you want to display from the choices in the **Level** area. Selecting logging levels in this area changes only the messages that are displayed, not the messages that are logged. To change the level of messages that are logged, you must change the logging level. See "Specify the Threshold Level for Service Logs" on page 194.

The numbers that are next to the entries in the **Level** list indicate the number of messages at each level.

6 To display messages only from a particular SAS Viya component, select the component in the **Source** area.

The numbers that are next to the entries in the **Source** list indicate the number of messages from each source, subject to any filters you have specified.

Note: The names that are used in the **Source** area are not the same as the names that are used when specifying logging levels. See "Specify the Threshold Level for Service Logs" on page 194.

7 If you make any selections in the **Logs Filter** area (including **Time**, **Message**, **Level**, or **Source**), you must click **Apply** to apply the filters and to update the table.

The **Messages** table and the message counts that are beside the entries in the **Level** and **Sources** area change to reflect the filters that you have applied. This list provides details about what causes them to change. The **By Level** and **By Source** charts change based on any filter that you select. The **Time Series** chart changes based on the filters for time, message, and source.

Level area counts
Change based on the selected time range and the **Message** filter.

The counts on the **By Level** graph and the **By Source** graph
Change based on all filters.

Source area counts
Change based on the selected time range, and the **Message** and **Level** filters.

Time Series graph
Displays the top five sources of ERROR-level and FATAL-level messages. The messages are subject to the selected time range, and the **Message** and **Source** filters.

Messages table
Displays the changes in the content of the table and the counts in the table header based on all filters (selected time range, and the **Message**, **Level**, and **Source** filters).

Save the Log Messages Table

1 In SAS Environment Manager, select **Logs** from the left navigation menu to display the Logs page.

2 By default, the **Messages** table displays log messages from the previous 30 minutes from all sources and all threshold levels. Use the **Logs Filter** options to filter the messages that are displayed. See "Filter Log Messages" on page 192 for details.

3 Click 💾 at the top of the **Messages** table.

4 In the Save Table dialog box, select whether you want to save the log messages as a CAS table or as a CSV file.

- If you choose to save the log messages as a CAS table, select a library and specify a table name to which the log message table should be saved. Because of potentially sensitive data that might be contained in the log messages, you can choose to save the log message table to the SystemData library (on the system CAS server) or your personal library. By default, the table is also loaded into memory when it is saved.
- If you choose to save the log messages as a CSV file, specify the file name of the CSV file. Because the CSV file can contain only 1000 messages, use filters to reduce the number of messages so that the file contains only the messages that you need.

5 Click **Save**. If you saved to a CAS table, the table is created and loaded into memory. If you saved to a CSV file, the file is displayed in your browser window. From the browser, you can save the file to a location that you choose.

Specify the Threshold Level for Service Logs

1 In SAS Environment Manager, select **Configuration** from the left navigation menu.

2 In the **View** menu, select **Definitions**.

3 In the list of definitions, select **logging.level**.

4 Select **New Configuration**.

5 In the New Configuration - logging.level window, specify the following parameters:

Services

Select the services to which the logging level applies. Some services (such as SAS Environment Manager and SAS Message Broker) correspond to SAS Viya web applications. CAS servers are listed by server name. If you use a value of **Global**, the logging level applies to all services.

Level

Specify the lowest level of messages that you want to be included in the service's logs. Possible values are OFF, INFO, FATAL, WARN, ERROR, DEBUG, TRACE, and ALL. See "Logging Thresholds" on page 200 for detailed information.

Note: Because a logging level of TRACE generates an extremely large number of messages and can fill up available disk space, an additional configuration step is required in order to use this level of logging. Use TRACE-level logging only with the assistance of SAS Technical Support.

Name

Specify the logger name. Some microservices are associated with more than one logger. See "Microservice and Web Application Loggers" on page 201 for a list of valid loggers.

6 Click **Save**.

Note: You cannot delete the configuration after you create it. You can only edit the configuration.

Manage CAS Server Logging

The logging for a CAS server is specified in the server's logging configuration file. This file specifies the loggers and logging thresholds that are used by the server. However, there might be instances, such as when working with SAS Technical Support, when you need to temporarily change the threshold level of a logger, add a logger, or delete a logger. You can use the **Servers** area in SAS Environment Manager to make these changes on a temporary basis. Any change that you make remains in effect until you restart the server. When you restart the server, the loggers and threshold levels are reset to the values that are specified in the logging configuration file.

To manage CAS server logging:

1 In SAS Environment Manager, select **Servers** from the left navigation menu.

2 In the **Servers** table, select a server whose logging levels you want to view or edit.

3 If you want to modify the logging levels for a server, add a logger, or delete a logger, select **Assume the Superuser role**.

   Note: You can ignore this step if you want to only view the server logging levels.

4 Select **Settings**.

5 In the Server Settings dialog box, select the **Logging** tab.

   The **Logging** tab displays a list of the loggers that are specified for the selected server and the current threshold level for each logger. For more information about server loggers, see "CAS Server Loggers" on page 199 .

6 To change the threshold level for one or more loggers, select **Edit**. This icon is visible only if you have assumed the Superuser role for the server.

7 From the drop-down menu, select the threshold level that you want to use for each logger that you want to change. See "Logging Thresholds" on page 200 for more information.

8 To add a logger, select ✚ **Add**.

9 A blank entry appears in the table. Specify the name of the logger and the threshold level for the logger.

10 To delete a logger, select the logger and select **Delete**.

11 Select **Save** when you have finished making changes to the CAS server loggers.

12 Select Close to close the Server Settings window and return to the Servers table.

13 To relinquish the Superuser role, select **Relinquish**. From the pop-up menu, select the server for which you have assumed the Superuser role.

Note: In addition to its normal server logging, CAS can also log the processing of configuration and start-up files. For more information, see "CAS Configuration and Start-up Logging" on page 547.

Locate Archived Log Messages

Once a day, logs that are older than the specified retention period are archived in a ZIP file. The default retention period of the ZIP file is 30 days. A separate ZIP file is created each day for the entries that are removed on that day. The ZIP files are stored in the directory **`/var/log/sas/viya`** (UNIX) or **`C:\ProgramData\SAS\Viya\var\log`** (Windows). The files are named using the convention `log-YYYYMMDD090000Z.zip`.

These logs are not included in the archive process:

- all-services
- cas
- cas-consul-registration
- httpproxy
- rabbitmq-server
- vault
- watch-log

In UNIX, to search in all the ZIP files for the log for a specific service, issue this command from the `/var/log/sas/viya` directory.:

```
for f in `ls *.zip`; do echo "$f: "; unzip -l $f | grep service_name; done
```

Do not uncompress archived ZIP log files to their original location (**`/var/log/sas/viya/`*`service_name`*`/default`** for UNIX, **`C:\ProgramData\SAS\Viya\var\log\`*`service_name`*`\default`** for Windows). Uncompressed files in these directories will be compressed again during the next compression operation, which results in an old log file that has a new date. The date mismatch makes it difficult to locate the log for the original date. Uncompress ZIP log files to a different directory, and delete the uncompressed files when you are finished with them. The original ZIP files are retained. Deleting the uncompressed files saves disk space.

Note: This process does not allow the ZIP files to be saved anywhere other than the default location. You can develop your own process for moving the ZIP files to another directory in order to manage disk space and for managing the log files that are not included in the ZIP files.

Enable Application Response Measurement (ARM) Logging

Application Response Measurement (ARM) enables you to monitor the availability and performance of transactions within and across SAS Viya applications. You can enable logging of ARM messages from SAS Viya servers. After enabling ARM logging, the log messages are displayed in the SAS Environment Manager Logging area in the same way as other log messages.

For more information, see ARM - Application Response Measurement under the **Scalability and Performance** focus area of support.sas.com.

To enable ARM logging, modify the server configuration file to point to the logconfig.arm.xml file. Here are the modifications to make for each server type:

SAS/CONNECT server

Modify the file `/opt/sas/deployment_name/config/etc/connectserver/default/connectserver_usermods.sh`.

Add the line -`logconfigloc=/opt/sas/deployment-name/config/etc/connectserver/default/logconfig.arm.xml`.

SAS Workspace Server

Modify the file `/opt/sas/deployment_name/config/etc/workspaceserver/default/workspaceserver_usermods.sh`.

Add the line -`logconfigloc=/opt/sas/deployment-name/config/etc/workspaceserver/default/logconfig.arm.xml`.

Logging: Troubleshooting

Why Can I Not See Logs in SAS Environment Manager?

If you are a tenant administrator, you do not have the permissions to look at logs. Your provider-level administrator can access the log messages.

Why Is the Logged Issues Chart on the Dashboard Blank?

The **Logged Issues** chart on the Dashboard and the **Time Series** graph on the Logs page display only ERROR-level and FATAL-level messages from the top five sources of these messages over a selected time range (30 minutes by default). If no ERROR-level or FATAL-level messages are received during this time period, the chart is blank and the message **No information is available** appears.

Why Do Log Messages Not Appear in SAS Environment Manager?

The charts and tables on the Logs page and the Dashboard in SAS Environment Manager use information from the VIYALOGS and VIYALOGS_SOURCES tables. If the Logs page indicates that no log messages are present (for example, the **Messages** table is blank), perform these checks:

- Verify that the selected time range is valid. By default, log messages that are more than three days old are removed from the VIYALOGS table once a day. If you specify a time range that is greater than three days ago, no messages will match the time filter.
- Verify that the CAS tables that are used for logging exist and that they are being updated. Use the Data page in SAS Environment Manager to verify that the VIYALOGS and VIYALOGS_SOURCES tables exist and that they contain data. The tables are in the SystemData library. If the tables exist and contain data, verify that the number of rows in the VIYALOGS table changes over time as new messages are added. If the tables do not exist or are not being updated, verify that the CAS server is running.

  Note: The VIYALOGS and VIYALOGS_SOURCES tables are readable only by the SAS Environment Manager logging function. They cannot be read directly.
- If there is a problem with one of the logging CAS tables, use the **Availability** tile on the SAS Environment Manager Dashboard to check the status of the `stream-evdm` and `watch-log` services. Restart the services if needed.
- Verify that your SAS license is valid.

Why Are Some Messages in the Message Table Blank?

The process that is used to extract the log message information from the TSV file and to prepare the information for display in SAS Environment Manager parses the message and attempts to identify the most important part of the message. In only a few cases, none of the text of the original message remains after this parsing process has completed, so the **Message** column in the **Messages** table is blank. A blank message does not indicate a problem with SAS Viya logging.

Logging: Reference

Overview and Terminology

The SAS logging facility as used by SAS Viya is a flexible, configurable framework that you can use to collect, categorize, and filter events and to write them to a variety of output devices. The logging facility supports problem diagnosis and resolution, performance and capacity management, and auditing and regulatory compliance. The logging facility has the following features:

- Log events are categorized using a hierarchical naming system that enables you to configure logging at a broad level or a fine-grained level.
- Log events can be directed to multiple output destinations. For each output destination, you can specify the following logging facility components:
  - the categories and levels of log events to report
  - the message layout, including the types of data to be included, the order of the data, and the format of the data
  - filters based on criteria such as diagnostic levels and message content
- Logging levels can be adjusted dynamically without starting and stopping processes.

Here are the common terms that this document uses:

appender
: a named entity that represents a specific output destination for messages. Destinations include fixed files, rolling files, operating system facilities, client applications, database tables, message queues, and custom Java classes. You can configure appenders by specifying thresholds, filters, log directories and filenames, pattern layouts, and other parameters that control how messages are written to the destination.

filter
: a set of character strings or thresholds, or a combination of strings and thresholds that you specify. Log events are compared to the filter to determine whether they should be processed.

level
: the diagnostic level that is associated with a log event. The levels, from lowest to highest, are TRACE, DEBUG, INFO, WARN, ERROR, and FATAL.

log event
: an occurrence that is reported by a program for possible inclusion in a log.

logger
: a named entity that identifies a message category. Loggers are named using a hierarchical system that enables you to configure logging at a broad level or a fine-grained level.

  Loggers inherit settings from their higher-level (ancestor) loggers.

logging configuration
: an XML file that determines how log events are processed. You use the logging configuration to assign thresholds to loggers, to configure appenders, and to specify which categories and levels of log events are to be written to each appender.

message category
: a classification for messages that are produced by a SAS subsystem.

pattern layout
: a template that you create to format messages. The pattern layout identifies the types of data, the order of the data, and the format of the data that is generated in a log event. It is delivered as output.

threshold
: the lowest event level that is processed. Log events whose levels are below the threshold are ignored.

Understanding Loggers

A logger is a named entity that identifies a message category. A logger's attributes consist of a level and one or more appenders that process the log events for the message category. The level indicates the threshold, or lowest event level, that is processed for this message category.

Loggers are specified in log events to associate the log event with a message category. By categorizing log events, the logger can write messages of the same category to the same destinations. When a log event occurs, the log event message is processed by the appender that is associated with the logger that is named in the event log. However, the log event level must be the same or higher than the level that is specified for the logger.

Loggers are organized hierarchically and inherit the attributes of their ancestor logger. Hierarchical logger names are separated by a period (.) (for example, Admin.Meta.Security). The root logger is the highest level logger. All loggers inherit the root logger's attributes. The logging configuration file defines several message categories that are immediate descendants of the root logger. These high-level categories — Admin, App, Audit, IOM, and Perf — are used for SAS server logging and can be referenced by log events in SAS programs.

Understanding Appenders

An appender is a named entity that is referenced by a logger. An appender specifies the destination for the message and how the message is formatted. It also specifies the attributes for the appender class and provides additional filtering capabilities.

When a log event occurs, the logging facility processes the message by using the appender that is named. The appender is named in the logger's <appender-ref> element in a logging facility configuration file.

CAS Server Loggers

A logger identifies how different categories of messages are processed, including the level of messages that are processed and the appender to which messages are sent. For example, an Admin logger specifies how administration-related log messages are processed.

In the logging configuration file, a logger has this structure.

```
<logger name="logger-name">
      <level value=threshold/>
      <appender-ref ref="appender-name"?>
</logger>
```

The `level value=threshold` parameter for a logger specifies the lowest level of messages that are processed by the logger. For example, a level of WARN specifies that only messages with a level of WARN, ERROR, and FATAL are included in the log.

The `appender-ref ref=appender-name` parameter for a logger specifies the appender, or the destination, for log messages.

These loggers are defined with a default threshold of INFO in the default logging configuration file for the CAS server, located at **/opt/sas/*deployment_name*/config/etc/cas/default/logconfig.xml** (UNIX) or **C:\ProgramData\SAS\Viya\etc\cas\default\logconfig.xml** (Windows).

Admin
: processes administration events. The log messages are sent to the operating system log.

App
: processes events from applications.

App.cas.actions
processes events from CAS actions.

Audit
processes events used for auditing. These events include user authentication requests and administration of access controls.

Logging
processes events from the logging system. Log messages are sent to the operating system log.

These loggers are not defined by default.

App.cas
processes events from the CAS server.

App.cas.actions.*actionsetname.actionname*
processes events from a specified CAS action set and CAS action.

App.cas.driver
processes events from start-up of the CAS server.

App.cas.tkcasa
processes events from internal processes

App.cas.datastep
processes the output and events from DATA step PUT statements, as well as messages that are sent to the SAS log.

CAS Server Appenders

The logging configuration file for the CAS server at **`/opt/sas/`*`deployment_name`*`/config/etc/cas/default/logconfig.xml`** (UNIX) or **`C:\ProgramData\SAS\Viya\etc\cas\default\logconfig.xml`** (Windows) defines these appenders:

RollingFileAppender
writes log messages to a time-based rolling log file. By default, the file rolls over at midnight. By default, messages from the App, App.cas.actions, and Audit loggers are sent to this appender.

UNXFacilityAppender
writes log messages to the syslogd logging facility in the UNIX operating system. It discards messages that have already been logged by the appender. By default, messages from the Admin and Logging loggers are sent to this appender.

WinEventLog
writes log messages to the Windows event viewer in the Windows operating system. By default, messages from the Admin and Logging loggers are sent to this appender.

ARM Appender

If you enable ARM logging, logging is defined in the file logconfig.arm.xml. This file defines this appender:

ARMAppender
defines the format of the ARM log messages and writes the messages to a specified location.

Logging Thresholds

The logging configuration file sets a threshold for each logger. Messages at the threshold or higher level are processed by the logger, and messages at the level that are lower than the threshold are ignored. For example, if you set the threshold level for a logger to Info, the logger processes messages with the levels of Info, Warn, Error, or Fatal (because those levels are equal to or higher than Info), and ignores messages with the levels of Debug or Trace (because those levels are lower than Info). For information about changing the logging thresholds, see "Manage CAS Server Logging" on page 195.

Here are the available threshold levels (ordered lowest to highest):

Trace
: produces the most detailed log messages. This level might be useful when isolating the cause of a problem, but it produces too many messages for normal use.

Debug
: produces detailed log messages, although less detailed than the Trace threshold. This level might be useful when isolating the cause of a problem, but it produces too many messages for normal use.

Info
: produces messages that show an application's progress.

Warn
: produces messages that identify areas of potential problems.

Error
: produces messages when errors occur, although the application might continue to run.

Fatal
: produces messages when severe errors occur. The application will probably end.

Microservice and Web Application Loggers

These loggers are associated with SAS Viya microservices and web applications. Specify these loggers in the **Name** field when you are creating a logging level definition. See "Specify the Threshold Level for Service Logs" on page 194 for more information.

Service	Loggers	Usage
All	com.sas.authorization	Authorization decisions
All	com.sas.authorization.bootstrap	Authorization rule bootstrapping
All	con.sas.configuration.bootstrap	Configuration bootstrapping
All	com.sas.credentials.bootstrap	Credential domain bootstrapping
All	com.sas.event	Generated and received events
All	com.sas.folders.bootstrap	Folder definition bootstrapping
All	com.sas.security.oauth2	Authentication issues
All	com.sas.security.oauth2.bootstrap	Client token bootstrapping (allows services to talk to other services)
All	com.sas.typeregistry.bootstrap	Type definition bootstrapping
All	org.apache.http.header	Request and response headers
All	org.apache.http.wire	Full requests and responses
All	org.springframework.security	Authentication issues

Service	Loggers	Usage
appregistry	com.sas.appregistry	
appregistry	com.sas.homeshared	
audit	com.sas.audit	
authorization	com.sas.authorization	
authorization	org.springframework.security	
backup-agent	com.sas.backup.worker	
cas-access-management	com.sas.casconnection	The TRACE level logs the Lua-equivalent of all actions that are sent to the CAS server by cas-access-management.
cas-formats	com.sas.casconnection	The TRACE level logs the Lua-equivalent of all actions that are sent to the CAS server by cas-formats.
cas-management	com.sas.casmanagement	
cas-management	com.sas.casconnection	The TRACE level logs the Lua-equivalent of all actions that are sent to the CAS server by cas-management.
cas-management	com.sas.casmanagement.server	The DEBUG level shows CAS server usage by SAS Visual Analytics and SAS Visual Analytics Viewer.
cas-management	com.sas.casmanagement.session.service.CasSessionService	Tracks by user ID the creation of user CAS server sessions. Can be used to track an individual user or the log file that is processed to obtain counts of CAS user sessions.
casproxy	com.sas.casproxy	
casproxy	com.sas.casconnection	The TRACE level logs the Lua-equivalent of all actions that are sent to the CAS server by casproxy.
casrowsets	com.sas.casrowsets	
casrowsets	com.sas.casconnection	The TRACE level logs the Lua-equivalent of all actions that are sent to the CAS server by casrowsets.
collections	com.sas.collections	
collections	com.sas.homeshared	
comments	com.sas.comment	
configuration	com.sas.configuration	

Service	Loggers	Usage
credentials	com.sas.credentials	
data-preparation-plans	com.sas.data.preparation	The DEBUG level displays the contents of job requests and results.
deploymentBackup	com.sas.backup	
device-management	com.sas.devicemgmt	
files	com.sas.svcs.file	
folders	com.sas.folders	
geodelocator	com.sas.locator	
home	com.sas.home	
home	com.sas.homeshared	
identities	com.sas.identities	
identities	com.sas.identities.provider.ldap	SAS LDAP access
identities	org.springframework.security.ldap	Spring LDAP access
monitoring	com.sas.svcs.monitoring	
preferences	com.sas.preferences	
reportdata	com.sas.report.common	
reportdata	com.sas.report.bireport	
reportdata	com.sas.bicommon.export.office	
reportdata	com.sas.bidata	CAS data provider, CAS server resources, timing for CAS sessions
reportdata	com.sas.reportdata.utils.DataUtils	Cache management, CAS queries
reportdata	com.sas.reportdata	Reportdata service classes. Because it produces verbose logging, you should narrow the scope.
reportdata	com.sas.reportdata.DataServicesBase	SAS Report XML. It shows total elapsed time for report generation tasks. This logger is good for performance analysis of reports in SAS Visual Analytics Viewer. You can search the log statements by user ID.
reportdata	com.sas.reportcache	Caching of report content

Service	Loggers	Usage
reportdata	com.sas.reportcommon.utils.debug.Report	
reportdata	com.sas.cas	CAS resource utilization
reportdata	com.sas.reportcommon.utils.debug.Time	It provides detailed, step-by-step timing data for report generation: XML parsing, CAS query, result caching. This is the best logger for detailed performance analysis. Use with com.sas.reportdata.DataServicesBase at the DEBUG level to compare CAS query time to total report generation time. You can search the log statements by a specific report GUID and a user ID. This logger also provides an XML report.
saslogon	com.sas.logon	It produces a high volume of messages at the DEBUG level.
saslogon	com.sas.logon.authentication	SAS authentication events
saslogon	org.cloudfoundry.identity	
saslogon	org.springframework.security	
sasthemedesigner	com.sas.themedesigner	
sasvisualanalytics	com.sas.van	
search	com.sas.svcs.search	
searchindex	com.sas.svcs.search.index	
themes	com.sas.themedesigner.service	
themes	com.sas.themedesigner.service.rest	
themes	com.sas.themedesigner.service.publish	
themes	com.sas.themedesigner.service.servlet	
themes	com.sas.themedesigner.service.persistence.model	
transfer	com.sas.transfer	
types	com.sas.typeregistry	

15

Monitoring

Monitoring: Overview

SAS Viya provides monitoring functions through several facilities. Use the monitoring system that matches your needs and your environment:

- The SAS Viya operations infrastructure collects metrics from SAS Viya applications and services. See "Operations Infrastructure: Overview" on page 224 for more information. SAS Environment Manager uses the collected data to display metric information and status in these interfaces:
  - To quickly view the health and status of your SAS Viya environment, see "Use the SAS Environment Manager Dashboard for System Monitoring" on page 214.
  - To view metrics, status, and performance charts for the machines in your environment, see "Monitoring: How to (SAS Environment Manager)" on page 206.
  - To view detailed reports for the status and activity in your system, see "Use SAS Environment Manager Reports for System Monitoring" on page 208.

If you are using the SAS Viya programming-only interface, SAS Environment Manager is not deployed.

- If you are using the SAS Viya programming-only interface, you can use CAS Server Monitor. CAS Server Monitor is a graphical web application that is embedded in the CAS server. It provides system-level monitoring for the machines and processes that run on the CAS server.

  To view detailed information about the load and performance for the machines and processes running on a CAS server, see "Monitoring: How to (CAS Server Monitor)" on page 216. If you are not using the SAS Viya programming-only interface, CAS Server Monitor is not available.

- CAS start-up or session options can enable returning of performance metric information each time a CAS action runs. The data provided by the metrics enables you to monitor the CPU load on the CAS grid and to determine how efficiently the CAS grid is processing the actions. See "CAS Action Metrics" on page 220 for a list of the metrics that are returned.

  The CAS options are available in the SAS Viya programming-only environment.

Monitoring: Concepts

A metric is a measurement that describes the performance of a component or a subsystem of SAS Viya. Because metrics are valuable only when they are regularly collected and evaluated, the operations infrastructure is dedicated to collecting data about the state of SAS Viya resources and services. A set of collector components from the infrastructure then publishes the data as a message to a RabbitMQ exchange, where a publisher sends it to ETL processes and a data mart. SAS Environment Manager uses the collected data from the data mart to display in various interfaces such as reports, tables, and availability indicators. See Operations Infrastructure on page 223 for more information.

In a SAS Viya environment, CAS uses a controller node to distribute work to worker nodes. In this type of distributed environment, it is important to monitor the performance of each of the nodes in the environment, to ensure that nodes are not becoming overloaded and slowing down. You should also monitor session processes on the CAS nodes to ensure that individual processes are not consuming excessive resources.

Monitoring: How to (SAS Environment Manager)

Monitor Machines

Navigation

In SAS Environment Manager, select **Machines** from the left navigation menu to display the Machines page.

The Machines page displays a list of machines across the top of the page. An icon next to the machine name indicates the status of the machine (available, unavailable, or partially available). Select a machine from the list to display information about the machine on the charts and tables in the Machines page.

View the Status of a Machine

1 In SAS Environment Manager, select from the left navigation menu to display the Machines page.

2 On the Machines page, select a machine name from the list at the top of the page. An icon beside the machine name indicates whether the services on the machine are available ✓, partially available ⚠, or completely unavailable ⊗.

3 By default, the chart on the Machines page displays the percentage of total CPU utilization over the last hour. Click **Last hour** to change the display to the last 6, 12, or 24 hours. Place your pointer on a line on the graph to view detailed information about the CPU utilization, divided into User, System, Wait, and Stolen usage. Place your pointer in the chart and use the control wheel on your mouse to zoom in to the chart.

Note: The chart is updated every two minutes. The data that is displayed on the chart is updated every five minutes.

4 Click **Memory** above the chart to display the percentage of memory that is used over the selected time period. Place your pointer on a line on the graph to view detailed information about memory usage.

5 The **Machine Checks** table displays the results of these predefined system checks that are performed on the machine:

Disk utilization of SAS Config filesystem
The check passes if disk usage does not exceed 95%.

Memory percent free
The check passes if memory usage does not exceed 95%.

Serf Health Status
The check passes if the SAS Configuration Server is running.

The table is refreshed every 10 seconds.

6 The **Service Instances** table displays a list of the service instances that are running on the selected machine and the status, address, and port for each service instance. The data is refreshed every 10 seconds.

7 To display the server properties, click in the toolbar on the right side of the page. The **Properties** area displays information such as the host name, operating system, uptime, and total memory.

8 To display the collected metrics for the server, click in the toolbar on the right side of the page. The **System Metrics** area displays detailed information about memory usage and availability.

9 To display the SAS packages that are installed on the machine, click in the toolbar on the right side of the page. The **SAS Packages** area displays the name and version number of the packages that are installed on the machine.

Note: This information is not displayed if you are using Windows.

10 To display the system limits for the machine, click in the toolbar on the right side of the page. The **System Limits** area displays the resource limits for users on the machine.

Note: This information is not displayed if you are using Windows.

View CAS Server Metrics

1 In SAS Environment Manager, select **Server** from the left navigation menu to display the Servers page.

2 Select a CAS server in the table.

3 At the right side of the page, select **System Metrics**.

4 The System Metrics area displays these metrics for the selected server:

- Active sessions
- Date and time that the last session was created or destroyed
- Number of sessions created
- Server uptime

- User CPU time
- System CPU time
- Number of I/O operations performed
- Number of active threads compared to the maximum number of threads
- Amount of memory used
- Maximum amount of memory used
- Virtual machine size

Use SAS Environment Manager Reports for System Monitoring

Working with System Reports

SAS Environment Manager provides a set of predefined reports that provide a view of the most important metrics for monitoring a SAS Viya deployment. The Dashboard displays a thumbnail of each report, which you can use to access the full report in SAS Report Viewer. You must be an administrator in order to view system reports.

To display the report thumbnails, on the SAS Environment Manager Dashboard, select **Show Reports**.

Note: This option is available only if you have specified one or more reports in the **Public Dashboard Items** application setting. See “Personalizing Your Dashboard” on page 668 for more information. By default, all system reports are available to administrators.

The report thumbnails are not live views of the full reports but are snapshots of the report from the last time the thumbnail was generated. You must refresh the thumbnail in order to view the current state of the report. To refresh a report thumbnail, in the title bar for the report, select ⋮ and then select **Refresh**.

To open a report, in the title bar for the thumbnail report, select ⋮ and then select **Open**.

To return to SAS Environment Manager from the full view of a report, click your browser’s back button or select ≡ and select **Manage Environment**.

The CAS tables that are used to create these reports are refreshed every five minutes. In addition, during the deployment process, data might be delayed from appearing in these reports. The delay time varies and depends on how quickly SAS Infrastructure Data Server, RabbitMQ, CAS, and authentication services are operational and able to respond.

Note: In SAS Environment Manager 3.2, some system reports used data from the CAS_SYSTEM table. In SAS Environment Manager 3.3, these reports now use data that is collected by the operations infrastructure in order to provide a consistent view of SAS Viya metric data. The CAS_SYSTEM table will be removed from SAS Viya releases after SAS Viya 3.4. Therefore, if you have created any reports that use data from this table, they must be rewritten.

Monitor Application Activity

The Application Activity report provides detailed information about SAS applications and services running on your system. See “Working with System Reports” on page 208 for information about accessing and opening reports.

When you open the report, the machines in your environment are listed along the top of the report. Select a machine for which to display the report.

Select the report page to view. The report pages are organized into these tabs:

Main

Displays a chart of memory usage of the 10 applications or services that are consuming the most memory. The report displays the amount of used heap memory (HeapUsedMax) and of used non-heap memory (NonHeapUsedMax).

Place your cursor over a bar to view the name of the application and the memory usage values.

Application History

Displays the thumbnails of detailed reports for a selected service or application. Use the menu in the upper left corner of the page to select the service or application whose reports you want to view. If you do not select a service or application, the thumbnails display aggregate data for all services and applications. Use the slider control at the top of the page to select the time range for the reports. Click ⤢ in the upper right of any chart to view a full-size version of the chart, including legends and labels. Click ⤡ in the upper right of the full-size chart to return to the thumbnail view.

Place your cursor over a line in a graph to view detailed values.

Here are the available charts:

Heap usage
Displays the amount of heap memory that is used. The chart displays the metrics HeapCommitted, HeapUsed, NonHeapCommitted, and NonHeapUsed.

HTTP sessions
Displays the number of HTTP sessions that are used. The chart displays the metrics HTTPSessionsActive and HTTPSessionsMax.

Class Usage
Displays the number of classes that are used by the application or service. The chart displays the metrics Classes, ClassesLoaded, and ClassesUnloaded.

DataSource Usage
Displays the number of data sources that are used by the application or service. The chart displays the metrics DatasourcePrimaryActive and DatasourcePrimaryUsage.

Garbage Collection Time
Displays the amount of time that is used for garbage collection. The chart displays the metrics GcPsMarksweepTime and GcPsScavengeTime.

Threads
Displays the number of application threads that is used. The chart displays the metrics Threads, ThreadsDaemon, and ThreadsPeak.

Uptime
Displays the amount of time that the application or service has been running.

Garbage Collection Count
Displays the number of items that are collected during garbage collection. The chart displays the metrics GcPsMarksweepCount and GcPsScavengeCount.

System Session History

Displays a graph of the top 10 applications or services that have had the most active HTTP sessions over the previous eight hours.

Data collection status

Displays a chart of metric data points that are collected for each application.

Monitor CAS Activity

The CAS Activity report provides detailed information about CAS. See “Working with System Reports” on page 208 for information about accessing and opening reports.

When you open the report, the machines in your environment are listed along the top of the report. Select a machine for which to display the report.

Select the report page that you want to view. The report pages are organized into these tabs:

Main
Displays the **Memory Used**, **I/O**, and **Threads** charts.

CPU Usage
Displays the **CPU Usage** chart. The chart displays the metrics SystemCPU and UserCPU.

System Info
Displays the thumbnails of detailed reports for the CAS servers, including I/O wait time, IRQ time, open files, and free memory. This page is visible only if you are using SAS Visual Analytics.

System Details
Displays a table of detailed metric information about the CAS servers that is captured at one-minute intervals. The table includes data for load averages, free memory, idle time, and IRQ time. This page is visible only if you are using SAS Visual Analytics.

Node Details
Displays a table of information about the CAS server nodes.

CAS Details
Displays a table of detailed metrics for the CAS servers that are captured at one-minute intervals. The table includes these metrics:

- I/O count
- Maximum memory used (bytes)
- Maximum thread count
- Memory used (bytes)
- System CPU (seconds per second)
- System CPU count (seconds)
- Thread count
- Uptime (seconds per second)
- Uptime count (seconds)
- User CPU (seconds per second)
- User CPU count (seconds)

Monitor Disk Space

The Disk Space report provides detailed information about disk space and usage. See “Working with System Reports” on page 208 for information about accessing and opening reports.

The machines in your environment are listed in the **Machine** menu at top of the report. Select a machine for which to display the report.

Select the report page to view. The report pages are organized into these tabs:

Main
Displays a chart of the top 10 filesystems on SAS Viya machines that have the least amount of free space.

Storage Dashboard
Displays a chart of the total percentage of free disk space on each machine in the system. It also displays a series of charts of the top 10 file storage locations that have the least amount of available space.

In the **Bottom 10 paths by Percent available** charts, the black line represents the available space. The background of the chart is color-coded to indicate whether the available space is in the acceptable zone (cyan), the warning zone (yellow), or the danger zone (red). For example, the disk corresponding to this graph has 18% free space, which is in the danger zone.

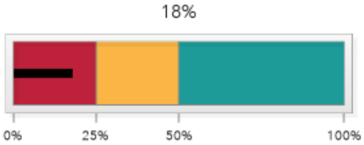

Machine disk usage over time

Displays a chart of the total percentage of free disk space on each machine in the system. It also displays a chart of the percentage of free space on all paths for each machine over the previous 48 hours.

Disk usage forecast

Displays a chart of the percentage of free space for a selected machine and the path over the previous 24 hours. It also includes a projection of the free space that will be available over the next 48 hours. Select a machine from the list above the chart, and then select a path from the list below the machine list.

Storage Map

Displays a visual representation of the size and available free space of all disks in all machines. Each disk is represented by a color-coded block. The size of the block represents the size of the disk. The color of the block represents the amount of free space. The color shifts from blue to red as the disk space decreases. Place your pointer on a block to view the size and percentage of free space for the disk.

Details

Displays a table of the size and free space (in bytes) for a selected machine and path, and that is recorded at one-minute intervals. Select a machine from the list above the table, and then select a path from the list below the machine list.

When you are monitoring CAS disk usage, keep in mind that owned disk space is the space used by files that are created in CAS_DISK_CACHE directories from in-memory blocks. These files cannot be shared with other server processes or session processes. Shared disk space is the space that is used by existing SASHDAT files from a co-located data source (PATH, HDFS, or DNFS). These files can be shared with other server processes or session processes.

Monitor SAS Infrastructure Data Server Tables

The Infrastructure Data Server Tables report provides detailed information about the table size and usage on the SAS Infrastructure Data Server. See "Working with System Reports" on page 208 for information about accessing and opening reports.

Select the report page to view. The report pages are organized into these tabs:

Main

Displays a chart of the five largest tables in the SAS Infrastructure Data Server. The chart displays the metrics TableSize Max, IndexSize Max, and ToastSize Max for each table.

Place your cursor over a bar to view the name of the table and the values for each metric.

Application Usage History

Displays an animated chart of the size of the largest SAS Infrastructure Data Server tables over the previous 36 hours. Click ▶ below the chart to start the animation. The chart displays the size of the tables at the time indicated on the slider control below the chart. You can use the slider control to view the size of the tables at a selected time. The chart separately displays the metrics TableSize, IndexSize, and ToastSize for each table.

Place your cursor over a bar to view the name of the table and the values for each metric.

Table Usage Trend
Displays a graph of the total size of all SAS Infrastructure Data Server tables over the past 36 hours. The chart separately displays the metrics TableSize, IndexSize, and ToastSize for all tables.

Table Size History
Displays an animated chart of the size of the largest SAS Infrastructure Data Server tables over the previous five hours. Click ▶ below the chart to start the animation. The chart displays the size of the tables at the time indicated on the slider control below the chart. You can use the slider control to view the size of the tables at a selected time. The chart separately displays the metrics TableSize, IndexSize, and ToastSize for each table.

Monitor Message Queue Activity

The Message Queue Activity report provides detailed information about traffic and activity on the RabbitMQ message queues that used by the operations infrastructure to provide log messages, metric data, notifications, and alerts to consumers such as SAS Environment Manager. See “Working with System Reports” on page 208 for information about accessing and opening reports.

Select the report page to view. The report pages are organized into these tabs:

Main
Displays a chart of the total amount of data that is published to and from each message queue. The chart displays the PublishInCount and PublishOutCount metrics for each message queue.

Messaging Activity
Displays a graph of the amount of data that is published to a selected message queue over the previous 48 hours. Select the queue name from the list at the top of the chart. The chart displays the PublishIn and PublishOut metrics.

Messaging Animation
Displays an animated chart of the amount of data that is published to message queues over the previous 36 hours. Click ▶ below the chart to start the animation. The chart displays the amount of data that is published to the queue at the time indicated on the slider control below the chart. You can use the slider control to view the amount of data that is published at a selected time. The chart displays the PublishIn and PublishOut metrics.

System info
Displays charts illustrating the number of RunQueue instructions, the amount of data written to queues, and the amount of memory that is used over the previous 48 hours. The charts display the RunQueue, IoWriteBytes, and MemUsed metrics.

Monitor System Activity

The System Activity report provides detailed information about CPU usage, memory usage, and network activity. See “Working with System Reports” on page 208 for information about accessing and opening reports.

The machines in your environment are listed in the **Machine** menu at top of the report. Select a machine for which to display the report.

Select the report page to view. The report pages are organized into these tabs:

Main
Displays a chart of the memory usage for a selected machine over a selected time range. Select the machine from the list at the top of the chart. Select the time range using the slider control at the top of the chart.

CPU history
Displays a chart of the CPU usage for a selected machine over a selected time range. Select the machine from the list at the top of the chart. Select the time range using the slider control at the top of the chart. The chart displays separate lines for the metrics System CPU% and User CPU%.

Memory Usage history
Displays a chart of the free memory and the used memory for a selected machine over a selected time range. The orange area at the top of the chart represents the free memory, and the green area at the bottom of the chart represents the used memory. The two values together always add up to the total memory. Select the machine from the list at the top of the chart. Select the time range using the slider control at the top of the chart. The chart displays the metrics Used Memory and Free Memory.

Network Activity history
Displays charts of the network activity and the cumulative network I/O for a selected machine and an interface over a selected time range. Select the machine and the interface from the lists at the top of the chart. Select the time range using the slider control at the top of the chart. **The Network Activity over time** chart displays the TransmitBytes and ReceiveBytes metrics. The **Cumulative Network I/O** chart displays the TransmitBytes_cnt and ReceiveBytes_cnt metrics.

Memory Animation
Displays an animated chart of the used memory and the free memory for all machines over the previous 36 hours. Click ▶ below the chart to start the animation. The chart displays the memory usage at the time indicated on the slider control below the chart. You can use the slider control to view the memory usage at a selected time. The chart separately displays the metrics Used Memory and Free Memory for each machine. The orange area at the top of the chart represents the free memory, and the green area at the bottom of the chart represents the used memory.

CPU Details Animation
Displays an animated chart of the CPU usage for all machines over the previous 36 hours. Click ▶ below the chart to start the animation. The chart displays the CPU usage at the time indicated on the slider control below the chart. You can use the slider control to view the CPU usage at a selected time. The chart separately displays the metrics UserCPU, IoWaitCPU, SystemCPU, and StolenCPU for each machine.

Network Activity Animation
Displays an animated chart of the network activity for all machines over the previous 36 hours. Click ▶ below the chart to start the animation. The chart displays the network activity at the time indicated on the slider control below the chart. You can use the slider control to view the activity at a selected time. The chart separately displays the metrics TransmitBytes_cnt and ReceiveBytes_cnt for each machine.

System Details
Displays a table of detailed system metrics for selected machines over a selected time period, which is captured at one-minute intervals. The table includes information about memory usage, CPU usage, and system load. Select a machine from the list at the top of the table. Select a time period by using the slider control at the top of the table.

Network Details
Displays a table of detailed network metrics for the selected machines and the interfaces over a selected time period, which is captured at one-minute intervals. The table includes information about received data, transmitted data, and transmit errors. Select a machine and an interface from the lists at the top of the table. Select a time period by using the slider control at the top of the table.

Monitor User Activity

The User Activity report provides a view of audit information. See “Working with System Reports” on page 208 for information about accessing and opening reports.

Select the report page to view. The report pages are organized into these tabs:

Main
Contains thumbnail graphs for the charts **Most active users**, **Activity counts**, **Most active data**, and **User Actions over time**.

Most Active Users
Displays the **Most Active Users** and **Activity Over Time** charts, and a table of the audit records that are ordered by level of user activity. The table does not display audit records from SAS internal users. Select a

bar in the **Most Active Users** chart to display the **Activity Over Time** chart for the selected user, and to list the audit records only for the selected user.

Application Usage
Displays the **Most used Applications** and **Application Activity** charts, and a table of the audit records that are ordered by level of application activity. Select a bar in the **Most used Applications** chart to display the **Application Activity** chart for the selected application, and to list the audit records only for the selected application.

Report Activity
Displays the **Top Report Usage** chart and a table of the audit records for report access. By default, the chart and the table display activity for all users. To view report usage and the audit records only for a specific user, select the user in the **Users** menu.

Data Activity
Displays the **Frequently Accessed Tables** chart and a table of the audit records for data table access. By default, the chart and the table display data table activity for all users. To view data table activity usage and the audit records only for a specific user, select the user in the **Users** menu.

Data Plan Activity
Displays the **Top Report Usage** chart and a table of the audit records for data plan access. By default, the chart and the table display activity for all users. To view data plan usage and the audit records only for a specific user, select the user in the **Users** menu.

Failures
Displays the **Failed Requests per Application** chart and the **Failed Activities** chart, and a table of the audit records only for failed requests. By default, the **Failed Activities** chart and the audit records table display failures for all applications. To view the **Failed Activities** chart and the audit records for a specific application, select the application's bar in the **Failed Requests per Application** chart.

Details
Displays a table of audit records. By default, the table displays all audit records. To filter the table, use the menus at the top of the table to display only those records that match your selected criteria. You can filter by user, application, action, and state. You can also filter using multiple criteria.

Note: Note: If the User Activity report is blank or displays the message `Cannot find the requested data source`, you must verify that the command-line interface (CLI) was deployed properly in your SAS Viya environment. See "Edit the Inventory File" in *SAS Viya for Linux: Deployment Guide* for more information.

Use the SAS Environment Manager Dashboard for System Monitoring

Monitor Availability of Machines and Services

The **Availability** tile displays grids of color-coded boxes, and each box displays the status of each machine, service, and service instance. A green box indicates that the item is available, a yellow box indicates that it is partially available, and a red box indicates that it is unavailable. The tile is updated every 10 seconds.

Selecting a box on one of the grids highlights the corresponding boxes on the other two grids. The box that you select is outlined with a solid line, and the associated boxes are outlined with a dashed line. Here are the associations between the selected boxes:

- When you click a box on the **Machines** grid, the services and the service instances that are running on that machine are highlighted on the **Services** grid and on the **Service Instance** grid.
- When you click a box on the **Services** grid, the machines on which that service is running are highlighted on the **Machines** grid, and the instances of the service are highlighted on the **Service instances** grid.
- When you click a box on the **Service instances** grid, the machines on which the service instance is running are highlighted on the **Machines** grid, and the service is highlighted on the **Services** grid.

Note: To deselect a box, hold down the Ctrl key and click the box. You can also hold down the Ctrl key and press the spacebar.

Place your cursor over a box to view the name of the machine, the service, or the service instance.

Double-click a box on the **Machine** grid (or right-click a box and select **View services**) to open the Machine Status dialog box, which lists the services that are running on that machine and their availability. Select **Machine Details** to open the Machines page for the selected machine.

Click a box on the **Service instances** grid to view the machine address and the port where the instance is running.

Note: Point to or click an instance of the postgres service to also identify whether the instance is a pgpool instance, whether it is a primary or standby data node, and whether SSL has been enabled for the node.

Use the **Search** field to display only certain machines, services, and service instances. When you enter characters in the **Search** field and click 🔍 or press **Enter**, the boxes that are displayed in the **Availability** area will change. The boxes that are displayed either match the filter that you specify or are associated with the boxes that are displayed. For example, entering the string `laun` in the **Search** field might cause the display of two **Services** boxes (for the Launcher service and the Launcher server), only the **Service instance** boxes that are associated with the displayed services, and only the **Machines** boxes that are associated with the displayed services.

Evaluate CAS Nodes

The **System Health** tile displays graphs that give you a quick view of the state of the nodes (machines) in your SAS Viya cluster for a selected CAS server. The data that is displayed on the graphs reflects all the work (not just the CAS operations) that occurs on the nodes. If you are using a UNIX system, the tile displays the **Node Memory Usage** and the **Load Average** graphs. Use the buttons at the top of the tile to select the graph to view. If you are using a Windows system, only the **Node Memory Usage** graph is displayed.

If your environment contains more than one CAS server, a menu above the graph enables you to select the server to view. When you display the dashboard, this functionality behind the tile attempts to connect to the default CAS server. If the default server cannot be found, the tile displays information for the first server to which it can connect. If it can connect to the default server, but the server does not respond within five seconds, the tile displays a message. You can then retry the server or choose another server. You specify the default server in the **default** ⇨ **casServer** property. This property is one of the `sas.casmanagement.global` properties for the CAS Management service. See “Introduction” on page 74 for information about setting this property.

Here are the graphs that are displayed in the **System Health** tile:

Load Average
Displays a graph of the 1-minute load average over the past five minutes for each node in your CAS cluster. This graph is displayed only on UNIX systems. The chart is updated every 10 seconds. Each node is represented by a separate line on the graph. The vertical scale of the graph changes, depending on the largest value that is displayed in the chart. Position your cursor over a line in the chart to identify both the node and the load average value.

Node Memory Usage
Displays a bar chart, which displays the percentage of memory usage for each node in your CAS cluster. This graph is displayed on both UNIX and Windows systems. Each bar represents a separate node. Bars for controllers use a different color than bars for workers. The colors that are displayed depend on the theme that you use. Point to a bar on the graph to view the name of the node, its type, and its memory usage. The chart is updated every 10 seconds.

Note: If your environment contains both a primary and a secondary controller node, this graph displays information only for the controller that is currently active.

Monitoring: How to (CAS Server Monitor)

Access CAS Server Monitor

CAS Server Monitor is available only if you are using a SAS Viya programming-only environment.

Enable CAS Server Monitor by setting the CAS_START_MONITOR_UI environment variable to 1.

To log on to CAS Server Monitor, open a web browser and enter the following URL in the address field:

```
https://http-proxy-machine-name/cas-tenant-name-deployment-instance-name-http
```

You must have an active CAS Server session in order to access CAS Server Monitor.

For more information, see “Using CAS Server Monitor” on page 509.

Monitor CAS Process Performance

The CAS processes you can monitor with these steps correspond to SAS server processes. You can separately monitor each session that is started from the CAS server.

1 In CAS Server Monitor, beneath the Cloud Analytic Services banner, click 📈.

2 Select **Add View** ⇨ **CAS Process CPU Usage**.

The **Process CPU Usage** panel on the window displays a set of histograms. There is one histogram for each machine and the corresponding CAS server process. The histogram in the upper left is the CAS controller node. If you are not an administrator, only the histogram for the CAS controller node is displayed.

Each histogram displays the percentage of CPU usage, from 0 to 100%.

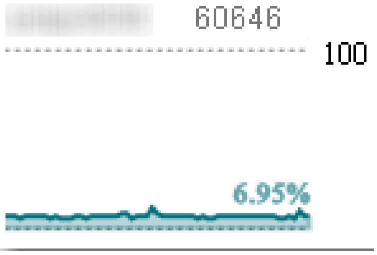

Use these histograms to note patterns of CPU usage among the CAS nodes.

3 Select **Add View** ⇨ **CAS Process Metrics**

The **CAS Process Metrics** panel on the window displays a set of histograms. There is one set of three histograms for each machine and the corresponding CAS server process. If you are not an administrator, only the set of histograms for the CAS controller node is displayed.

Each set of histograms displays the percentage of CPU used, amount of resident memory used, and amount of virtual memory used for the CAS process.

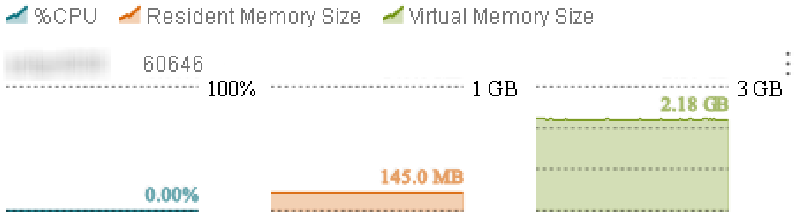

4 Click ■ if you want to stop metric collection. Click ▶ to resume collection.

Monitor CPU Usage for a Session

1 In CAS Server Monitor, select 📈 on the left side of the window.

2 Select **Add Session View** and select a session.

The panel for the session displays a set of histograms, with one histogram for each machine in the grid. If you are not an administrator, only the histogram for the CAS controller node is displayed. The top half of the histogram displays the percentage of CPU load used by the session, and the bottom displays the amount of resident memory used for the session.

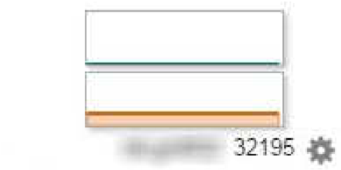

Monitor Host Performance

CAS Server Monitor displays histograms that enable you to view the CPU load and memory usage for all machines in your CAS server. Follow these steps:

1 In CAS Server Monitor, select 📈 on the left side of the window.

2 To view the CPU load, select **Add View** ⇨ **Host CPU Load Average**.

The **Host CPU Load Average** panel on the window displays a set of histograms. There is one histogram for each machine in the CAS grid. If you are not an administrator, only the histogram for the CAS controller node is displayed.

Each histogram displays the CPU load on the machine, using the same format as the Linux `xload` command. Each division on the histograms represents one load average point. The highest point on each histogram is displayed to the right of the histogram.

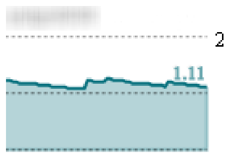

Use these histograms to note usage patterns among the CAS nodes. For example, if you notice that the load on a worker node machine is significantly and consistently higher than the load on other machines, you can use the **Show Processes** function to check for other running processes or defunct processes. See “Monitor Process Information” on page 218 for instructions on this function.

3 To view the memory usage, select **Add View** ⇨ **Host Memory Usage**.

The **Host Memory Usage** panel on the window displays a set of histograms. There is one histogram for each machine in the CAS grid. If you are not an administrator, only the histogram for the CAS controller node is displayed.

Each histogram displays the percentage of memory used on the machine, from 0 to 100%. The percentage of memory used is displayed in green, at the top of the histogram. The percentage of virtual memory used is displayed in orange, at the bottom of the histogram.

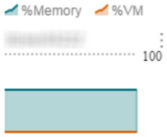

Use these histograms to note patterns of memory usage among the CAS nodes. For example, if the memory usage is consistently high on a machine, its memory might need to be increased.

4 Click ■ if you want to stop metric collection. Click ▶ to resume collection.

Monitor Process Information

1 Perform one of these actions in CAS Server Monitor:

- Select 📈 on the left side of the window and open one of the views from the **Add View** or **Add Session View** menus. Click ⋮ to the right of a histogram. Select **Show Processes**.
- Click 🖧 and select the **Nodes** tab. Click ⋮ on the right side of a node’s row and select **Show Processes**.

2 The Processes window appears. The window displays this information:

- Metrics for the selected node, including uptime, number of processes, memory usage, CPU load, and file usage
- A histogram of the CPU load for the node
- A table containing the output from the `top` command for the selected node. The output includes metrics such as CPU usage, time, and threads for each process. If you are a SAS administrator, the window displays information about all processes. If you are not a SAS administrator, you can view information only about your own processes.

If you are the process owner, you can open a terminal window to terminate processes that are causing problems. See “Open a Terminal Window on a Node” on page 218 for information.

Open a Terminal Window on a Node

After using the monitoring functions of CAS Server Monitor to identify problems with CAS nodes, you might want to issue commands to end processes on a node. If you are the process owner, you can launch a terminal window to manage processes on the node. Follow these steps.

1 Perform one of these actions in CAS Server Monitor:

- Select 📈 on the left side of the window and use the **Add View** menu to display the **Host CPU Load Average**, **Host Memory Usage**, **CAS Process CPU Usage**, or CAS **Process Metrics** views.

  Click ⋮ on the right side of the histogram for a node. Select **Launch Terminal**. This option is available only if you are an administrator.
- Click 🖥 and select the **Nodes** tab. Click ⋮ on the right side of a node's row and select **Launch Terminal**.

2 A terminal window appears on the selected machine. Use the window to manage processes on the machine.

3 Type `exit` to close the terminal window.

Change the Monitoring Display Options

When you are viewing the histograms in the **Grid Monitor** view in CAS Server Monitor, you can control how the histograms are displayed.

- To change how quickly the graph data is refreshed, move the slider next to the **Speed** label.
- To change the size of the histograms, move the slider next to the **Size** label.
- The default layout for a histogram view is a grid. To change to a single column, click the **column icon** ⁝ in the banner for a view. To return to a grid layout, click the **grid icon** ⊞.

To change the default view for the **Grid Monitor** view, select ***userid*** ⇨ **Settings** in the upper right of the CAS Server Monitor window. You can select a default monitor view and layout.

Monitoring: How to (CAS Options)

View Performance Metrics for a CAS Action

To view metric performance data when you execute a CAS action, start the CAS server with the -`metrics` option, or set the `cas.metrics` configuration option to **`true`**.

To start displaying performance metrics for a running server, set the `metrics` session option to **`true`**.

If you enable metric collection, a standard set of metric data is returned to the log each time that a CAS action completes. The same data is displayed by both the server and the client, although the names of the metrics are different. See "CAS Action Metrics" on page 220 for a list of the metrics that are displayed.

Here is an example of the metrics that are displayed for a CAS action:

```
NOTE: Executing action 'tkimstat.summary'
NOTE: Action 'tkimstat.summary' used (Total process time):
NOTE:       real time              0.024989 seconds
NOTE:       cpu time               0.165974 seconds (664.19%)
NOTE:       total nodes            4 (96 cores)
NOTE:       total memory           377.85G
NOTE:       memory                 22.53M (0.01%)
{
  disposition = { ...  },
  messages = { ... },
  results = { ... },
  performance = {
                  elapsedTime = 0.024989,
                  cpuUserTime = 0.132979,
                  systemCores = 96,
                  systemTotalMemory = 405711519744,
                  cpuSystemTime = 0.032995,
                  memoryOS = 45793280,
                  memory = 23621664,
                  memoryQuota = 47366144,
                  systemNodes = 4,
```

Evaluate CPU Utilization for an Action

If you specify that performance metrics are collected when CAS actions are executed, you can use these metrics to evaluate the utilization of your CAS environment.

The server metric CPU time is displayed in both the number of seconds and a percentage. Here is an example:

```
cpu time    0.165974 seconds (664.19%)
```

The percentage is calculated as `(cpuUserTime + cpuSystemTime)/elapsedTime`. On a single-threaded system, the maximum value for this metric is 100%. However, for a multi-core system, the maximum value is `100% * number of cores`. In this example, the system has 96 cores, so the maximum value is 9600%.

Monitoring:Troubleshooting

Why Can I Not See Machine Information in SAS Environment Manager?

If you are a tenant administrator, you do not have the permissions to look at machine health information or metric data. Your provider-level administrator can access this information.

If you are a provider-level administrator, verify that the CAS table that is used for machines information exists and that it is being updated. Use the Data page in SAS Environment Manager to verify that the SYSTEM table exists and contains data. The table is in the SystemData library. If the table exists and contains data, verify that the number of rows in the SYSTEM table changes over time as new messages are added. If the table does not exist or is not being updated, verify that the CAS server is running.

Monitoring: Reference

CAS Action Metrics

If you enable metric collection for CAS actions, a standard set of metric data is returned each time that a CAS action completes. The same data is displayed by both the server and the client. Here is the data that is displayed:

Server Metric Name	Client Metric Name	Description
real time	elapsedTime	The number of seconds in actual time required to run the action.

Server Metric Name	Client Metric Name	Description
	cpuUserTime	The total number of seconds taken by the action in user mode across all cores that were used to run the action.
	cpuSystemTime	The total number of seconds taken by the action in system mode across all cores that were used to run the action.
cpu time		CPU time is measured and displayed in these formats: ■ `cpuUserTime + cpuSystemTime`, displayed in seconds. ■ `(cpuUserTime + cpuSystemTime)/ elapsedTime`, displayed as a percentage.
total nodes	systemNodes	The number of nodes in the cluster (total nodes display both systemNodes and systemCores).
total nodes	systemCores	The number of cores in the cluster (total nodes display both systemNodes and systemCores).
total memory	systemTotalMemory	The total memory available to the system. Total memory is displayed in GB, and systemTotalMemory is displayed in bytes.
memory	memory	Memory used to execute the action.
	memoryOS	Operating system used by the action.
	contextVoluntary	The number of times a context switch occurred because a process relinquished its processor before its time slice had been completely used.
	contextInvoluntary	The number of times a context switch occurred because a higher priority process was present or because the current process exceeded its time slice.
	memoryQuota	The memory quota used by the action.
	dataMovementTime	The amount of time, in seconds, taken by the data that moved between the memory and the processors.
	dataMovementBytes	The number of bytes of data that moved between the memory and the processors.

See “View Performance Metrics for a CAS Action” on page 219 for information about displaying these metrics.

16

Operations Infrastructure

Operations Infrastructure: Overview

The operations infrastructure implements an event-driven architecture that underlies several areas of SAS Viya administration, most notably monitoring, auditing, and logging. An event in this architecture represents some unit of information, such as a metric reading (for example, CPU usage at a particular time) or a log message (for example, a microservice has failed). The architecture is flexible and adaptable, because it keeps the producer of the events and the consumers of the events separate. The producer of an event collects information such as a system metric or a log message, and publishes that information to a message exchange without knowledge about the consumer of the information. The information consumer looks for specified types of information and retrieves the information when it is found. Types of consumers include extract, transform, and load (ETL) processes or a data mart. Likewise, the consumer has no knowledge about the source of the information.

The operations infrastructure consists of these components:

Data collectors
: The sas-peek, sas-check, and sas-watch commands act as the producer by collecting metric data, log events, alert and notification messages. Then, the sas-ops-agent commands pass that data to the appropriate RabbitMQ exchange. See “Operations Infrastructure: Collectors” on page 225 for more information about collectors. See “Operations Infrastructure: Operations Agent” on page 226 for more information about the agent.

RabbitMQ exchanges
: The data that is collected in the RabbitMQ exchanges is available for the data mart and ETL processes. See “Operations Infrastructure: RabbitMQ Exchanges” on page 226 for more information.

Data mart and ETL processes
: The sas-stream commands act as the consumer by reading events from the RabbitMQ exchanges and writing them to the data mart. ETL processing runs periodically to load data from the data mart to the CAS server. See “How To: ETL and Data Mart Operations” on page 239 for more information.

The result of the operations infrastructure processes is to create or produce metric and log data in CAS tables. Applications, such as SAS Environment Manager or SAS Visual Analytics, can display or produce reports using the data from the CAS tables. For example, the Machines and Logging views and the system reports in SAS Environment Manager use metric data from the operations infrastructure.

The following diagram illustrates how the components of the operations infrastructure work together:

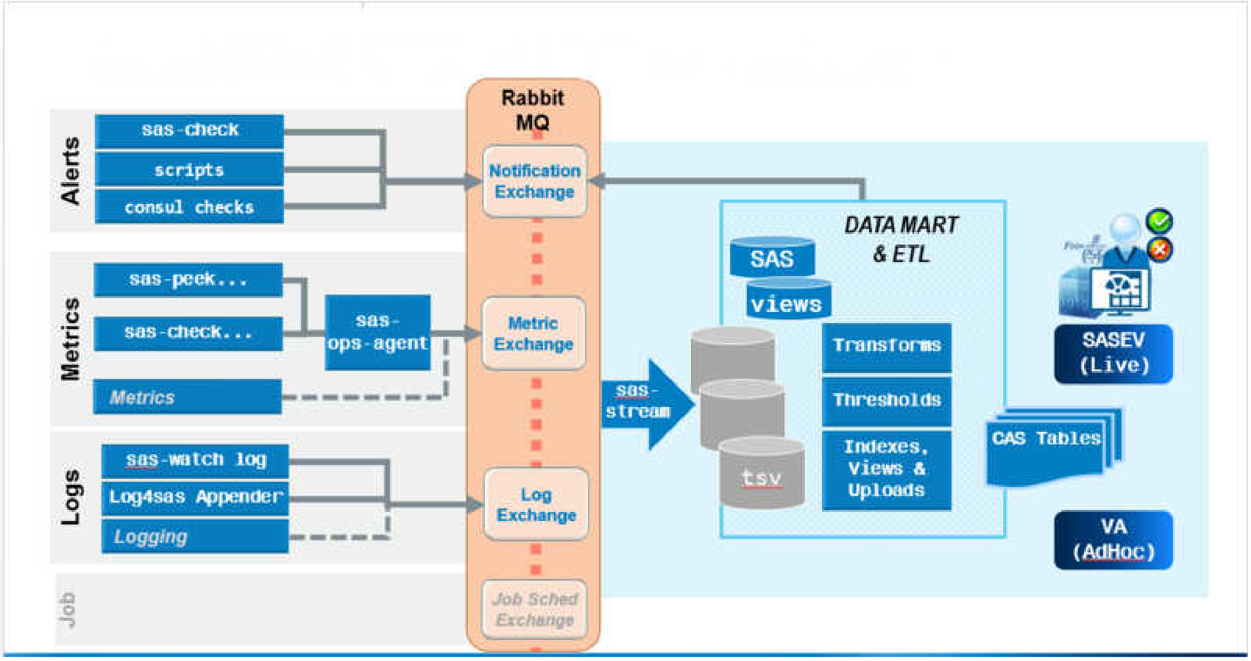

Operations Infrastructure: Collectors

Collectors

The operations infrastructure uses a set of collector commands to obtain metric and log information from a variety of sources.

The sas-peek collector periodically checks the value of a specified system resource or service metric. It uses GoSigar to capture system metrics. It uses REST APIs to capture metrics from CAS, RabbitMQ, and the Spring Boot microservice.

The sas-watch collector continuously monitors a data source such as a log file. Whenever an event is detected, it captures the data. It is used to monitor logs and to capture changes in Consul and CAS.

Metric Collection

Metric collection is performed by the sas-peek and sas-watch collectors. The sas-peek collectors periodically check the value of a predefined set of metrics and send the result as a metric event to the sas.metric RabbitMQ exchange. A set of predefined sas-peek collectors is installed on each machine in a SAS Viya deployment. These collectors capture metrics such as system resource usage, CAS usage, and RabbitMQ performance. For a list of metrics and associated collectors, see “Metrics in Data Mart Tables” on page 243 .

Log Collection

Log collection is performed by the sas-watch collector, which continuously monitors log files for new log messages. Here is the process used by the sas-watch log collector (the “log collector”) to collect and process log entries:

1 The log collector opens the ancestor directory of the directory tree that contains the log files. By default, the log collector looks for log files in the directory `/opt/sas/viya/config/var/log`.

2 The log collector locates and tracks log files that are in the log directory tree. The log collector scans the log directory tree and adds all subdirectories and log files to a watch list. Log files are identified by matching a regular expression that is specified in the sas-watch log command. The default regular expression is `*.log`.

   Note: To enable the sas-watch log command to watch a log, permissions must be specified for the log file so that the identity (under which the sas-watch command runs) can read the file.

3 After the log collector has identified all the log files in the directory tree, it opens the files and starts to watch for changes.

4 If new log files are created in the log directory, the log collector opens the files and starts to watch for changes.

5 As log messages are written to the log files, the log collector obtains the messages and transforms them into a stream of SAS log events.

6 The log collector publishes the log events to the sas.log RabbitMQ exchange.

After the log events are published to the RabbitMQ exchange, they are copied to the data mart and sent to SAS Environment Manager for display. See "Operations Infrastructure: Data Mart" on page 227.

Alert and Notification Collection

Collection of alert and notification messages is performed by a sas-check collector, which periodically monitors the notification service for messages. Collected messages are published to the sas.notification RabbitMQ exchange as notification events.

Operations Infrastructure: Operations Agent

The operations agent manages monitoring and operations activities in a SAS Viya deployment, including activities on all machines in the deployment.

Each machine in a deployment has a sas-ops-agent server that runs as a service. The role of the instance of the agent service on each machine is either as a client (ops-agent) or as a server (ops-agentsvr). Each role runs a set of tasks: client tasks and server tasks. All instances of the agents in the same role run the same set of tasks.

The client role (ops-agent) runs on every machine in a deployment. This role runs a schedule of tasks that use sas-peek commands to collect resource measurements on a regular basis. Other tasks use sas-check commands to query resources and check resource thresholds. The results of these tasks are published to a RabbitMQ exchange.

The server role (ops-agentsrv) runs on a single machine (the same machine that contains the data mart). It runs a schedule of tasks that gathers information from the RabbitMQ exchanges and runs processes to organize the data in the data mart.

Each task that the agent runs includes a command (such as a sas-peek or sas-check command) that is executed whenever the task runs. The task definition also includes information about how often the task runs (such as at a specific time of day or at a specified time interval) and the duration of the task. For a list of the default tasks, see "Summary of the Default Task List" on page 236.

Operations Infrastructure: RabbitMQ Exchanges

The RabbitMQ message exchanges are the component in the operations infrastructure that enables the information producers (such as metric data collectors) to operate independently of the information consumers (such as the data mart). The exchanges serve as an intermediate holding area between data collection and consumption.

The interaction between the collectors and the exchanges is simple. The collector places the collected data on the proper exchange and continues collecting data. Each type of collected data goes to a specific exchange. In some cases, data is written directly to the exchange, and in other cases, the ops-agent process writes the data. Because the collectors interact only with the exchanges, they do not need any knowledge of the consumers.See "Operations Infrastructure: Collectors" on page 225 for more information.

The interaction between the consumer and the exchanges is also simple. The consumer, which includes the data mart and its associated ETL process, collects whatever data is present from the exchanges at predetermined intervals. After the data is collected from the exchanges, the data mart and ETL processes prepare the data for storage and display. Because the collectors interact only with the exchanges, they do not need any knowledge of the collectors. See "Operations Infrastructure: Data Mart" on page 227 for more information.

The operations infrastructure uses these RabbitMQ exchanges:

sas.log
: Contains log message events from sas-watch log collectors

sas.metric
: Contains metric data from sas-peek collectors

sas.notification
: Contains notification and alert messages

sas.job
: Contains messages from job execution service jobs

sas.application
: Contains messages generated by the SAS middle-tier microservices

Operations Infrastructure: Data Mart

Metric, log, and notification events that are written to RabbitMQ exchanges are collected by the sas-stream process. The stream process collects events from these exchanges: sas.application, sas.log, sas.metric, and sas.notification. Then, the stream process writes the data to tab-separated value (TSV) files. There are different TSV file types for different types of event data. Here are examples of event data:

- metric data for each type of resource being monitored
- messages from the sas.application exchange
- raw log messages

Every five minutes, the etl_driver.sas process runs. For metric data, the etl driver process reads new data from the raw TSV file. It selects the fields that it needs to process and calculates rates for some metrics. Then, it writes the processed data to another TSV file and to a CAS table on the CAS server. For log data, the process reads new data from the raw TSV file, selects the fields that are needed to create a log record, and writes the data to a CAS search index on the CAS server.

SAS Environment Manager uses the processed data in the CAS table to display monitoring information, and it uses the processed data in the CAS search index to display logging information.

Each day in the early morning, old data is rolled off of the CAS table and the CAS search index. The existing table and the search index are deleted, and a new table and a search index are created using the next collection of metric and log data.

The TSV files that contain raw data and processed data store data for one day. At the end of each day, the ziptsv SAS job process compresses the TSV files that are older than a specified age (the default value is one day) in order to reduce the footprint that is occupied by the files. After the TSV files for the day are compressed, any new metric or log data is stored in a new TSV file. In addition, the ziptsv process removes compressed TSV files that are older than a specified number of days. The default is 10 days, but it can be changed with a configuration property in the SAS Infrastructure Data Server. You can use the `ops-dm-admin` command to modify the properties for the ziptsv job. See "Modify Property Values" on page 239 for more information. The process also deletes any log files that were produced by the ETL processes. After the process runs, the data mart contains the current day's uncompressed TSV file, the log files from the current day's processes, and a 10-day collection of compressed TSV files. As a result, after 10 days, the size of the data mart should have settled to a steady state.

How To: Operations Infrastructure Command Line

Overview

Operations infrastructure tasks are performed through the `sas-ops` command. Although you can run these commands manually, they are configured and run automatically as part of the operations infrastructure processes. The functions that are provided by this command follow:

- streaming of operations information, including notifications, alerts, metric data, and log messages
- validation of the SAS Viya environment and the operations infrastructure
- information about the SAS Viya environment, including the services, the machines, and the environment

Here is the format of the command:

```
sas-ops command --option
```

You must be the SAS install user (sas) in order to run the command.

Stream the Operations Information

Stream the Alert Messages

Use the `sas-ops alerts` command to stream the alert messages from SAS applications and components.

The default behavior is to stream alert messages until you stop the command. You can use the option `--timeout duration` to stream messages only for a specified time. The duration is specified using the format `0h0m0s0ms`, although subsets of this format are also allowed. See "Time Format" on page 241 for details and examples. For example, `--timeout 5m30s` specifies that alert messages are streamed for 5 minutes and 30 seconds.

Use the `--last n` option to display only a specified number (specified as *n*) of the most recent messages.

Use the `--format` option to specify the format for the data.

`--format json`
Streams the alerts in JSON format, with the data and parameters on one line.

`--format pretty`
Streams the alerts in JSON format, with the data and parameters on separate lines.

`--format line`
Streams the alerts on a single line. The timestamp is listed first.

`--format block`

`--format message`

Stream the Log Events

Use the `sas-ops logs` command to stream the log messages that are generated by SAS Viya applications and services.

The default behavior is to stream log messages until you stop the command. You can use the option `--timeout duration` to stream messages only for a specified time. The duration is specified using the format `0h0m0s0ms`, although subsets of this format are also allowed. See "Time Format" on page 241 for details and examples. For example, `--timeout 5m30s` specifies that messages are streamed for 5 minutes and 30 seconds.

The default behavior is to stream messages to the terminal on which the command was issued. However, you can use the `--format` option to specify a different format for the messages.

`--format json`
Streams the log messages in JSON format, with the message text and parameters on one line.

`--format pretty`
Streams the log messages in JSON format, with the message text and parameters on separate lines.

`--format line`
Streams the log message text and parameters on a single line. The timestamp is listed first.

`--format file`
Streams the log message text and parameters on a single line.

`--format term`
Streams the log message text and parameters in terminal format. The message level is listed first.

`--format plain`
Streams the log message text and parameters with no tagging to indicate different parts of the message. The abbreviated message level is listed first.

`--format logfmt`
Identifies the parts of the message using the format `message-part=string`. An example is `timestamp=2017-10-20T11:14:56.000000-04:00`. The messages are color-coded, depending on their level.

`--format template`
Streams the log message text and parameters with no tagging to indicate different parts of the message.

`--format event`
Streams the log messages in JSON format as used by the event service.

Because many log messages are produced in a typical environment, you can use these options to filter the message stream to include only those messages of interest.

`--match regular-expression`
Streams the messages that match the specified regular expression.

`--match-file file`
Streams the messages that match the regular expressions that are contained in the specified file.

`--min-level level`
Streams the messages that are at the specified level or a higher level.

Valid values for *level* are trace, debug, info, warn, error, fatal, and none.

`--source source`
Streams the messages only from the specified source.

Stream the Metric Data

Use the `sas-ops metrics` command to stream the metric data that is generated by SAS Viya applications and services.

The default behavior is to stream data until you stop the command. You can use the option `--timeout duration` to stream data only for a specified time. The duration is specified using the format `0h0m0s0ms`, although subsets of this format are also allowed. See "Time Format" on page 241 for details and examples. For example, `--timeout 5m30s` specifies that data is streamed for 5 minutes and 30 seconds.

The default behavior is to stream metric data to the terminal on which the command was issued. However, you can use the `--format` option to specify a different format for the data.

`--format json`
: Streams the metric data in JSON format, with the data and parameters on one line.

`--format pretty`
: Streams the metric data in JSON format, with the data and parameters on separate lines.

`--format line`
: Streams the metric data on a single line. The timestamp is listed first.

`--format property`
: Streams the metric data on a single line.

`--format event`
: Streams the metric data in JSON format as used by the event service.

Stream the Notification Messages

Use the `sas-ops notifications` command to stream notification messages from SAS applications and components.

The default behavior is to stream notifications until you stop the command. You can use the option `--timeout duration` to stream notifications only for a specified time. The duration is specified using the format `0h0m0s0ms`, although subsets of this format are also allowed. See "Time Format" on page 241 for details and examples. For example, `--timeout 5m30s` specifies that notifications are streamed for 5 minutes and 30 seconds.

Publish a Notification Message

Use the `sas-ops notify --level level message` command to publish a notification message. You can use the `--level` option to specify the level of the message. Supported level values are info, warn, or alert.

An example command is `sas-ops notify --level alert The server will be rebooted.`

Verify the Status of Your SAS Viya Environment

The Operations infrastructure provides the validate command to enable you to perform checks on your SAS Viya environment in order to locate problems. To use the command, SAS Viya does not have to be running. However, the SAS Configuration Server and RabbitMQ must be running.

The syntax of this command is
`sas-ops validate --level level --json --tags string --verbose`.

The value for *level* specifies the complexity of the validation checks. Each level performs the checks for its own level and the previous levels For example, specifying `--level 2` causes both the level 1 and level 2 checks to be performed.

Three levels of validation are available, in increasing order of complexity.

1

Verifies a connection to the SAS Configuration Server (Consul) and to the RabbitMQ exchanges sas.application, sas.log, sas.metric, and sas.notification. This level of validation ensures that you can perform validation checks at levels 2 and 3.

2

Verifies that the operation infrastructure is operating properly. These checks are performed:

- Verifies the SAS Configuration Server (Consul) services on each machine. The checks verify the following: the disk space used is not greater than 95%, the memory used is not greater than 95%, and the SAS Configuration Server is running.
- Verifies that the operations data mart ETL is running properly by performing checks on the standard regularly scheduled ETL jobs. The information displayed is the same information that is generated by the `sas-ops datamarts` command.
- Verifies the status of the operations services, including all instances of the ops-agent, alert-track, watch-log, ops-agentsrv, and stream-evdm services.

3

Verifies the status of CAS, the HTTP service, and the authorization service. These checks are performed:

- Verifies the status of the CAS servers by locating the specified servers and verifying that they are running.
- Verifies HTTP connectivity by attempting to connect to the base HTTP address.
- Verifies that the authorization service is working by attempting to obtain an OAuth token for the sas-ops command.

Display the Environment Information

Display Information about the Data Mart

Use the `sas-ops datamarts` command to display metric and status information about the data mart. Here is typical information that is returned:

```
evdm
   status
      etl_driver
         casLogLoad_SYSCCRC : 0
         casLogSearchLoad_SYSCCRC : 4
         casLogconnect_SYSCCRC : 0
         casMetricConnect_SYSCCRC : 0
         casMetricLoad_SYSCCRC : 0
         endtime : 2017-10-18T14:11:29.72-04:00
         jobExitRC : 0
         osLogRC : 1
         osMetricRC : 0
         readMetricTransform_SYSCCRC : 0
         starttime : 2017-10-18T14:11:05.881639-04:00
         status : ok
         statusRC : 4
      rolloff
         endtime : 2017-10-18T02:00:10.69-04:00
         jobExitRC : 42
         starttime : 2017-10-18T02:00:00.034636-04:00
         status : error
      ziptsv
         deleteOldZips_SYSCCRC : 0
```

```
        endtime : 2017-10-18T03:01:27.89-04:00
        jobExitRC : 0
        osZipTSVRC : 0
        starttime : 2017-10-18T03:00:00.023883-04:00
        status : ok
        statusRC : 0
        updateInventory_SYSCCRC : 0
        zipTSV_SYSCCRC : 0
```

The results include information about the following three jobs that are used by the data mart:

- The etl_driver job, which processes the metric data and the log data and loads the data into the data mart

  casLogLoad_SYSCCRC
  Return code from SAS for loading the log data into CAS for the log phase of the ETL job

  casLogSearchLoad_SYSCCRC
  Return code from SAS for updating the CAS search index for the log phase of the ETL job

  casLogconnect_SYSCCRC
  Return code from SAS for connecting to CAS for the log phase of the ETL job

  casMetricConnect_SYSCCRC
  Return code from SAS for connecting to CAS for the metric phase of the ETL job

  casMetricLoad_SYSCCRC
  Return code from SAS for loading the metric data into CAS for the metric phase of the ETL job

  endtime
  End time of the ETL job

  jobExitRC
  Return code from the operating system that is written by ops-runsas

  osLogRC
  Maximum return code (either 0,1, or 2) from SAS for the log phase

  osMetricRC
  Maximum return code (either 0,1, or 2) from SAS for the metric phase

  readMetricTransform_SYSCCRC
  Maximum return code from SAS for reading the raw TSV file for the metric phase

  starttime
  Start time of the ETL job

  status
  Status message (ok, error, or warning) that is written by ops-runsas at the end of the job

  statusRC
  Maximum return code from SAS

- Nightly rolloff job, which removes old data mart data from CAS

  endtime
  End time of the rolloff job

  jobExitRC
  Return code from the operating system that is written by ops-runsas

  starttime
  Start time of the rolloff job

  status
  Status message (ok, error, or warning) that is written by ops-runsas at the end of the job

- Nightly ZIPTSV job, which archives old TSV files into ZIP format, removes old archive files, and updates the data mart inventory

  deleteOldZips_SYSCCRC
  Return code from SAS for deleting old zipped TSV files

  endtime
  End time of the ZIPTSV job

  jobExitRC
  Return code from the operating system, written by ops-runsas

  osZipTSVRC
  Maximum return code (either 0,1, or 2) from SAS for archiving the TSV files

  starttime
  Start time of the ZIPTSV job

  status
  Status message (ok, error, or warning) that is written by ops-runsas at the end of the job

  statusRC
  Maximum return code from SAS

  updateInventory_SYSCCRC
  Maximum return code from SAS for updating all inventory files (one per resource type)

  zipTSV_SYSCCRC
  Maximum return code from SAS to Zip TSVs and delete old ZIP files, written by SAS

Access Information about Your Environment

Use the `sas-ops env` command to display information about the machine (on which you run the command), SAS environment variables, and the SAS Viya deployment.

Here is typical information that is returned:

```
Host Information:
  Full hostname              : full_hostname
  Short hostname             : short_hostname
  Consul node name           : full_hostname

SAS environment variables:
  CONSUL_HTTP_ADDR = https://localhost:8501

SAS Viya Deployment:
  Install user               : sas
  Deployment ID              : viya
  SAS root                   : /opt/sas
  Deployment root            : /opt/sas/viya
  Home directory             : /opt/sas/viya/home
  Config directory           : /opt/sas/viya/config
  Log directory              : /opt/sas/viya/config/var/log
  SPRE directory             : /opt/sas/spre
```

View Machine Information

Use the `sas-ops info` command to obtain information about each machine in your SAS Viya environment. For each machine in your environment, the command returns this information:

- machine identity

- packages installed on the machine
- system metrics
- system limits

The information returned by this command is the same information that is displayed on the Machines page in SAS Environment Manager.

View Information about Services

Use the sas-ops services command to view information about the services in your environment.

Run the command sas-ops services with no options to display a list of all SAS Viya services that are currently active in your environment.

Run the command sas-ops services --detail service-name to view detailed information about a specified service. Here is typical information that is returned:

```
{
  "ID": "e77ab2dc-b6c6-4a4d-af4e-bf3712de3c98",
  "Node": "vdmml-tue-17w47-ud.uda.sashq-r.openstack.sas.com",
  "Address": "10.104.29.192",
  "Datacenter": "",
  "TaggedAddresses": {
    "lan": "10.104.29.192",
    "wan": "10.104.29.192"
  },
  "NodeMeta": {},
  "ServiceID": "compute-10-104-29-192",
  "ServiceName": "compute",
  "ServiceAddress": "vdmml-tue-17w47-ud.uda.sashq-r.openstack.sas.com",
  "ServiceTags": [
    "proxy",
    "rest-commons",
    "https",
    "jobExecution-provider",
    "jobExecution-provider-Compute",
    "dataSources-provider",
    "dataSources-provider-Compute",
    "dataTables-provider",
    "dataTables-provider-Compute",
    "rowSets-provider",
    "rowSets-provider-Compute",
    "contextPath=/compute"
  ],
  "ServicePort": 39504,
  "ServiceEnableTagOverride": false,
  "CreateIndex": 16797,
  "ModifyIndex": 16797
}
```

Run the command sas-ops services --health service-name to perform the health checks on each instance of the specified service:

Disk utilization of SAS Config filesystem
    The check passes if disk usage does not exceed 95%.

Memory percent free
    The check passes if memory usage does not exceed 95%.

Serf Health Status
: The check passes if the SAS Configuration service is running.

Service '*service-name*' check
: The check passes if the service is running,

View Information about Metric Tasks

The operations agent (sas-ops-agent) runs a specified set of tasks to collect system metrics and to publish the metric data to RabbitMQ. Use the `sas-ops tasks` command to view a list of the tasks that are performed by the agent and the frequency of the task that is run. For more information about the agent and the tasks performed by the agent, see "How To: Operations Infrastructure Agent Command Line" on page 236.

Here is an example of the information that is returned by the `sas-ops tasks` command:

```
Task Name                    Description                                                         Frequency
---------                    -----------                                                         ---------
CASMetrics                   CAS performance metrics (level=2)                                   1m0s
CheckCpu                     Check CPU activity less than 95% busy                               1m0s
CheckFileSystemLinux         Check Linux file system space less than 90% used                    1m0s
CheckFileSystemWindows       Check Windows file system space less than 90% used                  1m0s
CheckMemory                  Check memory less than 95% used                                     1m0s
EmiSweeper                   Retry publishing any payloads that failed to publish earlier        1h0m0s
FileSystemMetrics            Host file system metrics (level=2)                                  1m0s
HostEnvSnapshot              Host environment snapshot                                           02:25
LogfileArchive               Archive daily                                                       04:00
NetworkInterfaceMetrics      Host network interface metrics (level=2)                            1m0s
OpsAgentActivity             Internal sas-ops-agent activity monitor                             2m0s
OpsAgentTaskStatistics       Internal sas-ops-agent task statistics activity monitor             4m0s
PostgresMetrics              Postgres metrics (level=2)                                          1m0s
RabbitmqMetrics              RabbitMQ performance metrics (level=2)                              1m0s
SpringBootMetrics            Spring Boot performance metrics (level=2)                           1m0s
SpringBootMetricsLevel3      Spring Boot performance metrics (level=3)                           4h0m0s
SystemMetrics                Host system metrics (level=2)                                       1m0s
TopProcessMetrics            Top CPU process consumers (level=2)                                 1m0s
registerOpsAgentServiceTask  Register Ops-Agent service task                                     5m0s
registerOpsServiceTask       Register Ops service task                                           5m0s
```

Use the `sas-ops tasks --name ops-agentsrv` command to view a list of the tasks that are performed by the agent server and the frequency of the task that is run.

Here is an example of the information that is returned by the `sas-ops tasks --name ops-agentsrv` command:

```
Task Name                       Description                                                         Frequency
---------                       -----------                                                         ---------
ARMEtl                          Application Response Measurement (ARM) ETL task                     01:18
DatamartEtl                     Datamart incremental etl driver                                     5m0s
DatamartRollOff                 Datamart daily rolloff task                                         02:18
DatamartzipTSV                  Datamart daily ZIP TSV task                                         03:43
EmiSweeper                      Retry publishing any payloads that failed to publish earlier        1h0m0s
FlushReqTask                    Flush task list to permanent storage on request                     0s
OpsAgentActivity                Internal sas-ops-agent activity monitor                             2m0s
OpsAgentTaskStatistics          Internal sas-ops-agent task statistics activity monitor             4m0s
genAudit                        Extract audit records. Generate a CSV files for given applications  2h0m0s
registerOpsAgentSvrServiceTask  Register Ops-AgentSvr service task                                  1h0m0s
updateInventory                 Update inventory                                                    01:23
```

How To: Operations Infrastructure Agent Command Line

Overview

The operations infrastructure agent runs a specified set of tasks. Each task is specified as a combination of a command to execute and information about how to publish the output of the command. Most tasks invoke the sas-peek or sas-check components and publish the output as an event to RabbitMQ.

The list of tasks for the agent to run are provided from the SAS Infrastructure Data Server or from a file to be read. The task definition also includes other attributes such as how often the task should be run.

Summary of the Default Task List

The following tasks are provided by default on the agent:

Task name	Command	Description
CASMetrics	`sas-peek cas —level 2`	Collects the CAS performance metrics (at level 2) every minute.
CheckCPU	`sas-check cpu -warning 95 -metric percentCpuBusy`	Checks the CPU activity and verifies that it is lower than 95% busy. It runs every minute.
CheckFileSystemLinux	`sas-check filesystem -warning 90 -metric percentUsedBytes -inctype 'xfs ext4' -min-level warn`	Checks the Linux file system and verifies that it is less than 90% used. It runs every minute.
CheckFileSystemWindows	`sas-check filesystem -warning 90 -metric percentUsedBytes -path 'c:' -min-level warn`	Checks the Windows file system and verifies that it is less than 90% used. It runs every minute.
CheckMemory	`sas-check memory -warning 95 -metric percentUsed`	Checks the system memory and verifies that it is less than 95% used. It runs every minute.
EmiSweeper (server)	`ops-event-sweep run -delete -verbose`	Attempts to publish any component outputs that were not published. It runs every hour.
FileSystemMetrics	`sas-peek filesystem -level 2`	Collects the host file system metrics (level 2). It runs every minute
EvdmDatamartEtl (server)	`ops-runsasjob -pgm etl_driver.sas -datamart evdm`	Specifies the ETL driver for the data mart. It runs every 5 minutes.
EvdmDatamartRollOff (server)	`ops-runsasjob -pgm rolloff.sas -datamart evdm`	Rolls off old data from the data mart. It runs at 2 AM every day.

Task name	Command	Description
evdmDatamartzipTSV (server)	`ops-runsasjob -pgm ziptsv.sas -datamart evdm`	Zips data from the data mart to a TSV file. It runs at 3 AM every day.
HostEnvSnapshot	`ops-sysinfo`	Takes a snapshot of the host environment at 2:25 AM every day.
LogfileArchive	`sas-archive`	Archives the logs from the previous day. It runs at 4 AM every day.
NetworkInterfaceMetrics	`sas-peek network -level 2`	Collects the host network interface metrics. It runs every minute.
OpsAgentActivity (server)	`sas-event-pub-exchange sas.metric`	Monitors the activity of the operations infrastructure agent. It runs every 2 minutes.
OpsAgentTaskStatistics (server)	`sas-event-pub-exchange sas.metric`	Monitors the activity of the operations infrastructure agent task statistics. It runs every 4 minutes.
PostgresMetrics	`sas-peek postgres -level 2`	Collects the metrics from the Infrastructure Data Server (level 2). It runs every minute.
RabbitmqMetrics	`sas-peek rabbitmq -level 2`	Collects the metrics from RabbitMQ (level 2), It runs every minute.
SpringBootMetrics	`sas-peek springboot -level 2`	Collects the metrics from Spring Boot (level 2). It runs every minute.
SpringBootMetricsLevel3	`sas-peek springboot -level 3`	Collects the metrics from Spring Boot (level 3). It runs every minute.
SystemMetrics	`sas-peek system -level 2`	Collects the host system metrics (level 2). It runs every minute.
TopProcessMetrics	`sas-peek top -level 2`	Collects the top consumers of CPU processes (level 2). It runs every minute.
registerOpsAgentSvrServiceTask (server)	`ops-util register -id sas.ops-agentsrv`	Registers the ops-agentSvr service task. It runs every 5 minutes.
genAudit	`genAudit.sh -a reports,folders,dataPlans,casManagement,casAccessManagement,--user-id -l 1000 -d 7`	Extracts the audit records for reports, folders, data plans, CAS management, and CAS access management, and generates a CSV file. It runs every 2 hours.

List the Tasks

To list the tasks that are in the current task list and that are loaded to the SAS Configuration Server, run the command `sas-ops-agent list`. The command returns a list of the tasks that are in the current list. The command also displays the frequency and a brief description for each task. Here is typical output:

```
Task name: CASMetrics
```

```
  Freq.........: 1m0s
  Description..: CAS performance metrics (level=2)
Task name: CheckCpu
  Freq.........: 1m0s
  Description..: Check CPU activity less than 95% busy
Task name: CheckFileSystemLinux
  Freq.........: 1m0s
  Description..: Check file system space on Linux system less than 90% used
Task name: CheckFileSystemWindows
  Freq.........: 1m0s
  Description..: Check file system space on Windows system is less than 90% used
Task name: CheckMemory
  Freq.........: 1m0s
  Description..: Check memory less than 95% used
Task name: EmiSweeper
  Freq.........: 1h0m0s
  Description..: Retry publishing any payloads that failed to publish earlier
Task name: FileSystemMetrics
  Freq.........: 1m0s
  Description..: Host file system metrics (level=2)
Task name: FlushReqTask
  Freq.........: 0s
  Description..: Flush task list to permanent storage on request
Task name: HostEnvSnapshot
  Freq.........: 02:25
  Description..: Host environment snapshot
Task name: LogfileArchive
  Freq.........: 04:00
  Description..: Archive daily
Task name: NetworkInterfaceMetrics
  Freq.........: 1m0s
  Description..: Host network interface metrics (level=2)
Task name: OpsAgentActivity
  Freq.........: 2m0s
  Description..: Internal sas-ops-agent activity monitor
Task name: OpsAgentTaskStatistics
  Freq.........: 4m0s
  Description..: Internal sas-ops-agent task statistics activity monitor
Task name: OpsArm
  Freq.........: 1h0m0s
  Description..: Runs Ops ARM
Task name: OpsValidate
  Freq.........: 0s
  Description..: Perform system validation when requested
Task name: PostgresMetrics
  Freq.........: 1m0s
  Description..: Postgres metrics (level=2)
Task name: RabbitmqMetrics
  Freq.........: 1m0s
  Description..: RabbitMQ performance metrics (level=2)
Task name: SpringBootMetrics
  Freq.........: 1m0s
  Description..: Spring Boot performance metrics (level=2)
Task name: SpringBootMetricsLevel3
  Freq.........: 4h0m0s
  Description..: Spring Boot performance metrics (level=3)
```

```
Task name: SystemMetrics
  Freq.........: 1m0s
  Description..: Host system metrics (level=2)
Task name: TopProcessMetrics
  Freq.........: 1m0s
  Description..: Top CPU process consumers (level=2)
Task name: registerOpsAgentServiceTask
  Freq.........: 1h0m0s
  Description..: Register Ops-Agent service task
Task name: registerOpsServiceTask
  Freq.........: 1h0m0s
  Description..: Register Ops service task
```

How To: ETL and Data Mart Operations

Obtain Information about the Data Mart

To obtain the resources used by the data mart:

```
ops-dm-admin usage -datamart datamart_name
```

Modify Property Values

You can manually change the property values that are used by the ETL and data mart processes.

To view a list of properties and their values:

```
ops-dm-admin show -datamart datamart_name
```

To modify property values:

```
ops-dm-admin set property-name=property-value -datamart datamart_name
```

Here are the properties that you can specify:

ARM_SUBSYS

Specifies the application response measurement (ARM) subsystem to use. Valid values are ARM_PROC or ARM_DSIO.

EMI_CAS_RETAIN_DAYS

Specifies the number of days that the metric and log data are retained in CAS. The default value is 3.

Note: Do not set this value to be greater than the value of the EMI_DELETE_TSVZIP_DAYS property, because the metric and log data will be deleted before the EMI_CAS_RETAIN_DAYS limit is reached. For example, if you set EMI_CAS_RETAIN_DAYS to 30 and EMI_DELETE_TSVZIP_DAYS to 10, the data is deleted after 10 days, and no data is ever available for days 11 through 30.

EMI_DELETE_TSVZIP_DAYS

Specifies the number of days to keep the metric and log data on disk in the data mart. The default value is 10.

EMI_ZIP_TSV_DAYS

Specifies the number of days of data in the data mart to collect before raw TSV files are compressed into a ZIP file. It also applies to the number of days to keep SAS log files that are generated by standard data mart batch jobs. The default value is 1. Do not use a value lower than 1.

Purge the Data Mart

You can delete files in the data mart and return the data mart to an empty status. This action is needed if you change memory or the hardware configuration on machines in your system. Purging the data mart ensures that all the stored metrics originate from the changed machines. An uncommon reason for purging the data mart is turning off the operations infrastructure while the rest of SAS Viya keep running. This scenario results in the accumulation of more log files than the operations infrastructure can process during its 10-minute processing window.

After the files are deleted, they cannot be recovered.

To delete all files in the data mart except for the tab-separated-value (TSV) input files:

```
ops-dm-admin purge -datamart datamart_name
```

You are prompted to confirm the action. To run the command without displaying the confirmation prompt, add the `-force` parameter to the end of the command.

To delete all files in the data mart including the tab-separated-value (TSV) input files, you must first stop the sas-viya-stream-evdm-default and sas-viya-ops-agentsrv-default services. After you have stopped the services, run this command:

```
ops-dm-admin purgeall -datamart datamart_name
```

You are prompted to confirm the action. To run the command without displaying the confirmation prompt, add the `-force` parameter to the end of the command.

Unlock the Data Mart

When an ETL job runs, it locks the data mart to prevent other jobs from modifying the contents of the data mart. If an error that keeps the data mart in a locked state occurs, you might need to unlock the data mart so that ETL jobs can resume operation.

To check whether the data mart is locked:

```
ops-dm-admin unlock -datamart datamart_name
```

If the data mart is already unlocked, this message is returned:

```
Datamart [datamart_name] is currently unlocked
```

If the data mart is locked, this output is returned:

```
datamart locked, process pid [pid_number] active

datamart lock user [user_id]

datamart lock process [active_process]

must use -force option to kill process then unlock

best practice is to wait until process finished
```

The recommended practice is to allow the process that has locked the data mart to complete. To end the process and unlock the data mart:

```
ops-dm-admin unlock -datamart datamart_name -force
```

Summary of ETL Return Codes

Here are possible return codes for ETL jobs:

Return code	Meaning
93	A required CAS table could not be initialized because the specified initialization job does not exist.
94	A required CAS table does not exist after the specified initialization job has been run.
95	A required SAS table could not be initialized because the specified initialization job does not exist.
96	A required SAS table does not exist after the specified initialization job has been run.
156	A lock on a data set could not be obtained within the specified time limit.
195	CAS configuration information has not been provided or located.
196	A CAS connection could not be established.
197	The contents of the data mart lock file could not be released when releasing the data mart lock.
198	The data mart lock could not be released.
199	The data mart lock could not be released because it is not locked by this process.
298	A data mart lock could not be obtained because the data mart is locked by another process.
299	A data mart lock could not be obtained because the data mart lock file cannot be renamed.

Operations Infrastructure Reference

Time Format

Time can be specified on the `sas-ops` commands in various formats. Here are the general forms of time specifications:

- An amount of time, specified in the format 0h0m0s0ms
- A specific time and date, specified in the format YYYY-MM-DDTHH:MM:SS.ssssssZhh:mm

Subsets and variants of each form are allowed. Here are some examples:

Specification	Time specified
"5m20s"	5 minutes and 20 seconds
"2h"	2 hours
"426s"	426 seconds

Specification	Time specified
"3:55:05.29754"	3:55 and 5.29754 seconds, in the current time zone, on any day Example: 0000-01-01 03:55:05.29754 -0500 EST
"14:18:34"	14:18 and 34 seconds, in the current time zone, on any day Example: 0000-01-01 14:18:34 -0500 EST
"14:18:34.38"	14:18 and 34.38 seconds, in the current time zone, on any day Example: 0000-01-01 14:18:34.38 -0500 EST
"14:28"	14:28, in the current time zone, on any day Example: 0000-01-01 14:28:00 -0500 EST
"5T14:25:04.54731"	14:25 and 4.54731 seconds, in the current time zone, on the fifth day of any month. Example: 0000-01-05 14:25:04.54731 -0500 EST
"05T09:37"	9:37, in the current time zone, on the fifth day of any month. Example: 0000-01-05 09:37:00 -0500 EST
"05T09"	9:00, in the current time zone, on the fifth day of any month. Example: 0000-01-05 09:00:00 -0500 EST
"05-23T17:26:09.237"	17:26 and 9.237 seconds, in the current time zone, on 23 May of any year. Example: 0000-05-23 17:26:09.237 -0500 EST
"04-05T09:37"	9:37, in the current time zone, on 5 April of any year. Example: 0000-04-05 09:37:00 -0500 EST
"04-16T21"	21:00, in the current time zone, on 16 April of any year. Example: 0000-04-16 21:00:00 -0500 EST
"03-15"	15 March, in the current time zone, of any year. Example: 0000-03-15 00:00:00 -0500 EST
"2015-04-05T14:18"	14:18, in the current time zone, on 5 April 2015. Example: 2015-04-05 14:18:00 -0400 EDT
"2015-04-05T18"	18:00, in the current time zone, on 5 April 2015. Example: 2015-04-05 18:00:00 -0400 EDT
"2018-01-16"	16 January, 2018, in the current time zone. Example: 2018-01-16 00:00:00 -0500 EST
"05:21:19Z"	5:21 and 19 seconds UTC, on any day. Example: 0000-01-01 05:21:19 +0000 UTC
"08:22:00.21Z"	8:22 and 0.21 seconds UTC, on any day. Example: 0000-01-01 08:22:00.21 +0000 UTC

Specification	Time specified
"14:25:04.54731Z"	14:25 and 4.54731 seconds UTC, on any day Example: 0000-01-01 14:25:04.54731 +0000 UTC
"3:55Z"	3:55 UTC, on any day Example: 0000-01-01 03:55:00 +0000 UTC
"17T14:25:04.54731Z"	14:25 and 4.54731 seconds UTC, on the 17th day of any month Example: 0000-01-17 14:25:04.54731 +0000 UTC
"05T09:37Z"	9:37 UTC, on the fifth day of any month Example: 0000-01-05 09:00:00 +0000 UTC
"05T09Z"	9:00 UTC, on the fifth day of any month Example: 0000-01-05 09:00:00 +0000 UTC
"05-23T17:26:09.23731Z"	17:26 and 9.23731 seconds UTC, on 23 May of any year Example: 0000-05-23 17:26:09.23731 +0000 UTC
"11-10T06:44Z"	6:44 UTC, on 10 November of any year Example: 0000-11-10 06:44:00 +0000 UTC
"09-21T16Z"	16:00 UTC, on 21 September of any year Example: 0000-09-21 16:00:00 +0000 UTC
"2019-11-07T13:21Z"	13:21 UTC, on 7 November, 2019 Example: 2019-11-07 13:21:00 +0000 UTC
"2019-11-07T11Z"	11:00 UTC, on 7 November, 2019 Example: 2019-11-07 11:00:00 +0000 UTC
"05T09+12"	9:00, UTC +12 hours, on any day Example: 0000-01-05 09:00:00 +1200 +1200
"2018-01-16+12:00"	16 January 2018, UTC +12 hours Example: 2018-01-16 00:00:00 +1200 +1200
"01-16+06:30"	16 January of any year, UTC +6 hours and 30 minutes Example: 0000-01-16 00:00:00 +0630 +0630

Metrics in Data Mart Tables

Overview

This topic provides information about the tables contained in the data mart and the metrics stored in each table. Each table in this topic lists the metrics that are stored in the associated data mart table, the metric names as returned by the sas-peek command, and the system reports in which the metrics are displayed, if applicable. See "Reports" on page 667 for information about the system reports. The sas-peek command that is used to return the metrics is listed prior to the table.

Understanding Calculated Metrics

Most metrics are stored in the data mart in the same format that they are collected in by the sas-peek command. However, some metrics are large integer values, and these values might not have much meaning by themselves. For such metrics, the ETL process creates a new rate-based metric in order to make the information easier to understand.

Here is an example of how a calculated metric is generated:

1 The sas-peek rabbitmq collector returns a metric named publish. This metric is a counter, and is a count of messages that are published on the RabbitMQ exchange.

2 The ETL process renames the metric as Publish_cnt and stores it in the RABBITMQ data mart table.

3 The ETL process creates a new metric named Publish that contains the rate of messages that are published on the RABBITMQ exchange at every data collection point. This new metric is calculated by determining the difference between the values of Publish_cnt at the present data collection point and the previous data collection point. The unit for the Publish metric is listed as count per second.

If the last time that the metric was collected is more than one hour from the current collection time, the calculated rate metric is set to the SAS missing value.

Calculated metrics can also be derived from time-based metrics. For example, the metric systemCPU in the CAS data mart table is calculated using the preceding steps. The ETL process uses the systemCPU metric, which is then stored in the CAS data mart table as systemCPU_cnt. Because the units of time for the original systemCpu metric are in seconds, and the collection interval is also in seconds, the units for the systemCPU metric are in seconds (or system CPU time) per second (of real time).

CAS Table

These metrics are captured by the sas-peek cas collector.

Note: This table contains calculated metrics. See "Understanding Calculated Metrics" on page 244 for information about calculations and units.

The following table lists the CAS system overview metrics:

Table A.1 Metrics in the CAS Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
IoCount	count per second	Calculated from the IoCount_cnt metric		Main page and CAS Details page of the CAS Activity report
IoCount_cnt	–	CAS process I/O count	ioCount	CAS Details page of the CAS Activity report
MaxMemoryUsed	bytes	CAS peak memory used	maxMemoryUsed	CAS Details page of the CAS Activity report

Metric	Units	Description	Associated sas-peek Metric	Used in Report
MaxThreadCount	–	CAS peak thread count	maxThreadCount	CAS Details page of the CAS Activity report
MemoryUsed	bytes	CAS memory used	memoryUsed	Main page and CAS Details page of the CAS Activity report
SystemCpu	seconds per second	Calculated from the SystemCpu_cnt metric		CPU Usage page and CAS Details page of the CAS Activity report
SystemCpu_cnt	seconds	CAS system CPU time	systemCpu	CAS Details page of the CAS Activity report
ThreadCount	–	CAS thread count	threadCount	Main page and CAS Details page of the CAS Activity report
Uptime	seconds per second	Calculated from the Uptime_cnt metric		CAS Details page of the CAS Activity report
Uptime_cnt	seconds	CAS uptime	uptime	CAS Details page of the CAS Activity report
UserCpu	seconds per second	Calculated from the UserCpu_cnt metric		CPU Usage page and CAS Details page of the CAS Activity report
UserCpu_cnt	seconds	CAS user CPU time	userCpu	CAS Details page of the CAS Activity report

CAS_NODE Table

These metrics are captured by the sas-peek cas collector.

The following table lists the metrics for the CAS nodes:

Table A.2 Metrics in the CAS_NODE Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
CpuLoad	–	CAS node CPU load	cpuLoad		Collected on UNIX only
IdleTime	ms		idleTime		
IoWaitTime	ms		ioWaitTime		
IrqTime	ms		irqTime		

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
MemoryFree	KB		memoryFree		
MemoryTotal	KB		memoryTotal		
Resident	bytes		resident		
SoftIrqTime	ms		softIrqTime		
SwapFree	KB		swapFree		
SwapTotal	KB		swapTotal		
SystemTime	ms		systemTime		
UserTime	ms		userTime		
VirtualTotal	KB		virtualTotal		
VirtualUsed	KB		virtualUsed		

CAS_SYSTEM Table

These metrics are captured by the sas-peek cas collector.

Note: This table contains calculated metrics. See “Understanding Calculated Metrics” on page 244 for information about calculations and units.

The following table lists metrics for the CAS system:

Table A.3 Metrics in the CAS_SYSTEM Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
Active	KB		active		
ActiveAnon	KB		activeAnon		
ActiveFile	KB		activeFile		
AnonHugePages	KB	Amount of memory used by transparent huge pages (2MB or 1GB)	anonHugePages		
AnonPages	KB		anonPages		
Bounce	KB		bounce		
Buffers	KB		buffers		
Cached	KB		cached		
CommitLimit	KB		commitLimit		

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
CommitedAs	KB		commitedAs		
DirectMap1G	KB		directMap1G		
DirectMap3M	KB		directMap3M		
DirectMap4K	KB		directMap4K		
Dirty	KB		dirty		
FifteenMinLoadAvg	none	15-minute load average	fifteenMinLoadAvg		Collected on UNIX only
FiveMinLoadAvg	none	Five-minute load average	fiveMinLoadAvg		Collected on UNIX only
Free	KB		free		
FreeFiles	–	Free files count	freeFiles		Collected on UNIX only
FreeInodes	–		freeInodes		Collected on UNIX only
HardwareCorrupted	KB		hardwareCorrupted		
HugePageSize	KB		hugePageSize		
HugePagesFree	KB		hugePagesFree		
HugePagesRsvd	KB		hugePagesRsvd		
HugePagesSurp	KB		hugePagesSurp		
HugePagesTotal	KB	total	hugePagesTotal		
Idle	seconds		idle		
IdleTime	seconds	System idle time	idleTime		
Inactive	KB		inactive		
InactiveAnon	KB		inactiveAnon		
InactiveFile	KB		inactiveFile		
Inodes	–		inodes		
IoWaitTime	seconds	I/O wait time	ioWaitTime		
IrqTime	seconds	IRQ time	irqTime		Collected on UNIX only
KernelStack	KB		kernelStack		
MLocked	KB		mLocked		
Mapped	KB		mapped		

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
MaxFiles	–	Maximum files count	maxFiles		Collected on UNIX only
MaxSystemThreads	–		maxSystemThreads		
NfsUnstable	KB		nfsUnstable		
OneMinLoadAvg	–	One-minute load average	oneMinLoadAvg		Collected on UNIX only
OpenFiles	–	Open files count	openFiles		
PageTables	KB		pageTables		
SReclaimable	KB		sReclaimable		
SUnreclaim	KB		sUnreclaim		
Shmem	KB		shmem		
Slab	KB		slab		
SoftIrqTime	seconds	Soft IRQ time	softIrqTime		Collected on UNIX only
SwapCached	KB		swapCached		
SwapFree	KB		swapFree		
SwapTotal	KB		swapTotal		
SystemTime	seconds	System time	systemTime		
Total	KB		total		
Unevictable	KB		unevictable		
Uptime	seconds	Calculated from the Uptime_cnt metric			
Uptime_cnt	seconds	System uptime	uptime		
UserTime	seconds	User time	userTime		
VmallocChunk	KB		vmallocChunk		
VmallocTotal	KB		vmallocTotal		
VmallocUsed	KB		vmallocUsed		
Writeback	KB		writeback		
WritebackTmp	KB		writebackTmp		

POSTGRES Table

These metrics are captured by the sas-peek postgres collector.

The following table lists overview metrics for the SAS Infrastructure Data Server (Postgres):

Table A.4 Metrics in POSTGRES Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
ConnectionsIdle	count	The number of idle connections	connectionsIdle	
ConnectionsMaximum	count	The maximum number of connections allowed	connectionsMaximum	
ConnectionUsage	percent	The percentage of connections that are being used	connectionsUsage	
ConnectionsUsed	count	The number of connections being used	connectionsUsed	
DatabaseSize	KB	The size of the Postgres database	databaseSize	
DirectorySize	KB	The size of the Postgres directory	directorySize	
FilesystemTotal	KB	The total size used by the file system	filesystemTotal	
FilesystemUsage	percent	The percentage of free space used by the file system	filesystemUsage	
FilesystemUsed	KB		filesystemUsed	
SASAuditEntryRows	count		sasAuditEntryRows	
SASAudirRows	count		sasAuditRows	

POSTGRES_TABLE_SIZE Table

These metrics are captured by the sas-peek postgres collector.

The following table lists the metrics for the size of tables in the SAS Infrastructure Data Server (Postgres):

Table A.5 Metrics in the POSTGRES_TABLE_SIZE Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
IndexSize	bytes		indexSize	Infrastructure Data Server Tables report

Metric	Units	Description	Associated sas-peek Metric	Used in Report
TableRank	–		tableRank	
TableSize	bytes		tableSize	Infrastructure Data Server Tables report
ToastSize	bytes		toastSize	Infrastructure Data Server Tables report
TotalSize	bytes		totalSize	

RABBITMQ Table

These metrics are captured by the sas-peek rabbitmq collector.

Note: This table contains calculated metrics. See "Understanding Calculated Metrics" on page 244 for information about calculations and units.

The following table lists the overview metrics for the RabbitMQ exchange bus:

Table A.6 Metrics in the RABBITMQ Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
Ack	messages per second	Calculated from the Ack_cnt metric		
Ack_cnt	–	Message_stats: raw ack	ack	
Confirm	messages per second	Calculated from the Confirm_cnt metric		
Confirm_cnt	–	Raw count of messages confirmed	confirm	
Consumers	–	Number of consumers	consumers	
Deliver	messages per second	Calculated from the Deliver_cnt metric	deliver	
Deliver_cnt	–	Raw count of messages delivered in acknowledgment mode to consumers	deliver	
DeliverGet	messages per second	Calculated from the DeliverGet_cnt metric		
DeliverGet_cnt	–	Raw count of messages delivered in acknowledgment mode in response to basic.get	deliverGet	
Exchanges	–	Number of exchanges	exchanges	

Metric	Units	Description	Associated sas-peek Metric	Used in Report
Messages	–	Sum of ready and unacknowledged messages (queue depth)	messages	
MessagesReady	–	Number of messages ready to be delivered to clients	messagesReady	
MessagesUnacknowledged	–	Number of messages delivered to clients but not yet acknowledged	messagesUnacknowledged	
Publish	messages per second	Calculated from the Publish metric		
Publish_cnt	–	Raw count of messages published	publish	
PublishIn	messages per second	Calculated from the PublishIn metric		
PublishIn_cnt	–	Raw count of messages published in to an exchange	publishIn	Message Queue Activity report
PublishOut	messages per second	Calculated from the PublishOut metric		
PublishOut_cnt	–	Raw count of messages published out of an exchange	publishOut	Message Queue Activity report
Queues	–	object_totals: queues	queues	
StatisticsDbEventQueue	messages per second	Calculated from the StatisticsDbEventQueue metric		
StatisticsDbEventQueue_cnt	–		statisticsDbEventQueue	

RABBITMQ_EXCHANGE Table

These metrics are captured by the `sas-peek rabbitmq` collector.

Note: This table contains calculated metrics. See “Understanding Calculated Metrics” on page 244 for information about calculations and units.

The following table lists the metrics for the RabbitMQ exchanges:

Table A.7 Metrics in the RABBITMQ_EXCHANGE Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
Ack	messages per second	Calculated from the Ack_cnt metric		
Ack_cnt	–	Message_stats: raw ack	ack	
Confirm	messages per second	Calculated from the Confirm_cnt metric		
Confirm_cnt	–	Raw count of messages confirmed	confirm	
Deliver	messages per second	Calculated from the Deliver_cnt metric		
Deliver_cnt	–	Raw count of messages delivered in acknowledgment mode to consumers	deliver	
DeliverGet	–	Calculated from the DeliverGet_cnt metric		
DeliverGet_cnt	–	Raw count of messages delivered in acknowledgment mode in response to basic.get	deliverGet	
Publish	messages per second	Calculated from the Publish_cnt metric		
Publish_cnt	–	Raw count of messages published	publish	
PublishIn	messages per second	Calculated from the PublishIn_cnt metric		
PublishIn_cnt	–	Raw count of messages published in to an exchange	publishIn	Message Queue Activity report
PublishOut	messages per second	Calculated from the PublishOut_cnt metric		
PublishOut_cnt	–	Raw count of messages published out of an exchange	publishOut	Message Queue Activity report

RABBITMQ_NODE Table

These metrics are captured by the sas-peek rabbitmq collector.

Note: This table contains calculated metrics. See “Understanding Calculated Metrics” on page 244 for information about calculations and units.

The following table lists the metrics for RabbitMQ nodes:

Table A.8 Metrics in the RABBITMQ_NODE Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
DiskFree	bytes	Disk free space	diskFree	
DiskFreeAlarm	–	Whether the disk alarm has gone off	diskFreeAlarm	
DiskFreeLimit	bytes	Point at which the disk alarm goes off	diskFreeLimit	
FdTotal	–	File descriptors available	fdTotal	
FdUsed	–	Used file descriptors	fdUsed	
IoReadAvgTime	ms	Average time for each disk Read operation	ioReadAvgTime	
IoReadBytes	bytes per second	Calculated IoReadBytes rate		
IoReadBytes_cnt	bytes	Total number of bytes read from disk by the persister	ioReadBytes	
IoReadCount	operations per second	Calculated from the IoReadCount_cnt metric		
IoReadCount_cnt	–	Total number of Read operations by the persister	ioReadCount	
IoReopenCount	operations per second	Calculated from the IoReopenCount_cnt metric		
IoReopenCount_cnt	–	Total number of times the persister has needed to recycle file handles between queues	ioReopenCount	
IoSeekAvgTime	ms	Average time for each Seek operation	ioSeekAvgTime	
IoSeekCount	operations per second	Calculated from the IoSeekCount_cnt metric		
IoSeekCount_cnt		Total number of Seek operations by the persister	ioSeekCount	
IoSyncAvgTime	ms	Average time for each fsync() operation	ioSyncAvgTime	

Metric	Units	Description	Associated sas-peek Metric	Used in Report
IoSyncCount	operations per second	Calculated from the IoSyncCount metric		
IoSyncCount_cnt	–	Total number of fsync() operations by the persister	ioSyncCount_cnt	
IoWriteAvgTime	ms	Average time for each disk Write operation	ioWriteAvgTime	
IoWriteBytes	bytes per second	Calculated from the IoWriteBytes_cnt metric		System Info page of the Message Queue Activity report
IoWriteBytes_cnt	bytes	Total number of bytes written to disk by the persister	ioWriteBytes	
IoWriteCount	operations per second	Calculated from the IoWriteCount_cnt metric		
IoWriteCount_cnt	–	Total number of Write operations by the persister	ioWriteCount	
MemLimit	–	Point at which the memory alarm goes off	memLimit	
MemUsed	bytes	Amount of memory used	memUsed	System Info page of the Message Queue Activity report
MnesiaDiskTxCount	transactions per second	Calculated from the MnesiaDiskCount_cnt metric		
MnesiaDiskTxCount_cnt	–	Number of Mnesia transactions that have been performed that required writes to disk	mnesiaDiskTxCount	
MnesiaRamTxCount	transactions per second	Calculated from the MnesiaRamTxCount_cnt metric		
MnesiaRamTxCount_cnt	–	Number of Mnesia transactions that have been performed that did not require writes to disk	mnesiaRamTxCount	
MsgStoreReadCount	messages per second	Calculated from the MsgStoreReadCount_cnt metric		
MsgStoreReadCount_cnt	–	Number of messages that have been read from the message store	msgStoreReadCount	

Metric	Units	Description	Associated sas-peek Metric	Used in Report
ProcTotal	–	Maximum number of Erlang processes	procTotal	
ProcUsed	–	Number of Erlang processes in use	procUsed	
Processors	–	Number of cores detected and usable by Erlang	processors	
QueueIndexJournalWriteCount	records per second	Calculated from the QueueIndexJournalWriteCount_cnt metric		
QueueIndexJournalWriteCount_cnt	–	Number of records written to the queue index journal	queueIndexJournalWriteCount	
QueueIndexReadCount	records per second	Calculated from the QueueIndexReadCount_cnt metric		
QueueIndexReadCount_cnt	–	Number of records read from the queue index	queueIndexReadCount	
QueueIndexWriteCount	records per second	Calculated from the QueueIndexWriteCount_cnt metric		
QueueIndexWriteCount_cnt	–	Number of records written to the queue index	queueIndexWriteCount	
RunQueue	–	Average number of Erlang processes waiting to run	runQueue	System Info page of the Message Queue Activity report
SocketsTotal	–	File descriptors available for use as sockets	socketsTotal	
SocketsUsed	–	File descriptors used as sockets	socketsUsed	
Uptime	ms per second	Calculated from the Uptime_cnt metric		
Uptime_cnt	ms	Time since the Erlang virtual machine started	uptime	

RABBITMQ_VHOST Table

These metrics are captured by the sas-peek rabbitmq collector.

Note: This table contains calculated metrics. See “Understanding Calculated Metrics” on page 244 for information about calculations and units.

The following table lists the metrics for the RabbitMQ virtual host:

Table A.9 Metrics in the RABBITMQ_VHOST Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
Ack	messages per second	Calculated from the Ack_cnt metric		
Ack_cnt		Message_stats: raw ack	ack	
Confirm	messages per second	Calculated from the Confirm_cnt metric		
Confirm_cnt	–	Raw count of messages confirmed	confirm	
Deliver	messages per second	Calculated from the Deliver_cnt metric		
Deliver_cnt	–	Raw count of messages delivered in acknowledgment mode to consumers	deliver	
DeliverGet	messages per second	Calculated from the DeliverGet_cnt metric		
DeliverGet_cnt	–	Raw count of messages delivered in acknowledgment mode in response to basic.get	deliverGet	
Messages	–	Sum of ready and unacknowledged messages (queue depth)	messages	
MessagesReady	–	Number of messages ready to be delivered to clients	messagesReady	
MessagesUnacknowledged	–	Number of messages delivered to clients but not yet acknowledged	messagesUnacknowledged	
Publish	messages per second	Calculated from the Publish_cnt metric		
Publish_cnt	–	Raw count of messages published	publish	
PublishIn	messages per second	Calculated from the PublishIn_cnt metric		
PublishIn_cnt	–	Raw count of messages published in to an exchange	publishIn	

Metric	Units	Description	Associated sas-peek Metric	Used in Report
PublishOut	messages per second	Calculated from the PublishOut_cnt metric		
PublishOut_cnt	–	Raw count of messages published out of an exchange	publishOut	
RecvOct	octets per second	Calculated from the RecvOct_cnt metric		
RecvOct_cnt	octets	Number of octets received	recvOct	
SendOct	octets per second	Calculated from the SendOct_cnt metric		
SendOct_cnt	octets	Number of octets sent	sendOct	

SPRINGBOOT Table

These metrics are captured by the sas-peek springboot collector.

Note: This table contains calculated metrics. See “Understanding Calculated Metrics” on page 244 for information about calculations and units.

The following table lists the metrics for the Spring Boot microservice:

Table A.10 *Metrics in SPRINGBOOT Table*

Metric	Units	Description	Associated sas-peek Metric	Used in Report
Classes	–		classes	Application History page of the Application Activity report
ClassesLoaded	–		classesLoaded	Application History page of the Application Activity report
ClassesUnloaded	–		classesUnloaded	Application History page of the Application Activity report
DatasourcePrimaryActive	–		datasourcePrimaryActive	Application History page of the Application Activity report
DatasourcePrimaryUsage	–		datasourcePrimaryUsage	Application History page of the Application Activity report

Metric	Units	Description	Associated sas-peek Metric	Used in Report
FreeMemory	KB		freeMemory	
GcPsMarksweepCount	operations per second	Calculated from the GcPsMarksweepCount_cnt metric		Application History page of the Application Activity report
GcPsMarksweepCount_cnt	–	Raw number of garbage collection marksweep operations	gcPsMarksweepCount	
GcPsMarksweepTime	ms per second	Calculated from the GcPsMarksweepTime_cnt metric		Application History page of the Application Activity report
GcPsMarksweepTime_cnt	ms	Raw time taken by garbage collection marksweep operations	gcPsMarksweepTime	
GcPsScavengeCount	operations per second	Calculated from the GcPsScavengeCount metric		Application History page of the Application Activity report
GcPsScavengeCount_cnt	–	Raw number of garbage collection scavenge operations	gcPsScavengeCount	
GcPsScavengeTime	ms per second	Calculated from the GcPsScavengeTime_cnt metric		Application History page of the Application Activity report
GcPsScavengeTime_cnt	ms	Raw time taken by garbage collection scavenge operations	gcPsScavengeTime	
Heap	KB		heap	
HeapCommitted	KB		heapCommitted	Application History page of the Application Activity report
HeapInit	KB		heapInit	
HeapUsed	KB		heapUsed	Application History page of the Application Activity report
HttpSessionsActive	–		httpSessionsActive	System Sessions History page of the Application Activity report
HttpSessionsMax	–		httpSessionsMax	

Metric	Units	Description	Associated sas-peek Metric	Used in Report
InstanceUptime	ms per second	Calculated from the InstanceUptime_cnt metric		
InstanceUptime_cnt	ms		instanceUptime	
LoadAverage	–		loadAverage	
NonHeap	KB		nonHeap	
NonHeapCommitted	KB		nonHeapCommitted	Application History page of the Application Activity report
NonHeapInit	KB		nonHeapInit	
NonHeapUsed	KB		nonHeapUsed	Application History page of the Application Activity report
Processors	–		processors	
Threads	–		threads	Application History page of the Application Activity report
ThreadsDaemon	–		threadsDaemon	Application History page of the Application Activity report
ThreadsPeak	–		threadsPeak	Application History page of the Application Activity report
ThreadsStarted	threads per second	Calculated from the ThreadsStarted_cnt metric		
ThreadsStarted_cnt	–		threadsStarted	
TotalMemory	KB		totalMemory	
Uptime	ms per second	Calculated from the Uptime_cnt metric		
Uptime_cnt	ms		uptime	Application History page of the Application Activity report

SYSTEM Table

These metrics are captured by the sas-peek system collector.

Note: This table contains calculated metrics. See "Understanding Calculated Metrics" on page 244 for information about calculations and units.

The following table lists metrics that are collected for the entire system:

Table A.11 Metrics in the SYSTEM Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
ActualFreeMemory	MB	Actual free memory	actualFreeMemory	System Details page of the System Activity report	
ActualUsedMemory	MB	Actual used memory	actualUsedMemory	System Details page of the System Activity report	
ContextSwitches	switches per second	Calculated from the ContextSwitches_cnt metric		System Details page of the System Activity report	
ContextSwitches_cnt	–	Raw value of system context switches	contextSwitches		
FreeMemory	MB	Free memory	freeMemory	System Details page and Memory Usage History page of the System Activity report	
FreeSwap	MB	Free swap space	freeSwap	System Details page of the System Activity report	
IdleCpu	ms per second	Calculated from the IdleCpu_cnt metric	idleCpu	System Details page of the System Activity report	
IdleCpu_cnt	ms	Raw value of idle CPU			Collected on UNIX only
IoWaitCPU	ms per second	Calculated from the IoWaitCPU_cnt metric		System Details page and CPU Details Animation page of the System Activity report	Collected on UNIX only
IoWaitCPU_cnt	ms	Raw value of I/O waiting CPU	ioWaitCPU		Collected on UNIX only
irqCpu	–	Calculated from the irqCpu_cnt metric		System Details page of the System Activity report	Collected on UNIX only

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
irqCpu_cnt	ms	Raw value of IRQ CPU	irqCpu		Collected on UNIX only
LoadAverage1	–	One-minute load average	loadAverage1	System Details page of the System Activity report	Collected on UNIX only
loadAverage5	–	Five-minute load average	loadAverage5	System Details page of the System Activity report	Collected on UNIX only
loadAverage15	–	15-minute load average	loadAverage15	System Details page of the System Activity report	Collected on UNIX only
NiceCpu	–	Calculated from the NiceCpu_cnt metric		System Details page of the System Activity report	Collected on UNIX only
NiceCpu_cnt	ms	Raw value of nice CPU	niceCpu		Collected on UNIX only
OpenFiles	files per second	Calculated from the OpenFiles_cnt metric			Collected on UNIX only
OpenFiles_cnt		Raw value of open files	openFiles		Collected on UNIX only
SoftIrqCpu	ms per second	Calculated from the SoftIrqCpu_cnt metric		System Details page of the System Activity report	Collected on UNIX only
SoftIrqCpu_cnt	ms	Raw value of soft IRQ CPU	softIrqCpu		Collected on UNIX only
StolenCpu	ms per second	Calculated from the StolenCpu_cnt metric		System Details page and CPU Details Animation page of the System Activity report	Collected on UNIX only
StolenCpu_cnt	ms	Raw value of stolen CPU	stolenCpu		Collected on UNIX only
systemCpu	ms per second	Calculated from the systemCpu_cnt metric		CPU Details Animation page of the System Activity report	
systemCpu_cnt	ms	Raw value of system CPU	systemCpu		

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
totalCpu	ms per second	Calculated from the totalCpu_cnt metric			
totalCpu_cnt	ms	Raw value of total CPU	totalCpu		
TotalMemory	MB	Total memory	totalMemory		
TotalSwap	MB	Total swap space	totalSwap	System Details page of the System Activity report	
Uptime	seconds per second	Calculated from the Uptime_cnt metric			
Uptime_cnt	seconds	Raw value of number of seconds since last boot	uptime	System Details page of the System Activity report	
UsedMemory	MB	Used memory	usedMemory	System Details page and Memory Usage History page of the System Activity report	
UsedSwap	MB	Used swap space	usedSwap	System Details page of the System Activity report	
UserCpu	ms per second	Calculated from the UserCpu_cnt metric		CPU Details Animation page of the System Activity report	
UserCpu_cnt	ms	Raw value of user CPU	userCpu	System Details page of the System Activity report	
idle_cpu_pct	percent				
system_cpu_pct	percent			CPU History page of the System Activity report	
total_cpu_pct	percent				
user_cpu_pct	percent			CPU History page of the System Activity report	

SYSTEM_CPU_USAGE Table

These metrics are captured by the sas-peek cpu collector.

The following table lists the metrics for the CPU usage for the system:

Table A.12 Metrics in the SYSTEM_CPU_USAGE Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
CpuDelta	ms	Total CPU available over interval	cpuDelta	
IdleUsage	percent	Idle CPU	idleUsage	
SystemUsage	percent	System CPU consumed	systemUsage	
UserUsage	percent	User CPU consumed	userUsage	

SYSTEM_FILESYSTEM Table

These metrics are captured by the sas-peek filesystem collector.

The following table lists the file system metrics for the system:

Table A.13 Metrics in SYSTEM_FILESYSTEM Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
Available	MB	Available space	available		
Files	count	Number of files (inodes)	files		Collected on UNIX only
Free	MB	Free space	free		
FreeFiles	count	Number of free files (inodes)	freeFiles		Collected on UNIX only
Size	MB	Total file system space	size		
Used	MB	Used space	used		

SYSTEM_NETWORK_INTERFACE Table

These metrics are captured by the sas-peek network collector.

Note: This table contains calculated metrics. See "Understanding Calculated Metrics" on page 244 for information about calculations and units.

The following table lists the metrics for system network performance:

Table A.14 *Metrics in the SYSTEM_NETWORK_INTERFACE Table*

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
ReceiveBytes	bytes per second	Calculated from the ReceiveBytes_cnt metric		Network Activity History page and Network Details page of the System Activity report	
ReceiveBytes_cnt	bytes	Raw value of bytes received in the interface	receiveBytes	Network Activity History page, Network Activity Animation page, and Network Details page of the System Activity report	
ReceiveCompressed	packets per second	Calculated from the ReceiveCompressed_cnt metric			Collected on UNIX only
ReceiveCompressed_cnt	–		receiveCompressed		Collected on UNIX only
ReceiveDropped	packets per second	Calculated from the ReceiveDropped_cnt metric			
ReceiveDropped_cnt	–	Raw value of number of received packets discarded	receiveDropped		
ReceiveErrors	packets per second	Calculated from the ReceiveErrors_cnt metric			
ReceiveErrors	–	Raw value of error packets received in the interface	receiveErrors		
ReceiveFrame	packets per second	Calculated from the ReceiveFrame_cnt metric			Collected on UNIX only
ReceiveFrame_cnt	–		receiveFrame		Collected on UNIX only
ReceiveMulticast	packets per second	Calculated from the ReceiveMulticast_cnt metric			
ReceiveMulticast_cnt	–	Raw value of multi-cast packets received	receiveMulticast		

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
ReceiveOverruns	packets per second	Calculated from the ReceiveOverruns_cnt metric			Collected on UNIX only
ReceiveOverruns_cnt	–		receiveOverruns		Collected on UNIX only
ReceivePackets	packets per second	Calculated from the ReceivePackets_cnt metric		Network Details page of the System Activity report	
ReceivePackets_cnt	–	Raw value of packets received in the interface	receivePackets		
TransmitBytes	bytes per second	Calculated from the TransmitBytes_cnt metric		Network Activity History page and Network Details page of the System Activity report	
TransmitBytes_cnt	bytes	Raw value of bytes sent in the interface	transmitBytes	Network Activity History page, Network Activity Animation page, and Network Details page of the System Activity report	
TransmitCarrier	packets per second	Calculated from the TransmitCarrier_cnt metric			Collected on UNIX only
TransmitCarrier_cnt	–		transmitCarrier		Collected on UNIX only
TransmitCollisions	packets per second	Calculated from the TransmitCollisions_cnt metric		Network Details page of the System Activity report	Collected on UNIX only
TransmitCollisions_cnt	–		transmitCollisions		Collected on UNIX only
TransmitCompressed	packets per second	Calculated from the TransmitCompressed_cnt metric			Collected on UNIX only
TransmitCompressed_cnt	–		transmitCompressed		Collected on UNIX only
TransmitDropped	packets per second	Calculated from the TransmitDropped_cnt metric		Network Details page of the System Activity report	

Metric	Units	Description	Associated sas-peek Metric	Used in Report	Operating System Support
TransmitDropped_cnt	–	Raw value of number of sent packets discarded	transmitDropped		
TransmitErrors	packets per second	Calculated from the TransmitErrors_cnt metric		Network Details page of the System Activity report	
TransmitErrors_cnt	–	Raw value of error packets occurring on transmission	transmitErrors		
TransmitMulticast	packets per second	Calculated from the TransmitMulticast_cnt metric			Collected on Windows only
TransmitMulticast_cnt	–	Raw value of multi-cast packets sent	transmitMulticast		Collected on Windows only
TransmitOverruns	packets per second	Calculated from the TransmitOverruns_cnt metric		Network Details page of the System Activity report	Collected on UNIX only
TransmitOverruns_cnt	–		transmitOverruns		Collected on UNIX only
TransmitPackets	packets per second	Calculated from the TransmitPackets_cnt metric		Network Details page of the System Activity report	
TransmitPackets_cnt	–	Raw value of packets sent in the interface	transmitPackets		

SYSTEM_PROCESS_USAGE Table

These metrics are captured by the sas-peek process collector.

Note: This table contains calculated metrics. See “Understanding Calculated Metrics” on page 244 for information about calculations and units.

The following table lists the system process metrics:

Table A.15 Metrics in the SYSTEM_PROCESS_USAGE Table

Metric	Units	Description	Associated sas-peek Metric	Used in Report
CpuDelta	ms	CPU consumed over sample interval	cpuDelta	

Metric	Units	Description	Associated sas-peek Metric	Used in Report
CpuTotal	ms per second	Calculated from the CpuTotal_cnt metric		
CpuTotal_cnt	ms	Total CPU consumed by process	cpuTotal	
CpuUsage	percent		cpuUsage	
DeltaRead	bytes	Calculated from the DeltaRead_cnt metric		
DeltaRead_cnt	operations per second	Read for process	deltaRead	
DeltaReadWrite	–	Calculated from the DeltaReadWrite_cnt metric		
DeltaReadWrite_cnt	bytes	Total Read/Write for process	deltaReadWrite	
DeltaWrite	operations per second	Calculated from the DeltaWrite_cnt metric		
DeltaWrite_cnt	bytes	Write for process	deltaWrite	
PageFaults	faults per second	Calculated from the PageFaults_cnt metric		
PageFaults_cnt	–	Page faults for process over sample interval	pageFaults	
ProcessRank	–	Current rank of process in resource consumption	processRank	
ResidentMemory	KB	Memory consumed by process	residentMemory	

17

Mobile

Mobile: Overview

The SAS Visual Analytics App (previously called SAS Mobile BI) enables mobile device users to view and interact with reports that can contain a variety of charts, graphs, gauges, tables, and other report objects. Supported mobile devices include iPads, iPhones, Android tablets and smartphones, and Windows 10 tablets and personal computers. For information about how to use the SAS Visual Analytics App, see the SAS Visual Analytics App Help.

As an administrator, you can control how a mobile device running the SAS Visual Analytics App can access reports and data located on a SAS Visual Analytics server. You can use the following features, rules, and properties (alone or in combination) to control access to the server data and reports from the app:

Note:

These settings and features can also be used to manage custom mobile apps that you create using Software Development Kits (SDKs), SAS SDK for Android or SAS SDK for iOS. For more information, see “Mobile: Software Development Kits” on page 281.

Blacklist and whitelist feature
: You can manage whether a device can access servers through the SAS Visual Analytics App, either by exclusion or inclusion.

Note: The device ID that is added to the blacklist or whitelist is validated, not the device itself.

Passcode properties and rule
: You can require SAS Visual Analytics App users to lock the app with a passcode. You can configure two properties that control the behavior of the passcode.

Offline access time-out property and rule
: If a user has not opened the SAS Visual Analytics App for a specified number of days, you can require the user to enter the user ID and password to access the server. The time-out is specified by a server property. You can use a rule to identify users who are exempt from the time-out.

Remote report data rule
: You can specify that when users view a report in the SAS Visual Analytics App, the mobile device must maintain a network connection to the server.

Limit functionality in the app
: You can limit the functionality of the SAS Visual Analytics App by applying one or more rules to a user or group of users. Functionality includes whether a user can subscribe to and view reports; share links to reports (and screen captures) by using email, text messaging, or other functionality; add or view comments; see and use the Favorites or Recent views; and view alerts.

Mobile: How To

Manage Mobile Devices

Navigate to the Mobile Devices Page

Note: This page is available only if you are a member of the SAS Administrators group.

1 Click ☰ and select **Manage Environment**.

2 In the navigation bar, click ▯ Mobile Devices.

Add a Device to a List from Last Access

You can add a device that has already connected (or attempted to connect) to the blacklist or whitelist by completing the following steps:

> **TIP** This option is disabled if the ID already exists on the respective list.

1 On the Mobile Devices page, click the **Last Access** tab.

2 Select the device and click +.

3 Select the list to which you want to add the device.

4 In the Add Device window, click **Yes**.

Add One or More Devices to the Blacklist or Whitelist

1 On the Mobile Devices page, click the **Blacklist** or **Whitelist** tab, depending on which list you want to add devices.

2 You can add one device or multiple devices to a list:

- To add one device to a list, click ✚.

  Enter the Device ID in the Add to Blacklist or Add to Whitelist window.

- To add multiple devices to a list, click ✚✚.

  In the Add to Blacklist or Add to Whitelist window, enter each Device ID to create a new line.

  Note: Validation is not performed on the device IDs as they are added to the list.

3 Click **Save**.

Move One or More Devices between Lists

You can move devices from one list to the other (for example, from the blacklist to the whitelist).

1 On the Mobile Devices page, click the tab that corresponds to the list from which you want to move a device.

2 Select one or more devices that you want to move, and click 📂.

3 In the Move Device window, click **Yes**.

Remove One or More Devices from a List

1 On the Mobile Devices page, click the tab that corresponds to the list from which you want to remove a device.

2 Select one or more devices that you want to remove, and click 🗑.

3 In the Confirm Remove window, click **Yes**.

View Logon Event Information

1 On the Mobile Devices page, click the **Last Access** tab.

2 View the device logon event information, including status. For more information, see "Device Logon Information".

> **TIP** Use the **Filter by** drop-down list to filter the information about the tab.

View Previous Logon Events

1 On the Mobile Devices page, click the **Last Access** tab.

2 To view records that were captured from devices on a prior application version or operating system version, select the **Include device history** option.

Determine Which List Is Enforced

There are several ways to determine whether the blacklist or whitelist is being enforced.

- On the **Mobile Devices** page, look for the following indicators:
  - The list that is being enforced has a ● next to the list name.
  - The list that is being enforced displays the following message above the **Device ID** table:

    "ⓘ This list is currently being enforced.".
  - The list that is not being enforced displays the following message above the **Device ID** table:

    "⚠ This list is not currently being enforced.".
- On the **Dashboard**, look on the **Mobile Devices** tile.
- Run the following command-line interface (CLI) command and look at the output:

  ```
  sas-admin devices enforcement status
  ```

  For more information about how to run the device management CLI, see "Command-Line Interface: Preliminary Instructions" on page 681.

Change How Devices Are Managed

CAUTION! These are deployment-level instructions that affect user access. Changing how devices are managed can disrupt existing users by changing which devices are eligible to connect to servers through the SAS Visual Analytics App.

1. Verify that the list that you intend to enforce is appropriately populated.
   - If you enforce the whitelist, the whitelist should contain all eligible devices. The blacklist is ignored.
   - If you enforce the blacklist, the blacklist should contain all excluded devices. The whitelist is ignored.
2. On the **Mobile Devices** page, click the tab that corresponds to the list that you want to enforce.
3. To change the list that is enabled, click **Enable**.

   Note:

   If the **Enable** button is disabled and you are working in a multi-tenant environment, contact the tenant administrator. You cannot alter the active list settings until the administrator sets the configuration properties on the individual tenant.
4. In the confirmation window, click **Yes** to enable the new list.

Limit Functionality

Initially, all authenticated users can access all functionality in the SAS Visual Analytics App. To change this behavior:

1. Locate the relevant authorization rule. For a list of the mobile-specific rules, see "Rules to Control Access to SAS Visual Analytics App Functionality" on page 278.
2. Choose one of the following:
   - Disable the rule.

- Change the principal type from its initial value (Authenticated Users) to a different value (for example, the group ID for a custom group).

  Note: Users who are within the scope of a revised rule have access to the functionality that the rule provides. Other users do not have access to the functionality that the rule provides.

For information about how to disable rules and change principal types, see "Example of the Basic Approach to Modifying Access to Functionality" on page 357.

Adjust Passcode Constraints

To adjust the passcode constraints, use the **passcodeAttempts** and **passcodeTimeoutMinutes** properties in the **sas.devicemanagement** configuration definition.

1 Click ≡ and select **Manage Environment**.

2 In the navigation bar, click 🔧.

3 In the **View** list, choose **Definitions**.

4 In the **Filter** field, type `device`.

5 Select **sas.devicemanagement** from the results. The configuration properties appear in the right pane.

6 Click ✎.

7 Edit the value in the **passcodeAttempts** field to configure the passcode lock-out behavior.

8 Edit the value in the **passcodeTimeoutMinutes** field to configure the passcode time-out behavior.

9 Click **Save**.

Adjust the Time-Out Interval

To adjust the time limit, set the **offlineLimitDays** property in the **sas.devicemanagement** configuration definition.

1 Click ≡ and select **Manage Environment**.

2 In the navigation bar, click 🔧.

3 In the **View** list, choose **Definitions**.

4 In the **Filter** field, type `device`.

5 Select **sas.devicemanagement** from the results. The configuration properties appear in the right pane.

6 Click ✎.

7 Edit the value in the **offlineLimitDays** field. Specify, in days, how many days a device can be offline without requiring the user log on to the server when opening SAS Visual Analytics App again.

8 Click **Save**.

Configure an Additional External URL

If you need to use a different URL prefix for external URLs (for example, in generated email messages or text messages), run the following command:

```
/opt/sas/viya/home/bin/sas-bootstrap-config  -token-file
/opt/sas/viya/config/etc/SASSecurityCertificateFramework/tokens/consul/default/client.token
 kv write config/application/sas.url.external.viya <userSuppliedValue>external-URL</userSuppliedValue> -force
```

Note: The previous command must be on one line. It is shown on more than one line for display purposes only.

For more information, see "Use the SAS Bootstrap Config CLI on Consul to Manage the KV Store and ACL Tokens " in *Encryption in SAS Viya: Data in Motion*.

Mobile: Concepts

SAS Viya provides ways to manage mobile devices and the security of reports and data. You can manage mobile devices by using a combination of configuration properties for the server and authorization rules that control the access of mobile device users to the server.

Prerequisites for Managing Mobile Devices

To manage mobile devices, you must be a SAS administrator, and your user ID must have the device management authorization rules for mobile devices. Initially, the SAS Administrators group has Read and Write access to the `/deviceManagement_capabilities/manageMobileDevices` object URI.

Blacklist and Whitelist Features

Overview

The *whitelist* manages the devices that can access servers by using the SAS Visual Analytics App. A device must be on the whitelist in order to use the SAS Visual Analytics App on your network. The whitelist affects devices, not users. If a device is lost, a SAS administrator can remove the device from the whitelist and prevent access to the reports and data.

The *blacklist* manages the devices that cannot access servers by using the SAS Visual Analytics App. All devices can use the SAS Visual Analytics App on your network except those that are on the blacklist. The blacklist affects devices, not users. If a device is lost, a SAS administrator can add the device to the blacklist and prevent access to the reports and data.

Important: To help ensure that unauthorized individuals do not gain access to your servers through the whitelist feature, ensure that a password policy is enforced with strict password controls.

Considerations

Here are the key points for managing mobile devices:

- You can manage devices either by exclusion or by inclusion.
  - If you manage by exclusion, all devices can access servers through the SAS Visual Analytics App, except those that are on the blacklist. A blacklist is a list of mobile devices that are not authorized to use the SAS Visual Analytics App.

  - If you manage by inclusion, only devices that are on the whitelist can access servers through the SAS Visual Analytics App. A whitelist is a list of mobile devices that are authorized to use the SAS Visual Analytics App.

- A deployment enforces only one list (either the blacklist or the whitelist) at a time.
- In a new deployment, the blacklist is enforced and contains no items. Therefore, all devices can access servers through the SAS Visual Analytics App.
- You can modify both lists. Making changes to a list that is not currently enforced can help accommodate a future change.
- The blacklist and whitelist affect devices, not users. As an administrator, you authorize what a particular user can see or do. For more information, see General Authorization on page 411.
- The device ID that is added to the blacklist or whitelist is validated, not the device itself.

Add Devices by User ID

The easiest way to add a device to the whitelist or blacklist is to add a device that has already connected (or attempted to connect) to the server. When the attempt is made, the **Last Access** tab logs the device owner’s user ID, device ID, device type, and other information. You can sort the **User ID** column to locate the user ID of the person whom you want to add.

Restricting and enabling devices by user ID is a best practice because users can have more than one device. By identifying the user ID, you can be sure to add all devices used by that person.

> **TIP** The only way to add a device running Windows 10 is by user ID.

Passcode Feature

Overview

The passcode feature locks the SAS Visual Analytics App. This feature is separate from and in addition to the passcode feature that is provided by mobile devices. There are two types of app passcodes: required and optional.

A *required passcode* is a passcode that is required by the server. When the app first connects to the affected server, the server forces the app to require that the app user create a passcode. Then, whenever the app user opens the app or views a report that is associated with that server, the user must enter the passcode.

Note: By using an additional rule, the SAS administrator can exempt app users from using a passcode. By using a combination of two rules, all mobile devices that access the server must use a passcode except for those separately exempted.

An *optional passcode* is a passcode that the app user can choose to use to lock the app. The passcode is not required to access the server. The app user can disable the passcode at any time.

Considerations

Here are some key points to remember when working with passcodes:

- The passcode should be known only to the app user. If the app user loses the mobile device, no one else should be able to guess the passcode and use it to open the app.
- The passcode has a time-out feature. The SAS administrator can customize the passcodeTimeoutMinutes setting to configure this feature. This setting specifies, in minutes, how long a user must wait before re-entering his or her passcode in the SAS Visual Analytics App. The default is 15.

If the app user (or another person) provides an incorrect passcode a specific number of times (passcodeAttempts), the app locks itself for a length of time (passcodeTimeoutMinutes). The app user can enter the passcode again after the time-out expires.

- The passcode has a lock-out feature. The SAS administrator can customize the passcodeAttempts setting to configure this feature. The setting limits the number of sequential, failed attempts to enter a passcode for the SAS Visual Analytics App. The default is 5.

  If a user reaches the specified limit (passcodeAttempts), the user is timed out of the app for 15 minutes (or the value set for passcodeTimeoutMinutes). After the time-out interval, the user can make one more attempt to enter his or her passcode. If the password fails again, all custom content (data, reports, settings, and connection information) is removed from the mobile device. The app is reset to its default settings.

- If the app user forgets the passcode, the app user must delete and re-install the app on the device. Doing so deletes the reports and data.

For information about how app users set a passcode, see the SAS Visual Analytics App Help. Be sure to view the Help for the platform (iOS, Android, or Windows 10) and release that you are using.

Cache Report Data Feature

Cache Report Data

When a user subscribes to a report, it appears in the **Subscriptions** view of the SAS Visual Analytics App. However, depending on the security assigned to the user ID, the report data might not exist on the mobile device. Report data can be local or remote:

- *Local* data is stored on the mobile device.
- *Remote* data exists on the mobile device only while the report is open and the device is connected to a Wi-Fi or cellular network. If a report uses remote data, the report tile in the **Subscriptions** view displays the cloud icon.

How Cache Report Data Works

Each time a user opens a cached report, the app connects to the server. The Prepare Data notification is displayed while the data is downloaded. The report opens when the data is available on the mobile device. The data is available only while the user views the report.

After the user closes the report, the data is removed from the device. The thumbnail image on the report tile no longer appears. If the user is not connected to a network and tries to open the report, it does not open.

This feature affects the user ID that is used to access the server. When the user accesses the server via the SAS Visual Analytics App using that user ID, all reports on that server use the caching report data feature.

Prevent Mobile Devices from Storing Report Data

The /SASMobileBI_capabilities/cacheMobileReportData rule specifies that a mobile device can store (or cache) report data on the device when it is not connected to a network. By default, all authenticated users' mobile devices can cache report data.

If you want to enforce additional security by preventing mobile devices from storing report data, then you must prohibit the authorization rule that is applied to a user or group of users.

Offline-Access Time-Out Feature

If a user has been offline for a specified number of days, he or she must sign in to the server used by the SAS Visual Analytics App. For example, if the user attempts to browse reports on the server or open a report in the report viewer, the app requires the user to enter the password for the requested server connection. If the user fails to sign in, then the app no longer downloads reports, updates subscribed reports, or opens reports for viewing.

This feature is not only useful when the device is missing. It also provides security when the employee leaves the organization but keeps the device. The blacklist and whitelist features require that the device must access the server before the list can look up the device to deny or permit access. The offline access time-out feature denies access by checking the employee's credentials, which the IT organization revokes when the employee leaves the organization.

Mobile: Reference

Device ID Criteria

To add one or more devices to the blacklist or whitelist, you must enter valid device IDs. If an invalid ID is entered, you cannot add devices to the lists. In order for an ID to be valid, the following conditions must be met:

- non-empty string
- length of 36 characters or less
- contains alphanumeric characters and hyphens
- is not a duplicate of an existing device ID

Device Logon Information

A device might appear multiple times in the blacklist or whitelist if a different user ID attempts to log on with a device that has already been captured. The following occurrences are logon events:

- a connection attempt that comes from a new source (a unique combination of device ID and user ID)
- a connection attempt that comes from an existing source (existing device ID and new user ID)
- a connection attempt that is accompanied by a device change (such as a new operating system version or application version)

The following table lists the device status icons that might be displayed:

Table A.1 Status Icons and Descriptions

Icon	Description
✓	Indicates that the authentication was a success.
⊗	Indicates that the authentication was a failure.
❗	Indicates that the device is in the blacklist.
⊘	Indicates that the device is not in the whitelist.

Rules to Control Access to SAS Visual Analytics App Functionality

You can limit the functionality of the SAS Visual Analytics App by applying one or more rules to a user or group of users.

The following table lists the rules that enable you to limit a user's access to functionality in the SAS Visual Analytics App:

CAUTION! Most of the rules in the table are not specific to the SAS Visual Analytics App. Therefore, careful consideration should be taken before making modifications to those rules that affect other applications. For more information, see "Identity Management: Access to Functionality" on page 355.

Table A.2 Rules and Descriptions

Object URI	Description	Rule Specific to SAS Visual Analytics App
/comments/comments	Create permission enables users to add comments to a report or its contents. Read permission enables users to view the comments that are associated with a report or its content.	No
/deviceManagement_capabilities/manageMobileDevices	Enables users to manage mobile device blacklist, whitelist, and device access history. Initially granted to SAS administrators.	Yes
/folders/folders/@myFavorites	Create permission enables users to add a report to the Favorites view. Delete permission enables users to remove a report from the Favorites view. Read permission enables users to view items for a particular server in the Favorites view. It also enables the user, in Add Reports, to see the contents of the Favorites folder for that server. Note: When Read permission is granted, it takes precedence over Create and Delete permissions. If a user is not authorized to view favorites, the user cannot add or remove favorites, even if the user was granted Create and Delete permission.	No
/folders/folders/@myHistory	Enables users to see items for a particular server in their Recent view. It also enables the user, in Add Reports, to see the contents of the Recent folder for that server.	No
/reportAlerts/*	Enables users to view alerts in a report and to subscribe to them.	No
/reportData_capabilities/exportData	Enables users to export data for a report object.	No
/reportData_capabilities/exportDetailData	Enables users to select the **Detailed data** option (if applicable) in the Export Data window. Note: The EXPORT_DATA rule takes precedence over this rule. If a user ID is not authorized to use the EXPORT_DATA rule, then authorizing the EXPORT_DETAIL_DATA rule to that user ID has no effect.	No
/reportRenderer/reports	Enables users to access PDF printing of reports.	No

Object URI	Description	Rule Specific to SAS Visual Analytics App
/reportViewerNaturalLanguageUnderstanding/interpretations	Enables users to process natural language on an iOS app. **Note:** Current support for this feature is provided only by the iOS version of SAS Visual Analytics App.	No
/SASMobileBI/**	Enables users to subscribe to and view reports. It also controls whether the user can define a connection to the server within the SAS Visual Analytics App.	Yes
/SASMobileBI_capabilities/allowWebContent	Enables users to render reports with web content.	Yes
/SASMobileBI_capabilities/cacheMobileReportData	Enables users to cache mobile report data from within the SAS Visual Analytics App. This is required for offline access to reports. For users who do not have this capability, report data is retained on the device only while the report is open. **Note:** For users who are within the scope of the revised rule, report data is cached. For users who are outside the scope of the revised rule, report data is downloaded when a report is open and purged when the report is closed. Offline access to reports is not supported for users who are outside the scope of the revised rule.	Yes
/SASMobileBI_capabilities/exemptFromOfflineTimeLimit	Enables users to be exempt from the SAS Visual Analytics App offline time-out. Initially, no users are subject to time-outs for offline access. **Note:** Users who are within the scope of the revised rule are exempt from time-outs. Users who are outside the scope of the revised rule are subject to time-outs.	Yes
/SASMobileBI_capabilities/exemptFromPasscodeRequirements	Enables users to be exempt from the requirement to enter a passcode to access the SAS Visual Analytics App. Initially, use of a passcode is not required. However, you can require passcodes or adjust passcode constraints. **Note:** If any of the mobile server connections require a passcode, then it is still required to access the application. This is true even if the exemption rule is in effect. In addition, users can enable a passcode even if the exemption rule is in effect.	Yes
/SASVisualAnalyticsCommon_capabilities/shareReport	Enables users to share links to reports (and screen captures) by using email, text messaging, or other functionality.	No
/webDataAccess/esri/user/token	Enables mobile users to view Esri maps.	No

Mobile: Troubleshooting

A user cannot open reports on an offline device.

Explanation:

The user ID might be required to use remote report data.

The user ID might be affected by the offline-access time out.

Resolution:

If the user ID is subject to the remote report data authorization rule, make sure the user understands that he or she must be connected to a network while viewing the report. See "Cache Report Data Feature".

If the user ID is subject to the offline-access time out authorization rule, make sure the user can log on to the server connection in the SAS Visual Analytics App. See "Offline-Access Time-Out Feature".

A user is prompted for an application passcode.

Explanation:

The user is required to secure the SAS Visual Analytics App with a passcode. See "Passcode Feature".

Resolution:

To learn how to create a required passcode in the SAS Visual Analytics App, see the SAS Visual Analytics App Help.

Note: Be sure to view the Help for the platform (iOS, Android, or Windows 10) and release the user is accessing.

On the Mobile Devices page, a message indicates that a list is not currently in use.

Explanation:

By design, only one list (either the blacklist or the whitelist) is in use.

As an administrator, you are unable to change the blacklist or whitelist using the Enable or Disable button.

Explanation:

If you are working in a multi-tenant environment, the device configuration properties are available only in the provider tenant configuration. You cannot alter the active list settings until the administrator sets the configuration properties on the individual tenant.

TIP For more troubleshooting information about the SAS Visual Analytics App, see the SAS Visual Analytics App Help. Be sure to view the Help for the platform (iOS, Android, or Windows 10) and release the user is accessing.

Mobile: Interfaces

There are multiple interfaces available to administer mobile devices. In the following table, the shaded part of each circle is an approximation of the amount of mobile functionality that a particular interface exposes. The shading indicates relative coverage. The shading does not indicate alignment of functional coverage across interfaces.

Table A.3 Interfaces to Mobile

Coverage Amount	Interface	Description
◕	Command-line interface	A simple scriptable interface that provides commands for accessing the Device Management command-line interface. For details, see "CLI Examples: Device Management" on page 713.
●	SAS Environment Manager	A graphical enterprise web application used to access Mobile Devices. For details, see "Mobile: How To" on page 270.

Mobile: Software Development Kits

The Software Development Kits (SDKs), SAS SDK for Android or SAS SDK for iOS, enable your mobile apps to include SAS Visual Analytics content. You can preconfigure, customize, and manage the app experience by doing the following:

- Creating custom mobile apps for viewing and interacting with SAS Visual Analytics report content.
- Substituting your organization's name and branding in the SAS Visual Analytics App.
- Displaying SAS Visual Analytics reports in a custom-designed app.
- Integrating the mobile app with your mobile device management (MDM) service.

Your customized apps can connect to SAS Viya servers and can be managed by your organization's SAS administrators. The SAS SDK is free and available for iOS and Android operating systems. It can be downloaded from https://developer.sas.com.

18

Auditing

Auditing: Overview

An audit record is generated whenever these types of events occur:

- an action is performed on a resource (such as a folder or a job). Actions include access to the resource and any changes made to the resource (such as updating, creation, or deletion).
- a security-related action occurred, such as logging on to an application or changing an authorization rule

By default, these actions generate audit records:

- resource read failure
- resource created, updated, or deleted
- security actions (logon attempts, logoff attempts, accessing authorization rules, updating authorization rules)

See “Change Auditing Configuration” on page 287 for information about changing the actions that generate an audit entry.

The audit records are stored in the SAS Infrastructure Data Server and, by default, are retained for seven days. Records older than seven days can be archived to a local storage location. See “Change Auditing Configuration” on page 287 for information about changing the archiving behavior.

All audit records contain this information:

ID
generated identifier of the audit record

Description
description of the action that is recorded (for example, authorization rule access)

Time Stamp
the date and time that the action occurred

Type
the type of action (such as security or resource)

Action
: the action that was performed (such as read, create, or update)

State
: the outcome of the action (success or failure)

User ID
: the user, application, or service that initiated the action

Trace ID
: the trace ID of the record

Properties
: information unique to the type of record

Application
: the application or service that performed the action

In addition, other fields might be included depending on the type of audit record.

In order to access the information in the audit records, commands are provided to list all of the audit records or to list records based on criteria such as date, application name, and user ID. The command also enables you to view details about a specific audit record. See "List Audit Records" on page 285 for more information.

SAS Viya operations infrastructure also includes a predefined task to process the audit records, create a CSV file of the extracted records, and then create a CAS table with the records. Predefined reports enable you to view detailed information about access to reports, applications data, and data plans; about access by user' and about access failures. See "View Audit Record Reports and Tables" on page 284 for more information.

Auditing: How To

View Audit Record Reports and Tables

The User Activity report is available from the SAS Environment Manager Dashboard. You can use it to view graphs and tables of the collected audit record data.

The genAudit task, which runs every two hours (by default), collects information from the audit records that is then used to create the User Activity report. Because the task runs using the credentials of the SAS install user (sas), it collects only those records to which the SAS install user has access. The SAS install user is not a SAS administrator ID.

Note: This report is not available if you are a tenant administrator.

Follow these steps to view the reports.

1. On the SAS Environment Manager Dashboard, select **Show Reports**. A gallery of available reports is displayed at the bottom of the Dashboard.

2. Click in the **User Activity** report and select **Open**. Use the control to navigate through the report gallery to locate the **User Activity** report.

3. The **User Activity** report contains pages that display the audit information based on different criteria, such as user activity, report access, and data table access. Audit records are retained for seven days, so by default, the report displays information from all of the past seven days. Use the slider on each report page to view information only for a selected time range.

   Select the page of the report that contains the type of information that you want to view. These pages are available:

Main
contains thumbnail graphs for the charts **Most active users**, **Activity counts**, **Most active data**, and **User Actions over time**.

Most Active Users
displays the **Most Active Users** and **Activity Over Time** charts, and a table of the audit records ordered by level of user activity. The table does not display audit records from SAS internal users. Select a bar in the **Most Active Users** chart to display the **Activity Over Time** chart for the selected user, and to list the audit records only for the selected user.

Application Usage
displays the **Most used Applications** and **Application Activity** charts, and a table of the audit records orders by level of application activity. Select a bar in the **Most used Applications** chart to display the **Application Activity** chart for the selected application, and to list the audit records only for the selected application.

Report Activity
displays the **Top Report Usage** chart and a table of the audit records for report access. By default, the chart and table display report activity for all users. To view the report usage and audit records only for a specific user, select the user in the **Users** menu.

Data Activity
displays the **Frequently Accessed Tables** chart and a table of the audit records for data table access. By default, the chart and table display data table activity for all users. To view the data table usage and audit records only for a specific user, select the user in the **Users** menu.

Data Plan Activity
displays the **Top Report Usage** chart and a table of the audit records for data plan access. By default, the chart and table display data plan activity for all users. To view the data plan usage and audit records only for a specific user, select the user in the **Users** menu.

Failures
displays the **Failed Requests per Application** and **Failed Activities** charts, and a table of the audit records only for failed requests. By default, the **Failed Activities** chart and the audit records table display failures for all applications. To view the **Failed Activities** chart and audit records for a specific application, select the application’s bar in the **Failed Requests per Application** chart.

Details
displays a table of audit records. By default, the table displays all audit records. To filter the table, use the menus at the top of the table to display only those records matching your selected criteria. You can filter by user, application, action, and state, and multiple criteria are allowed

Note: If the User Activity report is blank or displays the message `Cannot find the requested data source`, you must verify that the command-line interface (CLI) was deployed properly in your SAS Viya environment. See “Edit the Inventory File” in *SAS Viya for Linux: Deployment Guide* for more information.

List Audit Records

Use the command `sas-admin audit list` to list all of the audit records that have been collected. Because the list of records that are returned can be long, you can use these options to manage the records that are returned and more easily locate the records that you want to see:

`sas-admin audit list --limit "number_of_records"`
returns only the specified number of audit records. The default value is 50.

`sas-admin audit list --action action_name`
returns only audit records that contain the specified action

`sas-admin audit list --after YYYY-MM-DDTHH:MM:SS.ssssssZhh:mm`
returns only audit records that occur after the specified date and time

`sas-admin audit list --application application_name`
: returns only audit records that contain the specified application name

`sas-admin audit list --application—contains application_string`
: returns only audit records whose application name contains the string *application_string*

`sas-admin audit list --before YYYY-MM-DDTHH:MM:SS.ssssssZhh:mm`
: returns only audit records that occur before the specified date and time

`sas-admin audit list --description description`
: returns only audit records that contain the specified description

`sas-admin audit list --description—contains description_string`
: returns only audit records whose description contains the string *description_string*

`sas-admin audit list --remote—address address`
: returns only audit records that contain the specified remote address

`sas-admin audit list --remote—address—contains address_string`
: returns only audit records whose remote address contains the string *address-string*

`sas-admin audit list --state state`
: returns only audit records that contain the specified state

`sas-admin audit list --type type`
: returns only audit records that contain the specified type

`sas-admin audit list --user—id user_ID`
: returns only audit records that contain the specified user ID

`sas-admin audit list --user—id—contains user_ID_string`
: returns only audit records whose user ID contains the string *user_ID_string*

`sas-admin audit list --user—id—starts—with user_ID_string`
: returns only audit records whose user ID starts with the string *user_ID_string*

View a Detailed Audit Record

Use the `sas-admin audit show-info --id` command to display detailed information about a single audit record. The information returned look like this

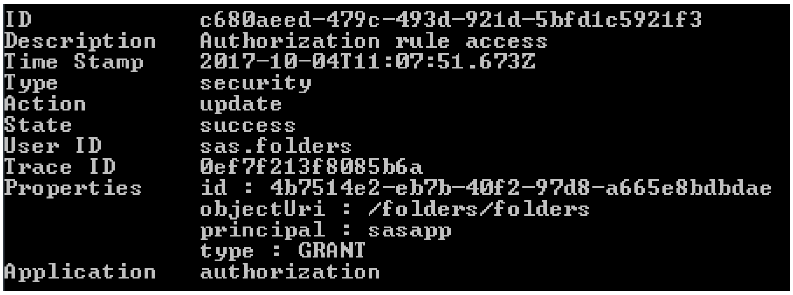

```
ID            c680aeed-479c-493d-921d-5bfd1c5921f3
Description   Authorization rule access
Time Stamp    2017-10-04T11:07:51.673Z
Type          security
Action        update
State         success
User ID       sas.folders
Trace ID      0ef7f213f8085b6a
Properties    id : 4b7514e2-eb7b-40f2-97d8-a665e8bdbdae
              objectUri : /folders/folders
              principal : sasapp
              type : GRANT
Application   authorization
```

View a File of Audit Records

The genAudit task is included in the default task list for the SAS Viya operations infrastructure agent. The task runs automatically every two hours and performs these functions:

- extract the audit records for reports, data plans, CAS management, and CAS access management
- write the extracted audit records to a CSV file in a cache location
- remove audit records in the CSV file from the eighth day of collection

- use the CSV file to create a table in the SystemData caslib called AUDIT

You can use the extracted audit data in the AUDIT table to perform analysis or create reports.

Reset Audit Record Extraction

If the data created by the audit record extraction process becomes corrupted or incorrect, you can reset the extraction process. This action does not alter or remove any of the original audit records. It deletes only the data in the CSV file that is extracted by the genAudit task.

This is an example of a scenario where you should reset the process. The CSV file is designed to hold seven days of audit records, so one step in the process is to remove records only from the eighth day of collection. It does not remove records that are older than the eighth day. If something prevents the genAudit task from running on a particular day, the eighth-day records are not removed, and they remain in the CSV file from that point forward.

You should reset the extraction process if any of these go down:

- sas-ops-agentsrv
- casManagement service
- audit service

To reset the record extraction process, delete all of the files in the directory **`/opt/sas/viya/config/var/cache/auditcli`** on the Operations host (as specified in the Ansible inventory.ini file). The genAudit task creates new extracted audit data when the task runs again after two hours.

Change Auditing Configuration

Follow these steps to edit configuration properties that specify aspects of the audit process:

1. In SAS Environment Manager, select **Configuration**.
2. In the **View** field, select **Definitions**.
3. To change the configuration for how audit records are archived, select **sas.audit.archive** in the definition list.
4. Click **Edit**. You can modify these properties:

   batchSize
   specifies the number of audit records archived at a time.

   enabled
   specifies whether to archive audit records. If enabled, records older than the specified retention period are removed from the archive table. If not enabled, the records are not archived and remain in the table.

   localRetention
   specifies the number of days that records are retained before they are archived.

   scanSchedule
   specifies the time at which the archive process starts (the default value is 000**?, which specifies midnight each day).

   storage.local.destination
   specifies the location where archived records are stored if the **storageType** property is set to **local**.

   storageType
   specifies whether records that are removed from the table are archived to a file (specify a value of **local** and specify a location in the **storage.local.destination** property) or discarded (specify a value of **none**).

5. Select **Save** to save your changes.

6 To change the configuration for the actions that generate an audit record, select **sas.audit.record** in the definition list.

7 Click **Edit**. You can modify these properties:

application

specifies recording of entries from a specified application or service. Select **+ Add property** to add a property for an application. In the Add Property window, specify the property in the **Name** field using the format `sas.audit.record.application.application.enabled`. Specify the value for the property in the **Value** field.

For example, specify `sas.audit.record.application.identities.enabled` in the **Name** field and `false` in the **Value** field to disable recording of entries from the Identities service.

type

specifies recording of entries for a specified action or audit type. Select **+ Add property** to add a property for an application. In the Add Property window, specify the property in the **Name** field.

For actions, use the format `sas.audit.record.type.resource.action.action_type.enabled`. Specify the value for the property in the **Value** field. For example, specify `sas.audit.record.type.resource.action.read.enabled` in the **Name** field and `false` in the **Value** field to disable recording of all read records.

For audit types, use the format `sas.audit.record.type.audit_type.enabled`. Specify the value for the property in the **Value** field. For example, specify `sas.audit.record.type.resource.enabled` in the **Name** field and `false` in the **Value** field to disable recording of all resource records.

8 Select **Save** to save your changes.

See “Configuration Properties: How to Configure Services” on page 74 for more information.

19

Authentication

Authentication: Overview

Authentication is the process of verifying the identity of a user that is attempting to log on to or access software.

In SAS Viya, authentication options vary, based on which interface and operating system are being used in your environment:

Table A.1 Authentication Options

Type of Deployment	Operating System	Authentication Mechanism
full deployment	Linux	The pluggable authentication module (PAM) validates the user’s credentials when accessing SAS Studio 4.x and CAS Server Monitor. Batch jobs submit credentials that require validation. Users can be authenticated through SAS Logon Manager, using an LDAP provider, Kerberos, Security Assertion Markup Language (SAML), or OAuth and OpenID Connect.
	Windows	Host authentication validates the user’s credentials when accessing SAS Studio 4.x and CAS Server Monitor. Batch jobs submit credentials that require validation. Kerberos is the only supported authentication mechanism for SAS Viya visual interfaces and configuration of the middle tier environment.
programming-only deployment	Linux	The only supported authentication mechanism is PAM.
	Windows	The only supported mechanism is Windows host authentication.

Authentication: How To

Authentication Mechanisms

Overview

Authentication mechanisms integrate SAS into your computing environment. External mechanisms include direct LDAP authentication (which is referred to as LDAP in this documentation), host authentication, Kerberos, Security Assertion Markup Language (SAML), and OAuth 2.0 with OpenID Connect. Pluggable authentication modules (PAM) extend UNIX host authentication.

The following sections are listed alphabetically. Configure the authentication mechanism that is appropriate for your environment. For more information, see “Authentication Mechanisms” on page 313.

Note: On Windows deployments, Windows host authentication validates the user's credentials when accessing SAS Studio 4.x and CAS Server Monitor, and for batch jobs. For SAS Viya visual interfaces and configuration of the middle-tier environment, Kerberos is the only supported authentication mechanism.

Configure Kerberos

To configure Kerberos, you must perform prerequisites to verify that certain conditions are met. Then, you must set up SAS Logon Manager, SAS Cloud Analytic Services, and SAS Launcher Server. Lastly, you must configure your web browser for Kerberos.

Verify Kerberos Prerequisites

Before configuring Kerberos, make sure that the following exists:

Note: These prerequisite components are usually configured by the Active Directory administrator.

1. A service account exists in Active Directory.

2. A service principal name (SPN) is mapped to the service account.

   a. Verify that there is a mapping already configured.

   ```
   setspn -F -Q HTTP/hostname.example.com
   ```

   Output A.1 *Sample SPN Query*

```
CN=user-logon-name,OU=Service Accounts,OU=Domain Controllers,OU=Servers,DC=EXAMPLE,DC=com
        HTTP/<hostname.example.com
        HTTP/HOSTNAME

Existing SPN found!
```

   If an SPN is not found, then contact your information technology support group for assistance with registering the machine.

   b. Verify that the service is linked to the service account.

   ```
   setspn -L user-logon-name
   ```

   Output A.2 *Sample Account Query*

```
Registered ServicePrincipalNames for CN=user-logon-name,OU=Service Accounts,OU=Servers,DC=EXAMPLE,DC=com:
        HTTP/<hostname>@<example>.com
        HTTP/<hostname>
```

   The value for **`user-logon-name`** is the same one identified in the common name (CN) from the previous command output, or as the sAMAccountName on the service account in Active Directory.

3. Verify that a keytab file has been generated by issuing one of the following commands:

   On Linux:

   ```
   ktutil
   rkt path-to-keytab-file.keytab
   list -e
   ```

   The following is sample output. Your keytab file will be different.

   Output A.3 *Sample Linux Output*

```

slot KVNO Principal
---- ---- ---------------------------------------------------------------------
   1    3    HTTP/<hostname>@<example>.com (arcfour-hmac)
```

   On Windows:

   ```
   ktab.exe -l -k FILE:path-to-keytab-file.keytab
   ```

   The following is sample output. Your keytab file will be different.

Output A.4 *Sample Windows Output*

```
Keytab name: <filename>.keytab
KVNO    Principal
-----   -------------------------------------
1       HTTP/<hostname>@<example>.com
```

For more information about the **ktutil** and **ktab**commands, see the vendor documentation.

4 On Windows, verify the following items:

a A CAS SPN exists following this convention: `sascas/hostname`. The host name specifies the fully qualified domain name of the machine on which CAS is running. The SPN must be registered on the service account that is running the CAS server. This service account must be trusted for delegation and can be the same service account that is used in step 2 above.

b The machine object must be trusted for delegation. SAS Launcher Server runs under the local system account on the machine it is deployed on and registers its own SPN. The server on which it is running must be marked in Active Directory as trusted for delegation. If the machine is not marked as trusted, it cannot use the user's Kerberos ticket to access remote file systems, nor can it launch CAS sessions under user identity.

For details about enabling a server to be trusted for delegation, see Enable computer and user accounts to be trusted for delegation.

Configure Kerberos for SAS Logon Manager

Note: This information does not apply to a programming-only deployment.

1 If you have not already done so, from SAS Environment Manager, add your user ID or an Active Directory group that contains the environment administrators, as a member of the SAS Administrators group. Then, log off from SAS Environment Manager. For more information, see “Add or Remove Custom Group Members” on page 351.

CAUTION! You must specify your personal user ID. Your user ID must be in your specified LDAP provider. It must match the user ID that you use to log on to your system. Also, your user ID must be added to the SAS Administrators group because once Kerberos is configured, you can no longer sign in as the sasboot user.

2 Make sure that the keytab file is saved to a directory that is accessible to the user account that runs the SAS services.

3 Verify that the SPN is mapped to the principal name.

```
setspn -F -Q HTTP/hostname.example.com
```

4 Configure the Kerberos authentication properties.

a Log on to SAS Environment Manager, using your user ID or the ID of a user who is a member of the SAS Administrators group.

b Navigate to the SAS Logon Manager configuration definitions. For more information, see “Edit Authentication Configuration Instances” on page 304.

c In the **Definitions** list, select **sas.logon.kerberos**.

d In the top right corner of the window, click **New Configuration**.

e In the New sas.logon.kerberos Configuration dialog box, enter the values for the following fields, based on your environment.

Table A.2 *Configuration Fields and Values*

Field	Value
debug	On
holdOnToGSSContext:	On
keyTabLocation	file:///*path-to-keytab-file* **Note:** You must use forward slashes, even on Windows systems (for example, file:///c:/*path-to-keytab-file*).

servicePrincipal	*principal-name-from-keytab*
	On Linux, issue the following command:
	`ktab -l -k FILE:path-to-keytab-file.keytab`
	On Windows, issue the following command from the directory where Java is installed on your machine:
	`ktab.exe -l -k FILE:path-to-keytab-file.keytab`

stripRealmForGss	On **Note:** When enabled, this option strips the realm from the user name.

Note: Contact your administrator for the keytab location and the host name of the service principal.

f Click **Save**.

5 Add Kerberos to the active profile.

a In the navigation pane, switch to the **All services** list and select **SAS Logon Manager**.

b In the **spring** instance, click ✎.

c In the Edit spring Configuration dialog box, add **`kerberos`** to the **profiles.active** field.

The following value should be specified for the **profiles.active** field:

```
ldap,postgresql,kerberos
```

d Click **Save**.

6 Restart the SAS Logon Manager service.

On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-saslogon-default stop
sudo service sas-viya-saslogon-default start
```

On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-saslogon-default
```

On Windows, in Windows Services Manager, right-click the **SAS Logon Manager service** and select **Restart**.

Note: It might take several minutes to restart SAS Logon Manager.

Note: Once Kerberos is enabled on Windows, a browser running on the same machine where the services are deployed cannot connect to SAS Viya visual interfaces.

Configure Kerberos for SAS Cloud Analytic Services

Note: This information does not apply to Windows systems.

1 Create a keytab file for CAS to use.

The file is used to validate incoming user Kerberos tickets and generate server identity Kerberos tickets for access to Kerberized resources, such as Hadoop. By default, the keytab file should reside in the **`/etc/sascas.keytab`** file and be readable by CAS. If you save the file in a different directory or use a different filename, set the **`KRB5_KTNAME`** environment variable (for example, `env.KRB5_KTNAME = 'fully-qualified-filename'`). For more information, see “ CAS Environment Variables” on page 535.

2 If you changed the default principal name, set the **`CAS_SERVER_PRINCIPAL`** environment variable (for example, `env.CAS_SERVER_PRINCIPAL = 'principal-name'`).

By default, CAS uses the following Kerberos principal name: sascas/*fully-qualified-DNSname*. CAS searches for this principal in the keytab file.

3 Add the 'kerb' option to the **`cas.PROVLIST`** configuration file option (for example, `cas.PROVLIST = 'oauth.ext.kerb'`).

For more information about the configuration file option, see “Configuration File Options” on page 512.

4 Enable the Kerberos option for authentication to SAS Compute Server.

a Log on to SAS Environment Manager, using your user ID or the ID of a user who is a member of the SAS Administrators group.

b Navigate to the Launcher service configuration instance. For more information, see “Edit Configuration Instances” on page 75.

c In the **sas.compute** instance, click ✎.

d In the Edit sas.compute Configuration dialog box, select the **kerberos.enabled** option.

e Click **Save**.

f Restart the CAS controller.

```
sas-viya-cascontroller-default restart
```

Configure Kerberos for SAS Launcher Server

Note: This information does not apply to Windows systems.

1 Source the consul.conf file to add configuration values that use the SAS Security framework certificate truststore.

```
sudo su sas
. /opt/sas/viya/config/consul.conf
```

Note: To source the consul.conf file, you must use the "sas" user account. For more information, see “User Accounts (Reference)” in *SAS Viya for Linux: Deployment Guide*.

2 Log on to SAS Configuration Server as a user with root or sudo privileges and run the *sas-bootstrap-config* script for the SAS Launcher Server keytab.

```
/opt/sas/viya/home/bin/sas-bootstrap-config --token-file
/opt/sas/viya/config/etc/SASSecurityCertificateFramework/tokens/consul/default/client.token
```

```
kv write --force --key config/launcher-server/global/keytab --value path-to-keytab-file
```

Note: The previous command must be on one line. It is shown on more than one line for display purposes only.

3 Restart SAS Launcher Server.

   On Red Hat Enterprise Linux 6.x, run the following commands:

   ```
   sudo service sas-viya-runlauncher-default stop
   sudo service sas-viya-runlauncher-default start
   ```

   On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

   ```
   sudo systemctl restart sas-viya-runlauncher-default
   ```

4 Enable the Kerberos option for authentication to SAS Compute Server.

   a Log on to SAS Environment Manager, using your user ID or the ID of a user who is a member of the SAS Administrators group.

   b Navigate to the Launcher service configuration instance. For more information, see “Edit Configuration Instances” on page 75.

   c In the **sas.compute** instance, click ✎.

   d In the Edit sas.compute Configuration dialog box, select the **kerberos.enabled** option.

   e Click **Save**.

Validate Kerberos Configuration

All users are authenticated using OAuth 2.0 and OpenID Connect. Complete the following steps to verify that Kerberos is configured correctly:

1 Check the CAS log to see how the non-delegated user authenticated to CAS by running the following command:

   ```
   cat /var/log/sas/viya/cas/default/* |grep non_delegated_user|grep authenticated|tail -1
   ```

   On Windows, navigate to the **C:\ProgramData\SAS\Viya\var\log\cas\default** directory and view the contents of the cas_*date_hostname* file.

2 Look for output similar to the following:

```
2018-06-12T11:03:35,376 INFO  [00002846] <non_delegated_user> local MAIN NoUser  [tkidentgss.c:741] - User
<non_delegated_user>@<domain_name> successfully authenticated using the OAuth authentication provider.
```

On Linux systems, delegation occurs only for users who are in the CASHostAccountRequired custom group. On Windows systems, users are automatically delegated. Users with delegated Kerberos credentials are also authenticated with the Kerberos authentication provider to delegate their identity to CAS. To validate Kerberos for the delegated user, complete the following steps:

1 Check the CAS log to see how the delegated user authenticated to CAS.

   On Linux, run the following command:

   ```
   cat /var/log/sas/viya/cas/default/* |grep delegated_user|grep kerberos|tail -1
   ```

   On Windows, navigate to the **C:\ProgramData\SAS\Viya\var\log\cas\default** directory and view the contents of the cas_*date_hostname* file.

2 Look for output similar to the following:

```
2018-06-12T11:03:35,376 INFO  [00002846] <delegated_user> local MAIN NoUser  [tkident.c:741] - User
<delegated_user> successfully authenticated using the Kerberos authentication provider.
```

Configure Microsoft Internet Explorer and Google Chrome to Use Kerberos

Configure Security Settings

1 In the Windows Control Panel, open Internet Options.

2 In the Internet Properties dialog box, select the **Security** tab.

3 Select **Local intranet**, and then click **Sites**.

4 In the Local intranet dialog box, configure the intranet domain settings.

a Verify that the check boxes for the following items are selected:

- **Include all local (Intranet) sites not listed in other zones**
- **Include all sites that bypass the proxy server**

b Click **Advanced** and add your domain name to the **Websites** list to ensure that Internet Explorer recognizes any site with your domain name as the intranet.

c Click **Close**, and then click **OK**.

5 Configure intranet authentication.

a In the **Security level for this zone** area, click **Custom level**.

b In the Security Settings - Local Intranet Zone dialog box, scroll to the **User Authentication** section, select **Automatic Logon only in Intranet Zone**, and click **OK**.

Configure Connection Settings

If your site uses a proxy server, follow these steps:

1 In the Internet Properties dialog box, select the **Connections** tab.

2 Click **LAN settings**.

3 In the Local Area Network (LAN) Settings dialog box, verify that the proxy server address and port number are correct.

4 Click **Advanced**.

5 In the Proxy Settings dialog box, verify that the correct domain names are entered in the **Exceptions** field. Then, click **OK**.

6 Click **OK**.

Configure Integrated Windows Authentication

1 In the Internet Properties dialog box, select the **Advanced** tab.

2 Scroll to the **Security** section, and verify that **Enable Integrated Windows Authentication** is selected.

3 Click **OK** and restart your computer to activate the changes.

Configure User Delegation for Microsoft Internet Explorer

Complete the following steps after configuring Integrated Windows Authentication:

1. In the Windows Control Panel, open Internet Options.
2. In the Internet Properties dialog box, select the **Security** tab.
3. Select **Trusted Sites**, and then click **Sites**.
4. In the Trusted sites dialog box, enter the middle-tier host name in the **Add this website to the zone** field and click **Add**.
5. Click **Close**, and then click **OK**.

   Note:

   For Internet Explorer to pass a forwardable ticket to the SAS Viya machine, the service account in Active Directory holding the SPNs must be trusted for delegation.

Configure User Delegation for Google Chrome

By default, Chrome disables the delegation of Kerberos credentials. The Windows registry must be updated. Microsoft recommends performing a system backup before editing the registry. Complete the following steps to enable Kerberos delegation after configuring Integrated Windows Authentication:

1. Open the Windows registry editor.
2. Add the following REG_SZ keys:

   \HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Google\Chrome\AuthServerWhitelist
   Specifies which servers should be whitelisted for integrated authentication. Set the value to the SAS Web Server host name: *hostname*.*example*.com.

   \HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Google\Chrome\AuthNegotiateDelegateWhitelist
   Specifies which servers Chrome can delegate to. Set the value to the SAS Web Server host name: *hostname*.*example*.com.

   Note: You might also need to add Google and Chrome under Policies.

Configure Mozilla Firefox to Use Kerberos

Configure Kerberos

1. From a browser window, navigate to `about:config`.
2. Click **I accept the risk!** to accept the security warning.
3. In the **Search** field, enter `network.negotiate`.
4. Double-click the **network.negotiate-auth.trusted-uris** Preference Name, enter `http://`*`hostname.example`*`.com,` in the **Enter string value** field, and then click **OK**.

   Note: The values in the **Enter string value** field are comma-separated.

Configure User Delegation

1. From a browser window, navigate to `about:config`.
2. Click **I accept the risk!** to accept the security warning.
3. In the **Search** field, enter `network.negotiate`.
4. Double-click the **network.negotiate-auth.delegation-uris** Preference Name, enter `http://`*`hostname.example`*`.com` in the **Enter string value** field, and then click **OK**.

Configure OpenID Connect

OpenID Connect uses a reverse proxy server as the single sign-on entry point for initial user authentication. To configure the OpenID Connect, complete the following:

Note: This information does not apply to a programming-only deployment. It also does not apply to Windows systems.

1 Log on to SAS Environment Manager.

2 Navigate to the SAS Logon Manager configuration definitions. For more information, see "Edit Authentication Configuration Instances" on page 304.

3 In the **Definitions** list, select **sas.logon.oauth.providers.external_oauth**.

4 In the top right corner of the window, click **New Configuration**.

5 In the New sas.logon.oauth.providers.external_oauth Configuration dialog box, enter values for the required fields, based on your environment. The following table provides guidance about the information needed for the listed fields:

Table A.3 OAuth Configuration Fields and Descriptions

Configuration Fields	Descriptions
attributeMapping.user_name	The attribute claim to use as the user name. By default, the value is *user_name*.
authUrl	The URL to the authorization endpoint.
emailDomain	The email address domain for users authenticating with the provider.
issuer	The principal that issued the token, specified as a case-sensitive string or URI. This value must match the issue claim in the token.
linkText	The text that should be displayed on the sign-in page for the provider.
relyingPartyId	The client ID that is registered with the provider.
relyingPartySecret	The secret that is registered with the provider for the client ID.
scopes	The comma-delimited list of scopes for the authorization request. The list should contain `openid`. **Note:** SAS Viya does not process any additional scopes that are returned in the token.
tokenUrl	The URL to the token endpoint.

Configuration Fields	Descriptions
type	The protocol type. By default, the value is **oidc1.0**. Note: SAS Viya requires an id_token in the authorization response from the provider. However, some providers return an id_token when the scope in the authorization request is *openid* and respose_type=*token*. For those providers, use type **`oauth2.0`**.

6 Click **Save**.

7 Restart the SAS Logon Manager Service.

On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-saslogon-default stop
sudo service sas-viya-saslogon-default start
```

On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-saslogon-default
```

On Windows, in Windows Services Manager, right-click the **SAS Logon Manager service** and select **Restart**.

Note: It might take several minutes to restart SAS Logon Manager.

See Also

"Authentication: OpenID Connect Scenario" on page 331

Configure PAM

Default pluggable authentication module (PAM) configuration files are installed for both the CAS server and SAS Studio.

Note: This information does not apply to Windows systems.

1 As a user with root authority, edit the ***`SAS-Viya-configuration-directory`*****`/etc/pam.d/`*****_`service`_** file. For the CAS server, *service* is *cas*. For SAS Studio, *service* is *sasauth*.

The following information is displayed for the CAS server:

```
$ vi /etc/pam.d/cas
#%PAM-1.0
auth        include      password-auth
account     include      password-auth
password    include      password-auth
session     include      password-auth
```

The following information is displayed for SAS Studio:

```
$ vi /etc/pam.d/sasauth
#%PAM-1.0
auth        include      password-auth
account     include      password-auth
```

2 Make any modifications to the file that are necessary for your environment.

3 Save the file and exit.

Configure SAML

Overview

Configuration for the Security Assertion Markup Language (SAML) typically follows this pattern:

Note: This information does not apply to a programming-only deployment. It also does not apply to Windows systems.

1 "Configure SAS Viya as a SAML Service Provider" on page 300

2 "Configure the SAML Identity Provider – Relying Party Configuration" on page 301

3 "Configure SAS Viya with Information about the SAML Identity Provider" on page 301

Note: By default, SAS Viya allows only same-origin requests. Authentication requests from the SAML identity provider might be seen as cross-origin. Therefore, the origin of the SAML provider might need to be added. For details, see "Configure Cross-Origin Resource Sharing" on page 305.

Configure SAS Viya as a SAML Service Provider

1 Log on to SAS Environment Manager.

2 Navigate to the SAS Logon Manager configuration definitions. For more information, see "Edit Authentication Configuration Instances" on page 304.

3 In the **Definitions** list, select **sas.logon.saml**.

Note: If you change any of the sas.logon.saml properties, the new metadata must be provided to the Relying Party in ADFS. If it is not, the SAML connections might fail.

4 In the top right corner of the window, click **New Configuration**.

5 In the New sas.logon.saml Configuration dialog box, enter values for the required fields, based on your environment. The following table provides guidance on what information needs to be provided for the listed fields:

Table A.4 SAML Configuration Fields and Descriptions

Field	Description
entityBaseURL	The external URL for the SAS Logon web application in SAS Viya (for example, https://*hostname.example.com*/SASLogon).
entityID	The unique ID that represents the service provider that is included in protocol messages between relying parties. Change from the default value that is pre-populated.
serviceProviderCertificate	Paste a copy of the PEM-encoded (base64) certificate, which is used by the service provider.
serviceProviderKey	Paste a copy of the PEM-encoded (base64) key, which is used by the service provider.
serviceProviderKeyPassword	Provide the password for the service provider, or leave blank if there is no password.

Field	Description
setProxyParams	**Note:** This field should not be modified. The value should remain **`false`**. Specifies whether to allow the base URL to reside behind an HTTP proxy.
signMetaData	Specifies whether the local service provider should sign the metadata.
signRequest	Specifies whether the local service provider should sign the SAML requests.
wantAssertionSigned	Specifies whether the assertions should be signed.

6 Click **Save**.

7 Restart the SAS Logon Manager Service.

On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-saslogon-default stop
sudo service sas-viya-saslogon-default start
```

On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-saslogon-default
```

On Windows, in Windows Services Manager, right-click the **SAS Logon Manager Service** and select **Restart**.

Note: It might take several minutes to restart SAS Logon Manager.

Configure the SAML Identity Provider – Relying Party Configuration

You can either configure the relying party trust or supply the required information to your information technology support group, in order for them to add the relying party trust. Here is an overview of the steps to perform, if you configure the relying party trust. The steps might vary, depending on which tool you use for configuration.

1 If the identity provider requires it, configure Transport Layer Security (TLS), if it has not already been configured. For more information, see "Update Apache HTTP Server TLS Certificates and Cryptography " in *Encryption in SAS Viya: Data in Motion*.

2 Download the application metadata.xml file, which contains information about the service provider, or provide the https://*hostname*/SASlogon/saml/metadata link to your information technology support group.

3 Request that your information technology support group configure a relying party in the identity provider.

Configure SAS Viya with Information about the SAML Identity Provider

1 Complete the following steps in SAS Environment Manager:

a In the **Definitions** list, select **sas.logon.saml.providers.external_saml**.

b In the top right corner of the window, click **New Configuration**.

c In the New sas.logon.saml.providers.external_saml Configuration dialog box, enter values for the required fields, based on your environment. The following table provides guidance on what information needs to be provided for the listed fields:

Table A.5 Configuration Fields and Descriptions

Field	Description
idpMetadata	The URL to the location of the identity provider metadata (for example, https://*hostname.example.com*/*filename*.xml). This information is provided by your information technology support group. Use the URL to configure Active Directory Federation Services (ADFS) or another service, along with endpoints. Note: In this document, ADFS is used for configuration.
metadataTrustCheck	Specify whether to trust the identity provider certificate.
nameID	The field is populated with a default value. Verify with your information technology support group that the value is correct.
showSamlLoginLink	Determines whether a link should be displayed on the logon page for this identity provider.

d Click **Save**.

2 Edit the ***SAS-Viya-configuration-directory*/etc/sysconfig/sas-javaesntl/sas-java-services** file, and add the following line where the truststore options are set:

```
[[ -f $truststore ]] && export
java_global_option_truststore_password="-Djavax.net.ssl.trustStorePassword=changeit"
```

3 Restart the SAS Logon Manager Service.

On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-saslogon-default stop
sudo service sas-viya-saslogon-default start
```

On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-saslogon-default
```

On Windows, in Windows Services Manager, right-click the **SAS Logon Manager Service** and select **Restart**.

Note: It might take several minutes to restart SAS Logon Manager.

Configure Authentication Options with SAS 9.4

Configure the SAS 9.4 Deployment

1 Log on to SAS Management Console and navigate to **Plug-ins** ⇨ **Application Management** ⇨ **Configuration Manager**.

2 Right-click **SAS Application Infrastructure** and select **Properties**.

3 Click **Advanced**, and then set the following property value:

ServiceUrl.Allowed
Specifies the address to where tickets should be sent on SAS Viya. The format of the address should be similar to the following: http://*hostname*/SASLogon/**.

Note: For SAS deployments prior to SAS 9.4M3, the *ServiceUrl.Allowed* property is not required.

4 Click **OK**.

5 Restart all instances of SASServer1 to pick-up the new property.

Configure the SAS Viya Deployment

1 Log on to SAS Environment Manager.

2 Navigate to the SAS Logon Manager configuration definitions. For more information, see "Edit Authentication Configuration Instances" on page 304.

3 In the **Definitions** list, select **sas.logon.sas9**.

4 In the top right corner of the window, click **New Configuration**.

5 In the New sas.logon.sas9 Configuration dialog box, enter values for the required fields, based on your environment. The following table provides guidance on what information needs to be provided for the listed fields:

Table A.6 SAS 9.4 Configuration Fields and Descriptions

Field	Description
autoLink	Specifies whether to automatically open the link to SAS 9.4 when the logon page is displayed. Note: If the *autoLink* property is enabled, then the SAS Logon Manager in SAS Viya form is not displayed. End users are automatically redirected to SAS Logon Manager in SAS 9.4 to authenticate. End users cannot use the LDAP provider.
enabled	Specifies whether to enable sign-ins using SAS 9.4 credentials.
linkText	Specifies the hyperlink to display on the sign-in page. Note: By default, the end user is presented with a link at the bottom of the standard SAS Logon Manager in SAS Viya form. The text of the link is controlled by the *linkText* property. This default behavior means that end users can choose to either use SAS 9.4 to authenticate or use the LDAP provider.
sas9LogonUrl	Specifies the URL of the SAS Logon Manager in SAS 9.4 (for example, https://*SAS9_hostname*/SASLogon).
showLinkText	Specifies whether to display the link text on the sign-in page.
single.signOn.enabled	Specifies whether to redirect to SAS 9.4 for single sign-on.
single.signOut.enabled	Specifies whether the local sign-out should also sign the user out of SAS 9.4.
viyaLogonUrl	Specifies the URL of the SAS Logon Manager in SAS Viya (for example, https://*SASViya_hostname*/SASLogon).

6 Click **Save**.

7 Restart the SAS Logon Manager Service.

On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-saslogon-default stop
sudo service sas-viya-saslogon-default start
```

On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-saslogon-default
```

On Windows, in Windows Services Manager, right-click the **SAS Logon Manager Service** and select **Restart**.

Note: It might take several minutes to restart SAS Logon Manager.

Session Management

Overview

The following sections provide information about customizing SAS Logon Manager and the user's session experience.

Note: This information does not apply to a programming-only deployment.

Edit Authentication Configuration Instances

1 From SAS Environment Manager, click ☰ and select **Manage Environment**.

2 In the navigation bar, click 🔧 Configuration.

3 In the top left corner of the window, select **Definitions** from the drop-down box.

See Also

- *SAS Viya Administration: Configuration Properties*

Customize Sign-in, Sign-out, and Session Time-out Content

You can configure customized content that is displayed when users of SAS web applications sign in, sign out, or the session reaches the time-out interval. To enable the display of customize content, follow these steps:

1 In the **Definitions** list, select **sas.logon.custom**.

2 In the top right corner of the window, click ✱.

3 In the New sas.logon.custom Configuration dialog box, specify the URI that contains the custom content that you want to display. Here are the available fields:

   - **login**
   - **logout**
   - **timedout**

   For a description of the properties, see "sas.logon.custom" on page 117.

4 Click **Save**.

Customize Concurrent Sign-in Sessions

1 In the **Definitions** list, select **sas.logon.sessions**.

2 In the top right corner of the window, click ✱.

3 In the New sas.logon.sessions Configuration dialog box, you can set the following properties:

maxConcurrentSessions
: Set this property to limit a user to a certain number of concurrent sessions.

rejectNewSessionsIfMaxExceeded
: When sessions are limited, the default behavior is to cause an existing session to expire and grant a new session to the user attempting to authenticate. To override this behavior and prevent a new session from being granted, set this property to *true*.

4 Click **Save**.

Configure the HTTP Session Time-out Interval

1 In the **Definitions** list, select **server**.

2 In the top right corner of the window, click ✱.

3 In the New server Configuration dialog box, complete the following:

a Select **SAS Logon Manager** from the **Services** drop-down list.

b Click +.

c In the **Name** field, specify `session.timeout`.

d In the **Value** field, specify the amount of time a session has to be idle before it times out, in seconds.

e Click **Save**.

4 Click **Save**.

5 Restart all services to reflect the new time-out interval. For more information, see "Start and Stop All Servers and Services" on page 462.

Disable Logins

As a SAS administrator, you can disable logins through operating system firewall rules or using LDAP. This disables new sessions, ends current sessions, and prevents others from using the deployment. For more information, see the appropriate documentation for your operating system.

Additional Authentication Topics

Configure Cross-Origin Resource Sharing

By default, SAS Viya allows only same-origin requests. If cross-origin requests are needed, complete the following steps:

1 In SAS Environment Manager, edit the CORS configuration instance. For details, see "Edit Authentication Configuration Instances" on page 304.

2 Select **sas.commons.web.security.cors**.

3 In the top right corner of the window, click **New Configuration**.

4 In the New sas.commons.web.security.cors Configuration dialog box, specify values that correspond to your environment. For a description of each field, see “sas.commons.web.security.cors” on page 127.

Note: The specified value for the **allowedOrigins** field must be a comma-delimited list of URIs or an asterisk ('*') to accept all origins. Partial wildcards are not supported. For example, **`https://*.example.com`** is not supported.

5 Click **Save**.

6 Restart the SAS Logon Manager Service.

On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-saslogon-default stop
sudo service sas-viya-saslogon-default start
```

On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-saslogon-default
```

On Windows, in Windows Services Manager, right-click the **SAS Logon Manager Service** and select **Restart**.

Note: It might take several minutes to restart SAS Logon Manager.

Obtain an Access Token Using Password Credentials

You can use the following commands to register a client ID and secret. You can also use the commands to obtain a token that can be used to call a SAS Viya API and to access SAS Viya credentials from SAS 9.4.

1 Register a new client ID and secret by completing the following steps:

Note: You must register a client ID once.

a Obtain a token to register a new client ID and secret. For more information, see “Obtain an OAuth Access Token to Register a New Client ID” in *Encryption in SAS Viya: Data in Motion*.

b Use the token to register the new client ID and secret by running the following curl command:

Note: The initial line of the curl command must be entered on one line. It is shown on more than one line for display purposes only.

```
curl -X POST http://localhost/SASLogon/oauth/clients -H "Content-Type: application/json"
-H "Authorization: Bearer token-from-previous-step"
    -d '{
         "client_id": "client-id",
         "client_secret": "client-secret",
         "scope": ["openid", "*"],
         "resource_ids": "none",
         "authorities": ["uaa.none"],
         "authorized_grant_types": ["password"]
         }'
```

Note: The value for the **scope** parameter can be a list of scopes and groups that you **might** request when obtaining a token. You might also specify the wildcard `"*"` to request all scopes always. Ensure that you specify the list correctly. SAS Viya treats group memberships as scopes. Therefore, the list of scopes is the list of group memberships that you might request when obtaining a token. The "openid" is a special scope that represents authentication only and should always be included.

2 A token can be used until it expires. By default, this is 12 hours. To acquire a token, run the following curl command:

```
curl http://localhost/SASLogon/oauth/token
          -H "Accept: application/json"
          -H "Content-Type: application/x-www-form-urlencoded"
          -d "grant_type=password&username=username&password=password"
          -u "client-id:client-secret"
```

Note: The values for *client-id* and *client-secret* should be the same as the values that were specified in Step 1b.

3 Retrieve the access token information from the results of the curl command in Step 2. This access token is used to perform the following tasks:

a Call a SAS Viya API by passing the HTTP Authorization header as a Bearer token: `Authorization: Bearer access_token`.

b Assign the access token to the SAS_VIYA_TOKEN environment variable. Setting this environment variable enables you to access SAS Viya credentials from SAS 9.4. For example, when the AUTHDOMAIN= option is set on a CAS or LIBNAME statement, an attempt is first made to retrieve credentials from the SAS Viya Credentials service before searching the metadata. For more information, SAS_VIYA_TOKEN Environment Variable.

Create an Authinfo File

The authinfo file supplies a user name and password that is sent to CAS for authentication. For information about how to create an authinfo file, see Create an Authinfo File.

Authentication: Concepts

Authentication Architecture

In a full deployment, authentication services are provided by SAS Logon Manager. SAS Logon Manager is based on the Cloud Foundry User Account and Authentication (UAA) server. The security architecture is built around Open Authorization (OAuth) and OpenID Connect. By default, authentication is performed via a Lightweight Directory Access Protocol (LDAP) provider. Authentication support is also available for Kerberos, OAuth 2.0 with OpenID Connect, and Security Assertion Markup Language (SAML).

Note: For Windows deployments, Kerberos is the only supported authentication mechanism for SAS Viya visual interfaces and configuration of the middle tier environment.

In a programming-only and full deployment, host authentication is supported on both Linux and Windows systems. On Linux systems, you can configure the host to use only pluggable authentication modules (PAM).

Authentication and SAS Viya Services

The following table lists the key services that are used in authentication in SAS Viya:

Table A.7 SAS Viya Services

Service Name	Description
SAS Logon Manager	Provides both an end-user interface for authentication and internal authentication to other services. Enables single sign-on within the SAS Viya environment between services. Enables single sign-on to the SAS Viya environment through configuration of third-party software.
Identities service	Provides the user and group information to other services. Reads user and group information from the LDAP provider.
Authorization service	Provides authorization information to other services.
Launcher Service	Provides the connection and authentication to the SAS Launcher Server. Resolves the credentials that are used when authenticating to the SAS Launcher Server.
SAS Cloud Analytic Services	Authenticates end users launching CAS sessions by way of the SAS Cloud Analytic Services controller.
SAS 9.4	Supports several mechanisms for coupling authentication with SAS 9.4.
SAS Studio 4.4	Leverages the SAS Object Spawner to authenticate users accessing SAS Studio 4.4.

In-bound and Out-bound Authentication

In-bound Authentication

In-bound authentication is the authentication of the end user to the environment. In-bound authentication provides an internal OAuth token and group membership information in the OAuth token. If Kerberos authentication is used, a delegated Kerberos credential is also stored.

Out-bound Authentication

Out-bound authentication is the authentication of the SAS process to a downstream process. Out-bound authentication occurs after the end user is initially authenticated to SAS Logon Manager. Out-bound authentication occurs to SAS Cloud Analytic Services, SAS Compute Server (through SAS Launcher Service), and then onto external resources, such as Secured Hadoop environments.

Authentication Options

Authentication for Visual Interfaces

With visual interfaces, users are authenticated through SAS Logon Manager. SAS Logon Manager is a web application that handles all authentication requests for SAS web applications and is accessed via the Apache HTTP Server.

The following figure shows how a user is authenticated on Linux to SAS Logon Manager and the supported authentication mechanisms.

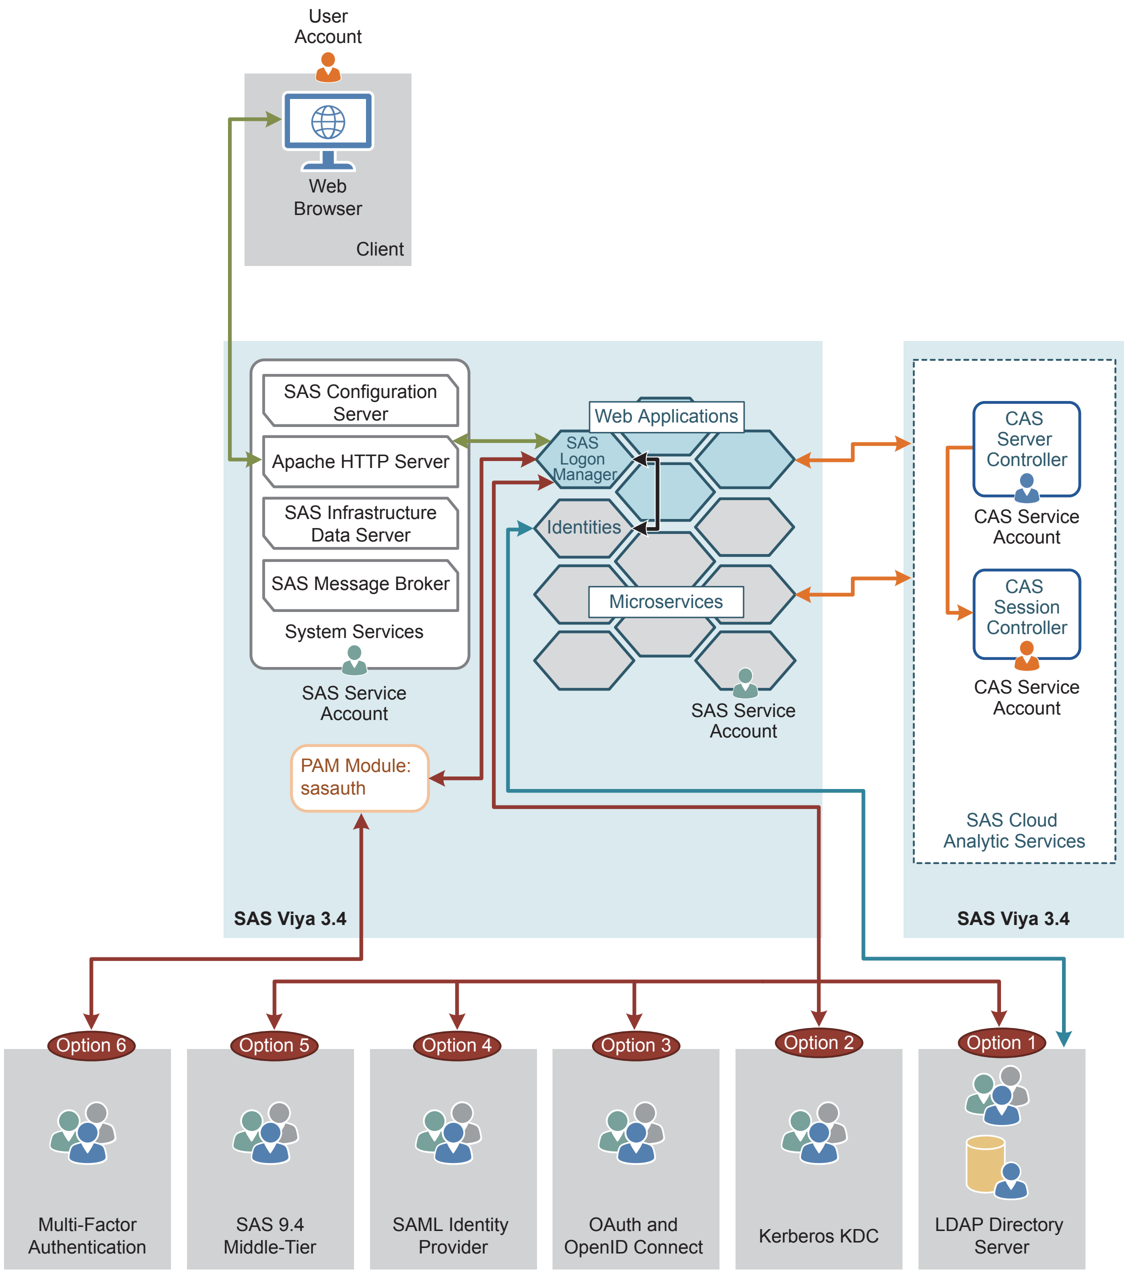

The following protocols are available for you to configure for authentication:

- The first option is a Lightweight Directory Access Protocol (LDAP) provider. This is the default configuration. In this configuration, SAS Logon Manager displays a logon form and submits the entered credentials to LDAP. The identity service verifies users in LDAP. For more information, see "LDAP Authentication" on page 313.
- The second option is Kerberos. In this configuration, SAS Logon Manager uses SPNEGO to authenticate users against the Kerberos Key Distribution Center (KDC). The identity service verifies users in LDAP. For more information, see "Kerberos Authentication" on page 313.

  Note: SAS Cloud Analytic Services sessions run as the end user only when using Kerberos delegation. On Linux systems, the user must be a member of the CASHostAccountRequired custom group. On Windows systems, users are automatically delegated.
- The third option is OAuth 2.0 and Open ID Connect. In this configuration, SAS Logon Manager uses OAuth 2.0 and OpenID Connect to authenticate users. The identity service verifies users in LDAP. For more information, see "OAuth and OpenID Connect Authentication" on page 321.

- The fourth option is Security Assertion Markup Language (SAML). In this configuration, SAS Logon Manager uses a SAML provider to authenticate users. The identity service verifies users in LDAP. For more information, see "SAML Authentication" on page 321.
- The fifth option is SAS 9.4. In this configuration, SAS Logon Manager supports single sign-on and single sign-off with SAS 9.4. The identity service verifies users in LDAP. For more information, see "SAS 9.4 Authentication" on page 322.
- The sixth option is pluggable authentication module (PAM) to support multi-factor authentication. In this configuration, SAS Logon Manager uses the operating system PAM stack. The identity service verifies users in LDAP. For more information, see "PAM Authentication" on page 324.

With all six options, the connection to SAS Cloud Analytic Services (CAS) environment is performed using internal OAuth tokens that are generated by SAS Logon Manager. In most cases, the session that is started by the CAS controller runs on the operating system as the same user who launched the CAS operating system service. This defaults to the cas account.

Authentication for Programming Interfaces

Overview of Programming Interfaces

The following figure shows how a user is authenticated on Linux while using programming interfaces.

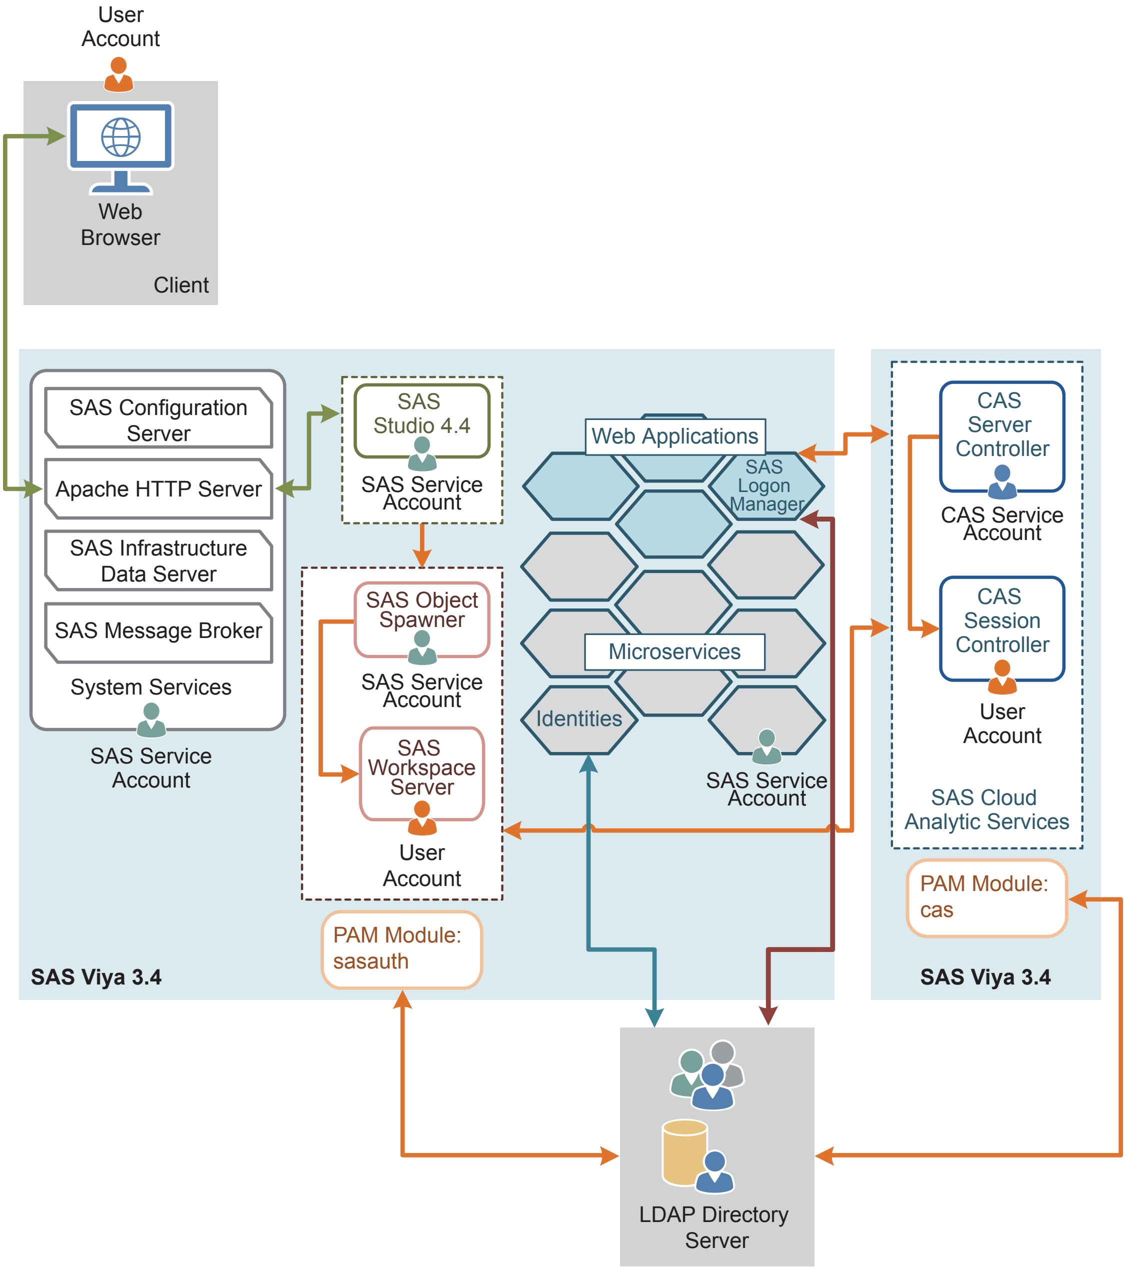

In a deployment with programming interfaces, the user's credentials are entered into SAS Studio via the Apache HTTP Server. Then SAS Object Spawner uses pluggable authentication module (PAM) configuration files on the host to validate the user ID and password. The user ID and password can be a local account on the host or, depending on the PAM configuration, an account in the LDAP provider. Once the user is authenticated, SAS Workspace Server is started. The PAM configuration file for SAS Studio is sasauth and includes the password module.

SAS Workspace Server connects to the CAS environment using the user ID and password that were used to start SAS Workspace Server. However, if the AUTHINFO= option is specified, it is used to find credentials to connect to CAS. For more information about the AUTHINFO= option, see AUTHINFO= SAS system option or AUTHINFO= CAS statement option.

The CAS controller uses its own PAM configuration to validate the user's credentials and launch the session process as the user. The PAM configuration file for CAS is cas and includes the password module.

The CAS controller uses the user ID and password to obtain an internal OAuth token from SAS Logon Manager. This requires the user ID and password to be valid in the LDAP provider that is configured for SAS Logon

Manager. Otherwise, CAS cannot obtain an OAuth token, and the session will fail. Therefore, PAM for SAS Studio (sasauth), PAM for CAS (cas), and SAS Logon Manager should all use the same or equivalent LDAP providers. If these three components are not sending the user ID and password that was entered into SAS Studio to the same provider, errors might be generated when trying to connect.

Programming Interfaces with Symmetric Multiprocessing CAS Server

In a symmetric multi-processing (SMP) environment, a CAS server consists of a controller and runs on a single machine. The following details the authentication process:

1. The end user connects to the SAS Studio 4.4 application and enters their user name and password in the logon form. SAS Studio is proxied by the Apache HTTP Server.
2. SAS Studio 4.4 passes the user name and password to SAS Object Spawner to start the SAS Workspace Server for the end user.
3. SAS Object Spawner uses the PAM configuration that is defined in /etc/pam.d/sasauth to validate the user name and password and launches SAS Workspace Server as the end user.
4. The end user enters code to start their CAS session. SAS Workspace Server passes the user name and password to the CAS controller.
5. The CAS controller connects to SAS Logon Manager to obtain an OAuth token presenting the end user's user name and password.
6. SAS Logon Manager validates the user name and password against the defined LDAP Provider. SAS Logon also connects to the Identities service to obtain group information to include in the OAuth token.
7. The Identities service connects to the LDAP provider with a simple BIND using stored credentials for a service account and regularly connecting to refresh the cache of users and groups, which is stored in SAS Infrastructure Data Server.
8. SAS Logon Manager sends the OAuth token back to the CAS controller.
9. The CAS controller uses the PAM configuration in /etc/pam.d/cas to validate the user name and password and launches the CAS controller as the end user.

See Also

- "Single-machine CAS Server" on page 496
- "Multiple CAS Servers" on page 498

Programming Interfaces with Massively Parallel Processing CAS Servers

In a massively parallel processing (MPP) environment, a distributed CAS server consists of one controller, one or more workers, and one backup controller (optional). Each component runs on a separate machine. The authentication process for MPP SAS Cloud Analytic Services is essentially the same as for SMP CAS, with the following key differences:

- Initial communication between the CAS controller and CAS workers is via Secure Socket Shell (SSH).
- On-going communication does not use SSH.
- A worker process is launched on each CAS worker as the end user.
  - The CAS controller authenticates the end user with PAM.
  - The CAS controller generates an internal identity token after authenticating the end user.
  - The internal identity token is used to launch the CAS worker processes.
  - PAM is not used on the CAS worker nodes.

See Also

"Multiple CAS Servers" on page 498

Authentication Mechanisms

LDAP Authentication

Overview of LDAP

Note: This information does not apply to a programming-only deployment.

In SAS Viya, LDAP is used for identifying and authenticating users. Third-party LDAP server implementations are supported, including Microsoft Active Directory and OpenLDAP.

How It Works in SAS Viya

LDAP is the default authentication mechanism. The Identities service always makes a direct connection to LDAP to obtain user and group information. By default, SAS Logon Manager authenticates users using a direct connection to the configured LDAP provider. To ensure that network connections are secure, the connection between the browser and the Apache HTTP Server can be secured with HTTPS. In addition, the connection between SAS Logon Manager and the LDAP provider can be secured with LDAPS.

For information about configuring LDAP, see Configure the Connection to Your Identity Provider.

Kerberos Authentication

Overview of Kerberos

Note: This information does not apply to a programming-only deployment.

Kerberos is a network authentication protocol that is used to verify user or host identity. The Kerberos protocol uses strong cryptography so that a client can prove its identity to a service (and vice versa) across an insecure network connection. During Kerberos authentication, a user's credentials (user ID and password) are not sent over the network. Instead, both the client and the service use the credentials that were supplied as a key in an encryption algorithm to encrypt the message that is sent between the client and the service. If the client sends an encrypted message, and the service uses the same key to decrypt the message, it is proven that the credential is known without having to transmit the credentials.

In SAS Viya, the visual interfaces are SAS Environment Manager and CAS Server Monitor. SAS Environment Manager can be enabled to support Kerberos authentication. Conversely, CAS Server Monitor does not support Kerberos authentication.

Key Terms

Table A.8 Term Definitions

Term	Definition
Client	An application that is attempting to connect to and access a resource, on behalf of a user. Resources include reports that are viewed, services that are accessed, and databases that are queried. In SAS Viya, the client is the web browser.
Service	A service, or server, that hosts a resource the user wants to connect to. The service must be able to validate the service tickets presented by the client.

Term	Definition
Key Distribution Center	A trusted third party within Kerberos that verifies the authenticity of the client and service. Both the client and service must trust the KDC. In addition, end users and services must register with the KDC.
Service Principal Name	A unique name that is used to identify a web service that is running on a server. Before a service principal name (SPN) can be used, it must be registered. Every web service that uses Kerberos authentication needs to have an SPN set for it so that clients can identify the server on the network. An SPN usually matches the pattern of HTTP/*hostname.example*.com.
Keytab File	A file containing pairs of Kerberos principals and encrypted keys. The keys are associated with a password for the principal. The principals are SPNs. Keys can use different encryption algorithms. For a single principal, you might have several entries that correspond to each encryption type.
Ticket-granting ticket	An encrypted identification file that is valid for a limited amount of time. After a user is authenticated, this file is granted to a user for data traffic protection by the KDC. The TGT file contains the session key, its expiration date, and the user's IP address.

How It Works in SAS Viya

In addition to using the LDAP provider to obtain user and group information, you can configure SAS Logon Manager for Kerberos authentication. This option replaces the option to use the default LDAP provider for authentication to SAS Logon Manager. Kerberos provides the user with single sign-on capabilities from the browser on their desktop. Single sign-on allows the user to access the SAS Viya visual interfaces without being prompted to enter their credentials.

For information, see “Configure Kerberos” on page 290.

Integrated Windows Authentication

Integrated Windows Authentication (IWA) uses Kerberos authentication and is a Microsoft technology that is used in an environment where users have Windows domain accounts. With IWA, the credentials are hashed before being sent across the network. The client browser proves its knowledge of the password through a cryptographic exchange with the web application server. When IWA is used in conjunction with Kerberos, IWA enables the delegation of security credentials. The Kerberos protocol uses strong cryptography so that a client can prove its identity to a server (and vice versa) across an insecure network connection.

Supported Kerberos Scenarios

There are six possible scenarios for accessing a Hadoop environment that is secured by Kerberos. The following table provides an overview of each use case and links to additional information.

Table A.9 *Supported Scenarios*

End-User Client	Connection to SAS Viya	Who Runs CAS Session	Connection to Hadoop
SAS Viya visual interface See “Kerberos in SAS Viya Visual Interface with Delegation” on page 315.	Kerberos (delegation)	End user	End user

End-User Client	Connection to SAS Viya	Who Runs CAS Session	Connection to Hadoop
SAS Viya visual interface See “Kerberos in SAS Viya Visual Interface without Delegation” on page 317.	Kerberos (no delegation)	Service account (cas)	Service account (sascas)
SAS Viya programming interface See “Kerberos in SAS Viya Programming Interface with User Credentials” on page 319.	User ID and password	End user	End user
SAS 9.4 See “Kerberos in SAS 9.4 with User Credentials” on page 320.	User ID and password	End user	End user
SAS 9.4 See “Kerberos in SAS 9.4 with Delegation” on page 320.	Kerberos (delegation)	End user	End user
SAS 9.4 See “Kerberos in SAS 9.4 with One-Time Password” on page 320.	One-time password	Service account (cas)	Service account (sascas)

Kerberos in SAS Viya Visual Interface with Delegation

The following figure illustrates the first scenario:

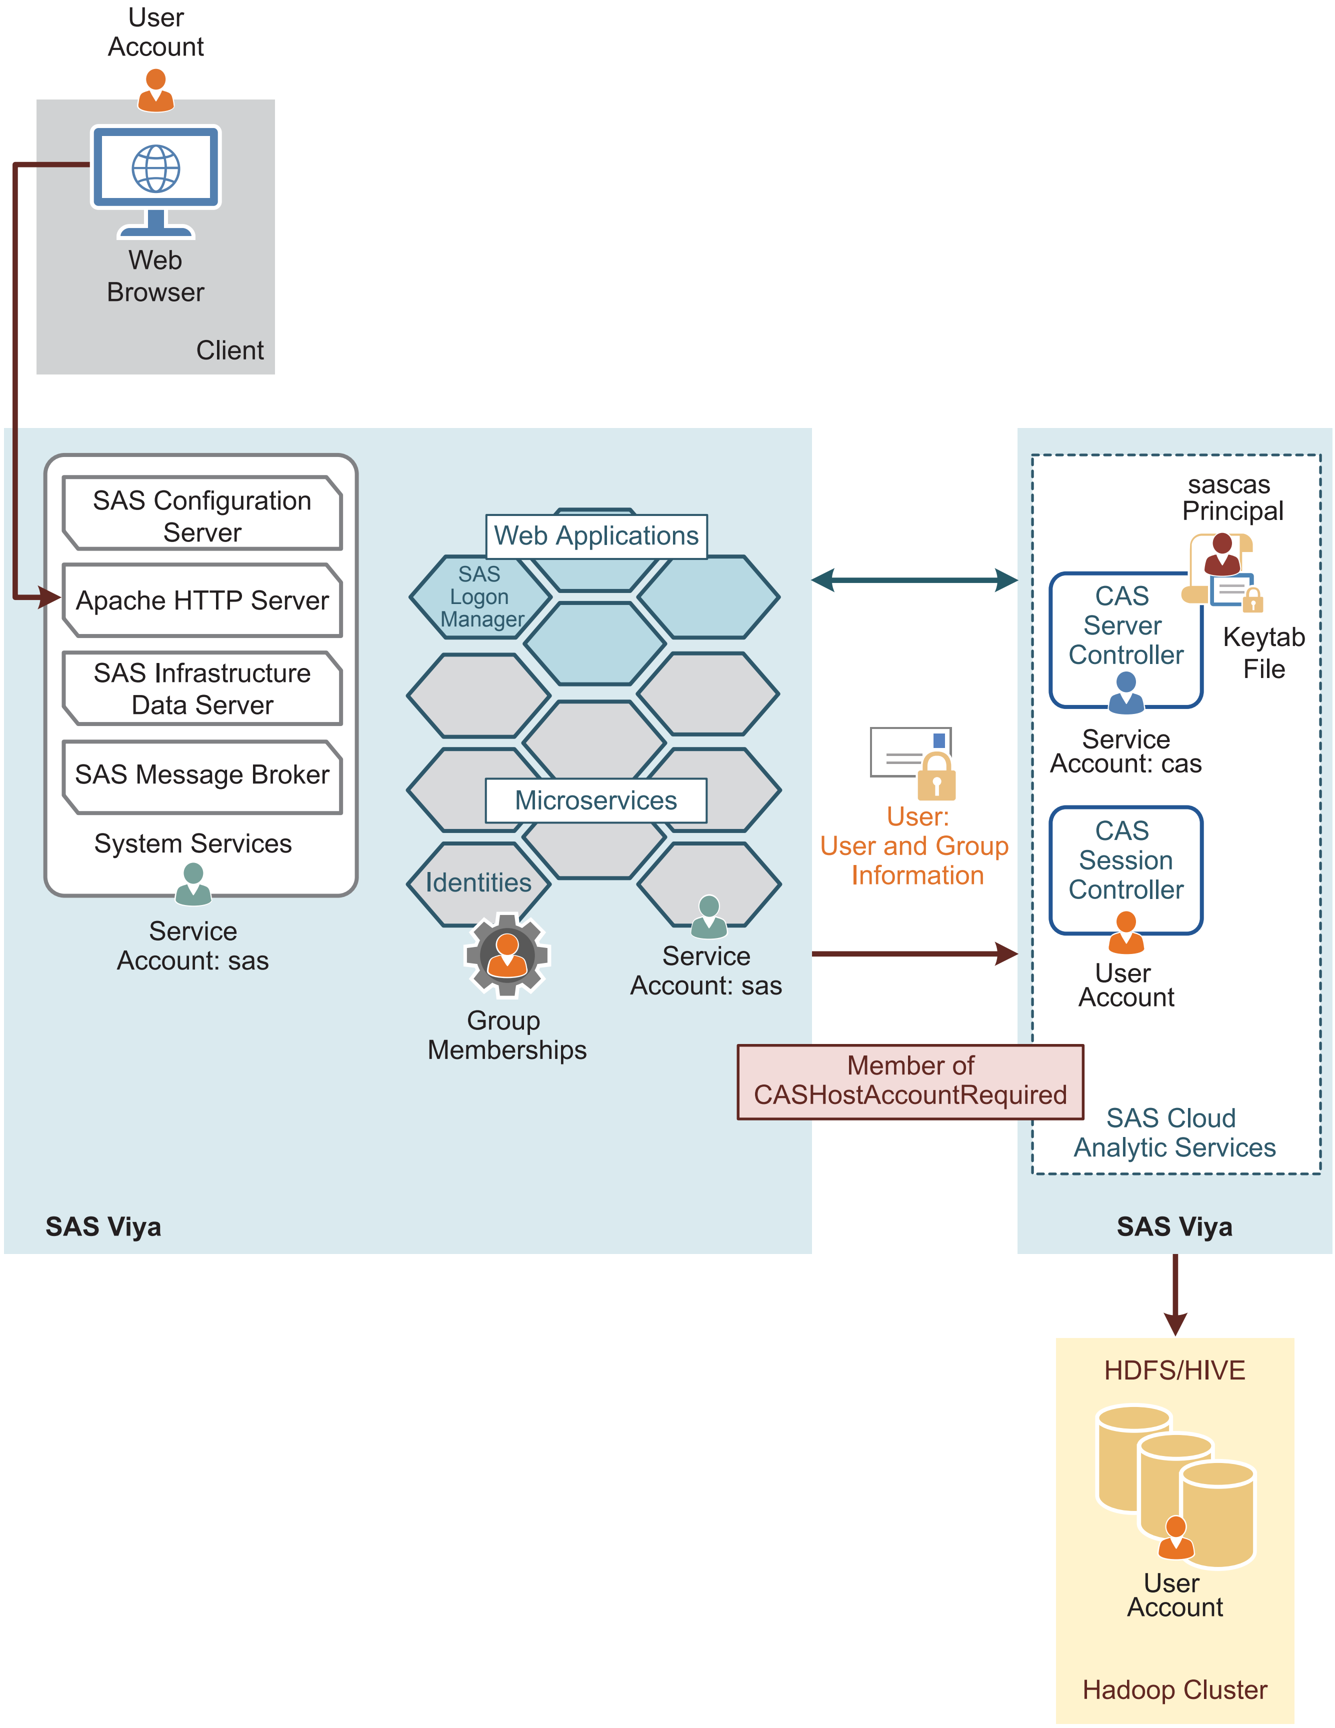

Kerberos delegation to CAS, or user delegation, is a feature that allows a SAS Viya application to reuse the end-user credentials to access Kerberized systems. Delegation allows a server to forward a user’s credentials to the

CAS server where they can be used to access other Kerberized services, such as Hadoop. By default, user delegation is not enabled and must be configured.

In this scenario, the CASHostAccountRequired custom group notifies CAS that the user session needs to be launched under the user's operating system account. Kerberos is used for authentication, and the user is a member of the CASHostAccountRequired group. Therefore, the user’s credentials are delegated to CAS, and the user can access Hadoop using Kerberos as himself or herself. For more information, see “The CASHostAccountRequired Custom Group” on page 365.

Note: On Windows, user sessions are launched under the user identity. However, there are times when this is not possible. A session can also be launched under the CAS service account if a session cannot be launched under the user identity and the user requesting the session is able to assume the Superuser role.

The service account that is associated with the Apache HTTP Server has to be trusted for delegation. An OAuth token is generated by SAS Logon Manager. It indicates to CAS that Kerberos was used to authenticate the user. If the user is a member of the CASHostAccountRequired group, CAS attempts to obtain a Kerberos ticket from SAS Logon Manager after the OAuth token is validated. The session is launched under the user’s operating system identity, and the user’s Kerberos ticket is used for secured Hadoop access.

See Also

“Configure Kerberos for SAS Cloud Analytic Services” on page 294

Kerberos in SAS Viya Visual Interface without Delegation

The following figure illustrates the second scenario:

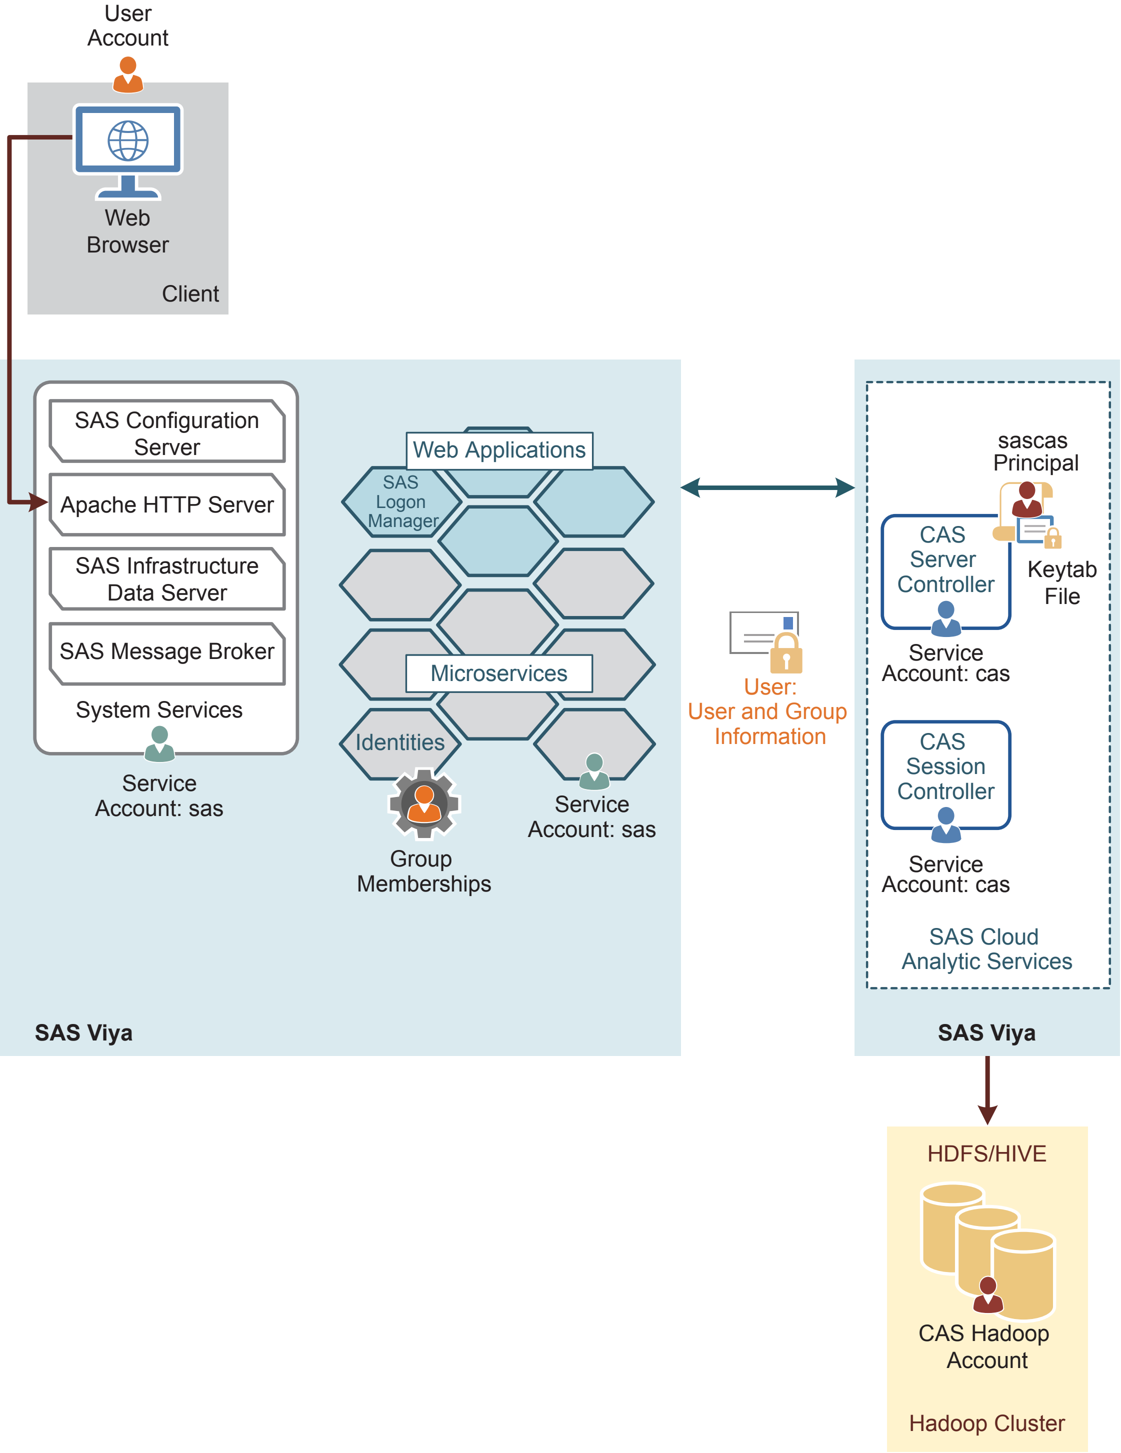

In this scenario, there is a full deployment and a user that relies on Kerberos to log on to SAS Logon Manager as part of his or her access to the visual interfaces. The user accesses the visual interfaces and connects to CAS and a secured Hadoop environment. The connection between the visual interfaces and CAS uses OAuth to authenticate the user. Users obtain an OAuth token from SAS Logon Manager as part of their initial authentication, and this token is used to authenticate to CAS. The OAuth token provides user and group information that enables CAS to provide integration with the authorization information that is stored in the SAS services.

The user is not a member of CASHostAccountRequired. Therefore, CAS does not try to obtain a delegated Kerberos ticket and does not attempt to launch the user's session under the user's operating system identity.

The CAS controller and CAS workers all run as the service account that launched the CAS operating system service. By default, this is the cas account.

If the user needs to access a secured Hadoop cluster, where Kerberos Service Tickets are required to access the data, then the CAS controller Kerberos credentials must be provided. The credentials are provided in the form of a Kerberos keytab file. This keytab file enables the CAS controller to initialize a set of Kerberos credentials at start-up. This set of Kerberos credentials does not have to be for the cas user. The credentials in the keytab can be for any user, but all access from CAS to Hadoop uses this single, shared credential. In the diagram, the shared credentials are listed as “CAS Hadoop Account”.

Kerberos in SAS Viya Programming Interface with User Credentials

The following figure illustrates the third scenario:

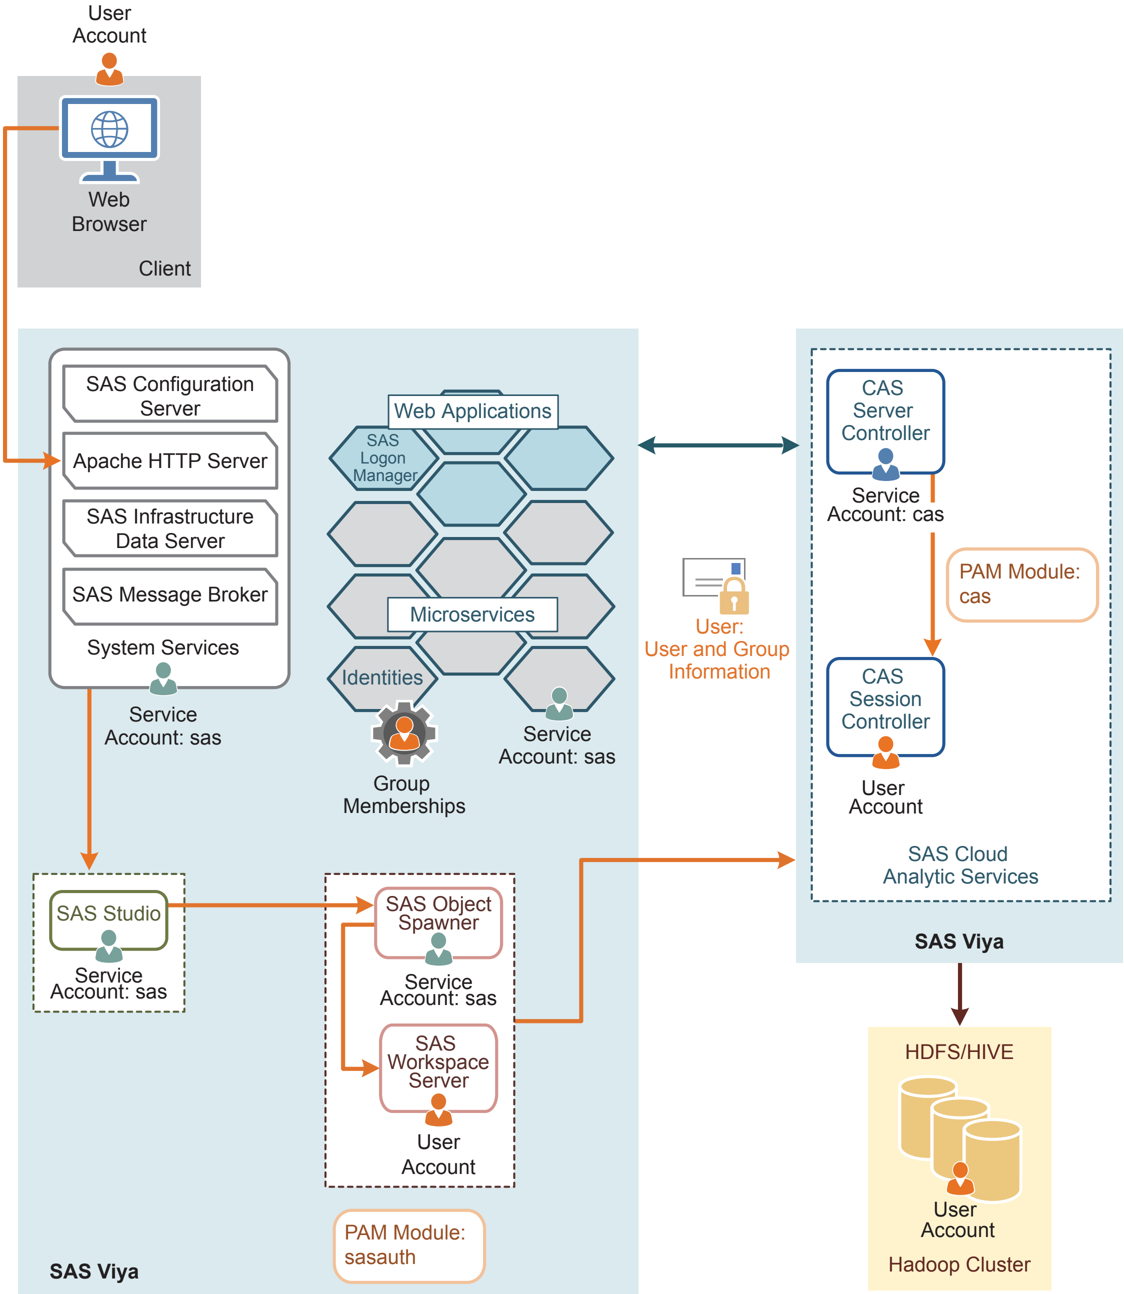

In this scenario, there is a full deployment and a user that provides his or her user ID and password to CAS. CAS uses its own pluggable authentication modules (PAM) configuration to validate the user’s credentials and launch the CAS controller process running as the user. In addition, the CAS controller also uses the user ID and password to obtain an OAuth token from SAS Logon Manager. The OAuth token provides the user's group memberships from the Identities service. These memberships are essential in enforcing access control.

The PAM stack is configured to generate a Kerberos credentials cache during authentication. The resulting cache can be used to access Hadoop as the user.

Depending on the deployment options that you chose, users who access both the programming interface and the visual interface might have different access to Hadoop.

Kerberos in SAS 9.4 with User Credentials

In the fourth scenario, end users provides their credentials to access SAS 9.4. SAS Workspace Server running SAS 9.4 is launched using a user ID and password, which are cached when SAS is launched. This enables SAS Workspace Server to use these cached credentials when connecting to CAS. The user credentials can also be provided by other sources in a SAS 9.4 environment, such as SAS Metadata Server or an authinfo file in the user’s home directory, because the process on the CAS controller is the same.

The user ID and password are validated through the PAM stack on the CAS controller and is used to generate an internal OAuth token from SAS Logon Manager running SAS Viya. The PAM stack is responsible for initializing the Kerberos credentials for the end user. These Kerberos credentials are placed into a Kerberos Ticket cache, which makes them available to the CAS session for the connection to the secured Hadoop environment. The different sessions within SAS 9.4, SAS Viya, and the secured Hadoop environment run as the end user.

Kerberos in SAS 9.4 with Delegation

In the fifth scenario, SAS 9.4 is configured for Kerberos authentication. SAS Workspace Server running SAS 9.4 is launched using Kerberos credentials and the service principal for SAS Object Spawner running SAS 9.4 must be trusted for delegation. A Kerberos credential for the end user is available to SAS Workspace Server, which can be used to request a service ticket for the connection to CAS. CAS is provided with a Kerberos keytab and principal that it can use to validate this service ticket. Validating the service ticket authenticates the SAS 9.4 end user to CAS. The principal for CAS must also be trusted for delegation. CAS session must have access to the Kerberos credentials of the SAS 9.4 end user.

The Kerberos credentials that are made available to CAS are used to make a Kerberized connection to SAS Logon Manager running SAS Viya to obtain the SAS Viya internal OAuth token. Therefore, SAS Logon Manager running SAS Viya must be configured to accept Kerberos connections. For information about the configuration property that must be configured, see “sas.logon.kerberos” on page 119. In addition, the Kerberos credentials for the SAS 9.4 end user are used to connect to the secure Hadoop environment.

Since all the principals are trusted for delegation, the SAS 9.4 end user can be authenticated using Kerberos with each component in the SAS Viya and SAS 9.4 integrated environment. Through the use of Kerberos authentication, the SAS 9.4 end user is authenticated in to CAS and out to the secure Hadoop environment.

Kerberos in SAS 9.4 with One-Time Password

In the sixth scenario, the SAS 9.4 session can be a SAS Stored Process Server, SAS Pooled Workspace Server, or SAS Workspace Server using server launch credentials, such as sassrv. The SAS 9.4 session is not running as the end user and does not have access to the end-user credentials. You can still connect to CAS and to the secured Hadoop environment by configuring one-time passwords generated by SAS Metadata Server running on SAS 9.4. SAS Metadata Server running on SAS 9.4 must be aware of CAS. This is done by creating a CAS server definition in SAS Metadata Server, using the AUTHDOMAIN= argument. For more information, see CAS Statement Arguments.

The SAS Viya environment must be able to validate the one-time password that is used to connect to CAS. When CAS receives the one-time password during the connection, it is sent to SAS Logon Manager running on SAS Viya for validation and to obtain a SAS Viya internal OAuth token. SAS Logon Manager running on SAS Viya must be configured to enable this validation. For information about the configuration property that must be configured, see “sas.logon.sas9” on page 122. SAS Logon Manager running on SAS Viya then passes the one-time password to SAS Web Infrastructure Platform running on SAS 9.4 to validate the password. After the one-time password is validated, a SAS Viya internal OAuth token is generated and passed back to CAS.

CAS does not have access to the end-user credentials. Therefore, the session that is created will be run using the account used to launch the controller process. By default, this account is cas. Since the end-user credentials are not available, the Kerberos credentials that are initialized for the session are from the Kerberos keytab provided to CAS. The connection to the secured Hadoop environment is made using those Kerberos credentials of the principal assigned to CAS.

OAuth and OpenID Connect Authentication

Overview of OAuth and OpenID Connect

Note: This information does not apply to a programming-only deployment. It also does not apply to Windows systems.

Open Authorization (OAuth) is a token-based authorization standard on the internet. OAuth 2.0 acts as an intermediary on behalf of the user, giving the third-party service an access token that authorizes specific account information. OpenID Connect is an extension to OAuth 2.0, which provides authentication support.

Key Terms

Table A.10 Term Definitions

Term	Definition
Access token	Specifies identifying information for a user, including the user’s credentials, groups, and privileges.
OpenID Connect	An authentication layer built on top of OAuth 2.0.
Flow	The process for obtaining an OAuth token.

How It Works in SAS Viya

An OAuth 2.0 and OpenID Connect provider can be internal to the customer’s environment, or it can be an external provider, such as Google Authenticator or Facebook. When the OAuth 2.0 option is configured, this does not completely replace the default LDAP provider. Instead, when users access SAS Logon Manager, they are presented with a link to authenticate using OAuth 2.0 and the standard logon form using the LDAP provider. Users can select which to use. The user identity and group membership information is looked up in LDAP. OAuth 2.0 can provide single sign-on from the OAuth 2.0 provider. For example, when a user signs in to his or her Google account, the user can access the visual interfaces of SAS Viya without being prompted any further for credentials.

SAML Authentication

Overview of SAML

Note: This information does not apply to a programming-only deployment. It also does not apply to Windows systems.

The Security Assertion Markup Language (SAML) standard defines a framework for exchanging security information about users between an identity provider and service provider. This security information is packaged in the form of portable XML assertions that applications working across security domain boundaries can trust. SAML allows for single sign-on to web browser applications.

Key Terms

Table A.11 Term Definitions

Term	Definition
Federation	Allows multiple identity management systems to work together and establish trust.
Assertion	A package of information, in the form of an XML document, that is created and sent during a federated access request.
Claims	Information that a federation member is asserting to be true.
Identity provider	A federation member that authenticates users and keeps track of their information. Creates assertions for the users, and sends them to service providers.
Service provider	A federation member that consumes assertions to make access control decisions for its applications.
Metadata	An XML document that is produced by a SAML provider to describe its service endpoint URLs, x.509 certificate, and other information in a standard way for consumption by partners in the federation.
Relying party	A server providing access to secure software.

How It Works in SAS Viya

SAML supports configuring SAS Logon Manager to be integrated with an external SAML identity provider. This identity provider can be internal or external to the customer’s environment. If it is internal, a tool similar to Oracle Access Manager can be used. If it is external, something like salesforce.com can be used. SAML does not completely replace the default LDAP provider. End-users accessing SAS Logon Manager can choose SAML authentication or the default LDAP provider. The user identity and group membership information is looked up in LDAP. This option also provides single sign-on with the third-party SAML provider.

When a user attempts to access a service URL, the service provider, which is SAS Logon Manager, initiates the exchange with an authentication request. The identity provider sends a response that contains the assertion. The SAML protocol defines the structure and content of these request and response messages. When the user logs on to a service or system, the service provider trusts the identity provider to validate the credentials, instead of providing credentials to the service provider. Therefore, users do not have to provide their credentials directly to anyone but the identity provider.

For configuration information, see “Configure SAML” on page 300.

SAS 9.4 Authentication

Overview of SAS 9.4 Authentication

This option enables integration between SAS Viya and an existing SAS 9.4 environment. The authentication to the SAS Viya visual interfaces is performed by the SAS Logon Manager in SAS 9.4. None of the authentication occurs with the SAS Logon Manager in SAS Viya. Any authentication mechanism supported by SAS 9.4 is supported by this configuration. For more information about the supported authentication mechanisms, see *SAS Intelligence Platform: Security Administration Guide*.

Note: All versions of SAS 9.4 support this configuration. The SAS 9.4 deployment does not have to be running the latest maintenance release.

How It Works in SAS Viya

Here is a sample scenario:

1 The client's web browser connects to SAS Logon Manager in SAS Viya.

   a If the request to SAS Logon Manager in SAS Viya does not have an existing session, the SAS Logon Manager in SAS Viya displays the logon form, which contains a link to perform SAS 9.4 authentication and the form to do LDAP authentication.

   b If the end user selects the link, SAS Logon Manager in SAS Viya constructs an authentication request and redirects the client's web browser to the SAS 9.4 middle tier.

2 The client authenticates to SAS 9.4, receives a service ticket, and is redirected to SAS Logon Manager on SAS Viya.

3 The client's web browser connects to SAS Logon Manager on SAS Viya, including the SAS 9.4 service ticket in the request.

4 SAS Logon Manager on SAS Viya connects to SAS 9.4 middle tier to validate the service ticket and the end user.

5 SAS Logon Manager on SAS Viya connects to the Identities service to get the custom and LDAP group information for the validated end user.

6 The Identities service either looks up the validated end user in its cache or connects to Active Directory using the LDAP service account to update the cache.

The SAS 9.4 authentication configuration impacts only the SAS Viya 3.4 visual interfaces using the SAS Logon Manager. An LDAP provider is still required by the Identities service. For authentication to SAS Logon Manager in SAS Viya that is not through a browser, the credentials are first passed to the SAS Logon Manager in SAS 9.4. If they fail, the credentials are tried against LDAP. Therefore, authentication with the administration command-line interface (CLI) and SAS Visual Analytics App (previously called SAS Mobile BI) is still authenticated against the SAS Logon Manager in SAS Viya first. SAS Studio 4.4 is not impacted by this configuration.

If you want to configure TLS for either the SAS 9.4 or SAS Viya deployment, the Apache HTTP server certificate must be trusted. You need to import the certificate of the one deployment into the SAS certificate framework of the other deployment. For more information, see "Configure SAS 9.4 Clients to Work with SAS Viya" in *Encryption in SAS Viya: Data in Motion*.

Compatibility of User Names

The Identities service must be able to take the authenticated user name from SAS 9.4 and correctly search for it in the SAS Viya LDAP provider. You can log on to SAS 9.4 using an internal account (which includes the @saspw suffix), but such accounts cannot exist in the LDAP provider. Therefore, these accounts do not work with SAS Viya.

Also, you can sign in to SAS 9.4 with an account that does not exist in any LDAP provider, such as a Google account. This does not work with SAS Viya unless the Google account is the accountId property that is used by the Identities service. For more information about the accountId property, see "sas.identities.providers.ldap.group (Field Mappings)" on page 103.

Finally, domain qualified user names cannot be used with SAS Viya. Even if the SAS 9.4 environment passed the domain qualified user name, the domain is stripped.

Single Sign-On and Single Sign-Out

Single sign-on and single sign-out is supported between SAS Viya and SAS 9.4. During single sign-on, a user with an active SAS 9.4 session can access SAS Viya applications without being required to sign-on to SAS Viya.

Single sign-out is initiated from SAS Viya. If a user has two browser tabs open, one with a SAS Viya web application and the other with a SAS 9.4 web application, selecting the sign-out option in SAS Viya also signs the user out of SAS 9.4. However, the reverse is not true. If the user signs out from the SAS 9.4 web application, he or she is not signed out from the SAS Viya web application.

PAM Authentication

Overview of PAM

Pluggable authentication modules (PAM) enable you to determine how applications use authentication to verify the identity of a user. It is an industry-standard technology that extends UNIX host authentication to recognize additional authentication providers. PAM uses *modules* or libraries to access multiple authentication methodologies. SAS Viya supports host authentication.

How It Works in SAS Viya

Default PAM configuration files, ***`SAS-Viya-configuration-directory`*****`/etc/pam.d/`*****`service`***, are installed as a part of the SAS Viya deployment process.

Note: For SAS Cloud Analytic Services (CAS) server, *service* is *cas*. For SAS Studio, *service* is *sasauth*.

For sasauth to perform authentication, entries must be made in the PAM configuration files that are provided by SAS. These entries describe the authentication services that are used when sasauth performs an authentication. This includes the account and auth module types. The session and password modules are not supported.

> **TIP** In a multi-machine deployment, configure PAM on the host with SAS Object Spawner and the host with CAS controller.

For configuration information, see “Configure PAM” on page 299.

Authinfo File

Authentication is used to control access to the CAS server and its resources. Your identity must be successfully authenticated before your session is created. SAS Studio authenticates the connection to CAS by using your user credentials. When password information is not available, an attempt is made to find an authinfo file (.authinfo is the default filename on Linux). The authinfo file provides a user name and password to CAS for host authentication. It is an alternative to including passwords in programs.

You can also force the use of the authinfo file by specifying authinfo= in the CAS statement. An alternative method is to use the CAS_AUTH_METHOD environment variable.

The authinfo file is required when you are using the command line to submit commands for the following tasks:

- Run programs in batch mode. The USER= option in the CAS statement or SAS system option CASUSER= can be specified.
- Perform limited server administration using the **casadmin** command.
- Run commands in line mode.
- Sign on to SAS/CONNECT and specify the casuser in the RSUBMIT block of code. This action is performed when the casuser is different from the SAS Viya user or when the user is the same for both SAS Viya and CASUSER, but the password is different.

Note: SAS Studio user credentials are used to authenticate your connection to CAS. SAS Studio does not use the authinfo file for authentication.

Typically, the authinfo file resides in the `$HOME` directory.

The authinfo file format is based on the .netrc file specification. The .netrc file format is an older format. You can see the file specification at Netrc Format. In addition to the standard .netrc file standards, the authinfo specification allows for putting commands in the file as well as using quoted strings for passwords. The quoted strings allow for spaces within passwords.

If the authinfo file contains values that match the host, port, or user name. The information contained in the authinfo file is used to connect to CAS.

The following system options and environment variables can be used to override the authinfo file. These options point to authinfo files that are located in a different directory or are named differently.

Here are the ways that the AUTHINFO system option, environment variable, and the statement option can be used to override the authinfo file:

- Environment variable AUTHINFO takes precedence over the authinfo file.
- SAS system option AUTHINFO= (alias CASAUTHINFO=) overrides the AUTHINFO environment variable as well as the authinfo file.
- AUTHINFO= option in the CAS statement overrides the AUTHINFO= system option, the AUTHINFO environment variable, and the authinfo file.

For more information, see the following documents:

- AUTHINFO= System Option
- CAS Statement
- CAS_AUTH_METHOD environment variable on page 335
- USER=user-ID argument
- Batch Mode in UNIX Environments

Multi-Factor Authentication

Multi-Factor authentication (MFA) is a security system that requires more than one method of authentication from independent categories of credentials to verify the user’s identity for a sign-on or other transaction.

MFA combines two or more of the following independent credentials:

- what the user knows – their password
- what the user has – a security token
- what the user is – biometric verification

The goal of MFA is to create a layered security defense, making it more difficult for an unauthorized person to access a target such as a physical location, computing device, network, or database.

Typical MFA scenarios include the following:

- swiping a card and entering a PIN
- logging on to a website and being requested to enter an additional one-time password (OTP) that the website's authentication server sends to the requester's phone or email address
- downloading a virtual private network (VPN) client with a valid digital certificate and logging on to the VPN before being granted access to a network
- swiping a card, scanning a fingerprint, and answering a security question
- attaching a universal serial bus (USB) hardware token to a desktop that generates a one-time passcode and using the one-time passcode to log on to a VPN client

Additional Authentication Topics

SAS/CONNECT Authentication

As an administrator, you might want to enable SAS Viya to accept connections for existing SAS 9 environments. SAS/CONNECT enables that connection, and passes credentials that can be used in the SAS Viya environment.

With SAS Viya, your credentials are used to authenticate to CAS when you are using SAS/CONNECT. When additional SAS/CONNECT servers are spawned, SAS/CONNECT forwards your credentials to the spawned SAS/CONNECT server session.

Here are the ways that SAS/CONNECT and CAS authenticate your user credentials:

- When the user is using any environment that is not a SAS Viya environment, and is connecting to SAS Viya via the SAS/CONNECT spawner, the spawner passes the SIGNON credentials to the SAS/CONNECT server where the credentials can be used to connect to CAS.
- When the user is in the SAS Viya environment using SAS Studio and starting SAS/CONNECT server sessions (using SASCMD SIGNON or the CONNECT Spawner), the CAS credentials (if they exist) are passed to the SAS/CONNECT server in SAS Viya.
- When running SAS Viya in batch or line mode, the authinfo file is used to authenticate to CAS. If you specified the USER= option in the CAS statement, CASUSER= system option, or if you specified the CAS_AUTH_METHOD environment variable, authinfo file authentication is used.

For more information, see the following documents:

- USER=user-ID
- CAS AUTH_METHOD environment variable on page 335.
- SAS/CONNECT 9.4 User’s Guide
- “Operate (Linux)” on page 578.

Single Sign-On

Note: This information does not apply to a programming-only deployment.

Single sign-on (SSO) is an authentication model that enables users to access a variety of computing resources without being repeatedly prompted for their user IDs and passwords. For example, SSO can enable a user to access SAS servers that run on different platforms without interactively providing the user's ID and password for each platform. SSO can also enable someone who is using one application to launch other applications based on the authentication that was performed when the user initially logged on.

SAS Logon Manager is the central point for handling changes to authentication mechanisms, such as the addition of third-party SSO products. SAS Viya supports the following SSO products:

- Kerberos on page 313
- “SAML Authentication” on page 321
- OAuth 2.0 and OpenID Connect on page 321

Dual Authentication

Linux

In a dual authentication environment on Linux, users are validated against the LDAP server and the host authentication mechanism. The following conditions exist:

- If PAM is configured to use local accounts and those users also log on to the visual components, then those local accounts must match the LDAP server used for SAS Logon Manager.
- If PAM is configured to use an LDAP server, SAS Logon Manager should be configured to use the same LDAP server.
- When directly connecting to the CAS server using SAS Studio or a batch job, the user ID and password that are supplied are authenticated against both the LDAP server and PAM.

Windows

In a dual authentication environment on Windows, users are validated against the LDAP server and the host authentication mechanism. The following conditions exist:

- The LDAP server should be configured to use the same Active Directory server that the Windows host is using.
- If host authentication is configured to use local accounts and those users also log on to the visual components, then those local accounts must match the LDAP server used for SAS Logon Manager.
- If host authentication is configured to use an LDAP server, SAS Logon Manager should be configured to use the same LDAP server.
- When directly connecting to the CAS server using SAS Studio or a batch job, the user ID and password that are supplied are authenticated against the LDAP server.

Authentication: Guest Access

About Guest Access

Note: This information does not apply to Windows systems.

Guest access is an optional feature that provides anonymous Read-Only access to a subset of resources and functionality in participating applications. Guest access is supported for viewing reports in the SAS Report Viewer and SAS Mobile BI.

For information about multi-tenancy, see “Guest Access in Multi-tenancy” in *SAS Viya Administration: Multi-tenancy*.

Enable Guest Access

Note: In a multi-tenancy environment, the following steps must be repeated for each tenant that supports guest access.

1 Set the sas.logon.provider.guest configuration property, using SAS Environment Manager:

   a In the applications menu (☰), select **Administration** ⇨ **Manage Environment**. In the navigation bar, select 🔧.

   b Create a new configuration instance for **sas.logon.provider.guest**, ensuring that you enable the guest access option. For more information, see “Create Configuration Instances” on page 75.

2 Add rules that provide the necessary access to functionality:

   a From the SAS Viya machine where the command line interfaces are installed, create a default profile, if you have not already created one, and sign on. For more information, see “Create at Least One Profile” on page 684.

b Modify the authorization rules.

- For a new SAS Viya 3.4 installation, run the following command:

```
sas-admin authorization facilitate-guest
```

- For an upgrade from SAS Viya 3.3 to SAS Viya 3.4 in which guest access was not previously configured, run the following command:

```
sas-admin authorization facilitate-guest
```

- For an upgrade from SAS Viya 3.3 to SAS Viya 3.4 in which guest access was previously configured, complete the following steps:

  □ Run the *facilitate-guest* command.

  ```
  sas-admin authorization facilitate-guest
  ```

  Output similar to the following will be displayed:

  ```
  The jsonPatch was not valid.
  Http Status: 400
  ErrorCode: 1177

  Detailed Messages:
          correlator: e607fd5d-c4c8-4548-ad2d-b9e608ccf41a
          traceId: 49f41d99e62595f2
          path: /authorization/rules
          FieldError: Rule [id=<defined_id>, type=GRANT, permissions=[READ], principal=null,
  principalType=guest, containerUri=null, objectUri=/identities/users/@currentUser, mediaType=null,
  condition=null, filter=null, reason=null, description=Guest Access: XXX, isEnabled=true,
  matchParams=false, isShare=false]:Provided authorization rule is a duplicate of this rule.
  ```

  □ Remove the rule ID that is specified in the output of the previous step:

  ```
  sas-admin authorization remove-rule --id=<defined_id>
  ```

  □ Run the *facilitate-guest* command again. If an error message is displayed stating `“Provided authorization rule is a duplicate of this rule”`, repeat the previous step to remove the rule ID.

  Repeat this step until the *facilitate-guest* command runs successfully.

- For an upgrade from a release prior to SAS Viya 3.3 to SAS Viya 3.4, run the following command:

```
sas-admin authorization facilitate-guest
```

c Modify the direct access controls for the predefined caslibs on the server, using the controls that are defined in the specified source file, run the following command:

Note: The following command must be executed by a user who is a member of the Superuser role.

```
sas-admin cas facilitate-guest --source-file path-to-controls-file --server CAS-server-name
--superuser
```

For more information about the controls file, see “Enable Guest Access” on page 706.

3 Add access controls that provide Read access to caslibs that should be accessible to guest users:

a From the SAS Viya machine, if you have not already signed in to SAS Viya, sign on using the default profile that was created in the previous step.

b Run the following commands as a user who is a member of the Superuser role:

```
sas-admin cas caslibs add-control --server server-name --caslib caslib-name --grant readInfo
--guest --superuser
```

```
sas-admin cas caslibs add-control --server server-name --caslib caslib-name --grant select
--guest --superuser

sas-admin cas caslibs add-control --server server-name --caslib caslib-name --grant limitedPromote
--guest --superuser
```

4 Use SAS Environment Manager to grant Read access to folders and reports that should be accessible to guest users:

   a From the **Content** page, identify the folder to which you want to grant Read access to guest users.

   b Right-click and select **Edit Authorization**.

   c Click + and select **Add Guest**. Grant Read and Read (convey) access. For more information, see "General Authorization: How to (Authorization Window)" on page 412.

   d Click **Save**.

   Note: From the **Content** page, you can move folders and objects into the **My Favorites** and **My Folder** folders for the guest user. You can also create and add folder and report shortcuts . For more information, see "Content Management: How To" on page 132.

Connect as Guest Users

Once guest access is enabled, guest users can view reports using SAS Report Viewer and SAS Mobile BI.

- SAS Report Viewer displays a guest login button.
- SAS Mobile BI displays a guest login button when a mobile connection is established.

See Also

- SAS Report Viewer Documentation
- SAS Mobile BI Documentation

Generate Custom Links to Reports

You can create a custom web link for guest users, allowing them to access a specific report. If guest access is enabled, the custom link is configured to bypass the logon page and automatically connect the user as guest. If guest access is disabled, a logon page is displayed, where users can choose to connect as a guest or log on with their credentials.

1. From SAS Report Viewer, open the report to which you want to generate a link.
2. Click ⋮ and then select **Share report** ⇨ **Link**.
3. In the Generate Link window, customize the link, if necessary, in the **Link** field.
4. Click **Copy Link**. You can paste the link and distribute to guest users.

See Also

SAS Report Viewer Documentation

Disable Guest Access

1 Set the sas.logon.provider.guest configuration property, using SAS Environment Manager:

  a In the applications menu (☰), select **Administration** ⇨ **Manage Environment**. In the navigation bar, select 🔧.

  b From the **Definitions** view, select **sas.logon.provider.guest**.

  c Click ✎. In the Edit sas.logon.provider.guest Configuration window, select the option to disable guest access.

  Note: The sas.logon.provider.guest option is tenant-specific and must be disabled for each tenant.

  d Click **Save**.

2 (Optional) Remove the rules that provide the necessary access to functionality:

  a From the SAS Viya machine, navigate to the ***SAS-Viya-installation-directory*/home/bin** directory.

  b At the command prompt, create a default profile and sign on by entering the following commands:

```
sas-admin profile init
sas-admin auth login
```

  c Modify the authorization rules by running the following command:

```
sas-admin authorization disable-guest-access
```

  Note: This command removes the rules that were automatically loaded by the *facilitate-guest* command. If you manually created any custom rules, using either SAS Environment Manager or the command-line interface, you must manually remove those rules. A list of the remaining guest rules can be viewed on the SAS Environment Manager **Rules** page.

3 (Optional) Run the following commands as a user who is a member of the Superuser role to remove CAS Access grants:

```
sas-admin cas sessions create --server server-name --name clisession --superuser

sas-admin cas caslibs remove-control --server server-name --caslib VAModels --grant readInfo
--guest --session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib VAModels --grant select
--guest --session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib VAModels --grant limitedPromote
--guest --session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib ReferenceData --grant readInfo
--guest --session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib ReferenceData --grant select
--guest --session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib ReferenceData --grant
limitedPromote --guest --session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib AppData --grant readInfo --guest
```

```
--session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib AppData --grant select --guest
--session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib AppData --grant
limitedPromote --guest --session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib Formats --grant readInfo --guest
--session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib Formats --grant select --guest
--session-id session-id

sas-admin cas caslibs remove-control --server server-name --caslib Formats --grant limitedPromote
--guest --session-id session-id

sas-admin cas sessions delete --server server-name --session-id session-id
```

Note: These commands remove the grants that were automatically defined by the facilitate-guest command. If you manually created any custom grants, using either SAS Environment Manager or the command-line interface, you must manually remove those grants.

Authentication: OpenID Connect Scenario

Note: This information does not apply to a programming-only deployment. It also does not apply to Windows systems.

In the following tasks, OpenID Connect uses IBM Security Access Manager (ISAM) WebSEAL reverse proxy server as the single sign-on entry point for initial user authentication. Other providers can be used, but configuration instructions are not provided here. To configure the OAuth and OpenID Connect, complete the following sections:

Configure OpenID Connect Provider Properties for IBM Security Access Manager

1 Log on to SAS Environment Manager.

2 Navigate to the SAS Logon Manager configuration definitions. For more information, see "Edit Authentication Configuration Instances" on page 304.

3 In the **Definitions** list, select **sas.logon.oauth.providers.external_oauth**.

4 In the top right corner of the window, click ✱.

5 In the New sas.logon.oauth.providers.external_oauth Configuration dialog box, enter values for the required fields, based on your environment. The following table provides guidance about the information needed for the listed fields:

Table A.12 OAuth Configuration Fields and Descriptions

Configuration Fields	Descriptions
attributeMapping.user_name	The attribute claim to use as the user name. For ISAM, use **sub**.
authUrl	The URL to the authorization endpoint (for example, https://*hostname.example.com*/isam/oidc/endpoint/amapp-runtime-ISAMOP/authorize).
emailDomain	The email address domain for users authenticating with the provider.
issuer	The principal that issued the token, specified as a case-sensitive string or URI. This is your WebSEAL instance (for example, the reverse proxy entry point, https://oidcidp.*example*.com).
linkText	The text that should be displayed on the sign-in page for the provider (for example, **OpenID Connect Login Using ISAM Reverse Proxy [WebSEAL]**).
relyingPartyId	The client ID that is registered with the provider.
relyingPartySecret	The secret that is registered with the provider for the client ID.
scopes	The comma-delimited list of scopes for the authorization request. The list should contain `openid`. **Note:** SAS Viya does not process any additional scopes that are returned in the token.
tokenUrl	The URL to the token endpoint.
type	The protocol type. By default, the value is **oidc1.0**. **Note:** SAS Viya requires an id_token in the authorization response from the provider. However, some providers return an id_token when the scope in the authorization request is *openid* and respose_type=*token*. For those providers, use type `oauth2.0`.

6 Click **Save**.

7 Restart the SAS Logon Manager Service.

On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-saslogon-default stop
sudo service sas-viya-saslogon-default start
```

On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-saslogon-default
```

On Windows, in Windows Services Manager, right-click the **SAS Logon Manager service** and select **Restart**.

Note: It might take several minutes to restart SAS Logon Manager.

Configure OpenID Connect Provider in IBM Security Access Manager

For basic steps to configure OpenID Connect in ISAM 9.0.3.1, see: IBM SECURITY ACCESS MANAGER Federation Cookbook 9.0.0.0 – 9.0.3.0 Installation, SAML 2.0, OpenID Connect, and Secure Token Service. To configure OpenID Connect Provider, complete the following steps:

1. In the ISAM 9.0.3.x admin console, create the WebSEAL reverse proxy instance as a single sign-on entry point.

2. Configure an OpenID Connect Provider and its partner.

   An OpenID Connect Provider on ISAM is a federation. First create a federation that represents the OpenID Connect Provider. Then, create a partner that represents the SAS Viya application under it.

3. Create a federation for OpenID Connect Provider. The following table shows values that you should provide while creating the new federation.

Table A.13 Create New Federation Values

Field Name	Value
Federation Name	ISAMOP
Protocol for this federation	OpenID Connect
Role	OpenID Connect Provider
Issuer Identifier	www.oidcidp.*example*.com **Note:** This is your WebSEAL instance.
Signature Algorithm	HS256
Grants	Authorization Code
Identity Mapping	Do not perform identity mapping . The same user name exists both in ISAM LDAP and SAS Viya LDAP.

4. Create an OpenID Connect Provider Partner for SAS Viya (SASLogon). The following table shows values that you should provide while creating the new partner.

Table A.14 Create New Partner Values

Field Name	Value
Name	ISAM-to-SASViya
Enabled	Yes
Connection Template	OIDC
Client ID	isamClientID

Field Name	Value
Client Secret	isamClientSecret
Client Display Name	SAS Viya Client
Response Types	code, id-token token, and token
Allow Refresh Token Grant	Enabled
Redirect URIs	https://*sas-viya-host*/SASLogon/login/callback/external_oauth
Scope	openid

5 Test your configuration.

OpenID Connect and IBM Security Access Manager

The following diagram depicts the IBM Security Access Manager reverse proxy components and process flow.

Figure A.1 IBM Security Access Manager Components and Flow

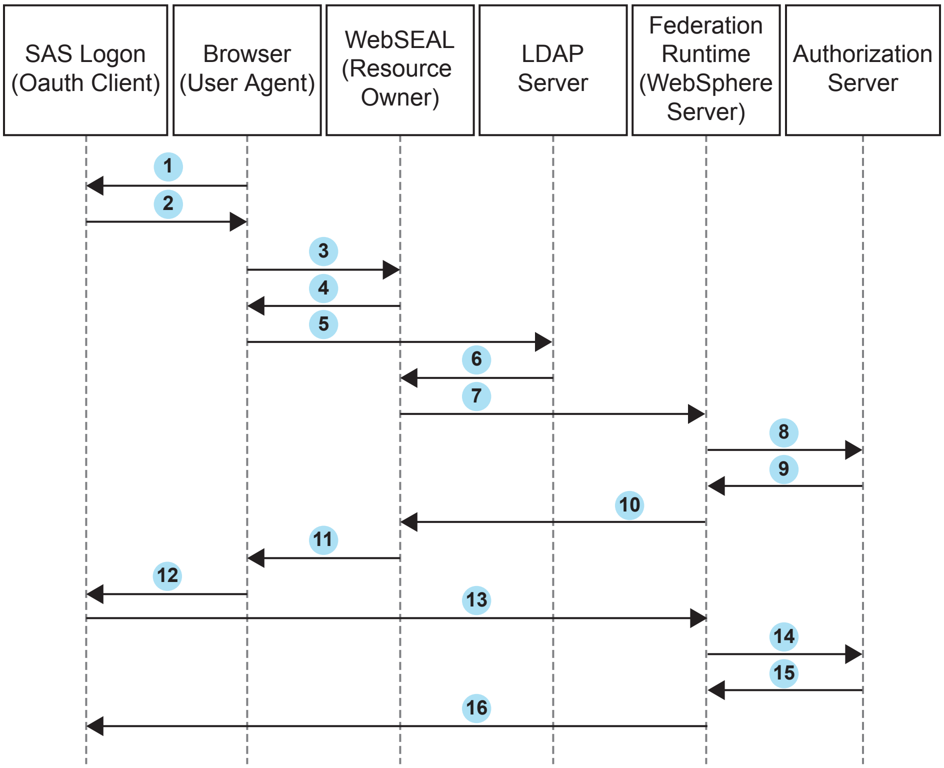

In this figure, the numbered arrows correspond to the following activities:

1 A client browser (user agent) accesses SAS Logon Manager (OAuth client)

2 SAS Logon Manager redirects the client browser to the SAS Logon Manager logon page. The end user clicks **OpenID Connect logon using ISAM Reverse Proxy (WebSEAL)**.

3. The client browser sends an authentication request to WebSEAL (resource owner).
4. WebSEAL redirects the client browser to the IBM Security Access Manager (ISAM) sign-in page. The end user provides their authentication information.
5. The client browser sends the authentication information to the Lightweight Directory Access Protocol (LDAP) server.
6. The LDAP server authenticates the user with IBM Security Access Manager.
7. WebSEAL sends an authorization request to the ISAM federation run time (WebSphere Application Server).
8. The ISAM federation run time sends the authorization request to the authorization server.
9. The authorization server sends an authorization code to the ISAM federation run time.
10. The ISAM federation run time sends the authorization code to WebSEAL.
11. WebSEAL sends the authorization code to the client browser.
12. The client browser sends the authorization code to SAS Logon Manager.
13. SAS Logon Manager sends a request to the ISAM federation run time to convert the authorization code to an access token.
14. The ISAM federation run time sends the request to the authorization server.
15. The authorization server sends the access token to the ISAM federation run time.
16. The ISAM federation run time sends the access token to SAS Logon Manager.

Authentication: Reference

CAS Environment Variables for Clients

The environment variables in this section are set on the client and affect how the client authenticates with the CAS server.

CAS_AUTH_METHOD=authinfo | kerberos
specifies the authentication method that CAS clients use.

Valid in	operating system command line
Category	Security
Operating environment	Environment variables on Linux are case-sensitive.
See	Authinfo File Authentication
Examples	In these examples, the CAS client is forced to authenticate using the credentials in the authinfo file (Kerberos authentication is not attempted). Here are two examples of specifying the command for Linux.
	export CAS_AUTH_METHOD=authinfo
	set CAS_AUTH_METHOD=kerberos

Authentication: Troubleshooting

The following error message is displayed while trying to log on to SAS Logon Manager: **“An error occurred while the system was verifying your credentials. Please enter your credentials again.”**

Resolution:
Change the operating system clock to the correct time.

After configuring Kerberos for SAS Logon Manager, you are unable to log on to a visual interface, such as SAS Environment Manager.

Resolution:
You must use a web browser on a different machine. Once Kerberos is enabled on Windows, a browser running on the same machine where the services are deployed cannot connect to SAS Viya visual interfaces.

The Kerberos authentication handshake fails and a session is not launched.

Resolution:
Users can store their credentials from the My Credentials page. Then, if the Kerberos handshake fails, authentication will fallback to the stored credentials in DefaultAuth. For more information, see “Add New Credentials” on page 347.

After configuring Kerberos for SAS Logon Manager, no one is able to log on to SAS Environment Manager.

Resolution:
If the information that you specified while adding Kerberos to the active profile, profiles.active, is incorrect or missing, the only way to change the information is by using the SAS Bootstrap Config CLI.

Run the following command:

```
/opt/sas/viya/home/bin/sas-bootstrap-config --token-file
$consul-token kv write --force config/SASLogon/spring/profiles.active ldap,postgresql
```

Note: The previous command must be on one line. It is shown on more than one line for display purposes only.

For more information, see “Use the SAS Bootstrap Config CLI on Consul to Manage the KV Store and ACL Tokens ” in *Encryption in SAS Viya: Data in Motion*.

20

External Credentials Management

External Credentials: Overview

In addition to logon credentials, users on SAS Viya systems might need external credentials for accessing databases and other third-party products. This document describes administrative tasks to manage external credentials using SAS Environment Manager.

To enable users to retrieve data from external data management systems (Oracle, Teradata, Facebook, Amazon, and so on), the business user must have the appropriate credentials and SAS Viya must be able to use those credentials.

This document assumes that you are familiar with the data and caslib concepts explained in Caslibs, Files, and Tables. For information about using SAS Environment Manager to manage caslib data see Caslibs on the Data Sources Tab and Import Tab.

External Credentials: How To

About the Domains Page

The Domainspage in SAS Environment Manager enables you to manage the following types of domains and credentials:

Note: The Domains area is available only if you are a member of the SAS Administrators group.

authentication domain	Makes stored credentials (user IDs and passwords) available to designated identities to facilitate connections to servers that require a password.
connection domain	Makes stored credentials (user IDs) available to designated identities to facilitate connections to servers that do not require a password.
encryption domain	Makes a stored credential (an encryption key) available to designated identities to facilitate loading of encrypted files. See Encryption for Data at Rest on page 443.

Navigation

The Domains area is available if you are a member of the *SAS Administrators* group and you have opted into your assumable groups. For more information, see Accessing SAS Environment Manager on page 665.

1 In the applications menu (≡), from the Administration section, select **Manage Environment**.

2 In the navigation bar, locate the **Security** section and click **Domains** .

3 You can select one of two views from the Domains page. The default view is Domains. From the **View** drop-down list, select one of the following views:

Domains	Lists all domains displayed. Domains is the default view. On the Domains page, you can view the information for each domain that is defined, or you can create a new domain. **Note:** This view is available only to SAS Administrators. **Note:** This document discusses only the Authentication domain and the Connection domain. For information about the Encryption domain, see Encryption for Data at Rest on page 443.
Credentials	Enables you to access external data sources and other third-party products requiring authentication. Credentials are associated with a specific domain for use with a specific data source type.

Pre-Defined Authentication Domains

The following Authentication domains are pre-defined and appear in the **Domains** view of SAS Environment Manager. For more information, see “Authentication Domains Pre-Defined in the Deployment” on page 346.

DefaultAuth

used to store user credentials for process launching. The DefaultAuth domain is used by CAS and the Launcher service to fetch user's stored credentials when the Kerberos handshake fails due to expired or non-delegated credentials or when CAS or the Launcher service requests an OS-launched session using OAuth. You can change the DefaultAuth domain by changing the configuration property sas.compute.domain.default. See "Change DefaultAuth Domain" on page 341.

EsriAuth

used to store credentials for ArcGis online server definition.

Note: Only User credentials should be added to the EsriAuth domain to access premium services. No groups should be added to the EsriAuth domain.

EsriPortalAuth

used to store credentials on-premises for an Esri server or portal.

Manage Domains

The Domains area is available if you are a member of the *SAS Administrators* group and you have opted into your assumable groups.

Create a New Domain

Create an Authentication or Connection domain.

1 In the **Domains** view, click ✱.

2 In the New Domain window, you can specify the following settings:

ID	Create an ID name. Required for the Authentication domain and the Connection domain.
Type	There are three domain types: Authentication, Connection, and Encryption. Select **Authentication** or **Connection** from the list of domains.
Description	Add a description.
Identities	From the Select Identities window, you can select from users, groups, and custom groups. When you add an identity, you are prompted to enter a user ID and password. **Note:** In the New Domain window, entering identities is optional. You can simply create the domain name. Users can add their own credentials using the My Credentials page.
User ID	Enter the user ID that has access to the external data. All identities selected connect using this user ID.
Password	Enter the password for the user ID that can connect to the external data source. This is not needed for the Connection domain.
Confirm password	Confirm the password. This is not needed for the Connection domain. **Note:** If the passwords do not match, you cannot save the domain.

3 For additional information about identities, from the New Domain window, click ?.

4 Add **Identities**. From the New Domain window, click 👤.

Note: In the New Domain window, entering identities is optional. You can simply create the domain name. Users can add their own credentials using the My Credentials area. See “Manage My Credentials” on page 346.

a In the left pane of the Choose Identities window, you can filter by Users, Groups, and Custom Groups. From the drop-down menu, select 👤**Users**, 👥**Groups**, or 👥 **Custom Groups**.

You can also filter using the search 🔍.

Note: A best practice is to use a custom group. Then you can add additional users to this custom group as needed to grant access to the external data. Be sure to assign correct permission for this custom group in the associated caslib Authorization.

Note: Only User credentials can be added to the EsriAuth domain to access premium services. No groups should be added to the EsriAuth domain.

b Move the selected user, group, or custom group to the **Selected Identities** pane. Click ➡.

c Click **OK**.

5 Add a user ID.

6 Add a password. Click **Confirm password**.

7 After you enter all of the parameter settings needed, click **Save**.

View Properties of the Domain

1 In the **Domains** view, select an ID of type Authentication or Connection.

2 Right-click, and select **Properties**. Or select ☰ from the taskbar. In the Domain Properties window, properties pertaining to that domain are displayed. Examples of properties displayed are the ID, Type, Description, Date created, Date modified, Created by, and Modified By.

3 Click **Close**.

View Credentials by Domain

1 From the Domains page, select an ID of type Authentication or Connection to view the credentials associated with that specific domain.

2 Right-click, and select **Credentials**. Or select ** from the taskbar. In the Credentials for Domain view, the credentials associated with your domain. The user ID and Identity are displayed along with the Modified By, Date Modified, Created By, and Date Created information.

3 Manage the columns that are displayed.

a At the right edge of the table, click the **Options** icon ⋮ , and select **Manage Columns**.

b From the Manage columns window, select items to move from the **Hidden columns** pane to the **Displayed columns** pane. After selecting the **Hidden columns** to display, click the **Add** arrow ➡

After selecting the **Displayed columns** to hide, click the **Remove** arrow ⬅

TIP To select more than one item, use the Shift key.

c Click **OK**.

4 From the **Credentials by Domain** view, you can also add new credentials to the existing domain, edit existing credentials, or delete a credential. See “Manage Credentials” on page 342 for the details.

Edit the Description of a Domain

1 From the Domains page, select an ID that is of domain type Connection or Authentication.

2 Right-click, and select **Edit**. Or select ✎ from the taskbar.

3 Edit the description of the Authentication Domain.

4 Click **Save**.

Refresh a Domain

1 In the **Domains** view, you can refresh the view.

2 Select ↻ from the taskbar.

A toast message is displayed to indicate when the Refresh has completed.

Delete Credentials from a Domain

1 In the **Domains** view, select an ID that is of domain type Connection or Authentication.

2 Right-click, and select **Credentials** . Or select ** from the taskbar.

3 From the Credentials for Domain window, right-click, and select **Delete**. Or select 🗑 from the taskbar.

4 In the Caution window, this message is displayed: `“Are you sure you want to delete the credential for the identity 'identity' with user ID 'userid'?”`

5 Click **Yes**.

Delete a Domain

1 From the Domains page, in the **Domains** view, select the ID of a domain and right-click to select **Delete**. Or select 🗑 from the taskbar.

2 In the Delete window, this message is displayed for an authentication domain that was being deleted: `“If a library is associated with this domain, you will not be able to access the data in the library after you delete the domain. Are you sure you want to delete the authentication domain named '`*your-domain*`' and all credentials associated with the domain?”`

3 Click **Yes**.

Change DefaultAuth Domain

The Compute and CAS servers can be configured to use a different DefaultAuth domain by changing the value of sas.compute.domain.default property.

1 In the applications menu (≡), from the Administration section, select **Manage Environment**.

2 In the navigation bar, locate the **System** section and click **Configuration** 🔧.

3 From the Configuration window, select **Definitions** from the drop-down menu.

4 In the search box, type sas.compute. Double-click on sas.compute to see the configuration properties.

5 Edit the sas.compute properties. Click ✎.

6 From the Edit sas.compute Configuration window, change the domain.default value from DefaultAuth to another domain.

7 Click **Save**.

For information about the sas.compute properties, see “Compute Service” on page 98.

Manage Credentials

The Domains area is available if you are a member of the *SAS Administrators* group, and you have opted into your assumable groups.

Add New Credentials

From the Domains page, select the **Credentials** view from the drop-down list.

Note: You are creating new credentials for an existing domain.

1 To add new credentials, click ✱.

2 From the New Credential window, select an existing *Authentication* or *Connection* domain from the drop-down menu.

3 After the Domain has been selected, specify the other settings in the New Credential window.

Domain	Select an existing *Authentication* or *Connection* domain.
Identities	In the Choose Identities window, you can select from users, groups, and custom groups. See below for instructions on selecting identities.
User ID	Enter the user ID and password required to access the external data.
Password	Enter the password associated with the Identities. Note: If a Connection domain is selected, a password is not required.
Confirm Password	Enter the same password as above.

4 For additional information about identities, in the New Credential window, click ?.

5 Add **Identities**. From the New Domain window, click 👤.

a In the left pane of the Choose Identities window, you can filter by Users, Groups, and Custom Groups. From the drop-down menu, select 👤**Users**, 👥**Groups**, or 👥 **Custom Groups**.

You can also filter using the search 🔍.

Note: A best practice is to use a custom group. Then you can add additional users to this custom group as needed to grant access to the external data. Be sure to assign correct permission for this custom group in the associated caslib Authorization.

Only User credentials can be added to the EsriAuth domain to access premium services. No groups should be added to the EsriAuth domain.

b Move the selected user, group, or custom group to the **Selected Identities** pane. Click ➡.

c Click **OK**.

6 Add a user ID.

7 Add a password. Confirm your password by entering it again.

Note: If a Connection domain is selected, a password is not required.

8 After you have entered all of the parameter settings needed, click **Save**. These credentials are now listed in the Credentials view.

View Credential Properties

1 In the **Credentials** view of the Domains page, right-click a user ID that you want to view, and select **Properties** or select ☰ from the taskbar. The properties are displayed in the Credential Properties window.

Only the credentials for the selected domain are displayed.

Note: Properties that can be displayed can include User ID, Identity ID, Identity type, Domain ID, Domain type, Date created, Date modified, Created by, and Modified by.

2 Manage which columns are displayed:

a At the right edge of the table, click the **Options** icon ⋮, and select **Manage columns**.

b From the Manage columns window, select items to move from the **Hidden columns** pane to the **Displayed columns** pane. After selecting the **Hidden columns** to display, click the **Add** arrow ➡

After selecting the **Displayed columns** to hide, click the **Remove** arrow ⬅

> **TIP** To select more than one item, use the Shift key.

c Click **OK**.

3 (Optional) You can also add a new credential using the **Credentials by Domain** view. You can delete a credential from this view also. Follow the steps in “Add New Credentials” on page 342 and “Delete Credentials” on page 344 for details.

Edit a Credential

If you are a member of the group or custom group, or a user associated with the selected credentials, you can add Identities and remove Identities (users, groups, and custom groups) from an existing credential. You must supply the credentials to edit.

1 In the **Credentials** view of the Domains page, select a credential entry.

2 Right-click, and select **Edit**. Or select ✎ from the taskbar.

3 To add and delete identities, in the Edit Credential window, click 👤.

a In the left pane of the Edit Credential window, you can filter by Users, Groups, and Custom Groups. From the drop-down menu, select 👤**Users**, 👥**Groups**, or 👥 **Custom Groups**.

   You can also filter using the search 🔍.

   Note: Only User credentials can be added to the EsriAuth domain to access premium services. No groups should be added to the EsriAuth domain.

b To add a selected user, group, or custom group to the **Selected Identities** pane, click ➡.

   To remove a selected user, group, or custom group from the **Selected Identities** pane, click the **Remove** arrow ⬅.

c Click **OK**.

4 Edit the user ID if you want to change it.

   Note: If a Connection domain is selected, a password is not required.

5 Edit the password if you want to change it. Enter the password.

6 Confirm the password. Enter the password again to confirm the password.

7 After you have edited all of the parameters, click **Save**.

Delete Credentials

Perform the following tasks in the **Credentials** view of the Domains page.

1 Right-click the user ID of the credential that you want to delete, and select **Delete**. Or select 🗑 on the taskbar.

2 In the Delete window, this message is displayed: `“Are you sure you want to delete the credential for the identity "name" with user ID 'userid'?”`

3 Click **Yes**.

External Credentials: Concepts

What Are Credentials?

Credentials are associated with identities. Identities can be individual users, groups, or custom groups. A credential enables you to assign a user ID for external data to one or more identities. Passwords are optional. For example, you can use a custom group called OracleUsers as an Identity and assign an Oracle user ID and password. The individual users or groups in this OracleUsers custom group do not need to know the Oracle credentials. These individual users have access through this OracleUsers custom group credential definition.

In SAS Viya, SAS Environment Manager uses the credentials service to store credentials.

What Is a Domain?

Overview

Domains are used to store both the credentials required to access external data sources (for example, databases such as Oracle, Teradata, and other data sources such as Facebook and Amazon) and the identities that are allowed to use those credentials. A domain contains one or more references to identities (users or groups) who have access to the credentials in the domain. Users can access the credentials either directly with their user ID or indirectly as a member of a group that is defined as an identity.

The ID, or name, of a domain is used in the definition of a non-path-based caslib to access and load tables from external data sources. A domain is associated with a caslib to provide access. Examples of external sources include SAS LASR, Oracle, Teradata, Hadoop, Postgres, and Impala. Users of a caslib with an associated domain do not have to know or enter data source credentials to access or load external data.

There are three domain types: authentication, connection, and encryption.

Authentication Domain

An authentication domain is a name that facilitates the matching of logons with the servers for which they are valid. Authentication domains are used to store credentials that are used to access an external source (for example, an Oracle database) that can then be associated with a caslib of the appropriate type. They can also be used for batch processing and scheduling where you store your credentials in a domain to run jobs in batch mode.

Each user ID and password is valid within a specific scope. For example, the user ID and password that you use to log on to your computer at work are probably not the same as the user ID and password that you use to log on to a personal computer at home. It is also common for database servers and web servers to have their own authentication mechanisms, which require yet another, different, user ID and password.

The software attempts to use only the credentials that it expects to be valid for a particular resource or system. The software’s knowledge of which credentials are likely to be valid is based entirely on authentication domain assignments. For this reason, you must correctly assign an authentication domain to each set of resources that uses a particular authentication provider, and also assign that same authentication domain to any stored credentials that are valid for that provider.

For example, assume that the user wants to define an Oracle caslib name "oralib" and to allow the Oracle users to access this caslib. From SAS Environment Manager, the administrator first creates a custom group of users called **orausers**, and then defines a domain called "OracleAuth." For the domain, they add **orausers** to the list of identities that have access to the credentials. They then provide a set of credentials (user ID and password) that give access to the Oracle database and schema that they plan to use in the caslib. Note that in this case, all users are accessing the schema using a shared set of credentials. Finally, when defining the caslib, the administrator associates the caslib with the "OracleAuth" domain.

Note: If the intent is for users to access the server with their own credentials, simply define the domain and instruct users to use the My Credentials window to enter their own credentials. See “Manage My Credentials” on page 346.

Authenticating to SAS Viya can be done through SAS Logon Manager. See Authentication: Overview on page 289.

Connection Domain

A Connection domain is used when the external source has been set up to require a user ID but no password. For example, a third-party database like Hadoop might be configured with accounts for authentication that do not require a password.

Encryption Domain

An encryption domain is used to store an encryption key that is required to read data at rest in a path assigned to a caslib. The Identities selected in this encryption domain have access to the key. When you create a path-based caslib, you can choose to enable encryption. You then select an encryption domain to assign an encryption key. Tables imported to this caslib are now encrypted. If the path contains preexisting tables, those tables are not encrypted. Users who are not defined in Identities as individuals or as members of a group cannot load data from this caslib. See Encryption for Data at Rest on page 443 for information about how to use encryption domains.

Authentication Domains Pre-Defined in the Deployment

The following three default authentication domains are present in a full deployment of SAS Viya:

EsriAuth
: authentication domain to store credentials for ArcGis online server definition.

  Note: Only User credentials can be added to the EsriAuth domain to access premium services. No groups should be added to the EsriAuth domain.

EsriPortalAuth
: authentication domain to store credentials on-premises for an Esri server or portal.

DefaultAuth
: used to store user credentials for process launching. The DefaultAuth domain is used by CAS and the Launcher service to fetch user’s stored credentials when the Kerberos handshake fails due to expired or non-delegated credentials or when CAS or the Launcher service requests an OS-launched session using OAuth.

Geo Enrichment Service from Environmental Systems Research Institute (Esri) provides a large collection of data sets, including population, income, housing, consumer behavior, and the natural environment. You can use SAS Visual Analytics to gain access to Esri premium services.

In order for the credentials to be validated in SAS Visual Analytics, you must be a member of the ESRI Users custom group and have valid credentials to access Esri premium services. Users can be added to this group using SAS Environment Manager. The user ID and password combination of this identity is entered in the Domain (the user ID and password combination that has access to the Esri server). You can provide single identities that all users share or each user can have their own identities. See “Custom Groups” on page 364.

If your login has appropriate privilege, you can Import Esri data into a caslib. See “Importing Esri Data for Geo-enrichment and Geocoding” in *SAS Data Explorer: User’s Guide*.

Manage My Credentials

Introduction to the My Credentials Page

These tasks explain how to manage your stored credentials for authentication and connection domains using SAS Environment Manager. These domains are used to store credentials for authenticating and connecting to third-party providers and to the CAS Server and the SAS Compute Server. The credentials are associated with users, groups, or both users and groups.

The My Credentials page is visible to all users, not just administrators. If you access the page while assuming administrator privileges, the functionality on the My Credentials page is the same as if you had not assumed administrator privileges. You can perform the following tasks from the My Credentials page:

- You can view, edit, and delete only the credentials for the currently logged-in user, and only for Authentication or Connection domains.

- You can add credentials to an authentication or connection domain, but only to domains for which a credential is not already associated.

  Note: Only a user with Administrator privileges can manage Identities, Domains, and User IDs using the Domains page.

Navigation

The My Credentials page is available when you log on to SAS, and you opt in (yes) or opt out (no) of assumable groups to manage your credentials. For more information, see Accessing SAS Environment Manager on page 665.

1 In the applications menu (≡), in the navigation bar, select **Administration** ⇨ **Manage Environment**.

2 In the navigation bar, locate **Security** and click the **My Credentials** icon 🔑.

Add New Credentials

You can add credentials to an existing domain. Only an administrator can add a domain. See "Manage Domains" on page 339.

1 To add new credentials, click ✱.

2 From the New Credential window, select an existing *Authentication* or *Connection* domain from the drop-down menu.

  Note: This list of domains displayed in the drop-down menu are domain names with which your personal credential is not associated.

3 Add a user ID.

4 Add a password and confirm your password by entering it again.

  Note: If a Connection domain is selected, a password is not required.

5 Click **Save**.

View Credentials and Manage Columns Displayed

You can view only the credentials for the currently logged-in user, and only for Authentication or Connection domains. All columns are displayed by default. Properties displayed are User ID, Domain ID, Modified By, Date Modified, Created By, and Date Created.

1 Select a row to enable the **Properties**, **Edit**, and **Delete** icons on the toolbar.

2 To view your Credential Properties, right-click a user ID that you want to view, and select **Properties** or select the **Properties** icon ☰.

3 Click **Close**.

4 To manage the columns that are displayed on the My Credentials page, at the right edge of the table, click the **Options** icon ⋮ , and select **Manage Columns**.

5 From the Manage columns window, select items to move between the **Displayed columns** pane and the **Hidden columns** pane. After selecting the items to display or hide, click the **Add** arrow ➡ or the **Remove** arrow ⬅

TIP To select more than one item, use the Shift key.

6 Click **OK**.

Edit Your Credentials

1 To change your credentials, select the **Edit** icon ✎.

2 From the Edit Credential window, you can change your user ID and password. For this example, select **Password** and enter a new password. Confirm your password by entering it again.

3 Click **Save**.

Refresh Password Properties

Manually refresh the password properties shown on the My Credentials page by selecting the **Refresh** icon ↻. A toast message is shown indicating that the page has been refreshed.

Delete Your Credentials

1 Right-click the user ID of the credential that you want to delete, and select **Delete**. Or select 🗑 on the taskbar.

2 A Delete window appears and asks “Are you sure you want to delete the credential with user ID "*yourName*". Click **Yes**.

21

Identity Management

Identity Management: Overview

SAS identity management includes the following:

- managing the membership of custom groups and CAS roles

- giving users, groups, and custom groups access to SAS functionality

You can use SAS Environment Manager for most identity management tasks in full deployments. SAS Environment Manager is not available in programming-only deployments.

Familiarize yourself with the predefined custom groups and CAS server roles. Based on this information, determine which of your users to add to roles and each predefined custom group.

Identity Management: How To (SAS Environment Manager)

View User and Group Information

1 In the applications menu (≡), select **Administration** ⇨ **Manage Environment**. In the navigation bar, select 👥.

2 On the **Users** page, you can do the following:

   - Select **Users**, **Groups**, or **Custom groups** from the drop-down list in the toolbar. Custom groups are displayed when you first open the page.

     Note: A custom group is a group that exists in SAS but not in your identity provider.

   - Enter a string in the **Filter** field to search for identities within the category that you selected (Users, Groups, or Custom groups). To restore the complete list of identities, clear the filter field.

   - Click an identity in the left pane to see its properties in the right pane. An identity's properties include the following:

     - profile picture (avatar) that is associated with the identity
     - basic properties including name, ID, title, and description
     - contact information (for users only)
     - a list of members (for groups and custom groups only)
     - a list of groups that the identity is a member of. 👥 indicates custom groups, and 👥 indicates groups from your identity provider.

     Note: Properties for users and groups (other than memberships in custom groups) are retrieved from your directory service and are read-only. Properties for custom groups are stored in SAS and can be edited using SAS Environment Manager.

   - Access recently viewed identities by using the drop-down box at the top of the right pane.

Note: To add, edit, or delete users and groups (other than custom groups), use your organization's identity provider (for example, Microsoft Active Directory) to which SAS Viya is connected.

Manage Custom Groups

A custom group is a group that exists in SAS Viya but not in your identity provider. Your deployment includes a set of predefined custom groups. You can also create your own custom groups, which are useful if you do not want to (or do not have permission to) create groups in your identity provider.

Add or Remove Custom Group Members

1. On the **Users** page in SAS Environment Manager, select **Custom groups** from the drop-down list in the toolbar.
2. In the left pane, click the name of the group whose members you want to update.
3. In the **Members** section of the right pane, click ✎.

   The Edit Members window displays the custom group's current members in the right pane.
4. To add a member, do the following:
   a. In the left pane of the Edit Members window, select **Users**, **Groups**, or **Custom groups** from the drop-down box.
   b. In the left pane, click the name of a user, group, or custom group identity. The identity's properties are displayed in the far right pane.
   c. Click ➕ or double-click the identity.
5. To remove a member, do the following in the Edit Members window:
   a. In the **Select Identities** list, click the user, group, or custom group identity that you want to remove. The identity's properties are displayed in the right pane.
   b. Click ⬅ or double-click the identity.
6. When you are finished adding and removing members, click **OK**.

Note: If you add or remove a user, the change takes effect the next time that this user logs on. If the user is currently logged on, his or her previous memberships continue to apply.

Create a New Custom Group

Create custom groups to give members similar permissions.

1. On the **Users** page in SAS Environment Manager, select **Custom groups** from the drop-down list in the toolbar.
2. Click ✱ in the toolbar.
3. In the New Custom Group window, enter a unique name and ID for the group. You can also enter a description.
   - Do not assign a custom group the ID of sasapp. SAS Viya reserves the group identifier `sasapp` for internal use by services.
   - Do not assign a custom group the ID of CASHostAccountRequired. CASHostAccountRequired is a reserved custom group name.
   - Do not use an apostrophe (') in a custom group ID. The use of an apostrophe (') interferes with the use of that group's identity on the **Users** page in SAS Environment Manager as well as accessing that group's identity when working with authorization.
   - Create an ID that is easily recognizable. For example, for the group "Report Testers", you could use "ReportTesters" as the ID.
4. Click **Save**.

TIP You can also create a custom group by copying a custom group. To do so, click the existing group (or custom group) and select ⧉. Then you can edit the properties and members of the new custom group as needed.

Edit a Custom Group's Basic Properties

1 On the **Users** page in SAS Environment Manager, select **Custom groups** from the drop-down list in the toolbar.

2 In the left pane, click the name of the group whose properties you want to edit.

3 In the basic properties section of the right pane, click ✎.

4 In the Edit Custom Group window, enter your changes to the name or description.

Note: You cannot edit the ID of a custom group.

5 Click **Save**.

Delete a Custom Group

1 On the **Users** page in SAS Environment Manager, select **Custom groups** from the drop-down list in the toolbar.

2 Click the custom group that you want to delete. The group's properties are displayed in the right pane.

3 Click 🗑, and then click **Delete** in the confirmation window.

Manage CAS Role Memberships

For each CAS server, be sure to designate at least one user (other than the server's process owner) to the Superuser role. In the initial deployment, users that you add to the SAS Administrators predefined custom group have membership in the Superusers role. If you want to designate a user to the role without providing the extra privileges of SAS Administrators, follow these instructions.

Manage Direct Membership in the CAS Superuser Role

1 In the applications menu (≡), select **Administration** ⇨ **Manage Environment**. In the navigation bar, select 🖥.

2 Right-click a CAS server, and select **Assume the Superuser role**.

3 Right-click the server again, and select **Settings**.

4 In the Superuser Role Membership section of the Server Settings window, click ✎.

5 To add a member, do the following in the Select Identities window:

a In the left pane, select **Users**.

b In the left pane, click the name of a user. The user's properties are displayed in the far right pane.

c Click ➕.

6 To remove a member, do the following in the Select Identities window:

   a In the **Select Identities** list, click the user that you want to remove. The identity’s properties are displayed in the right pane.

   Note: You cannot change or remove the account that starts the server.

   b Click ⬅.

7 Click **OK**.

8 Click **Relinquish** in the status bar to relinquish the Superuser role.

Assume the Superuser Role

In SAS Environment Manager, you become a Superuser only after you explicitly assume that role. For example, you might assume the role to troubleshoot and resolve an access issue or to manage format libraries. To assume the Superuser role:

1 In the applications menu (≡), select **Administration** ⇨ **Manage Environment**. In the navigation bar, select 🖥.

2 In the list of servers, right-click the name of the server for which you want to assume the role, and select **Assume the Superuser role**.

   The status message reminds you that you have assumed the role.

3 After you perform the task that required the role, click **Relinquish** in the status bar.

Note: Use the Superuser role only when it is required for a specific task. Be sure to relinquish the role when you are finished.

Manage Profile Pictures (Avatars)

1 In the applications menu (≡), select **Administration** ⇨ **Manage Environment**. In the navigation bar, select 👥.

2 On the **Users** page, select **Users**, **Groups**, or **Custom groups** from the drop-down list in the toolbar.

3 On the **Users** page, enter a string in the **Filter** field to search for identities within the category that you selected (Users, Groups, or Custom groups). To restore the complete list of identities, clear the **Filter** field.

4 Click an identity in the left pane to see its properties in the right pane. The profile picture appears below the identity name.

5 To add a profile picture:

   a Click on the profile picture.

   b From the Edit Profile Picture window, select **Choose Picture** from the drop-down list.

   c Navigate to the image that you want to use for the profile picture, and click **Open**.

6 To remove a profile picture:

   a Click on the profile picture.

   b From the Edit Profile Picture window, select **Remove** from the drop-down list.

   c In the Remove Profile Picture confirmation window, click **Remove**.

7 Click **Save**.

Note: When you copy a group or custom group, the profile picture is not copied to the new group. If desired, you should assign a new profile picture to the new group.

Reload Identities Cache

1 In the applications menu (≡), select **Administration** ⇨ **Manage Environment**. In the navigation bar, select 👥.

2 On the **Users** page, select ⋮ from the toolbar, and click **Reload Identities**.

3 A Reload Identities confirmation window appears with a warning that reloading the users and groups can take several minutes. Click **Yes** if you want to continue.

4 Monitor the identities log on page 191 for information about the reload. Messages about when the reload starts and completes are produced at the Info level.

The identities log file is located here:

Table A.1 Identities Log File

Linux	`/var/log/sas/viya/identities/default`
Windows	`\ProgramData\SAS\Viya\var\log\identities\default`

See Also

"Identities Synchronization" on page 366

Identity Management: How To (CAS Server Monitor)

Add or Remove CAS Role Members

Note: Starting with SAS Viya 3.4, CAS Server Monitor is available exclusively in programming-only deployments.

1 Sign in to CAS Server Monitor with an account that is already a CAS (Superuser).

2 In the left navigation bar, select 🔧.

3 On the Configuration page, select the **Administrators** tab.

4 To add a member:

   a Click **Add**.

   Note: If the **Add** button is not present, you are not signed in as a CAS administrator (Superuser).

   b In the Add Administrator window, enter a user or group name, select the appropriate identity type, and select the **CAS** or **Data** radio button.

> **TIP** The user and group names that you enter are not validated. You can enter any user or group name from your identity provider.

  c Click **OK** to save your changes.

5 To change a role assignment:

  a Click ⋮ in the appropriate row, and select **Modify**.

  Note: You cannot change the assignment for the account that starts the server.

  b In the Edit Administrator window, select **Data** or **CAS**, and click **OK**.

6 To remove a role assignment, click ⋮ in the appropriate row, and select **Delete**.

  Note: You cannot remove the account that starts the server.

7 Under **Administrators**, review the results.

8 Verify that full administrative privileges are available when designated users sign in to CAS Server Monitor. For example, any user who sees the **Add** button on the **Administrators** tab is a CAS administrator (Superuser).

See Also

CAS Server Roles on page 362

Identity Management: Access to Functionality

Overview

This section is about access to applications, features, services, and service endpoints.

Initially, access to functionality is distributed as follows:

SAS Administrators group	Provides access to all applications and features.
Other predefined groups	Provide access to certain specialized applications and features.
Authenticated Users (includes anyone who signs in)	Provides access to most applications and features.

If the initial distribution is appropriate, the only task is to assign users who need specialized or administrative access to the appropriate predefined groups.

If the initial distribution is not appropriate, you must expand or reduce access by working with authorization rules on page 425 as follows:

- The basic approach is to make only limited changes to only the documented rules.
- The advanced approach can include making broader changes to the documented rules, modifying undocumented rules, and adding new rules.

  CAUTION! Managing access to functionality can be an extremely complex task. Use the advanced approach only if you have a thorough understanding of target URIs on page 425, the functionality that you want to limit, the effect of each permission, and the interactions with any related rules. Make sure you have a current backup before you begin. Test your changes to make sure they do not have unintended effects.

Basic Approach: Planning

Before making changes, make a plan.

- Determine how many levels or categories of access you need.
- Decide which object URIs fit each level or category. See documented rules.
- Identify or create a group for each access level or category.

Here are tips to help with planning:

- Make appropriate use of the Authenticated Users principal. Unless it is unacceptable for all authenticated users to have at least the lowest level of access, leave Authenticated Users as the principal in the rule (or rules) for the lowest level of access. Here are examples:
  - All users are implicitly members of Authenticated Users, so they can access SAS Report Viewer. If this is acceptable, you can leave the predefined grant of /SASReportViewer/** assigned to Authenticated Users. Users can access the report viewer without any administrative intervention.
  - All users are implicitly members of Authenticated Users, so they can access SAS Visual Analytics. If this is not acceptable, change the principal in the relevant grant rule (/SASVisualAnalytics/**) from Authenticated Users to a designated group.
  - If you want to hide other applications from Authenticated Users, you must change the principal in each relevant rule, so that access is no longer granted to Authenticated Users.

    CAUTION! **`Never prohibit Authenticated Users. Prohibiting Authenticated Users blocks access for all authenticated users. That block has absolute precedence. It cannot be mitigated by more specific grants.`** Instead, make sure that Authenticated Users is not granted access. Any access that is not granted is implicitly denied.
- Take advantage of the group structure. To establish cumulative levels of access, make each higher-privilege group a member of the next-most-privileged group.

  For example, if your intent is to make applications menu items available as follows:

Menu Item	Report Viewer Group	Visual Analytics Group	Model Studio Group
View Reports (SAS Report Viewer)	✓	✓	✓
Explore and Visualize Data (SAS Visual Analytics)		✓	✓
Build Models (Model Studio)			✓

Then specify principals in authorization rules as follows:

Principal	Rule’s Target objectURI
Report Viewer group	/SASReportViewers/**
Visual Analytics group	/SASVisualAnalytics/**

Principal	Rule’s Target objectURI
Model Studio group	/ModelStudio/**

Or, if it is acceptable for all authenticated users to access SAS Report Viewer, omit the Report Viewers group.

The following figure depicts the membership structure for each approach:

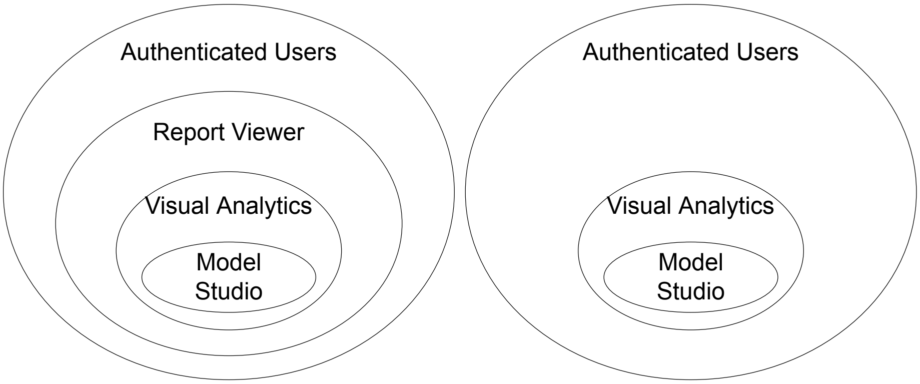

- Be aware that modifying access to applications does not affect access to underlying services. For example, a user who cannot access SAS Environment Manager might still be able to access the folders service through another interface.
- To grant the same access to two distinct groups, make a copy of the original rule. Specify one group as the principal in the original rule and the other group as the principal in the new rule.
- Avoid use of prohibit rules. Instead, use selective grants to provide selective access.
- You do not have to add rules that grant access to the SAS Administrators group. That group has a universal grant.

Example of the Basic Approach to Modifying Access to Functionality

The basic approach to modifying access to functionality involves making limited changes to a documented rule. In this example, you reduce the availability of the SAS Visual Analytics web application (Explore and Visualize) so that it is available to only members of a designated group. The process involves creating or identifying an appropriate group and making that group the principal in the relevant rule.

Here are instructions:

1 In the documented rules list, find the URI for the rule that affects the functionality that you want to work with. In this example, you want to limit the ability to access SAS Visual Analytics, so the URI that you will be working with is /SASVisualAnalytics/**.

2 On the Users page, create or identify an appropriate group. Add the users for whom you want to grant the ability to access SAS Visual Analytics.

3 On the Rules page, locate and edit the rule (principal type, principal, description).

   a In the applications menu (☰), select **Administration** ⇨ **Manage Environment**. In the navigation bar, select ▤.

   b Enter /SASVisualAnalytics/** in the field under **Object URI**, and click **Apply**.

c Select the rule, and click ✎.

d In the Edit Rule window, select `group` in the **Principal Type** field. In the **Principal** field, select the group that you just selected.

e In the **Description** field, update the description for the group for which you provided the ability to access SAS Visual Analytics.

Click **Save**.

4 Verify that users who are in the group are able to access SAS Visual Analytics, and that users who are not in the group are not able to access SAS Visual Analytics.

If a user has more access than is expected, make sure that there are no other rules that grant the same functionality (the same object URI) to another principal.

> **TIP** Add yourself as a member of the new group, and sign in without assuming your membership in the SAS Administrator's group.

Rules Reference

Here are rules for use with the basic approach to managing access to functionality. The following access rules are grouped by type of SAS Viya functionality. Except where otherwise specified, the functionality is initially granted to Authenticated Users. In most cases, the relevant permission is Read. For details and exceptions, see the rule properties in SAS Environment Manager.

Access Applications

Table A.2 Object URIs for Accessing Applications

Application or Function	Object URI
SAS Environment Manager (all functionality)[1]	/SASEnvironmentManager/**
SAS Environment Manager (gatekeeper)	/SASEnvironmentManager/
SAS Graph Builder	/SASGraphBuilder/**
SAS Visual Analytics App	/SASMobileBI/**
SAS Report Viewer	/SASReportViewer/**
SAS Theme Designer[2]	/SASThemeDesigner/**
SAS Visual Analytics	/SASVisualAnalytics/**
SAS Data Explorer	/SASDataExplorer/**
Authorization dialog from SAS Drive[1]	/authorizationDialog
SAS Drive	/SASDrive/**
Model Studio	/ModelStudio/**

[1] Initially granted to SAS Administrators.

[2] Initially granted to SAS Application Administrators.

Table A.3 Object URIs for Page-Level Rules in SAS Environment Manager

Application or Function	Object URI	Access Rule Exists
Dashboard page in SAS Environment Manager	/SASEnvironmentManager/dashboard	✓
Data page in SAS Environment Manager	/SASEnvironmentManager/data	✓
Servers page in SAS Environment Manager	/SASEnvironmentManager/servers	✓
Content page in SAS Environment Manager	/SASEnvironmentManager/content	✓
Users page in SAS Environment Manager	/SASEnvironmentManager/identities	
Licensed Products page in SAS Environment Manager(only present in multi-tenant environments)	/SASEnvironmentManager/licenses	
Tenants page in SAS Environment Manager	/SASEnvironmentManager/tenants	
Configuration page in SAS Environment Manager	/SASEnvironmentManager/configuration	
Contexts page in SAS Environment Manager	/SASEnvironmentManager/contexts	
User Defined Formats page in SAS Environment Manager	/SASEnvironmentManager/udf	
Logs page in SAS Environment Manager	/SASEnvironmentManager/logs	
Machines page in SAS Environment Manager	/SASEnvironmentManager/machines	
Jobs page in SAS Environment Manager	/SASEnvironmentManager/jobs	✓
Domains page in SAS Environment Manager	/SASEnvironmentManager/domains	
My Credentials page in SAS Environment Manager	/SASEnvironmentManager/credentials	✓
Mobile Devices page in SAS Environment Manager	/SASEnvironmentManager/devices	

Application or Function	Object URI	Access Rule Exists
Rules page in SAS Environment Manager	/SASEnvironmentManager/rules	
Publishing Destinations page in SAS Environment Manager	/SASEnvironmentManager/destinations	

Interact with Reports and Data

For additional information about how to adjust rules associated with accessing data that is imported from social media, see "Controlling Access to Features" in *SAS Data Explorer: User's Guide*.

Action	Object URI
Import reports.	/importVASpk/**
Upload data files (via casManagement service).	/casManagement/servers/*/caslibs/*/tables
Access the **Import** window.	/casManagement_capabilities/importData
Export reports as PDF.	/reportRenderer/reports/**
Export data from reports.	/reportData_capabilities/exportData
Export detail data from reports.	/reportData_capabilities/exportDetailData
Create jobs to obtain report images (for example, thumbnails and section images).	/reportImages/jobs/**
Export report images from SAS Visual Analytics and web or mobile viewers.	/SASVisualAnalyticsCommon_capabilities/exportImage
Email or share report images from SAS Visual Analytics and web or mobile viewers.	/SASVisualAnalyticsCommon_capabilities/shareReport
Subscribe to report alerts.	/reportAlerts/**
Evaluate text templates.	/reportImages/textTemplateOutput
Manage report states for reports.	/reports/reports/*/states
Access imported Facebook data.	/webDataAccess_capabilities/facebookImport
Access imported Google Drive data.	/webDataAccess_capabilities/googledriveImport
Access imported Google Analytics data.	/webDataAccess_capabilities/googleanalyticsImport
Access imported YouTube data.	/webDataAccess_capabilities/youtubeImport
Access imported Twitter data.	/webDataAccess_capabilities/twitterImport
Manage comments.	/comments/**

Action	Object URI
Create models in SAS Visual Analytics.	/SASVisualAnalytics_capabilities/buildAnalyticalModel

Manage Jobs

Action	Object URI
Schedule jobs.[1]	/jobExecution/jobRequests/* /jobExecution/jobRequests/*/
Edit scheduled jobs.	/scheduler/jobs/**
Monitor jobs.	/jobExecution/jobs/**

[1] Both of the object URIs are required.

Manage Geo

Action	Object URI
Add a custom map provider.[1]	/maps/providers
Manage custom map providers (update and delete).[1]	/maps/providers/*
Use the ESRI service.	/webDataAccess/esri/user/token

[1] Initially granted to SAS Administrators.

Manage Mobile

The following table lists only the object URIs that are specific to mobile devices. You might need to work with other object URIs to manage mobile devices. For additional information about rules that affect the SAS Visual Analytics App and the SAS Software Development Kits (SDKs), see Table 17.45 on page 278.

Action	Object URI
Cache mobile report data.	/SASMobileBI_capabilities/cacheMobileReportData
Exempt from offline time-out.	/SASMobileBI_capabilities/exemptFromOfflineTimeLimit
Exempt from requirement to enter passcode.	/SASMobileBI_capabilities/exemptFromPasscodeRequirements
Manage mobile device blacklist, whitelist, and device access history.[1]	/deviceManagement_capabilities/manageMobileDevices
Enable natural language processing.	/reportViewerNaturalLanguageUnderstanding/interpretations

Action	Object URI
Render reports with web content.	/SASMobileBI_capabilities/allowWebContent

[1] Initially granted to SAS Administrators.

Personal Folders and Preferences

Item	Object URI
Set preferences.	/preferences/preferences/@currentUser/**
Access personal history folder.	/folders/folders/@myHistory
Manage personal favorites folder.	/folders/folders/@myFavorites

Identity Management: CAS Roles

Superusers have permission-exempt access to CAS (with the exception of access to data) and are exempt from all CAS authorization requirements.

In SAS Environment Manager, the Superuser role is never initially or automatically assumed. If you are a member of a CAS server's Superuser role, you can become a Superuser by explicitly assuming the role for that server. For example, you might assume the role to troubleshoot and resolve an access issue. After the issue is resolved, you relinquish the role.

The account that starts a CAS server is automatically assigned to that server's Superuser role.

Note: The following built-ins actions for SAS Cloud Analytics Services require a user ID that can assume the Superuser role:

- addNode
- installActionSet
- refreshLicense
- removeNode
- shutdown

For more information about the built-ins actions for SAS Cloud Analytics Services, see Builtins Action Set: Details.

Role	Description	Initial Members
Superuser	Provides permission-exempt access to a CAS server. Only a Superuser can perform the following tasks: - Stop the server. - Add and remove nodes. - Manage role membership. - See and manage the paths list. The account under which a CAS server runs is an implicit member of that server's Superuser role. Make sure each CAS server has at least one other designated Superuser. **Note:** By default, the users that are assigned this role have permission-exempt access to metadata. However, they do not have permission-exempt access to data (CAS libraries). To give users with this role permission-exempt access to data, you must modify access controls to explicitly grant them access.	SAS Administrators Process owner for the server Backup Administrator (sas.deploymentBackup) Report Distribution Service administrator account (sas.reportDistribution) Report Images Service administrator account (sas.reportImages) Scheduler Service administrator account (sas.scheduler) Report Alerts Service administrator account (sas.reportAlertsEval) VSD Service administrator account (sas.svi-vsd-service)
Data	Assign members to this role only if you have users who should have permission-exempt access to data but should not be able to perform all administrative tasks. Not all interfaces support the Data role. **Note:** By default, the users that are assigned this role have permission-exempt access to metadata. However, they do not have permission-exempt access to data (CAS libraries). To give users with this role permission-exempt access to data, you must modify access controls to explicitly grant them access.	None
Action	Do not use this role. Not all interfaces support the Action role.	None

Note: The Data role provides a subset of the abilities of the Superuser role. You cannot be a member of both the Superuser role and the Data role in the same session.

See Also

- Manage CAS Role Memberships in SAS Environment Manager
- Add or Remove CAS Role Members in CAS Server Monitor

Identity Management: Reference

Initial Users

sasboot Account

The sasboot account is an internal user account that is created during the deployment process. For details, see the deployment guide for Linux, or the deployment guide for Windows.

Note:

The sasboot account exists only in a full deployment. A full deployment includes all of the software to which you are entitled, whereas a programming-only deployment excludes SAS Drive, SAS Environment Manager, and most graphical user interfaces, and most services.

Operating System Accounts

Some user and service accounts are required on the operating system during deployment. For details, see Linux accounts and Windows accounts.

Custom Groups

What Is a Custom Group?

A custom group is a group that exists in SAS Viya but not in your identity provider. These groups are persisted in a SAS database.

Your deployment includes a set of predefined custom groups. can also create your own custom groups. This feature is useful for creating new groups of SAS users if you do not want to (or do not have permission to) create groups in your identity provider.

Note: These groups are not supplied in a programming-only deployment.

Assumable Custom Groups

The SAS Administrators group is a predefined custom group. This group is *assumable*. When a user in this group signs in to SAS Viya, a prompt appears asking `Do you want to opt in to all of your assumable groups`?

If the user selects **Yes**, the user gets the extra permissions that are associated with the assumable group. If the user selects **No**, the user does not get the extra permissions. The selection remains in effect until the user signs out.

As a best practice, users should select **Yes** only when they need to perform tasks that require the extra permissions

Predefined Custom Groups

Certain custom groups are provided with your deployment. These groups provide an easy way to give users and groups access to the appropriate data, content, or functionality.

- The custom groups effectively implement a role within SAS Viya. The members of the custom groups have access to privileges associated with the role. Search on the Rules page by the group name to see all privileges that are associated with the role.
- The predefined groups below are a part of a deployment that contains SAS Visual Analytics, SAS Visual Statistics, and SAS Visual Data Mining and Machine Learning. Some products and solutions have additional predefined groups. See the documentation for these products and solutions for information about other predefined groups.

  For example, if you have SAS Data Studio, then you have a predefined group called Data Builders. This group is not assumable, and there are no initial members.

The custom groups are as follows:

SAS Administrators

Have access to the following:

- all functionality that is controllable through authorization rules.
- all folders and all objects that the folders contain (for example, plans and reports).

Is an assumable group.

Members can assume the CAS Superuser role.

Note: Access to data (caslibs) is not included. For example, users in this group can create, run, and view reports only if they have explicitly been granted access to the underlying data.

Esri Users

Can access Esri systems for geo map access.

Is not an assumable group.

Has no initial members.

Note: Esri requires that organizations pay for tokens to use the Esri geographic mapping services. You can add a user or group of users to the Esri Users group to control who has access to these tokens. Therefore, you can control the cost of using Esri geographic services.

Application Administrators

Can access selected administrative functions within applications.

Is not an assumable group.

Has no initial members.

Note: An additional custom group is predefined, but not created. If you create a group with ID: *CASHostAccountRequired*, members of this group automatically run their CAS sessions under their own host account. By default, CAS sessions run using the `cas` account. For more information, see CASHostAccountRequired custom group on page 365.

The CASHostAccountRequired Custom Group

Note: The CASHostAccountRequired custom group is not applicable in a Windows deployment.

The CASHostAccountRequired custom group is predefined, but not created. If you create a group with ID: CASHostAccountRequired, members of this group automatically run their CAS sessions under their own host account. By default CAS sessions run using the `cas` account.

Therefore, members of this group must have host accounts.

Note: If a user is a member of the CASHostAccountRequired custom group, but has no host account, then SAS Environment Manager cannot access information about the CAS Server. You might observe the following behavior:

- From SAS Environment Manager, the CAS server appears to be down even though it is not. No libraries or tables are displayed.
- From SAS Data Studio, you receive a `connection refused` or `access denied` error message when you attempt to select a CAS server.

When you modify the membership of this group, the users that have been added or removed must log off from their sessions before the changes can take effect.

If a user has previously created SASHDAT files and is then added to the CASHostAccountRequired custom group, the user can continue to work with data in memory. However, if certain triggering events occur, such as a CAS server restart, the same user can no longer see the SASHDAT files because the location of these files is different for members of this group. Users in this situation should copy the SASHDAT files from the default location to the host CAS user path.

The original default location is **`/opt/sas/viya/config/data/cas/default/casuserlibraries/`** ***`username`*** where user name is the user's host account.

The host CAS user location is **`~/casuser/`** where the ~ represents the users home directory.

Note that files should be copied in the opposite direction for users that are removed from the CASHostAccountRequired group.

See Also

"Manage Custom Groups" on page 350

Identities

Identity Filtering

When configuring the connection to your identity provider, you should specify a filter to limit the identities that SAS Viya returns. For example, you can create a filter to exclude identities whose accounts are disabled or expired, or to exclude objects that represent computer resources rather than actual users or groups. You can modify this filter at any time.

If you have a large number of users, using a filter can improve performance and reduce memory requirements. In addition, user management tasks can be performed more efficiently if only relevant identities are listed in SAS Environment Manager.

A default filter is provided for sites that use Active Directory. If you use another identity provider such as openLDAP, then you might need to modify the default filter. For more information about the default filter, see "Identities Service" on page 100.

Note: Identity filtering does not apply in a programming-only deployment.

Identity Caching

Identity caching is available for enhanced performance. Search requests go to the cache, reducing the number of direct requests to the identity provider. You can configure the cache refresh interval, and enable or disable the cache. The cache is enabled by default. See "Identities Service" on page 100.

Note: Identity caching does not apply in a programming-only deployment.

Identities Synchronization

Information about LDAP identities is available in SAS Environment Manager, and is synchronized with the SAS Infrastructure Data Server (PostgreSQL) periodically. The amount of time between each synchronization is determined by a configuration option on page 101, which is set to 12 hours by default.

If you want to manually synchronize the identities at any time, you can reload the identities on page 354. Note that reloading the set of users and groups can take several minutes to complete.

See Also

- "Identities Service" on page 100

Identity Management: Interfaces

In the following table, the shaded part of each circle is an approximation of the amount of user management functionality that a particular interface exposes. The shading indicates relative coverage. The shading does not indicate alignment of functional coverage across interfaces.

●	SAS Environment Manager	A graphical enterprise web application. See "Identity Management: How To (SAS Environment Manager)".
◔	CAS Server Monitor	A graphical web application that is embedded in the CAS server. See "Identity Management: How To (CAS Server Monitor)" on page 354.
◔	Access Control action set	A programmatic interface for SAS (the CAS procedure), Python, R, and Lua. See Access Control Action Set.
◕	Command-line interface	A simple scriptable interface that provides commands for managing identities. See "CLI Examples: Identities" on page 720.

Identity Management: Troubleshooting

Cannot Sign In to SAS Studio

- Ensure that the user's account is known to the host of the SAS Studio web application. See Authentication on page 289.
- Examine the object spawner log. See Logging on page 191.
- If users cannot make a secure connection, see Encryption in SAS Viya: Data in Motion.

Cannot Access Cloud Analytic Services

- If the user cannot start a CAS session, ensure that the user's account meets all applicable requirements. See Authentication on page 289.
- If an error message in the CAS log states that the user "failed mid-tier authentication", the user's credentials are not valid for your direct LDAP provider. See the discussion of dual authentication in Authentication on page 289.
- Ensure that users have a host account before adding them to the CASHostAccountRequired group. A member of the CASHostAccountRequired group without a host account cannot start the necessary CAS session.

Cannot Sign In to CAS Server Monitor

Note: Starting with SAS Viya 3.4, CAS Server Monitor is available exclusively in programming-only deployments.

- Ensure that the user's account meets all applicable requirements. See Authentication on page 289.
- If an error message in the CAS log states that the user "failed mid-tier authentication", the user's credentials are not valid for your direct LDAP provider. See the discussion of dual authentication in Authentication on page 289.
- If users cannot make a secure connection, see Encryption in SAS Viya: Data in Motion.

Cannot View Users and Group Members

If you receive the following error while viewing users, groups, or their memberships from SAS Environment Manager or any other client, then a referral might have been encountered. SAS Viya does not process LDAP referrals.

Here is an example of this error message:

```
Load Users
An error occurred loading the list of users.
exception:
org.springframework.ldap.PartialResultException
Caused by: javax.naming.PartialResultException: Unprocessed Continuation
Reference(s); remaining name 'DC=COMPANY,DC=COM'
```

This occurs because LDAP is initialized based only on what the Identities service itself configures. Therefore, any environment variables that are set will not be processed. Connecting to the global catalog might be a viable solution.

Cannot Access Esri Geographic Mapping Resources

Ensure that the user is a member of the Esri Users group. Users that are members of the Esri Users group have access to tokens for which there is a fee. See "Esri Users" on page 365.

Cannot Retrieve List of Users or Groups

If the following error occurs while attempting to retrieve a list of users or groups that are defined in your environment, then this is due to a failed LDAP search by the Identities service:

```
[LDAP: error code 12- Unavailable Critical Extension]
```

The Identities service attempted an LDAP search for a collection of users or groups, but the request failed because the LDAP server does not support paged queries.

To resolve this, follow these steps to change the value of the **`sas.identities.providers.ldap.pagedResults`** configuration property:

1. Log on to SAS Environment Manager as an administrator.
2. Navigate to the Configuration page. From the **View** drop-down list, select **Definitions**. In the **Filter** field, enter *sas.identities.providers.ldap*.
3. From the **Identities service** drop-down list on the right pane, click ✎. Change the **`pagedResults`** property to **`Off`**.

4 Click **Save**.

Cannot Update User Information

Any service or application that uses the Identities service pulls the associated user information from the LDAP server directly. Therefore, user information such as phone number, work address, and email address cannot be updated in SAS Viya, but must be updated in LDAP.

For example, to specify a different email address to receive SAS Visual Analytics alerts, the email address field must be updated directly in LDAP.

Cannot Log In to SAS

Ensure that no two users have the same email address in LDAP. Users might have problems logging in to SAS if another user has the same email address.

Membership in a sasapp LDAP, Custom, or Host Group Is Ignored

SAS Viya reserves the group identifier **`sasapp`** for internal use by services. Only services are members of the privileged internal group **`sasapp`**. If you also have a **`sasapp`** host group, LDAP group, or custom group, unintended results can occur. Descriptive information (such as the Authorization window and the Users page in SAS Environment Manager) reflects membership in the group. However, actual access does not reflect membership in the group. Here are details:

- When a user who is a member of a **`sasapp`** LDAP or custom group signs in to SAS Viya, SAS Logon Manager discards the user’s **`sasapp`** membership information, excluding it from the user’s OAuth token. By discarding that membership information, SAS Logon Manager ensures that the privileges of the **`sasapp`** internal group are not made available to users.
- If you specify the **`sasapp`** group as the principal in a general authorization rule, that rule affects only the **`sasapp`** internal group.
- When a user who is a member of a **`sasapp`** host group authenticates to CAS, CAS alters its copy of the user’s membership information, replacing the **`sasapp`** group name with the group ID. By altering that membership information, CAS ensures that the privileges of the internal group **`sasapp`** are not made available to users.
- If you specify the **`sasapp`** group as the principal in a CAS access control, that access control affects only the **`sasapp`** internal group.

  Note: SAS Viya does not currently prevent the creation of a **`sasapp`** custom group. SAS Viya cannot prevent the creation of a **`sasapp`** host or LDAP group.

Identity Management: Guidelines

The following basic guidelines contribute to simplicity and security:

- Limit membership in administrative roles and groups.
- Assume administrative group memberships only when you need to perform tasks that require the extra permissions.
- Assume a CAS administrative role only when you need to perform tasks that require the extra permissions, and relinquish the role when you are finished.

- If you delete a custom group, any custom rules that you created still exist. Manually delete such rules.
- As you plan your group structure, remember that you can use a group for either or both of these purposes:
  - To make shared resources available to multiple users. For example, you might use a group as the principal in an authorization rule, or you might store shared credentials for a group.
  - As a parent to other groups. For example, if groupA and groupB should have identical access to multiple resources, you might assign both groups to a parent group, and grant access to the parent group.

Note: When you design a group structure, consider both clarity (Will others be able to interpret the structure?) and conciseness (Is the structure as minimal as possible?). In a complex authorization model, you might prioritize clarity, using more than the strict minimal number of groups, and giving each group a name that describes its purpose.

22

Orientation to Authorization

Two Authorization Systems

Introduction

Authorization is the aspect of security that determines which resources are available to which users. The SAS Viya authorization layer consists of two authorization systems:

- Cloud Analytic Services (CAS) authorization system
- general authorization system

Each system uses a distinct model to protect a distinct class of resources.

Similarities

- Both systems can share the same identity provider.
- Both systems implicitly disallow any access that is not granted.
- Both systems can be administered using SAS Environment Manager or a command-line interface.

Differences

Characteristic	CAS Authorization System	General Authorization System
Basis:	DBMS-style access control.	Attribute-based access control.
Targets:	CAS objects, such as caslibs and tables.	Most other objects, such as folders and reports.
Inheritance:	Through a hierarchy of objects (for example, from a caslib to its tables).	Through a hierarchy of containers (for example, from a folder to its members).
Precedence:	By object hierarchy (closest wins), then by identity type (user wins), and then by type of setting (denial wins).	By type of setting (Prohibit always wins).
Row-level access:	You can attach a filter to a grant of the Select permission on a table.	(Not applicable).
Conditional access:	(Not applicable).	You can attach a Boolean expression to any rule.
Highest privileges:	An assumable role (Superuser) is exempt from authorization requirements throughout a CAS server, except for data access requests.	An assumable group (SAS Administrators) is granted broad access throughout the general authorization system.*
Applicable to:	All deployments of SAS Viya.	Full deployments of SAS Viya.

* The SAS Administrators group is not unrestricted (exempt from authorization requirements). Access is provided by a predefined rule.

Influences

In the CAS authorization system, memberships, inheritance, and row-level filters can influence access.

In the general authorization system, information about the requesting user, the target resource, and the environment can influence access. Each access request has a context that includes environmental data such as time and device type. Environmental constraints can be incorporated using conditions.

Key Terms

Term	Definition
Access control or rule	A composite of authorization elements. CAS example: An access control grants the ReadInfo permission to groupA on caslibA. General example: A rule grants the Add permission to groupA on folderA.
Setting	An indication of whether (and to what extent) access is provided. CAS values: Grant, Row-Level Grant, Deny General values: Grant, Conditional Grant, Prohibit, Conditional Prohibit

Term	Definition
Permission	A type of access. CAS values: ReadInfo, Select, LimitedPromote, Promote, CreateTable, DropTable, DeleteSource, Insert, Update, Delete, AlterTable, AlterCaslib, ManageAccess General values: Create, Read, Update, Delete, Secure, Add, Remove
Principal	The user, group, or construct to which an access control or rule is assigned. Examples: UserA, GroupA, Authenticated Users
Target	A resource or set of resources. CAS examples: tableA, caslibA General examples: folderA, reportA
Condition	In a conditional rule, the constraint expression. General example: `currentUser() == #preferenceOwner`
Filter	In a row-level grant, the constraint expression. CAS examples: `User='SUB::SAS.Userid'`, `sales>1000`
Effective access	A context-neutral description of the net result of all relevant access controls or rules. Effective access does not incorporate evaluation of conditions. CAS values: Authorized, Not Authorized, Row-Level General values: Authorized, Not Authorized, Conditional
Access outcome	The authorization decision for a specific access request. CAS values: Authorized, Not Authorized, Row-Level Authorization General values: Authorized, Not Authorized

Demonstration

In this demonstration, you enable a set of users to access a caslib and a folder.

- Users can read and write data in the caslib and objects in the folder.
- Administrators can manage the caslib and the folder.

Note: This demonstration assumes that you are a member of the SAS Administrators group, and that the initial settings and memberships for that group are in place.

1. Sign in to SAS, and opt in to your assumable groups.
2. To represent the set of users, create a group.
   - a. In the vertical navigation bar for SAS Environment Manager, click 👥.

     > **TIP** If you are not already in SAS Environment Manager, select **Manage Environment** from the applications menu (≡).

   - b. Make sure that **Custom groups** is selected in the **View** drop-down list.
   - c. Click ✱.

d In the New Custom Group window, enter `groupA` as the name and as the ID. Click **Save**.

e In the right pane, click ✎ on the **Members** tab.

f In the Edit Members window, assign members by moving them to the **Selected Identities** pane. Click **OK**.

Note: For this demonstration, adding members is optional.

3 Give groupA appropriate access to a new global caslib.

a In the vertical navigation bar, select ▦.

b On the **Data** page, select the **Data Sources** tab.

c On the **Data Sources** tab, click ⚡.

d In the Connection Settings window, create a path-based global caslib that has the following properties:

Property	Instruction
Name	Enter `caslibA`.
Server	Select a CAS server.
Type	Select **File system**.
Source type	Select **PATH**.
Persist this connection...	Select the check box.
Path	Enter a path that is accessible by and relative to the selected CAS server.

TIP If you need more information, see "Connecting to Remote File Systems" in *SAS Data Explorer: User's Guide*.

e Click **Save**.

f On the **Data Sources** tab, right-click **caslibA**, and select **Edit authorization**.

TIP If the **Edit authorization** action is not available, make sure the caslib that you created has global scope. See Caslibs on the Data Sources Tab and Import Tab in *SAS Data Explorer: User's Guide*.

g In the Edit Authorization window, click 👤, and add groupA to the window.

TIP In the Add Identities window, make sure **Custom groups** is selected. Move groupA to the **Selected Identities** list. Click **OK**.

h In the Edit Authorization window, notice that groupA has an effective access value of ⊘ (Not Authorized) for all permissions. In the **Access Level** column, adjust groupA's gauge to the **Write** access level.

i Click **Save**.

4 Give groupA appropriate access to a new top-level folder.

a In the vertical navigation bar for SAS Environment Manager, click 📄.

b Make sure you are at the top level of the folder structure, and then click 📁 to create a new top-level folder. Enter the name `folderA`, and press the Enter key.

c Grant access to folderA and conveyed access to folderA’s members.

i Select **folderA**, right-click, and select **Edit authorization**.

ii In the Edit Authorization window, click 👤, and add groupA.

> **TIP** In the Add Identities window, make sure **Custom groups** is selected in the **Filter by** drop-down list. Move groupA to the **Selected Identities** list. Click **OK**.

iii In folderA’s Edit Authorization window, click the effective access icon for groupA’s Read permission. In the pop-up window, select **Grant** as the direct setting. Repeat that process for groupA’s Add and Remove permissions.

These settings target folderA’s object URI and affect access to the folder.

iv In folderA’s Edit Authorization window, click the effective access icon in groupA’s **Read (convey)** column. In the pop-up window, select **Grant** as the direct setting. Repeat that process in groupA’s **Update (convey)** and **Delete (convey)** columns.

These settings target folderA’s container URI and affect access to the folder’s members. See Inheritance in *SAS Viya Administration: General Authorization*.

v Click **Save**.

5 (Advanced) View direct rules that affect access to folderA.

a In the vertical navigation bar, click ☰.

b In the **Rules Filter** pane, under **Object URI**, select **URI** from the drop-down list.

c In the Choose an Item window, select **folderA**. Click **OK**.

d In the **Rules Filter** pane, click **Apply**.

e To ensure that you are seeing all available information, click ↻. Notice that there are two rules that target folderA’s object URI:

- The rule that grants all permissions to you was automatically generated because you added folderA as a top-level folder. Lower level folders do not have automatically generated rules.
- The first set of changes that you made in the Authorization window created the rule that grants the Read, Add, and Remove permissions to groupA.

Note: For groups and users, the **Principal** column on the **Rules** page contains IDs, not display names.

6 (Advanced) View direct rules that affect access that folderA conveys to its members.

a At the right edge of the table, click ⋮ , and select **Manage columns**.

b In the Manage Columns window, move **Container URI** to the **Displayed columns** pane.

c In the **Displayed columns** pane, select **Container URI**, and click ⤒. Click **OK**.

d At the top of the **Rules Filter** pane, click the **Reset all** link to clear all filters that are currently in effect.

e In the **Rules Filter** pane, under **Container URI**, select **URI** from the drop-down list.

f In the Choose a Location window, select **folderA**. Click **OK**.

g In the **Rules Filter** pane, click **Apply**.

h To ensure that you are seeing all available information, click ↻. Notice that there are two rules that target folderA’s container URI:

  - The rule that grants all permissions to you was automatically generated because you added folderA as a top-level folder. Notice that the generated rule targets both folderA’s object URI and folderA’s container URI.
  - The second set of changes that you made in the Authorization window created the rule that targets folderA’s container URI, granting the Delete, Read, and Update permissions to groupA. That rule provides conveyed access to the members of folderA.

Note: The **Rules** page does not display CAS access controls. You can use a command-line interface to view the direct access controls for a CAS object (such as a caslib or table). See CLI Examples: CAS Authorization in *SAS Viya Administration: Using the Command-Line Interfaces*.

See Also

- General Authorization
- Cloud Analytic Services Authorization

Impact of Assumable Memberships

Introduction

Most memberships are always in effect. For example, if UserA is a member of GroupA, that membership affects UserA all of the time. UserA cannot temporarily opt in or opt out of experiencing the effects of his membership in GroupA.

The most highly privileged memberships are assumable. Assumable memberships are in effect in only certain circumstances. Here are examples:

- In a programming interface or SAS Environment Manager, members of a CAS role can temporarily experience that role’s elevated privileges by assuming that role at any time.
- In most visual interfaces, members of the SAS Administrators group can temporarily experience that group’s elevated privileges by opting in to that group at sign-in time.

Effective Access

When you examine effective (net) access for a user who has assumable memberships, information about whether those memberships are currently in effect is, in most cases, unavailable.

- In general authorization, effective access information presumes that all assumable memberships are in effect.
- In CAS authorization, effective access information presumes that no assumable memberships are in effect.

Access Outcomes

When a user who has assumable memberships makes an access request, the outcome of that request is affected by whether those memberships are currently in effect.

Demonstration

This demonstration uses SAS Environment Manager to explore the availability of access that you get exclusively through your assumable memberships.

Note: This demonstration assumes that you are a member of the SAS Administrators group, and that the initial settings and memberships for that group are in place.

1 Examine the availability of your elevated privileges in the general authorization system.

   a If you are currently signed in to SAS, click your user name in the banner, and select **Sign out**.

   b In the Sign in to SAS window, click **Sign In**, and sign back in. In the Assumable Groups window, click **No**.

   c In the vertical navigation bar for SAS Environment Manager, notice that there is no 👥 icon.

   > **TIP** If you are not already in SAS Environment Manager, select **Manage Environment** from the applications menu (≡).

   d Sign out, and then sign in again. In the Assumable Groups window, click **Yes**.

   e In the vertical navigation bar for SAS Environment Manager, notice that additional items are present, including a 👥 icon.

2 Examine the availability of your elevated privileges in the CAS authorization system.

   a In the vertical navigation bar for SAS Environment Manager, click ▤.

   b On the **Servers** page, select a CAS server, right-click, and select **Settings**.

   c In the Server Settings window, notice that none of the tabs are editable.

   On the **Superuser Role Membership** tab, verify that you are an indirect member of the Superuser role through your membership in the SAS Administrators group. Close the window.

   d Right-click the CAS server again, and select **Assume the Superuser role**. Notice that a message at the top of the page indicates your elevated status.

   e Right-click the CAS server again, and select **Settings.** Notice that an additional tab is present (**Paths List**) and all of the tabs are editable. (They have an edit icon ✎ in their upper right corner.) Close the window.

   f Right-click the CAS server again, and select **Relinquish the Superuser role**.

3 Examine the relationship between your assumable memberships in the SAS Administrators group and the Superuser role.

   a Click your user name in the banner, and select **Sign out**.

   b Sign back in. In the Assumable Groups window, click **No**, so that your membership in the SAS Administrators group is not in effect.

   c Navigate back to the **Servers** page in SAS Environment Manager.

   d Select the CAS server that you used in the preceding steps, right-click, and select **Assume the Superuser role**. A message indicates that you cannot assume the Superuser role in your current session.

Note: Initially, your membership in the Superuser role is indirect, through the SAS Administrators group. If you opt out of your assumable membership in SAS Administrators, you do not experience any of the privileges that you obtain exclusively from that membership.

TIP If you want to always be able to assume the Superuser role, add yourself to that role in a way that does not involve the SAS Administrators group. For example, make yourself a direct member of the Superuser role.

See Also

Identity Management

Comparison to SAS 9 Authorization

About This Topic

This topic is intended for SAS Viya administrators who are already familiar with SAS 9 authorization.

Here are some possible uses for this topic:

- You can use this topic to transfer some of your SAS 9 knowledge to SAS Viya.
- You can use this topic as an input to the creation of your new SAS Viya authorization model.
- If you discover that access in SAS Viya differs from access in SAS 9, you can use this topic to help determine the causes of the differences.

Note: This topic provides conceptual information. It does not describe the scope or effects of using the promotion tools. For information about the promotion tools, see “Promotion from SAS 9: Reference” in *SAS Viya Administration: Promotion (Import and Export)*.

Considerations

Here are the primary differences that SAS 9 administrators must consider when they begin to manage access to SAS Viya resources:

CAUTION! For content, be aware of a radical difference in precedence. In general authorization, denials have absolute precedence and are referred to as Prohibit settings. Any relevant Prohibit setting prevents access, regardless of where it is set or who it is assigned to. For example, if you assign a Prohibit rule to Authenticated Users, you cannot provide access for any authenticated user, even if you explicitly grant access to a specific user. Avoid use of Prohibit rules. Instead, grant access selectively.

CAUTION! For content, be aware of a significant difference in inheritance. In general authorization, permissions that affect access to a folder are fully independent from permissions that affect access to the folder’s members. In SAS 9 authorization, only the WriteMemberMetadata permission has different effects on a folder and its members.

CAUTION! For data, be aware of a significant difference in inheritance. In CAS authorization, inheritance begins at the library (caslib) level. This differs from inheritance in SAS 9, which flows to libraries from folders.

CAUTION! For data, be aware of a minor difference in precedence. In CAS authorization, all non-implicit memberships have equal precedence, regardless of any nesting of groups. If your SAS 9 libraries (or tables) have conflicting direct access controls for nested groups, initial SAS Viya access to data might be more limited than in SAS 9.

Here are additional comparative details:

- In SAS Viya, there are no access control templates.

- In SAS Viya, the equivalents to the SAS 9 implicit groups are as follows:

	SAS Viya **Equivalent**	
SAS 9 Implicit Group	**General Authorization**	**CAS Authorization**
PUBLIC	Everyone	(none)*
SASUSERS	Authenticated Users	Authenticated Users (group *)

* There is no single corresponding construct in CAS. Together, Authenticated Users (group *) and Guest correspond to PUBLIC.

- In SAS Viya, identity-driven substitution for row-level security can be based on only user IDs and memberships. (Other forms of identity-driven substitution that are available in SAS 9 are not available in SAS Viya.)
- In SAS Viya, access to functionality is primarily driven by authorization rules that are granted to groups. The only roles in SAS Viya are the administrative roles for the CAS server. See Access to Functionality in *SAS Viya Administration: Identity Management*.

About Permission Mappings

If you use permission mappings to help translate authorization information between SAS 9 and SAS Viya, make sure those mappings fit your security goals and authorization model. Here are key points:

- Some permissions are not mapped, because not all SAS 9 permissions have an exact counterpart in SAS Viya.

  For example, SAS Viya does not have an exact counterpart for the CheckInMetadata permission. The mappings in this topic exclude SAS 9 permissions that do not map to a SAS Viya permission.
- Some mappings are approximate, because not all activities in SAS 9 have a precise counterpart in SAS Viya.

  For example, inserting rows in a metadata-bound table directly affects that physical table, while inserting rows in a CAS table occurs in memory, affecting the corresponding physical table only after a separate save action. Adding rows to the in-memory table requires only the Insert permission. Persisting that change to the corresponding physical table requires additional permissions.
- Some mappings are discretionary, because not all mappings are optimal for all deployments or for all objects within a deployment. There is no universally appropriate formula for translating SAS 9 authorization information into SAS Viya authorization information.

  For example, in SAS 9 a single permission (WriteMetadata) governs the ability to update, delete, and secure content objects. In SAS Viya, you can define access more precisely, enabling one group to only update a content object and another group to both update and secure that content object.

Permission Mappings for Data Objects

SAS LASR Analytic Server (LASR)

Here are permission mappings for LASR libraries and tables:

SAS 9: LASR	SAS Viya**: CAS Authorization**
ReadMetadata	ReadInfo

SAS 9: LASR	SAS Viya**: CAS Authorization**
WriteMetadata	AlterCaslib, AlterTable, DropTable, DeleteSource, ManageAccess
Administer	Promote*
Read	Select, LimitedPromote**
Write	AlterTable, DropTable, Update, Insert, Delete

* In LASR, the Administer permission is required to load or import a table. In CAS, the Promote (or LimitedPromote) permission is required to move a table to global scope.

** LimitedPromote is optional. It supports just-in-time loading of data.

> **TIP** In SAS 9, the ability to stop a SAS LASR Analytic Server is provided by the Administer permission on the target server's definition. In SAS Viya, the ability to stop a CAS server is provided by the target server's Superuser role.

For SAS 9 details, see Permission Definitions in *SAS Visual Analytics: Administration Guide*.

For SAS Viya details, see Permissions by Task in *SAS Viya Administration: Cloud Analytic Services Authorization*.

Metadata LIBNAME Engine (MLE)

Here are permission mappings for MLE libraries and tables:

SAS 9: MLE	SAS Viya**: CAS Authorization**
ReadMetadata	ReadInfo
WriteMetadata	AlterCaslib, AlterTable, CreateTable, DeleteSource, ManageAccess
Read	Select*
Write	Update, DeleteSource, CreateTable, (ManageAccess)
Create	Insert, DeleteSource, CreateTable, (ManageAccess)
Delete	Delete, DropTable, DeleteSource, CreateTable, (ManageAccess)

* Select in CAS is broader than Read in MLE. Select supports the loadTable, upload, and addTable actions, in addition to providing the ability to read data.

Here are details about the CAS side of the mappings in the preceding table:

- The DeleteSource and CreateTable permissions are included in the mappings for Write, Create, and Delete because CAS does not directly modify physical (on-disk) data. CAS modifies an in-memory table and then deletes and re-creates the corresponding physical file.
- The ManageAccess permission is included in parentheses where it supports deletion or replacement of files that have direct access controls.
- For some permissions, CAS offers more levels of control than does MLE. For example, in CAS you can prevent table creation at the caslib level or on an individual (existing) table.

For SAS 9 details, see Permissions That Affect Data Access through the Metadata Engine in *SAS Language Interfaces to Metadata*.

For SAS Viya details, see Permissions by Task in *SAS Viya Administration: Cloud Analytic Services Authorization*.

Metadata-Bound Libraries

CAUTION! In terms of limiting host access by SAS, CAS does not currently offer an exact equivalent to SAS 9 metadata-bound libraries. CAS does prevent creation of multiple caslibs that reference the same host path. That is a per-server constraint; it does not extend across CAS servers. For greater protection, you can encrypt path-based caslibs, limit the ability to create caslibs, and limit the host paths against which non-administrators can create caslibs.

Here are permission mappings for metadata-bound libraries and tables:

SAS 9: Metadata-Bound	SAS Viya**: CAS Authorization**
ReadMetadata	ReadInfo
WriteMetadata	AlterCaslib, ManageAccess
Delete	Delete, DropTable, DeleteSource, CreateTable, (ManageAccess)
Insert	Insert, DeleteSource, CreateTable, (ManageAccess)
Update	Update, DeleteSource, CreateTable, (ManageAccess)
Select	Select
Create Table	CreateTable, Promote (or LimitedPromote)
Drop Table	DropTable, DeleteSource, (ManageAccess)
Alter Table	AlterTable, CreateTable, DeleteSource, (ManageAccess)

Here are details about the CAS side of the mappings in the preceding table:

- The DeleteSource and CreateTable permissions are included in the mappings for Insert, Update, and Delete because CAS does not directly modify physical (on-disk) data. CAS modifies an in-memory table and then deletes and re-creates the corresponding physical file.
- The ManageAccess permission is included in parentheses where it supports deletion or replacement of files that have direct access controls.
- The mapping for the Delete permission is included only for completeness. Currently, there is no CAS action for deleting rows.

For SAS 9 details, see Permissions for Metadata-Bound Data in *SAS Guide to Metadata-Bound Libraries*.

For SAS Viya details, see Permissions by Task in *SAS Viya Administration: Cloud Analytic Services Authorization*.

Other SAS 9 Data

Here are permission mappings for other SAS 9 libraries and tables:

SAS 9: Other Data	SAS Viya**: CAS Authorization**
ReadMetadata	ReadInfo, Select

SAS 9: Other Data	SAS Viya**: CAS Authorization**
WriteMetadata	(all other caslib and table permissions)

For SAS 9 details, see Working with Tables in *SAS Intelligence Platform: Security Administration Guide*.

For SAS Viya details, see Permissions by Task in *SAS Viya Administration: Cloud Analytic Services Authorization*.

Permission Mappings for Content Objects

Folders

> **TIP** Because authorization precedence in SAS Viya is very different from authorization precedence in SAS 9, it is a good idea to review and rethink your folder structure, rather than automatically replicating your SAS 9 folder structure in SAS Viya. The general authorization system is flexible enough to accommodate almost any folder structure, but preserving an existing folder structure is not always the cleanest and most efficient approach.

Here are permission mappings for folders:

SAS 9	SAS Viya**: General Authorization**	**Target**
ReadMetadata	Read	objectURI, containerURI
WriteMetadata	Update, Delete, Secure	objectURI, containerURI
WriteMemberMetadata	Add, Remove	objectURI
	Update, Delete, Secure	containerURI

For SAS 9 details, see Working with Folders in *SAS Intelligence Platform: Security Administration Guide*.

For SAS Viya details, see Permissions by Task in *SAS Viya Administration: General Authorization*.

Reports

Here are permission mappings for reports:

SAS 9	SAS Viya**: General Authorization**
ReadMetadata	Read
WriteMetadata	Update, Delete, Secure

For SAS 9 details, see Working with Reports in *SAS Intelligence Platform: Security Administration Guide*.

For SAS Viya details, see Permissions by Task in *SAS Viya Administration: General Authorization*.

23

CAS Authorization

CAS Authorization: Overview

To learn about the Cloud Analytic Services (CAS) authorization system, see Concepts.

To manage access, use the interface that best meets your needs. Here are suggestions:

- To interactively manage access, use the Authorization window (or, in a programming-only deployment, use CAS Server Monitor).
- To script management of CAS access controls in a full deployment, use the command-line interface. See CLI Examples: CAS Authorization in *SAS Viya Administration: Using the Command-Line Interfaces*.
- For comprehensive programmatic management of CAS access controls, use CAS actions. See the Access Control action set in *SAS Viya: System Programming Guide*.

CAUTION! Do not rely exclusively on CAS access controls to protect data. You must also consider direct host access. See Host Access Considerations.

CAS Authorization: How to (Authorization Window)

Introduction

These instructions explain how to manage access to caslibs, tables, and rows using SAS Environment Manager.

Navigation

1 In SAS Data Explorer, locate and select a global caslib or table.

> **TIP** You can access SAS Data Explorer by selecting **Manage Environment** from the applications menu (☰) and then clicking ▦ in the navigation bar. For other access methods, see SAS Data Explorer and the Choose Data Window in *SAS Data Explorer: User's Guide*.

2 Right-click, and select **View authorization** or **Edit authorization**.

Note: If **Edit authorization** is not enabled, you are not authorized to modify access to the selected object.

Examine Access

Scope

The scope of the display is as follows:

- There is always a row for Authenticated Users.
- There is always a row for you, the currently connected user who is using the display.
- There is a row for each principal that is assigned to an access control that affects access to the current object.
- If you add an identity and do not give that identity at least one direct setting, that identity is automatically removed from the display.
- You cannot directly remove a row. If you remove all direct settings for an identity and there is no other reason for that identity to be displayed, that identity is automatically removed from the display.
- Only the permissions that are relevant for an object (directly or for inheritance purposes) are displayed for that object.
- The display does not reflect the impact of CAS role membership or status.

Permissions

For each principal and permission, the following icons describe effective (net) access to the current caslib or table.

Icon	Meaning
✓	Authorized
⊘	Not Authorized
◐	Row-Level
●	Unknown
◆	Direct*

* Indicates that a permission is directly assigned to the specified principal on the current object.

Access Levels

The **Access Level** column provides an alternative to interacting with individual permissions.

- When you manage access, each access level is a shortcut for adding a set of direct access controls.
- When you view access information, each access level is a shorthand description of a set of effective permissions.

Table A.1 Access Levels

Access Level	Permissions
No access	None
Read	Only ReadInfo and Select
Write	All except ManageAccess and AlterCaslib
Full control	All
Custom	Any other combination, including any row-level access

Note: Access levels exist only in the presentation layer in SAS Environment Manager. CAS stores and evaluates individual permissions, not cumulative access levels.

Set an Access Level

1 Open the Edit Authorization window for a CAS object.

2 If the principal that you want to work with is not already listed, click 👤. In the Add Identities window, move the user or group to the right pane, and click **OK**.

Note: If guest access is enabled, you must select **Add Identities** after you click 👤. Or, if you need to add Guest to the display, select **Add Guest** after you click 👤. See Guest Access in *SAS Viya Administration: Authentication*.

3 In the **Access Level** column, click and drag a gauge to adjust access.

Note: If a gauge is not displayed, select an access level other than **(custom)** from the drop-down list.

Note: Each time you change an access level, direct access controls are added as needed to meet the definition of the new access level. If you want to discard all unsaved changes, click **Cancel**.

CAUTION! Reducing the access level for a group that you belong to might block your access. To preserve your access, make sure you have a higher precedence (offsetting) direct grant. If you are a Superuser, this precaution is not strictly necessary.

4 If you modified access for a group, click **Preview**. Examine the impact of the change on other principals. For example, increasing the access level for Authenticated Users from **No access** to **Full control** affects all authenticated users who do not have a more specific denial.

5 Click **Save**.

Add a Direct Access Control

1 Open the Edit Authorization window for a CAS object.

2 If individual permissions are not already displayed, select the **Show individual permissions** check box.

3 If the principal that you want to work with is not already listed, click 👤. In the Add Identities window, move the user or group to the right pane, and click **OK**.

Note: If guest access is enabled, you must select **Add Identities** after you click 👤. Or, if you need to add Guest to the display, select **Add Guest** after you click 👤. See Guest Access.

4 Click the effective access icon (for example, ✓) for the principal and permission that you want to modify.

5 In the pop-up window, select **Grant** or **Deny** in the **Direct Setting** drop-down list.

CAUTION! Before you deny access for a group that you belong to, make sure you have a higher precedence (offsetting) direct grant. If you are a Superuser, this precaution is not strictly necessary.

6 If you modified access for a group, click **Preview** in the Edit Authorization window.

- Notice that a diamond is displayed in the cell that you modified. The diamond indicates that effective access comes from a direct setting.
- Examine the impact on other principals. For example, a direct denial for GroupA affects all members of GroupA who do not have their own direct settings.

7 Click **Save**.

Remove a Direct Access Control

1 Open the Edit Authorization window for a CAS object.

2 In a cell that includes a diamond, click the effective access icon.

3 In the pop-up window, select **(none)** from the **Direct settings** drop-down list.

4 In the Edit Authorization window, notice that the new effective access value is unknown (●). Click **Preview**.

- Notice that the new effective access value is known. Or, if you removed the only setting that made the associated user or group a relevant principal for the current object, the user or group is no longer included in the display.
- If you modified access for a group, examine the impact on other principals.

5 Click **Save**.

Remove Multiple Direct Access Controls

1 Open the Edit Authorization window for a CAS object.

2 In a row that includes at least one direct setting, click the first cell. The row is selected.

3 Click ✖ to remove all direct access controls for the selected identity.

4 Notice that effective access for any affected cells is unknown (●). Click **Preview**.

- Notice that all effective access values are known. Or, if the associated user or group is no longer a relevant principal for the current object, the user or group is no longer included in the display.
- If you modified access for a group, examine the impact on other principals.

5 Click **Save**.

Examples: Manage Access to a Caslib

Provide Public Access to a Caslib

To give all users Read access to a new global caslib that you added:

1. Open the caslib’s Edit Authorization window. (See Navigation.)
2. In the row for Authenticated Users, increase the **Access Level** to **Read**.
3. Click **Save**.

Provide Selective Access to a Caslib

To give a particular user Read and Write access to a new global caslib that you added:

1. Open the caslib’s Edit Authorization window. (See Navigation.)
2. Click 👤 in the table toolbar.
3. In the left pane of the Add Identities window, locate the user. Move the user to the right pane, and click **OK**.
4. In the Edit Authorization window, increase the user’s **Access Level** to **Write**.
5. Click **Save**.

Block All Access to a Caslib

To block all access for a particular identity:

1. Open the caslib’s Edit Authorization window. (See Navigation.)
2. If the identity is not already listed, click 👤 in the table toolbar.

   In the left pane of the Add Identities window, locate the user, group, or custom group that you want to block. Move that identity to the right pane, and click **OK**.
3. In the Edit Authorization window, decrease the identity’s **Access Level** to **None**.
4. If the identity is not an individual user, click **Preview**. Examine the impact of your change on other listed identities.
5. Click **Save**.

Limit Write Access to a Caslib

To allow only Read access for a particular identity:

1. Open the caslib’s Edit Authorization window. (See Navigation.)
2. If the identity is not already listed, click 👤 in the table toolbar.

   In the left pane of the Add Identities window, locate the user, group, or custom group that you want to block. Move that identity to the right pane, and click **OK**.
3. In the Edit Authorization window, decrease the identity’s **Access Level** to **Read**.
4. If the identity is not an individual user, click **Preview**. Examine the impact of your change on other listed identities.
5. Click **Save**.

Provide Row-Level (Filtered) Access

To make different subsets of rows available to different identities, set one or more row-level grants. Each row-level grant includes a filter that limits the available rows.

1 Open the Edit Authorization window for a CAS table.

2 If the principal that you want to work with is not already listed, click 👤. In the Add Identities window, move the user or group to the right pane, and click **OK**.

   Note: If guest access is enabled, you must select **Add Identities** after you click 👤. Or, if you need to add Guest to the display, select **Add Guest** after you click 👤. See Guest Access in *SAS Viya Administration: Authentication*.

3 If individual permissions are not displayed, select the **Show individual permissions** check box.

4 In the **Select** column, click an effective access icon.

5 In the pop-up window, select **Row-level Grant** from the **Direct setting** drop-down list.

6 In the **Row-Level Filter** window:

   a Specify an expression that includes only the rows that the principal should be able to access. The basic format is: **`column-name operator value`**. Here are basic examples:

Type of Filter	Example
Numeric	`sales<1000`
Character	`Make='Ford'`
Dynamic	`user='SUB::SAS.Userid'`

   For details, see Row-Level Access.

   Note: If you view or edit a filter that was initially created programmatically, you might see escape characters and a different pattern of quotation marks.

   b Click **OK**.

7 In the Edit Authorization window, next to the new setting, notice that a diamond is displayed. The diamond indicates that effective access comes from a direct setting.

8 If you modified access for a group, click **Preview**. Examine the impact on other principals.

9 Click **Save**.

Identify the Source of Effective Access

To determine which access control causes a particular effective access result, examine the origins information for that result.

1 Open the View Authorization window for the target CAS object.

2 Click the effective access icon for which you want origins information.

3. In the pop-up window, next to the **Effective Access** value, click ⓘ.

   Note: The icon is disabled if you have changes that you have neither saved nor previewed.

4. In the Origins window, review the displayed information.

   - The **Source object** field indicates where the determinative access control is set.
   - The **Principals** field indicates which identity the determinative access control are assigned to.

     Note: If multiple access controls of equal precedence cause the result, multiple principals are listed.

   For more information, see Origins of Effective Access.

CAS Authorization: How to (CAS Server Monitor)

Introduction

These instructions explain how to manage access to global caslibs using CAS Server Monitor in a programming-only deployment.

> **TIP** To manage access at the table, column, or row level, use another interface.

Navigation

1. In CAS Server Monitor, beneath the **SAS Cloud Analytic Services** banner, click 🔧.
2. On the **Configuration** page, select **Access Controls**.
3. In the **Caslibs** list, select the caslib.

Here are details:

- All of the global caslibs that you are authorized to see are listed.
- To see any new global caslibs, click ↻. The **Caslibs** list is not updated automatically.
- Session and personal caslibs are not listed because you cannot set access controls on them. Resources in session and personal caslibs are not sharable with other users.
- The **Global Caslib Creation** and **Session Caslib Creation** caslibs do not contain data. These special caslibs determine which non-administrators can add and delete caslibs. See Caslib Management Privileges in *SAS Viya Administration: SAS Cloud Analytic Services*.
- Any predefined caslibs have appropriate access controls.

Examine Access to a Caslib

On the **Configuration** page, under **Access Controls**, the columns for the selected caslib are populated as follows:

Column	Content
Applies To	Specifies a type of access control principal.

Column	Content
Identity	Specifies the name (unique identifier) for an access control principal.
Grant	A check mark indicates that access to the selected caslib is authorized for the specified principal and permission.
Deny	A check mark indicates that access to the selected caslib is not authorized for the specified principal and permission.
Activity	Specifies a permission, such as Select.

The rows for a new custom caslib are initially populated as follows:

- For the user who added the caslib, there is initially a row for every permission, because that user has a direct setting for each permission.
- For Authenticated Users, there is always a row for every permission, because Authenticated Users has an inherited setting for each permission.
- For other identities, there is a row for each direct setting. For example, if UserA has a direct grant of the Select permission, there is a Select row for UserA.

Provide Public Access to a Caslib

To give all users access to a new global caslib that you added:

1. On the **Configuration** page, select **Access Controls**.
2. In the **Caslibs** pane, select the caslib.
3. Click **Edit**.
4. In the Edit Access Controls window, adjust settings as follows:

   a. In the **Authenticated Users** row for **Read Info**, select the **Grant** radio button.

   b. Click **Add Row**. In the new row at the end of the page, select **Authenticated Users**, the **Grant** radio button, and the **Select** activity.

   c. If you want to also provide Write access, add rows that grant the following additional permissions to Authenticated Users: **Insert**, **Update**, **Delete**, **Create Table**, **Drop Table**, **Delete Source**, **Alter Table**, **Limited Promote**, **Promote Table** (Promote).

   > **TIP** As an alternative to adding each row individually, click **Add Set**, and select **Add Set** from the drop-down list. In the new **Authenticated Users** rows for **Alter Caslib** and **Manage Access**, click 🗑 to delete those direct access controls. In the remaining new rows, select the **Grant** radio button.

   Note: If you want to provide access to guest users, you must grant access to Guest. Guest is not affected by access controls that are assigned to Authenticated Users. By default, guest access is not enabled. See Guest Access in *SAS Viya Administration: Authentication*.

5. Click **OK** to save your changes.
6. Under **Access Controls**, review the results of your changes.

Selectively Grant Access to a Caslib

This example gives UserA Read and Write access to a new global caslib that you added:

1. On the **Configuration** page, under **Access Controls**, select the caslib.
2. Click **Edit**.
3. In the **Edit Access Controls** window, click **Add Set**, and select **Add User Set** from the drop-down list.
4. In the next window, enter *UserA* in the **User name** field. Click **OK**.
5. In the **Edit Access Controls** window, select the **Grant** radio button in each **UserA** row, except the **Alter CASLib** and **Manage Access** rows.
6. At the end of UserA’s **Alter CASLib** and **Manage Access** rows, click 🗑.
7. Click **OK** to save your changes.
8. Under **Access Controls**, review the results of your changes.

   Note: UserA does not have rows for **Alter CASLib** and **Manage Access**, because you did not give UserA direct access controls for those permissions. Unless UserA is a member of a group that has direct access controls for those permissions, UserA’s effective access for those permissions comes from Authenticated Users.

Here are some additional details about editing access controls:

- To add a single row, click **Add Row**.
- To provide a complete set of editable rows for an identity, click **Add Set** and then select the appropriate item:
  - For a group, select **Add Group Set**.
  - For a user, select **Add User Set**.
  - For Authenticated Users, select **Add Set**.
  - For guest users, select **Add Guest Set**.

    Note: By default, guest access is not enabled. See Guest Access in *SAS Viya Administration: Authentication*.
- To delete a direct access control, click 🗑.
- You cannot change or delete inherited settings. You can add a direct access control that has precedence over an inherited setting.

Selectively Limit Access to a Caslib

About Setting Direct Denials

CAUTION! Identity names that you enter are not validated. For sensitive data, do not grant access to Authenticated Users and then rely on selective direct denials. The safer practice is to verify that access is broadly denied, and then grant access selectively.

CAUTION! Do not block your own access. Before you add a direct denial for a group that you belong to, make sure you have a higher precedence (offsetting) direct grant. If you are a CAS administrator (Superuser) or Data administrator, this precaution is not strictly necessary.

Block All Access for an Identity

To block all access for an identity:

1. On the **Configuration** page, under **Access Controls**, select the caslib.
2. Click **Edit**.

   Note: If the **Edit** button is disabled, you are not authorized to set permissions for the selected caslib.
3. In the **Edit Access Controls** window, click **Add Row**.
4. In the new row, select an identity type, enter a name (unique identifier), make sure the **Deny** radio button is selected, and select the **Read Info** activity.
5. Click **OK** to save your changes.
6. Under **Access Controls**, review the results of your changes.

Block Write Access for an Identity

To block Write access for an identity:

1. On the **Configuration** page, under **Access Controls**, select the caslib.
2. Click **Edit**.

   Note: If the **Edit** button is disabled, you are not authorized to set permissions for the selected caslib.
3. In the **Edit Access Controls** window, click **Add Set**, and select **Add Group Set** or **Add User Set**.
4. In the next window, enter the user or group name (unique identifier). Click **OK**.
5. In the identity's **Read Info** and **Select** rows, click 🗑.
6. In the identity's remaining rows, make sure the **Deny** radio button is selected.
7. Click **OK** to save your changes.
8. Under **Access Controls**, review the results of your changes.

Resolve Duplicate Access Controls

You cannot save more than one direct setting for a particular caslib, principal, and permission. Here are examples:

- You cannot save two direct grants of the Update permission for UserA on caslibA.
- You cannot save both a direct denial and a direct grant of the ReadInfo permission for UserA on caslibA.

If the error message `Access control has duplicates` is displayed twice in the Edit Access Controls window, there is one duplicate direct setting. Delete one of the settings that have the error message. One error message remains. You can now save your changes by clicking **OK** again.

CAS Authorization: Concepts

Scope

CAS authorization manages access to the following CAS objects:

- caslibs
- CAS tables and columns
- CAS action sets and actions

CAS authorization requirements do not apply in the following circumstances:

- The requesting user has assumed a role that is exempt from all applicable authorization requirements. For example, the user has assumed the Superuser role and the request is to add a caslib.
- The target object is not potentially sharable. For example, the target is a table in a personal caslib, a session caslib, or the session scope of a global caslib.

Note: Not all interfaces expose all aspects of CAS authorization.

Key Terms

Term	Definition
Access control	A composite of authorization elements. Example: An access control grants ReadInfo to groupA on caslibA.
Target	A resource. Examples: tableA, caslibA
Principal	The user, group, or construct to which an access control is assigned. Examples: UserA, GroupA, Authenticated Users
Permission	A type of access. Values: ReadInfo, Select, LimitedPromote, Promote, CreateTable, DropTable, DeleteSource, Insert, Update, Delete, AlterTable, AlterCaslib, ManageAccess
Setting	An indication of whether (and to what extent) access is provided. Values: Grant, Row-Level Grant, Deny
Filter	In a row-level grant of the Select permission, the constraint expression. Example: `User='SUB::SAS.Userid',sales>1000`
Effective access	A context-neutral description of the net result of all relevant access controls. Values: Authorized, Not Authorized, Row-Level
Access outcome	In an access request, the authorization decision. Values: Authorized, Not Authorized, Row-Level Authorization

Principals

The principal in an access control is the user, group, or construct to which the access control is assigned. The CAS authorization system supports the following principals:

- an individual authenticated user
- a user group (a custom group or a group in your authentication provider)
- Authenticated Users (the construct that represents all authenticated users)

  Note: In some programmatic contexts, this construct corresponds to the group that is named `*`.
- Guest (the identity type that facilitates guest access)

  Note: Guest is not part of Authenticated Users.

Administrators

CAS roles provide per-server assumable access to administrative functionality. For example, the Superuser role is exempt from authorization requirements throughout a CAS server, except for data access requests. See CAS Server Roles in *SAS Viya Administration: Identity Management*.

Inheritance

Access flows through a hierarchy of objects. Each parent object conveys settings to its child objects. Each child object inherits settings from its parent object.

Here are the inheritance relationships:

- Access flows from a caslib to its tables.
- Access flows from a table to its columns.
- Access flows from an action set to its actions.

Note: Each caslib always has inherited denials of all permissions for Authenticated Users. Those inherited denials prevent access if there are no higher precedence grants.

Permissions

	Data Enforcement Levels*			
Permission	**Caslib**	**Table**	**Column**	**Affected Activities**
ManageAccess	✓	✓	✓	Set access controls.
ReadInfo	✓	✓	✓	View and traverse objects.
LimitedPromote		✓		Promote from source in the same caslib.**
Promote	✓			Promote from any caslib.
CreateTable	✓	✓		Save (persist) a table.
DeleteSource		✓		Delete a physical source table.***
DropTable		✓		Remove a table from global scope.***

	Data Enforcement Levels*			
Permission	**Caslib**	**Table**	**Column**	**Affected Activities**
Select		✓	✓	Read data values.
AlterCaslib	✓			Change the properties of a caslib.
AlterTable		✓		Change the attributes or structure of a table.
Insert		✓		Add rows.
Delete		✓		Delete rows.
Update		✓		Change data values.
Execute				Run an action.
Load				Load an action set.

* You can set permissions at or above the level where they are enforced. See also Access to Actions.

** If no table is specified in the request, LimitedPromote is checked at the caslib level.

*** To delete any direct access controls, the ManageAccess permission is required.

Permissions by Task

To complete a task, you must have sufficient access to all relevant data objects. The following tables document permissions that are required for selected tasks.

Table A.2 Simple Tasks

Task (CAS Action)	Caslib	Table	Column
Set caslib permissions	ReadInfo ManageAccess	-	-
Modify caslib properties	ReadInfo AlterCaslib	-	-
Set table permissions	ReadInfo	ReadInfo ManageAccess	-
Modify table properties*	ReadInfo	ReadInfo Select AlterTable	-
Load a table from a caslib's data source (loadTable)	ReadInfo	ReadInfo Select	-
Transfer and load an entire file (upload)	ReadInfo	ReadInfo	-
Transfer rows to the server (addTable)	ReadInfo	ReadInfo Select	-

Task (CAS Action)	Caslib	Table	Column
Move a table to global scope (promote)	ReadInfo Promote**	ReadInfo	-
Remove a table from global scope (dropTable)	ReadInfo	ReadInfo DropTable	-
Delete a file (deleteSource)	ReadInfo	ReadInfo DeleteSource (ManageAccess)***	-
Persist a file (save)	ReadInfo CreateTable	ReadInfo CreateTable (DeleteSource)†	-
Read data	ReadInfo	ReadInfo Select	ReadInfo Select
Insert rows	ReadInfo	ReadInfo Insert	-
Update rows	ReadInfo	ReadInfo Select Update	-
Delete rows††	ReadInfo	ReadInfo Select Delete	-

* You cannot alter a table that has fine-grained constraints (row-level filters or column-level denials) that apply to you.

** For promotion of a table within the same caslib, LimitedPromote for the table (instead of Promote for the caslib) is sufficient.

*** ManageAccess is required if the request involves deletion of direct access controls.

† DeleteSource is required to replace a source table.

†† This entry is included only for completeness. Currently, there is no CAS action for deleting rows.

Table A.3 Compound Tasks

Task (CAS Actions)	Required Permissions
Import (upload + save + dropTable + loadTable + promote)	Caslib: ReadInfo, CreateTable, Promote* Table: ReadInfo, Select, (DeleteSource, DropTable)**
Just-in-time load (loadTable + promote)	Caslib: ReadInfo, Promote* Table: ReadInfo, Select
Delete a global-scope table (dropTable + deleteSource)	Caslib: ReadInfo Table: ReadInfo, DeleteSource, DropTable

* LimitedPromote for the table (instead of Promote for the caslib) is sufficient.

** DeleteSource is required to replace a source table. DropTable is required to replace a global-scope table.

For information about who can add and delete caslibs, see Caslib Management Privileges in *SAS Viya Administration: SAS Cloud Analytic Services*.

Row-Level Access

Overview of Row-Level Access

A row-level grant includes a filter that limits the Select permission on a table. A user who has row-level access to a table can view only those rows that are within the associated filter. See also Application and Persistence.

For example, you can use a row-level grant to enable groupA to see only those rows in tableA where the value in the Toy_Price column is 25. Here is an overview of the process:

1. On tableA, give groupA a row-level grant of Select permission.

   Specify the following filter: `Toy_Price=25`

   Note: For detailed instructions, see “Provide Row-Level (Filtered) Access”.

2. Make sure that groupA has ReadInfo access to tableA and its parent caslib.

3. Make sure that groupA is not a member of another group that has a grant or denial of the Select permission on tableA.

4. Verify that when a member of groupA accesses tableA, the expected rows are returned.

Syntax for Row-Level Filters

Operator (Alias)	Example Filter
Contains (?) Not Contains	Toy_Type Contains 'cars'
In Not In	Toy_Type In ('dolls' 'cars' 'animals')
Between -inclusive Not Between -inclusive	Toy_Price Between 20 AND 30
Like	Toy_Type Like 'd%'
= (EQ) > (GT) < (LT) <> >= (GE) <= (LE) ^= (NE, ~=)	Toy_Price=25

Operator (Alias)	Example Filter

`+ -addition`
`- -subtraction`
`/ -division`
`* -multiplication`
`** -exponentiation`
`() -parentheses`
`|| -string concatenation`
	`Profit > (Sales * .5)`

`AND (&)`
`OR (|, !)`
`NOT`
	`Toy_Type='cars' OR Toy_Type='dolls'`

`Is Missing`
`Is Not Missing`
`Is Null`
`Is Not Null`
	`Toy_Type Is Not Null`

Identity-Based Substitution

Identity-based substitution is a powerful and concise technique for defining row-level access. You can use substitution to implement any number of per-user access distinctions with a single row-level filter.

Identity-based substitution parameters map a user's authenticated ID or group memberships to values in a specified column in your data. Values are dynamically substituted into the filter at run time, as appropriate for each requesting user. Here are the supported substitution parameters:

Substitution Parameter	Description and Example
`SUB::SAS.Userid`	Determines whether a data value is the same as the requesting user's authenticated ID. *empID*='SUB::SAS.Userid' **Note:** If the casing of your user IDs might not match the casing of corresponding values in your data, use the upcase function on both sides of the expression. For example: `upcase(empID)=upcase(':SUB::SAS.Userid:')`
`SUB::SAS.IdentityGroups`	Determines whether a data value matches any of the requesting user's group memberships. *FacilityRegion* In ('SUB::SAS.IdentityGroups') **Note:** The comparison is against each group's unique name, so your data must contain unique group names. In a group definition in SAS Environment Manager, the **Group ID** field contains the group's unique name.

Example: User ID Substitution

If a tableB has an empID column with values that match the user IDs with which users authenticate, you might assign this filter to Authenticated Users:

```
empID='SUB::SAS.Userid'
```

At request time, each user's ID is substituted into the right side of the expression. In a request from userA, the expression resolves as:

```
empID='userA'
```

As a result, userA gets only those rows where the value in the empID column is **`userA`**.

Example: Membership Substitution

If tableC has a FacilityRegion column with values that match the unique names for user groups, you might assign this filter to an AllRegions group:

```
FacilityRegion In ('SUB::SAS.IdentityGroups')
```

At request time, each affected user's list of group memberships is substituted into the right side of the expression. In a request from user13 (who is a member of the grp7, grp9, and AllRegions groups), the expression resolves as:

```
FacilityRegion In ('grp7','grp9','AllRegions')
```

As a result, user13 gets only those rows where the value in the FacilityRegion column is **`grp7`**, **`grp9`**, or **`AllRegions`**.

Note: **`Authenticated Users`** is not one of the listed memberships, because it is an access control principal, not a user group.

Multiple Filters and Cumulative Access

If multiple row-level filters are applicable to a user, only the highest precedence filter provides access. If there is an identity precedence tie (the user is a member of multiple groups, each of which has a filter), the user can access any row that meets any of the filters.

Here are details:

The filters for multiple row-level grants provide cumulative access only if all of the following circumstances exist:

- The requesting user does not have a direct access control for the Select permission.
- None of the requesting user's groups have a direct grant or denial for the Select permission.
- Two or more of the requesting user's groups have row-level grants.

  Note: All custom and LDAP groups have equal precedence, regardless of any nested memberships.

  Note: A filter for a row-level grant that is assigned to Authenticated Users is never cumulative (joined with other filters by OR). Authenticated Users is a construct that has lower precedence than any group.

Column-Level Access

CAS supports column-level permissions, where a user can access some (but not all) columns in a table. You can use the Access Control action set to set column-level permissions. See also Application and Persistence.

To prevent a user from accessing a column, deny the user both the ReadInfo permission and the Select permission for that column. Denying both permissions ensures that the user cannot access the column through any CAS action or interface.

Note: Do not rely exclusively on a denial of the ReadInfo permission on a column to hide that column. Not all CAS actions require the ReadInfo permission at the column level.

CAUTION! Not all interfaces can successfully interact with tables that have column-level permissions. Before you provide a production implementation of column-level permissions, verify that results in all applicable interfaces are acceptable.

For example, in SAS Visual Analytics, column-level access is not supported and can yield unexpected results. If userA lacks access to any column that is included in a SAS Visual Analytics report object, userA cannot see any data in that report object.

Access to Actions

Action sets and actions that have no access controls are available to all authenticated users. As a result, almost all action sets and actions are available to all users. In general, the ability to perform a particular task is managed by access controls on the target data, not by access controls on actions.

An exception is actions for adding nodes and stopping the server. The initial configuration denies Authenticated Users the Execute permission for those actions. Initially, only Superusers can add nodes or stop the server.

Here are additional details:

- The ReadInfo and Execute permissions are enforced for actions.
- The ReadInfo and Load permissions are enforced for action sets.
- Unregistered action sets are subject to access constraints that are defined on the _UNREGISTERED action set. The initial configuration denies Authenticated Users the Load permission on the _UNREGISTERED action set.

  Note: An unregistered action set is an action set that is not listed in the database of action sets that SAS provides. SAS solutions use only registered action sets.

  Note: An attempt to load an action set that does not exist generates an `access denied` error message, because no such action set is known and registered. For example, if you do not correctly specify the action set name in a load request, an `access denied` error message is generated.
- The identity type Guest has the same access to non-administrative actions as Authenticated Users. That access is relevant only for sites that choose to enable guest access. See Guest Access in *SAS Viya Administration: Authentication*.

Authorization Decisions

Precedence

In the CAS authorization system, precedence is determined by where an access control is set and who an access control is assigned to.

- Direct access controls have precedence over inherited settings.
- The principal precedence hierarchy is relatively flat. It consists of only the following three levels: 1) individual users, 2) user groups, and 3) the construct Authenticated Users.

  Note: All user group memberships are at the same level of precedence, even if groups are nested.

Direct access controls have precedence over inherited access controls, regardless of who the principal is. For example, if only the following access controls exist, then UserA cannot access TableA:

- UserA has a direct grant of ReadInfo on caslibA.
- Authenticated Users has a direct denial of ReadInfo on TableA, which is in caslibA.

Note: One way to enable UserA to access TableA is to add a direct grant of ReadInfo for UserA on TableA. UserA's direct grant has precedence over the direct denial for Authenticated Users.

How Access Is Evaluated

Each access request initiates an authorization decision process. That process terminates when an outcome is reached. For example, here is the authorization decision process for the Select permission in a request to access data in a CAS table:

1. If there are relevant direct access controls on the table, those access controls determine the outcome as follows:

   a. If there is a setting that is specifically assigned to the requesting user, that setting determines the outcome.

   b. If there is a denial from a group, the outcome is Not Authorized.

   c. If there is a grant from a group, the outcome is Authorized.

   d. If there is exactly one row-level grant from a group, the outcome is Row-Level Authorization (authorized for rows within the applicable filter).

   e. If there are two or more row-level grants from groups, the outcome is Row-Level Authorization (authorized for any row that is within any of the applicable filters). See Multiple Filters and Cumulative Access.

   f. If there is a setting for Authenticated Users, that setting determines the outcome.

2. If there are no relevant direct access controls on the table, direct access controls on the parent caslib determine the outcome as follows:

   a. If there is a setting that is specifically assigned to the requesting user, that setting determines the outcome.

   b. If there is a denial from a group, the outcome is Not Authorized.

   c. If there is a grant from a group, the outcome is Authorized.

   d. If there is a setting for Authenticated Users, that setting determines the outcome.

3. If there are no relevant direct access controls on the table or the parent caslib, the outcome is Not Authorized.

Access Control Transactions

If you want to preview the results of changes to CAS access controls before you save those changes, use an access control transaction. Here are examples:

- When you have unsaved changes in a CAS object's Authorization window, you can click **Preview**. Review the results, and then save or cancel your changes.
- When you manage CAS authorization using the command-line interface, you can choose to check out an object, modify its access controls, and then commit or roll back the changes. In the command-line interface, a transaction is automatically started when you check out an object.
- When you manage CAS authorization programmatically, you can choose to start a transaction and check out one or more objects before you make changes. When the transaction is open, any whatIsEffective actions that you run incorporate the effects of your unsaved changes. Review the results, and then commit or roll back your changes.

Here are key points about access control transactions:

- This feature is for only changes to access controls. This feature does not provide transaction support for interactions with data or with any other aspect of CAS objects.
- In SAS Environment Manager, access control transactions are used when CAS access controls are managed.
- In programmatic and command-line interfaces, use of access control transactions is optional. You do not have to use a transaction in order to modify CAS access controls. Use access control transactions if you want to preview the results of your changes or ensure that nobody else is modifying access controls for the same objects at the same time.

Origins of Effective Access

Origins information explains effective access by answering the question: Why does this identity have this effective access to this object?

Origins information identifies the highest precedence access control that causes the access outcome for a particular identity, object, and permission. If there are multiple tied highest precedence access controls, origins information includes all of them. Additional, lower precedence controls are not included.

The following table provides simple examples of origins information. Each row in the table is for a different (independent) scenario. In each example, we are looking at why UserA has an effective access value of Not Authorized for the ReadInfo permission on TableA. UserA is a member of GroupA and GroupB. TableA is in CaslibA.

Table A.4 Origins: Examples

	Origins Information	
Highest-Precedence Access Control (or Controls)	**Object**	**Principals**
On TableA, a direct denial for UserA.	TableA	UserA
On TableA, a direct denial for GroupA.	TableA	GroupA
On TableA, direct denials for GroupA and GroupB.	TableA	GroupA, GroupB
On TableA, a direct denial for Authenticated Users.	TableA	Authenticated Users
On CaslibA, a direct denial for UserA.	CaslibA	UserA
On CaslibA, a direct denial for GroupA.	CaslibA	GroupA
On CaslibA, direct denials for GroupA and GroupB.	CaslibA	GroupA, GroupB
On CaslibA, a direct denial for Authenticated Users.	CaslibA	Authenticated Users
There are no relevant access controls.	Caslib default	

TIP To obtain origins information in the Authorization window, see Identify the Source of Effective Access.

Reduced Visibility: Hidden Caslibs

A hidden caslib is omitted from most lists of caslibs. Tables in a hidden caslib are omitted from most lists of tables. A caslib is hidden if its `hidden` parameter is set to **`true`**. You can set the `hidden` parameter in the tables.addCaslib action.

Hiding a caslib does not protect the caslib or limit access to the caslib's data. Hiding a caslib just reduces visibility by preventing the caslib and its tables from being listed in certain contexts. In those contexts, hidden caslibs and their tables are unlisted for all users and administrators, regardless of roles and access controls. Hiding a caslib affects everyone equally.

For example, two of the predefined caslibs (AppData and ReferenceData) are hidden. Users can use the data in those caslibs because appropriate predefined access controls are in place. However, those caslibs (and their tables) are not listed in most contexts. They are excluded from selection lists in SAS Visual Analytics.

> **TIP** Hide a caslib only if you have users who must be able to use that caslib's data but should not see that caslib or its tables in most lists.

Here are examples of contexts in which hidden caslibs are listed:

- On the **Data** page in SAS Environment Manager, hidden caslibs are listed.
- In the tables.caslibinfo action, hidden caslibs are listed if you specify the value **`true`** for the `showHidden` parameter.

Application and Persistence of CAS Access Controls

Protection of In-Memory Data

CAS access controls on a source file protect any in-memory table that is loaded from that source file (and is in the same caslib as that source file).

Persistence of Access Controls

When You Save Data

When you save data, existing direct access controls are managed in one of the following ways:

- If the original source file is replaced, direct access controls are preserved.

  > **TIP** The save action always produces a SASHDAT file, so a true replacement can occur only if the original source file is a SASHDAT file.

- If a new source file is created, direct access controls are not replicated. Access to the new source file is determined by access controls on its parent caslib.
- If a different existing source file is overwritten, direct access controls are not replicated. Access to the existing source file is unaffected.

Here is an example:

1. You add access controls to the source file CARS.csv.
2. You load the source file CARS.csv to global scope.
3. You modify the in-memory table CARS (for example, you add a calculated column).

4 You save the in-memory table CARS, specifying to replace any existing same-named source file.

5 A new source file CARS.sashdat is created. The access controls that you set on the source file CARS.csv are not replicated on the new source file.

   Note: If you could save the in-memory table without creating a new source file, the problem would not occur. In this example, a true replacement cannot occur because the in-memory table was loaded from a CSV source file.

In workflows that involve saving data, the best practice is to establish and maintain any direct access controls on the target (post-preparation) files. By default, such access controls survive activities such as deletion and replacement.

When You Save Access Controls

In almost all cases, access controls that you save are immediately persisted. Even if you specify or select an in-memory table, access controls that you save are immediately replicated on that in-memory table's corresponding source file.

The default behavior and best practice is to set access controls on persisted data (data that has been saved as a source file within a caslib), not on in-memory data. That default behavior and best practice applies to table-level, column-level, and row-level access controls.

Note: The exception is for an in-memory table that does not have a corresponding source file or backing store. If you set access controls on such a table, and then unload the table (without first saving the table), those access controls are discarded. To make sure you are referencing a source file when you set CAS access controls, include the source file extension in the table parameter.

The Ambiguity Problem

If you specify or select an in-memory table when you set access controls, the results might be unexpected. In a caslib that contains multiple same-named source files, it might not be obvious which of those source files corresponds to the table that is currently in memory.

Here is an example:

1 caslibA includes the source files CARS.csv and CARS.sashdat.

2 Someone loads CARS.csv to global scope.

3 In SAS Data Explorer, someone else specifies or selects the in-memory table CARS when setting access controls, incorrectly assuming that they are protecting the source file CARS.sashdat.

4 The new access controls are replicated, but only to the source file CARS.csv, not to the source file CARS.sashdat.

Access Controls for Rows and Columns

Row-level filters and column-level access controls are applied to requests to access or save data, not to requests to load data. (There is an exception for loading of cross-caslib data.)

For example, if tableA has a row-level filter that enables userA to see only those rows where the value in the MAKE column is Ford, the filter is applied as follows:

- If userA loads tableA, all rows are loaded.
- If userA accesses tableA, he sees only those rows where the value for MAKE is Ford.

As explained in the preceding sections, row-level filters and column-level access controls are not replicated when a user saves a table as a new or different table. The following information continues the preceding example:

- If userA saves tableA in place, the replaced table contains only those rows where the value for MAKE is Ford. This reduction in scope affects all subsequent loads of tableA and all access by subsequent users. Some users might access fewer rows than intended. The row-level filter still exists on the source file and can further reduce access, so no users access more rows than intended.
- If userA saves tableA to a different caslib or with a different name, the new source file contains only those rows where the value for MAKE is Ford. This reduction in scope affects all subsequent loads of tableA and all access by subsequent users. Some users might access fewer rows than intended. Because the new source file does not have the row-level filter, some users might access more rows than intended.
- If userA deletes tableA, the row-level filter is (by default) persisted.

Additional Considerations for Cross-Caslib Data

In most workflows, data does not move across caslibs. For example, when you load and promote data, you are usually performing a same-caslib activity, moving data from session scope to global scope in the same caslib.

In a few workflows, data does move across caslibs. For example, if you promote data from one caslib to another, or copy data from one caslib to another, you are performing a cross-caslib activity. The target data that results from a cross-caslib activity is referred to as cross-caslib data.

Cross-caslib data has no ongoing relationship to the source (original) caslib and file. Changes to access controls on the source caslib and file have no effect on cross-caslib data.

Access to cross-caslib data is subject to access controls on the *target* caslib and file. Avoid performing cross-caslib activities on sensitive data. If you must perform cross-caslib activities on sensitive data, make sure the target caslib and file have appropriate protections.

> **TIP** If the *source* data for a cross-caslib activity has fine-grained constraints (row-level filters or column-level denials), make sure the cross-caslib activity is performed by an identity that has sufficient access to the source data. Any constraints on the source data are applied during the cross-caslib activity, and are based on the identity that performs the cross-caslib activity. Subsequent access to the cross-loaded data cannot be expanded to include data that was unavailable to the identity that performed the cross-caslib activity.

Physical Storage of CAS Authorization Information

This topic refers to access controls being *replicated on*, *set on*, or *persisted to* a source file. That phrasing is an abstraction. On disk, access controls are stored in an item store. Each table-level access control in the item store is associated with a source file (physical table).

For more information, see cas.PERMSTORE on page 527.

CAS Authorization: Guidelines

The following guidelines can contribute to simplicity and security:

- Limit membership in administrative roles.
- Minimize the use of individual tables as targets.
- Minimize the use of individual users as principals.
- Remember that any access that is not granted is implicitly denied. Do not set unnecessary denials.
- If you deny someone access to part of a table (using a column-level or row-level access control), make sure that identity cannot update or insert rows in that table.

- Perform a backup before and after you make significant changes to your system.

CAS Authorization: Troubleshooting

Unrecognized Principals

If the Authorization window displays a warning icon next to a principal's name, that principal does not exist in the identities service.

- If the principal is a host account (for example, **`cas`**) that does not exist in your LDAP provider, you can ignore the warning icon.
- If the principal is an internal service account (for example, **`sas.ops-agentsrv`**, **`sas.searchIndex`** or **`sas.search`**), you can ignore the warning icon.
- If the principal should still exist in the identities service, make sure the identities service can still contact your LDAP provider.
- Otherwise, consider deleting the access controls that are assigned to the principal that no longer exists.

  Note: Deletion of a custom group does not cause automatic deletion of rules in which that custom group is the principal.

Inability to Modify Access

- Make sure the target object is global-scope (not in a personal caslib or session scope).
- Make sure you have the ManageAccess permission for the target object. Or, assume a role that is exempt from that permission requirement. See "Assume the Superuser Role" on page 353.
- If you get the following error: `"The access controls for table {name} are being blocked by a loaded table of the same name"`, select or specify a source file (for example, tableA.sashdat), not an in-memory table. See "Application and Persistence of CAS Access Controls" on page 404.
- If you get the following error: `"The object {name} {type} is currently locked, so you cannot modify its access controls"`, the target object is already participating in an active access control transaction. Either wait for the lock to be released, or terminate the CAS session in which the object is locked. See "Access Control Transactions" on page 402.

Unintended Loss of Access

Reinstate Access: Instructions for Users

If you inadvertently block your own access to an object, contact an administrator for assistance.

Note: Anyone who has the ReadInfo and ManageAccess permissions for the object can reinstate your access.

Reinstate Access: Instructions for Administrators

1 In the applications menu (≡), under **Administration**, select **Manage Environment**.

  Note: These instructions use SAS Environment Manager. You can instead use an alternate interface.

2 In the navigation bar, click 🗎.

3 On the **Servers** page, right-click on a server, and select **Assume the Superuser role**.

Note: If the **Assume the Superuser role** action is not available, you are not a member of the Superuser role for the selected server. If your Superuser role membership is exclusively through the SAS Administrators group, make sure you have opted in to that membership in your current session.

CAS Authorization: Interfaces

All CAS authorization requirements and constraints are always fully enforced. However, not all interfaces expose all CAS authorization features.

In the following table, the shaded part of each circle is an approximation of the amount of CAS authorization functionality that a particular interface exposes. The shading indicates relative coverage. The shading does not indicate alignment of coverage across interfaces.

Table A.5 Interfaces to CAS Authorization

Interface	Description
● Access Control action set	Programmatic interfaces for CASL (the CAS procedure), Python, Lua, and R.
● REST API	The REST interface for CAS.
● SAS Java Client Interface for SAS Viya	The Java programming interface for CAS actions.
◕ Command-line interface	A simple, scriptable interface that includes commands for managing access at the caslib, table, and row levels in a full deployment.
◕ SAS Environment Manager	A graphical web application for managing access at the caslib, table, and row levels in a full deployment.
◒ CAS Server Monitor	A graphical web application for managing access at the caslib level in a programming-only deployment.

CAS Authorization: Host Access Considerations

Why Host Access Matters

If host-layer access requirements are not met, grants in the CAS authorization layer do not provide access.

If host-layer access protections are not in place, denials in the CAS authorization layer do not fully prevent access.

Which Host Account Matters

The account under which a CAS server process runs must have appropriate host-layer access to target directories and files. Depending on context, that account is either an individual user's personal host account or the CAS server's shared service account. The following table provides examples:

Context	Identity of the CAS Process
A request from the SAS Visual Analytics web application to a CAS server on Linux.	The CAS server's shared service account.*
A request from the SAS Visual Analytics web application to a CAS server on Windows.	The requesting user's individual host account.
A request in a programming-only deployment.	The requesting user's individual host account.

* Unless the requesting user is a member of the custom group CASHostAccountRequired.

Host Access in a Programming-Only Deployment

In a programming-only deployment, every user accesses a CAS server's host files and directories using his or her individual host account. Users must have host access, so it is possible for them to access the back-end machine directly, bypassing the CAS authorization layer.

If you have sensitive data, ensure that all CAS access distinctions are mirrored in the host authorization layer. For example, if you use the CAS authorization system to deny userA Read access to a path-based caslib called caslibA, you must also set up host access controls that prevent userA from accessing that caslib (directory). Without such protection, userA could use a host command to copy files from the caslibA directory to the directory for a caslib that userA can access from CAS.

Host Access in a Full Deployment

Note: This section applies to a CAS server on Linux. Host access from a CAS server on Windows is always under each user's individual identity.

In a full deployment, host access from CAS is sometimes under individual identity and sometimes under a shared identity. Here are details:

- For users who access CAS only from a programming interface such as SAS Studio, all host access from CAS is under each user's individual identity. For such users, you must mirror CAS layer access distinctions in the host layer.
- For users who access CAS only from a visual interface such as SAS Visual Analytics or SAS Environment Manager, all host access from CAS is under a shared identity. Such users need CAS layer access to data, but they do not need host access to data. Only the shared identity needs host access to data.

  For such users, there is no reason to create host access controls that mirror your CAS access controls. Of course, you should always host protect your resources in accordance with your security requirements.
- For users who access CAS from both types of interface, host access and experience vary depending on the type of interface that is used. For example, the personal caslib that they use from programming interfaces is not automatically accessible from the visual interfaces.

  You can align access and experience for such users by assigning them to the CASHostAccountRequired group. For members of that group, host access from CAS is always under individual identity. Before you use this approach, review the associated limitations. See The CASHostAccountRequired Custom Group in *SAS Viya Administration: Identity Management*.

Using CAS to Modify Host Access

You can use CAS to add host access controls to a new directory or file in the following circumstances:

- You add a caslib of the type PATH or DNFS and specify to create a directory (on Linux).
- You save a table to a file (on Linux).

This functionality is provided by the permission parameter in the CAS actions tables.addCaslib and tables.save.

Here are the available fixed values:

Value	Octal*	Description
Private	700	Grants Read and Write access to only the owner.
GroupRead	750	Grants Read and Write access to the owner. Grants Read access to the owning group.
GroupWrite	770	Grants Read and Write access to the owner and the owning group.
GroupWritePublicRead	775	Grants Read and Write access to the owner and the owning group. Grants Read access to everyone.
PublicRead	755	Grants Read and Write access to the owner. Grants Read access to everyone.
PublicWrite	777	Grants Read and Write access to everyone.

* These octal values are for directories. For saved files (SASHDAT and CSV), the executable bit is not set.

Here are additional details:

- The owner is the host account that creates the directory or file. the owner always gets full access, regardless of whether you use the permission parameter. Use the permission parameter to further refine access.
- The owning group is the host group that is the primary group for the owner.
- You can specify an octal. You are not limited to the fixed values that are listed in the preceding table.

Access to Files in the Public Caslib

Initially, all users have host-layer Write access to the directory for the Public caslib. To limit host-layer Write access to that directory, adjust host-layer access controls.

For example, on a Linux directory you can set a user ownership access flag that is called the sticky bit. If the sticky bit is set for the Public caslib’s directory, only the host account that creates a file in that directory can remove that file. If the sticky bit is not set for the Public caslib’s directory, any user with Write access to the directory can remove files from that directory.

Note: For information about the Public caslib, see Predefined Caslibs in *SAS Viya Administration: Data*.

24

General Authorization

General Authorization: Overview

To learn about the general authorization system, see Concepts.

To manage access, use the interface that best meets your needs. Here are suggestions:

- To adjust access to content (such as folders and reports), use the Authorization window.
- To manipulate rules directly, use the Rules page or the command-line interface.

General Authorization: How to (Authorization Window)

Introduction

These instructions explain how to set permissions on content (such as folders and reports) using the Authorization window.

Navigation

To access the Authorization window from SAS Environment Manager:

1. In the applications menu (☰), under **Administration**, select **Manage Environment**.
2. In the vertical navigation bar, click 🗎.
3. Locate and select the object.
4. Right-click, and select **View authorization** or **Edit authorization**.

Note: If the **Edit authorization** item is not available, you are not authorized to modify access to the object.

> **TIP** The Authorization window is also available in SAS Drive to users who assume their membership in the SAS Administrators group. To expand or reduce that access, see Access to Functionality on page 355. (That access is provided by the Read permission on the following object URI: /authorizationDialog.)

Examine Access

For each principal and permission, the following icons describe effective (net) access to the current object:

Icon	Meaning
✓	Authorized
⊡	Conditional
⊘	Not Authorized
●	Unknown
◆	Direct (indicates that effective access comes from a direct setting)*
⋖	Not Authorized (but can share)**

* This icon indicates that effective access comes from a permission that is directly assigned to the specified principal on the current object. If a direct setting exists but does not win (does not determine effective access), a diamond is not displayed.

** This icon is applicable to only the **Secure** and **Secure (convey)** columns. This icon is displayed only if Secure access is not granted and sharing is possible. If Secure access is authorized, only the Authorized icon is displayed, because the ability to share is inherent in Secure access.

The scope of the display is as follows:

- There is always a row for Authenticated Users.
- There is always a row for you, the currently connected user who is using the display.
- There is a row for each principal that is assigned to a rule that affects access to the current object. The exception is that internal service principals (for example, sasapp or sas.folders) are not displayed in the Authorization window.
- If you add an identity and do not give that identity at least one direct setting, that identity is automatically removed from the display.
- You cannot directly remove a row. If you remove all direct settings for an identity and there is no other reason for that identity to be displayed, the identity is automatically removed from the display.
- For a non-container object (such as a report), only the Read, Update, Delete, and Secure permissions are displayed. The Create permission is not applicable to an individual content object.
- For a container (such as a folder), two sets of permissions are displayed:
  - The first set of permissions affects access to the object, including the ability to add members to and remove members from the object. This set of permissions has no effect on the folder’s members.
  - The second set of permissions affects the access that this object conveys to its child members. See Inheritance.

> **TIP** An effective access value of Not Authorized on the conveyed side of a folder’s Authorization window does not guarantee that access to child members is not authorized. A direct setting on the child or another influencing rule might provide access to the child member. For example, when you create a top-level folder, the effective access values on the conveyed side for SAS Administrators are all Not Authorized. However, SAS Administrators does have effective access to a folder that you add below the top-level folder. That access comes from a predefined rule that gives SAS Administrators access to all folders.

Provide Access

1 Open the Edit Authorization window for the target object.

2 If the principal that you want to work with is not already listed, click 👤.

Note: If guest access is enabled, you must select either **Add Identities** or **Add Guest** after you click 👤.

3 Click an effective access icon (for example, 🛇).

4 In the pop-up window, select **Grant** as the direct setting.

Note: If you cannot change a direct setting, you do not have Secure permission for the current object.

5 In the Edit Authorization window, click **Preview**. Examine the impact of your unsaved change.

Note: If there is a relevant prohibit setting anywhere in the system, that setting has precedence over the direct grant that you added. In that case, the effective (net) result is **Not Authorized** (🛇), and a diamond is not displayed.

6 Click **Save**.

Limit Access

Any access that is not granted is implicitly denied. The preferred approach is to grant selectively and to avoid use of prohibit settings.

If you must add a prohibit setting, make sure that you do not inadvertently block your own access, particularly for the Read and Secure permissions. If you do block your own access, see Troubleshooting.

CAUTION! A prohibit setting has absolute precedence, even if a more specific grant setting exists.

Add a Condition

To provide access within a particular scope or set of circumstances, add a condition.

1 Open the Edit Authorization window for the target object.

2 If the principal that you want to work with is not already listed, click 👤.

Note: If guest access is enabled, you must select either **Add Identities** or **Add Guest** after you click 👤.

3 Click an effective access icon (for example, 🛇).

4 In the pop-up window, select **Conditional Grant**.

Note: A conditional prohibit setting does not provide access. A conditional prohibit setting blocks all access within its scope, regardless of any more specific grant settings. A conditional prohibit setting can limit access that is provided by a grant or conditional grant setting.

5 In the Condition window, create an expression that specifies the scope and circumstances in which access is granted. Your syntax is validated when you click **OK**. See Rule Conditions.

6 In the Edit Authorization window, click **Preview**. Examine the impact of your unsaved change.

7 Click **Save**.

Edit a Condition

1 In the Edit Authorization window for an object, click the effective access icon for the direct conditional setting that you want to modify.

2 In the pop-up window, next to the **Conditional Grant** or **Conditional Prohibit** direct setting, click <=≠>.

3 In the Condition window, edit the expression. Your syntax is validated when you click **OK**. See Rule Conditions.

4 In the Edit Authorization window, click **Preview**. Examine the impact of your unsaved change.

5 Click **Save**.

Delete a Condition

1 In the Edit Authorization window for an object, click the effective access icon for the direct conditional setting that you want to delete.

2 In the pop-up window, next to the **Conditional Grant** or **Conditional Prohibit** direct setting, click <=≠>.

3 In the Condition window, delete the expression. Click **OK**.

4 In the Edit Authorization window, click **Preview**. Examine the impact of your unsaved change.

5 Click **Save**.

Remove a Direct Setting

1 Open the Edit Authorization window for an object.

2 In the cell that has the direct setting that you want to remove, click the effective access icon. In the pop-up window, select **(none)** as the direct setting.

Note: If you cannot change the direct setting, you do not have Secure permission for the current object.

3 In the Edit Authorization window, click **Preview**. Examine the impact of your unsaved change.

Note: Any identities that are no longer principals are automatically removed.

4 Click **Save**.

Identify the Source of Effective Access

To determine which rules and shares contribute to a particular effective access result, examine the origins information for that result.

1 Open the View Authorization window for the target content object.

2. Click the effective access icon for which you want origins information.

3. In the pop-up window, the **Contributing Rules** tab provides a read-only display of all applicable rules, except share-based rules. Here are tips:

   - If you opened the Edit Authorization window and have unsaved (and un-previewed) changes, the **Contributing Rules** tab is disabled.
   - To view additional details, add columns to the table. Click ▤ and select **Manage columns**.
   - To directly modify a rule, use the Rules page.

4. In the pop-up window, the **Contributing Shares** tab provides a read-only display of all relevant shares. Here are tips:

   - If sharing is disabled, the **Contributing Shares** tab is not displayed.
   - To view additional details, add columns to the table.
   - To directly modify a rule, use the Rules page.
   - To manage shares, see “Administrative Oversight on the Rules Page”.

General Authorization: How to (Rules Page)

Introduction

These instructions explain how to directly manage general authorization rules using SAS Environment Manager.

Navigation

In the applications menu (☰), under **Administration**, select **Manage Environment**. In the vertical navigation bar, select ☷.

The **Rules** page is an advanced interface. It is available to only SAS Administrators. You can use a simpler interface to set permissions on content such as folders and reports.

Rules Page

Use the **Rules** page to manage authorization rules directly. Here are examples:

- View and filter rules.
- Enable and disable rules.
- Replace the principal in a rule.
- View and edit a rule’s description or reason.
- Use an existing rule as the basis for a new rule.
- Work with rules that affect access to functionality.
- View conditions for multiple rules at the same time.

Here are additional details:

- To ensure that all rules that should be visible to you are displayed, refresh the display and click **Reset all** in the **Rules Filter** pane.

- On the **Rules** page, you cannot see rules that are directly assigned to objects for which you lack the Secure permission.
- The search field searches only the **Object URI**, **Description**, **Reason**, and **Condition** columns.
- To add, remove, or reorder columns, click ☰ , and select **Manage columns**.
- You cannot sort the values within a column.
- You can use the **Rules Filter** pane to view a subset of rules. Some filters take effect immediately, other filters take effect after you click **Apply**.
- You can clear a filter by clicking its **Reset** link. You can clear all filters by clicking **Reset all** at the top of the pane.
- The **Guest** principal type is always listed, regardless of whether guest access is enabled.
- Display names for users and groups are not available on the **Rules** page.
- For details about rule attributes, see Rule Attributes.

Add a Rule

1. On the **Rules** page, click ✱.
2. In the New Rule window, provide values for at least the required attributes. Here are tips:
   - In some fields, you can click 📁 or 👤 to browse instead of directly entering a value.
   - If a warning indicates that the **Principal** value cannot be validated, make sure the value is an ID, not a display name. If the principal is a service account (such as sasapp or sas.folders), you can ignore the warning.
   - To populate the list of permissions, use the **Clear All**, **Select All**, and **Choose** buttons.
3. Click **Save**.
4. On the **Rules** page, right-click the new rule, and select **Properties**. Verify that the attributes of the new rule are as you intended.
5. If the rule affects a content object (such as a folder or report), use the Authorization window to verify that the results are as you intended.

Edit a Rule

1. On the **Rules** page, select a rule, and then click ✎.

   Note: You cannot edit a share-based rule. See “Sharing: Details for Administrators”.
2. In the Edit Rule window, modify attributes as needed. Here are tips:
   - If a warning indicates that the **Principal** value cannot be validated, make sure the value is an ID, not a display name. If the principal is a service account (such as sasapp or sas.folders), you can ignore the warning.
   - If the rule does not have a condition, an **Add Condition** button is present. If the rule has a condition, an **Edit Condition** button is displayed.
3. Click **Save**.
4. On the **Rules** page, right-click the rule, and select **Properties**. Verify that the attributes of the new rule are as you intended.

5 If the rule affects a content object (such as a folder or report), use the Authorization window to verify that the results are as you intended.

Copy a Rule

1 On the **Rules** page, select a rule, and then click ⧉.

Note: You cannot copy a share-based rule. See "Sharing: Details for Administrators".

2 In the New Rule window, modify attributes as needed.

3 Click **Save**.

4 On the **Rules** page, right-click the new rule, and select **Properties**. Verify that the attributes of the new rule are as you intended.

5 If the rule affects a content object (such as a folder or report), use the Authorization window to verify that the results are as you intended.

Delete a Rule

1 On the **Rules** page, select a rule, and then click 🗑.

2 In the confirmation window, click **Delete**.

3 If the rule affects a content object (such as a folder or report), use the Authorization window to verify that the results are as you intended.

Edit a Condition

1 On the **Rules** page, select a rule, and then click ✎.

2 In the Edit Rule window, click **Edit Condition**.

Note: If a rule does not have a condition, an **Add Condition** button is present. If a rule has a condition, an **Edit Condition** button is displayed.

3 In the Edit Condition window, edit the expression. Your syntax is validated when you click **OK**.

4 Click **Save**.

5 If the rule affects a content object (such as a folder or report), use the Authorization window to verify that the results are as you intended.

Delete a Condition

1 On the **Rules** page, select a rule, and then click ✎.

2 In the Edit Rule window, click **Edit Condition**.

Note: If a rule does not have a condition, an **Add Condition** button is present. If a rule has a condition, an **Edit Condition** button is displayed.

3 In the Edit Condition window, delete the expression.

4. Click **Save**.

5. On the **Rules** page, verify that the condition no longer exists. Right-click on the rule, select **Properties**, and verify that the **Condition** field is blank.

6. If the rule affects a content object (such as a folder or report), use the Authorization window to verify that the results are as you intended.

Locate a Particular Rule

Here are general tips:

- Filter requirements are cumulative. For example, if you set two filters, only rules that meet both criteria are displayed.
- Remember to click **Apply** after you set or modify certain filters.
- You can search for rules that contain specified text in the ObjectUri, Description, Reason, or Condition field. The search is not case-sensitive. The search looks for the specified text in any of the four supported fields.

Here are tips for locating a rule by date:

- In the **Rules Filter** pane, under **Date Modified**, click 📅 to select a date or date range.
- Each rule's **Date Modified** value indicates when the rule was created or most recently modified.
- To add the **Date Modified** column to the display, click ▤ , and select **Manage columns**.

Here are tips for locating a rule by identity:

- The **Modified By** filter is based on who created or last updated a rule.
- To add the **Modified By** column to the display, click ▤ , and select **Manage columns**.
- The **Principal** filter is based on who a rule is assigned to.
- Both the **Modified By** and the **Principal** filters use ID values, not display name values. For example, to display only rules that were created by userA, specify `usera` in the **Modified By** filter. To display only rules that are assigned to the SAS Administrators group, specify **SASAdministrators** in the **Principal** filter.
- Rules that are predefined or generated can have a **Modified By** value that does not correspond to a user or group that is known to the identities service.

Here are tips for locating a rule by the URI of the rule's target:

- To view only those rules that target a specific content object, use the technique that is appropriate for the type of URI, as follows:
  - Browse content objects for the object URI for a rule target. In the drop-down list under **Object URI**, select **URI**.
  - Browse container objects for the container URI for a rule target. In the drop-down list under **Container URI**, select **URI**.
- To view only those rules that do not specify an object URI, select **(blank URI)** from the drop-down list under **Object URI**.
- To view rules that either specify `/**` as the object URI or do not specify an object URI, select **(global URI)** from the drop-down list under **Object URI**.
- To view only those rules that do not specify a container URI, select **(blank URI)** from the drop-down list under **Container URI**.

For tips for locating share-based rules, see "Administrative Oversight on the Rules Page" on page 441.

Extend the Ability to Create Top-Level Folders

Initially, only SAS Administrators can create top-level folders. Here are examples of how you can use the **Rules** page to extend that ability to other users:

- To enable all authenticated users to create top-level folders, locate the rule that targets the object URI `/folders/folders` and grants the `Add` and `Read` permissions to `Authenticated Users`. Edit that rule so that it also grants the `Create` permission.
- To enable a group that has the ID `groupA` to create top-level folders, add a new rule that targets the object URI `/folders/folders` and grants the `Create` permission to `groupA`.

General Authorization: Concepts

Scope

General authorization manages access to the following resources:

- content, such as folders and reports
- functionality, such as applications, features, and services

> **TIP** For information about the sharing endpoint within the authorization service, see “Sharing: Details for Administrators”.

Key Terms

Term	Definition
Rule	A composite of authorization elements. Example: A rule grants groupA the Read permission for folderA.
Target	The affected resource, such as an individual object, a set of objects, a service, or a service endpoint. Examples: folderA, reportA
Principal	The user, group, or construct to which a rule is assigned. Examples: UserA, GroupA, Authenticated Users
Permission	A type of access. Values: Add, Create, Delete, Read, Remove, Secure, Update
Setting	In a rule, the indication of whether (and to what extent) access is provided. Values: Grant, Conditional Grant, Prohibit, Conditional Prohibit
Condition	In a rule, the constraint expression. Most rules do not include a condition. Example: `currentUser() == #preferenceOwner`

Term	Definition
Effective access	A context-neutral description of the net result of all relevant rules. Effective access does not incorporate evaluation of any conditions. Values: Authorized, Not Authorized, Conditional
Access outcome	In a context-aware access request, the authorization decision. Values: Authorized, Not Authorized

Principals

The principal in an authorization rule is the user, group, or construct to which the rule is assigned. The general authorization system supports the following principals:

- A user is either an individual authenticated user or a service account.
- A user group is either a custom group or a group in your authentication provider.
- Authenticated Users is the principal type that represents all authenticated users.
- Everyone is the principal type that represents all principals.
- Guest is the principal type that facilitates guest access. Guest is not part of Authenticated Users, but is part of Everyone.

Note: When a principal is deleted, rules that are assigned to that principal are not automatically deleted. Such rules are reused if a new principal of the same type and ID is created. The general authorization system does not have an automated cleanup process for orphaned rules.

Administrators

The SAS Administrators group provides access throughout the general authorization system. A predefined rule grants all permissions throughout the general authorization system to the SAS Administrators group. However, the SAS Administrators group is not unrestricted or exempt from authorization requirements.

For more information, see Predefined Custom Groups in *SAS Viya Administration: Identity Management*.

Inheritance

Access flows through a hierarchy of containers. Each container conveys settings to its child members. Each child member inherits settings from its parent container. For example, a folder’s child members might include reports and other folders.

Note: A reference member (such as a shortcut) does not inherit access from its parent folder.

You can manage access that a container conveys independently from access to the container. Here are examples of that separation:

- In a folder’s Authorization window, the first set of permissions depicts access to the folder, and the second set of permissions depicts access that the folder conveys.
- On the **Rules** page, a rule that targets a folder can affect either or both types of access, depending on which fields (**Object URI**, **Container URI**, or both) are populated.

A rule that targets the object aspect of a container (the container’s objectUri attribute) has different effects than a rule that targets the container aspect of a container (the container’s containerUri attribute). Here are details, using folderA as an example container.

Rule Target	Potential Impact of the Rule
folderA (as an object)	Can affect the ability to read, update, or delete folderA. Can affect the ability to add or remove members for folderA. Settings are not conveyed to the objects within folderA.
folderA (as a container)	Settings are conveyed to folderA's child members.
folderA (as an object and as a container)	Can affect the ability to read, update, or delete folderA. Can affect the ability to add or remove members for folderA. Container settings are conveyed to folderA's child members.

Some interfaces enable you to create rules that target both aspects of access to a container. However, containerUri settings are never derived from or implicitly matched to objectUri settings. This separation enables you to provide Write access to the objects in a container without providing Write access to the container itself.

Permissions

Permission	Affected Activity
Create*	Create a new object.
Read	Read an object.
Update	Update or edit an object.
Delete	Delete an object.
Secure	Set permissions on an object (manipulate the object's direct rules).
Add**	Put an object into a container.
Remove***	Move an object out of a container.

* Applicable for a service, service endpoint, or media type.

** Applicable for a container, such as a folder.

*** Applicable for a container, such as a folder. In SAS Environment Manager, also affects the ability to delete a child member from a folder.

Table A.1 Permission Settings

Setting	Description
Prohibit	Prevents access.
Conditional Prohibit	Prevents access in specified circumstances and scope.
Grant	Provides access, unless there is a relevant Prohibit or Conditional Prohibit setting.
Conditional Grant	Provides access in specified circumstances and scope, unless there is a relevant Prohibit or Conditional Prohibit setting.

Note: Setting is a compound of rule type and whether a condition is present. Setting is a client-layer convenience construct, not a service-layer rule attribute.

Permissions by Task

Introduction

To provide sufficient access to complete a task, you must consider both the availability of functionality and the availability of individual objects. For example, here are the primary requirements for creating and saving a new report:

- the ability to create reports (for example, the Create permission for the service endpoint that controls the ability to create reports). For more information, see Access to Functionality in *SAS Viya Administration: Identity Management*.
- the ability to add members to the target folder (the Add permission for the object aspect of the target folder).
- the ability to see and update the new report (for example, the Read and Update permissions for the target folder's containerUri). See Inheritance.

Note: In addition to the requirements that are documented in this topic, most interfaces enable you to interact with only those resources for which you have Read access.

Details for Top-Level Folders

Here are the required permissions for managing top-level folders:

Task	Service URI	Top-Level Folder
Add a top-level folder*	Create	-
Delete a top-level folder	-	Delete
Rename a top-level folder	-	Update
Manage access to a top-level folder	-	Secure

* Initially, only members of the SAS Administrators group can add top-level folders. See Extend the Ability to Create Top-Level Folders.

Initial access to a new top-level folder is as follows:

- The user who creates a new top-level folder has full access to that folder. Automatically generated direct grants provide that access.
- SAS Administrators has full access to every new top-level folder. The predefined rule that grants SAS Administrators permissions for all objects provides that access. See Examine Access.

See also "Content Management: How To".

Details for Child Members

A child member is an object that is in a folder and is not a reference member. For example, folders (other than top-level folders) and reports are child members.

Here are the required permissions for managing child members:

Task	Service URI	Child Member	Current Parent Folder	New Parent Folder
Add	Create	(Read, Update)*	Add	-
Delete	-	Delete	Remove**	-
Update	-	Update	-	-
Rename	-	Update	-	-
Move	-	Update	Remove	Add
Manage access	-	Secure	-	-

* These permissions are required for only objects that are updated during their creation process. For example, the process of creating and saving a new report includes an internal update to the content of the new report. The necessary access is usually conveyed from the parent folder.

** This requirement applies only in SAS Environment Manager.

Initial access to a new child member is determined by inheritance and any other influencing rules, including any automatically generated direct settings.

See also "Content Management: How To".

Details for Reference Members

A reference member is a pointer to another resource. For example, an item in a list of favorite or recent objects is a reference member.

Access to a reference member is independent from access to the referenced resource. For example, having Read access to a favorite that points to ReportA does not equate to having Read access to ReportA. Conversely, having Read access to ReportA does not equate to having Read access to all reference members that point at ReportA.

Access to reference members is as follows:

- Anyone who has Read access to a folder can see all reference members in that folder.
- Anyone who has Remove access to a folder can delete all reference members in that folder.
- You cannot set permissions on a reference member.
- A reference member does not inherit permissions.

See also "Content Management: Concepts".

Details for Authorization Rules

Anyone who has the Secure permission for a resource can add, modify, and delete direct rules for that resource.

In the initial configuration, the SAS Administrators group can add, modify, and delete all authorization rules.

Note: Sharing can enable users who do not have the Secure permission for an object to share that object, indirectly creating a share-based authorization rule.

Details for Models

Registered models are stored beneath the **Model Repositories** node in the **SAS Content** tree. Initially, all authenticated users have Read and Write access to all registered models. That access is provided by predefined and generated rules that grant broad access to model repositories. See "Default Permissions" in *SAS Model Manager: Administrator's Guide*.

Unregistered models are stored in Model Studio projects beneath a user’s **My Folder**. Unregistered models are initially private. See “Sharing a Project” in *SAS Visual Data Mining and Machine Learning: User’s Guide*.

HTTP Mapping

Here are the standard mappings of HTTP verbs to permissions:

HTTP Verb	Permission
POST	Create
DELETE	Delete
GET, OPTIONS, HEAD	Read
PUT, PATCH	Update

Some actions override the default mappings and instead require a different permission.

Rule Targets

About Rule Targets

Each rule affects a target resource (or set of resources), as identified in the rule’s target-related attributes (objectUri, containerUri, and the media type properties). A rule can specify any, all, or none of the target-related attributes.

CAUTION! A rule that does not specify at least one target-related attribute affects access to all resources throughout the general authorization system.

About Target URIs

> **TIP** Do not confuse uniform resource identifiers (URIs) and folder locations. Although both constructs use paths to reference resources, the two constructs are entirely independent and distinct. For example, a report might have a URI of `/reports/reports/qwe3429ryjw12567` and a folder location of `/SAS Content/Public/reportA`.

A target URI is the value in an authorization rule’s objectUri or containerUri attribute. A target URI references a resource such as a service, service endpoint, or content object. The authorization service sees target URIs as strings. The authorization service does not know whether a particular URI represents an application, a feature, an object, or a collection of objects.

Authorization decisions for a particular resource evaluate all rules that have target URIs that match the requested URI. To compare a requested URI to the target URIs in authorization rules, the authorization service uses Ant-style pattern matching. In that pattern matching, the authorization service supports a wildcard character, so that the target URI in a rule can be more general than the target URI in an individual request. Here are details:

- A single wildcard (`*`) matches a single element in a URI path. For example, it matches the name of a service or the ID of an object.
- A double wildcard (`**`) matches any number of consecutive elements in a URI path. For example, it matches a forward slash (`/`), the name of a service, the ID of an object, a multi-element path (`/lmn/xyz/`), or the absence of any element (null).

The following tables provide examples. The first table uses fictional URIs to demonstrate the principles. The other tables use selected URIs to illustrate how the principles apply in context.

Table A.2 Scope of a Target URI: Principles

Rule's Target URI	Rule's Relevance
`/**`	Relevant to all requests.
`/abc`	Relevant to requests for exactly this URI: `/abc`.
`/abc/`	Relevant to requests for exactly this URI: `/abc/`.
`/abc/*`	Relevant to requests for URIs such as `/abc/lmn` (not `/abc`, `/abc/`, or URIs such as `/abc/lmn/` and `/abc/lmn/xyz`).
`/abc/*/xyz`	Relevant to requests for URIs such as `/abc/lmn/xyz` (not `/abc/xyz`).
`/abc/**`	Relevant to requests for `/abc`, `/abc/`, and URIs such as `/abc/lmn`, `/abc/lmn/`, and `/abc/lmn/xyz`.
`/abc/**/xyz`	Relevant to requests for `/abc/xyz`, and URIs such as `/abc/lmn/xyz` and `/abc/lmn/rst/xyz`.

Table A.3 Scope of a Target URI: Examples from Reports Service

Rule's Target URI	Rule's Scope
`/reports`	The root of the reports service. Relevant to requests that omit a trailing slash.
`/reports/`	The root of the reports service. Relevant to requests that include a trailing slash.
`/reports/*`	All first-level endpoints within the service (not the service root, lower-level endpoints, or individual reports).
`/reports/*/`*report-ID*	All first-level endpoints within the service (for the specified report).
`/reports/**`	All reports, all endpoints within the service, and the service root.
`/reports/**/`*report-ID*	All endpoints within the service (for the specified report).
`/reports/reports/*`	All reports, and all second-level endpoints beneath the reports endpoint.
`/reports/reports/`*report-ID*	The specified report.

Table A.4 *Scope of a Target URI: Examples from Model Repository Service*

Rule’s Target URI	Rule’s Scope
`/modelRepository`	The root of the model repository service. Relevant to requests that omit a trailing slash.
`/modelRepository/`	The root of the model repository service. Relevant to requests that include a trailing slash.
`/modelRepository/*`	All first-level endpoints within the service (not the service root, lower-level endpoints, or individual objects).
`/modelRepository/**`	All objects that are managed by the service (repositories, projects, and models), all endpoints within the service, and the service root.
`/modelRepository/repositories/*`	All repositories, and all second-level endpoints that are beneath the repositories endpoint.
`/modelRepository/projects/*`	All projects, and all second-level endpoints that are beneath the projects endpoint.
`/modelRepository/models/*`	All models, and all second-level endpoints that are beneath the models endpoint.
`/modelRepository/models/`*model-ID*	The specified model.

TIP As the preceding examples make clear, not all services have identical structures. Before you create or customize a rule that uses a wildcard in its target URI, make sure you understand the structure of the service, and the resulting scope of the rule. (For services that have published APIs, you can use developer.sas.com to discover structure.)

objectUri

A rule that targets an objectUri affects access to the referenced resource. Here are examples:

- A rule that targets a folder’s objectUri affects access to that folder.
- A rule that targets a report’s objectUri affects access to that report.
- A rule that targets a service’s objectUri (and does not target a specific object instance) affects access to functionality. See Access to Functionality in *SAS Viya Administration: Identity Management*.

In general, rules that specify the objectUri for a content object (such as a folder or report) should include the `/**` suffix. Inadvertently omitting the suffix narrows the effects of a rule and can yield unintended results due to insufficient access. Here are details:

- The Authorization window appends the `/**` suffix to new rules that target objectUris. You cannot see or modify URIs in the Authorization window.
- The New Rule window appends the `/**` suffix when you populate the **Object URI** field by clicking 📂 and selecting a content object in the Choose an Item window. You can see and modify URIs in the New Rule and Edit Rule windows.
- In other interfaces and contexts, you must remember to include a suffix, when appropriate.

containerUri

A rule that targets a folder's container URI affects access that the folder conveys to its child members.

Note: You cannot append the `/**` suffix to a containerUri. The `/**` suffix does not reference contained objects (such as reports within a folder). The `/**` suffix has nothing to do with container-based object inheritance.

mediaType

This property has no effect and is deprecated. It is replaced by the more specific properties acceptType, acceptItemType, and contentType.

acceptType, acceptItemType, and contentType

The three media type properties (acceptType, acceptItemType, and contentType) provide specialized, advanced support for refining the scope of a rule. Most rules do not specify these attributes. To view or set these attributes, use the REST API.

You can use media type properties to limit a rule's applicability so that it targets only a particular media type in only a specified context.

- If a rule specifies one media type property, the rule applies only to requests that specify a corresponding value for the corresponding attribute.

  For example, if a rule specifies `contentType = 'application/pdf'` then the rule applies only to requests in which the following query parameter is specified: `(contentType == application/pdf).`

- If a rule specifies multiple media type properties, the rule applies only to requests that specify corresponding values for all specified media type properties.

  For example, if a rule specifies the value `application/pdf` in both the contentType and the acceptType properties, then the rule applies to only requests for which the following compound set of parameters is specified: `(contentType == application/pdf and acceptType == application/pdf).`

These properties are processed as follows:

- In response to an authorization request that specifies one or more of the media type query parameters (contentType, acceptType, or acceptItemType), only matching rules are applied. A troubleshooting item provides an example.
- In response to an authorization request that does not specify a media type query parameter (contentType, acceptType, or acceptItemType), any rules that specify one or more media type properties are ignored.

Rule Attributes

Attribute	Description
Target-related attributes:	
objectUri	A relative URI that represents a resource such as a report, a folder, a service, or a service endpoint. Character limit: 500
containerUri	A relative URI that represents the container aspect of a container, such as a folder. Rules that specify a containerUri affect access that a container conveys to its child members. Character limit: 500

Attribute	Description
mediaType	This property is deprecated.
acceptType acceptItemType contentType	See "acceptType, acceptItemType, and contentType".
Principal-related attributes:	
principalType	Three of the values (Authenticated Users, Everyone, and Guest) represent classes of users. Everyone includes all authenticated users and any guest users. - Assign broad grant rules to Authenticated Users. - Do not assign prohibit rules to Everyone or to Authenticated Users. Such rules block access for all users, including yourself.
principal	The unique string that identifies a particular user or group by its ID. If principalType is **`user`** or **`group`**, you must specify a value for this attribute. Character limit: 100
Access-related attributes:	
type	The indication of whether a rule blocks (prohibit) or attempts to provide (grant) access. Prohibit rules have absolute precedence.
condition	An expression that limits the scope or applicability of a rule. Character limit: 5120
permissions	A list of access types. At least one permission is required.
enabled	The indication of whether a rule is enabled. By default, rules are enabled. To temporarily prevent a rule from being enforced, disable the rule.*
Documentation-related attributes:	
description	Text that documents a rule for administrative purposes. Character limit: 1000
reason	Text that provides information for end users, where supported by a client. For example, a prohibit rule's reason could be displayed to an end user as part of an `access denied` message. Character limit: 1000

* In SAS Environment Manager, this attribute is labeled **Rule Status**, and has values of **Enabled** and **Disabled**.

Rule Conditions

Overview

A condition is a Boolean expression that limits the scope of a rule.

- A rule that has no condition is always applied.

- A rule that has a condition that evaluates to **`true`** for a particular access request is applied in the authorization decision process for that access request.
- A rule that has a condition that evaluates to **`false`** for a particular access request is ignored in the authorization decision process for that access request.
- If a rule has an invalid condition, an error is logged and access is restricted as follows:
  - If a grant rule has an invalid condition, the rule is always ignored.
  - If a prohibit rule has an invalid condition, the rule is always applied.

You can specify a condition in inclusive or exclusive terms. Here are two examples:

- A rule grants the Read permission to GroupA for folderA, with a condition that the rule applies only on weekdays. A request from a member of GroupA to access folderA on Sunday is outside the condition. For that access request, the condition is false, so the rule is not applicable (it does not provide access).

  Note: A conditional grant rule provides access in specified circumstances, but it does not prevent access outside of those circumstances.

- A rule prohibits the Read permission for GroupA for folderA, with a condition that the rule does not apply on weekends. A request from a member of GroupA to access the folder on Sunday is inside the condition. For that access request, the condition is true, so the rule is applicable (access is blocked).

  Note: A conditional prohibit rule prevents access in specified circumstances, but it does not provide access outside of those circumstances.

Condition Syntax

- Conditions are written and stored in Spring Expression Language (SpEL).
- Boolean operators (AND, OR, and NOT) and parentheses are supported. For example, the following condition always evaluates to **`true`**:

  ```
  (4 < 6) and (5 > 3)
  ```

- Built-in functions correspond to attributes of the requestor or the environment. You must append **`()`** to each built-in function (for example, **`currentUser()`**).

  Note: You can use a constant instead of a function. However, functions are often more useful because they are dynamic. At request time, actual context-specific values are dynamically substituted into each function.

- Variables correspond to attributes of the target. You must prepend **`#`** to each variable (for example, **`#userId`**).

  Note: The available condition variables for a particular type of object are designated in the service for that object type. For example, the preference service designates the **`userId`** attribute on preference objects as available for use as a condition variable.

Built-In Functions

Table A.5 Location-Based Functions

Function	Description	Type
locale()	Locale of the client that made the request (for example, **`en_US`**).	String
remoteHost()	Name of the client machine that made the request.	String
remoteIp()	IP address of the client machine that made the request.	String

Function	Description	Type
serverIp()	IP address of the middle-tier server that received the request.	String
serverName()	Machine name of the middle-tier server that received the request.	String
serverPort()	Port of the middle-tier server that received the request.	int

Table A.6 Target-Based Functions

Function	Description	Type
contentType()	Content type of the target (for example, `application/vnd.sas.credential.domain+json`).	String
contentLength()	Length of the request.	long
uri()	URI of the target.	String

Table A.7 Time-Based Functions

Function	Description	Type
timestamp	Coordinated Universal Time (UTC) timestamp.	ZonedDateTime
timestamp(zoneId)	Timestamp of the request, based on a specified zoneid.*	ZonedDateTime
localTime(zoneId)	Time of the request, based on a specified zoneid.*	LocalTime
localDate(zoneId)	Date of the request, based on a specified zoneid.*	LocalDate
localDateTime(zoneId)	Date and time of the request, based on a specified zoneid.*	LocalDateTime

* A time zone ID that is valid for java.time.ZoneId.

Table A.8 Other Functions

Function	Description	Type
currentUser()	Identifier for the currently connected user.	String
groupsForCurrentUser()	Identifier for each group to which the current user belongs. Nested memberships are included. Unassumed memberships are not included.	List
method()	Method that the request invoked (for example, `GET`).	String
protocol()	Protocol of the request (for example, `HTTP/1.1`).	String
header(headerName)	Headers for a specified headerName.	List

Examples of Conditions

This condition makes its associated rule applicable only for weekday requests (in the US Eastern time zone):

```
localDate('US/Eastern').dayOfWeek != T(java.time.DayOfWeek).SUNDAY and
localDate('US/Eastern').dayOfWeek != T(java.time.DayOfWeek).SATURDAY
```

This condition makes its associated rule applicable only if the target's user ID is the same as the requesting user's ID:

```
#userId == currentUser()
```

Evaluation of Conditions

Here are key points about evaluation of conditions:

- If a request does not meet the criteria in a condition, the request is outside that condition. If a request meets the criteria in a condition, the request is inside that condition.
- In a description of effective access, there is no request context, so conditions are not evaluated. Even an atypical condition that is always true (1=1) or never true (1>2) yields an effective access result of Conditional in certain scenarios. A condition is evaluated only in the context of a specific request.
- In an actual request, there is a request context. Any relevant conditions are evaluated, and a definitive answer is provided (Authorized or Not Authorized).

Authorization Decisions

Precedence

In the general authorization system, precedence is extremely flat. The *only* factor that affects precedence is the type of rule (grant or prohibit). Prohibit rules have *absolute* precedence. If there is a relevant prohibit rule, effective access is always Not Authorized.

Neither object inheritance nor identity hierarchy has precedence implications. Here are examples:

- A grant setting that is assigned to you has *less* precedence than a prohibit setting that is assigned to Authenticated Users.
- A direct grant on a report has *less* precedence than a prohibit setting that the report inherits from its parent folder.

Cheat Sheet

In the following table, each row indicates the effective access answer for a separate, independent scenario. For example, if the only relevant rule is a Conditional Prohibit, the effective access answer is Not Authorized (because there is no relevant grant setting).

All Relevant Rules	Effective Access and Explanation	
(none)	🚫	Not Authorized (implicit). Any access that is not granted is implicitly denied.*
Prohibit	🚫	Not Authorized. A relevant prohibit setting blocks access.**

All Relevant Rules	Effective Access and Explanation	
Prohibit + (any other rules)	⊘	Not Authorized. A relevant prohibit setting has absolute precedence.
Conditional Prohibits	⊘	Not Authorized. No relevant grants, no access.
Grant	✓	Authorized. A relevant grant provides access, if there are no relevant prohibit settings.
Grant + Conditional Grants	✓	Authorized. Relevant grants provide cumulative access, if there are no relevant prohibit settings.
Grant + Conditional Prohibits	▣	Conditional. Authorized for requests that are outside all of the prohibit conditions. Prohibit wins, but only within its scope.
Conditional Grants	▣	Conditional. Authorized for requests that are inside any of the grant conditions.
Conditional Grants + Conditional Prohibits	▣	Conditional. Authorized for requests that are outside all of the prohibit conditions *and* inside at least one grant condition.

* This result is due to the lack of a grant, so you can override it by adding a grant.

** This result is due to a prohibit setting, so you cannot override it by adding grants. As long as the prohibit setting exists and is relevant, effective access is not authorized.

For details about conditional access, see Evaluation of Conditions.

Origins of Effective Access

An explanation of effective access answers the question, Why does this identity have this effective access result for this permission and object?

In general authorization, the explanation consists of these items:

- A list of contributing rules, which includes all relevant rules, except share-based rules. Rules that are relevant but not determinative are included.
- A list of contributing shares, which includes all relevant shares.

To view origins information in the Authorization window, see Identify the Source of Effective Access.

General Authorization: Guidelines

The following basic guidelines contribute to simplicity and security.

- Minimize use of prohibit rules.
- Limit membership in administrative groups.

- Use groups, not individual users, as principals.
- Use folders, not individual objects, as targets.
- Use conditions only if you cannot efficiently express your authorization requirements another way.
- Perform a backup before and after you make significant changes to your system.

General Authorization: Troubleshooting

Unexpected Outcomes

Here are tips for troubleshooting an authorization outcome that differs from what you expect:

- Make sure all relevant rules are enabled. On the Rules page, right-click a rule to view all of its properties. Or, add the **Rule Status** column to the display.
- Make sure you understand the precedence model. See Authorization Decisions.
- Examine the origins information for the unexpected outcome. See Identify the Source of Effective Access.
- If the unexpected outcome relates to inheritance from a folder to objects in that folder, make sure you are using the second set of permissions in the folder's Authorization window to convey access to the folder's child members. See Inheritance.
- If the unexpected outcome relates to your access, and you have changed your memberships in the current session, sign out and then sign back in.
- If the unexpected outcome relates to your access, and you are a member of the SAS Administrators group, sign out and then sign back in, indicating whether you want that membership to be in effect. See Impact of Assumable Memberships.
- If the unexpected outcome is for access to a caslib or table, see Cloud Analytic Services Authorization.

Unavailable Principals

To grant access to a principal that is not in the identities service, use the **Rules** page in SAS Environment Manager or use the command-line interface.

Unrecognized Principals

If the **Rules** page or the Authorization window displays a warning icon next to a principal's name, that principal does not exist in the identities service.

- If the principal is a service account (for example, sas.folders or sasapp), you can ignore the warning icon.
- If the principal is a non-existent group that has **Administrators** as its name and ID, the rule has no effect. Determine whether the rule should be modified or deleted.

  For example, a predefined wildcard rule targets the objectURI /modelRepository/** and grants full access to **Administrators**. The rule is unnecessary, because the SAS Administrators group has a universal grant. You can delete the rule or leave it in place.
- If you are using the New Rule or Edit Rule window, make sure that the correct value is selected in the **Principal type** field and the principal's unique identifier (not display name) is specified.
- If you are using the Authorization window, make sure the identity still exists.

Note: Deletion of a custom group does not cause automatic deletion of rules in which that custom group is the principal.

Unintended Loss of Access

Reinstate Access: Instructions for Users

If you inadvertently block your own access to a resource, contact an administrator for assistance.

Note: Anyone who still has Secure access to the blocked resource can reinstate your access.

Reinstate Access: Instructions for Administrators

To reinstate access that is blocked by a prohibit rule, complete the following steps:

1 Opt in to your assumable membership in the SAS Administrators group.

2 Try to reinstate access by disabling, modifying, or deleting the prohibit rule. Here are some tips:

- If the resource is a content object (such as a folder or report) and you cannot see the resource on the **Content** page, you lack Read access to the resource. Use the Rules page.
- If the resource is a content object and you cannot make changes in the resource's Authorization window, you lack Secure access to the resource. Either delete the resource (if you have Delete access and the resource is not already in use) or proceed to the next step.
- If the resource is not a content object, use the **Rules** page.
- If you know who created (or last modified) a problematic rule or when a problematic rule was created (or last modified), use the **Modified By** or **Date Modified** filter on the **Rules** page to locate the problematic rule.

If you cannot reinstate access, proceed to the next step.

3 To temporarily prevent users other than yourself from using the deployment, close current sessions for users other than yourself, and disallow new sessions. See Disable Logins in *SAS Viya Administration: Authentication*.

4 Temporarily disable self-enforcement of authorization requirements for the authorization service.

a In the configuration definition for the authorization service, add a supplemental property named `remote` with a value of `false`.

Note: To learn how to set configuration properties, see How To Configure Services in *SAS Viya Administration: Configuration Properties*.

b Restart the authorization service.

Note: To learn how to restart services, see Operate in *SAS Viya Administration: General Servers and Services*.

5 Disable, modify, or delete the problematic rule or rules.

6 Enable self-enforcement of authorization requirements for the authorization service.

a In the configuration definition for the authorization service, remove the supplemental property named `remote`.

b Restart the authorization service.

7 Verify that access is reinstated.

8 Make the deployment available again by allowing new user sessions.

If you cannot reinstate access, contact SAS Technical Support for assistance.

A Deleted Rule Reappears

Some of the predefined rules are bootstrapped by their associated service. If you delete one of those rules, it reappears the next time the service starts. Modifications that you make to such rules are preserved. If you are sure you do not want one of those rules to be in effect, disable that rule (instead of deleting it).

Inability to Browse to an Object

If you have access to a content object but cannot browse to that object, you lack Read access to at least one of the object's parent folders.

To access the object, use an alternate navigation method. Here are examples:

- Access the object from a direct link or reference.
- Search for the object, and open the object from the search results list.
- If the object was shared with you, access the object from the **Shared with me** pane in SAS Drive. (On the **All** tab, select **Shared**. In the center pane, make sure **Shared with me** is selected.)

Not All Contributing Rules are Visible

If the number that is displayed on the **Contributing Rules** tab exceeds the number of visible listed rules, you are not currently authorized to see all contributing rules.

For the most comprehensive view, assume membership in the SAS Administrators group when you sign in.

See "Identify the Source of Effective Access".

A Pair of Predefined Rules is Difficult to Interpret

The search service makes use of two predefined rules that might appear to block Read access to all resources for most authenticated users. No administrative action or adjustment is needed, because the rules do not actually have that effect.

Each of the two rules includes a specialized property that is not displayed in SAS Environment Manager. Those specialized properties, acceptType and acceptItemType, limit the applicability of the rules. The rules affect only requests that involve usage of the search service's index. The rules ensure that users are not offered search results that include resources that they are unable to access.

Here are the key properties of one of the rules:

Property	**Value**
objectUri	/**
principalType	authenticatedUsers
type	prohibit
permissions	read, create

Property	Value
condition	!(groupsForCurrentUser().contains('SASAdministrators') || clientId() == 'sas.searchIndex')
description	Prevent ordinary users from getting a collection of indexable data when requesting a collection of vnd.sas.search.indexable.data.
acceptItemType*	application/vnd.sas.search.indexable.data+json

* This advanced property is not displayed in SAS Environment Manager. It limits the applicability of the predefined rule. See "acceptType, acceptItemType, and contentType".

General Authorization: Interfaces

All general authorization requirements and constraints are always fully enforced. However, not all interfaces expose all general authorization features.

In the following table, the shaded part of each circle is an approximation of the amount of general authorization functionality that a particular interface exposes. The shading indicates relative coverage. The shading does not indicate alignment of coverage across interfaces.

Table A.9 Interfaces to General Authorization

Interface		Description
●	REST API	The REST interface for general authorization.
◕	Rules page	The advanced enterprise graphical interface for managing rules directly.
◒	Authorization window	The basic enterprise graphical interface for managing access to content such as folders and reports.
◒	Command-line interface	A simple, scriptable interface for managing access to objects and resources.
◔	Share Window	A SAS Drive interface for simple sharing of content objects.

Sharing: Details for Administrators

Introduction

The purpose of sharing is to help users make content available to one another.

This topic documents administrative aspects of the implementation of sharing that is provided by the authorization service. For usage information, see "Share" in *SAS Drive: Getting Started*.

The authorization service's implementation of sharing is separate and distinct from all of the following features:

- sharing of Model Studio projects. See "Sharing a Project" in *SAS Visual Data Mining and Machine Learning: User's Guide*.
- sharing of generic (formerly SAS Data Management) projects, which is achieved by assigning owners and members. See "Projects" in *SAS Drive: Getting Started*.
- sharing of reports in SAS Visual Analytics, which does not affect access. See "Sharing Reports and Objects with Other Users" in *SAS Visual Analytics: Designing Reports*.

Who Can Share?

In the initial configuration, sharing is available as follows:

- Any user who has Secure access to an object can share that object with other users and groups.
- Any user with whom an object has been shared can further share that object, passing along some or all of the access that they received. Only the initiator of a chain of sharing has to have the Secure permission on the shared object.

  For example, after UserA shares reportA with UserB, UserB can share reportA with other users, even if there is no authorization rule that grants UserB the Secure permission for reportA. (However, if an authorization rule explicitly prohibits the Secure permission for UserB on reportA, UserB cannot further share reportA.)

For a more restrictive configuration, see "How to Prevent All Sharing" or "How to Limit Re-sharing".

What is a Share?

A share is of a set of attributes that is backed by a corresponding authorization rule.

When users share objects in SAS Drive, shares and corresponding share-based rules are generated. Each share specifies a particular type of access, a recipient user or group, and a target object. Each corresponding share-based rule translates the share information into an authorization rule. In authorization decisions, share-based rules are evaluated in the same way as other authorization rules.

Here are details about shares:

- Shares can only expand access. Users can remove shares, but they cannot use shares to prevent access.
- Shares do not provide guaranteed or comprehensive access.
  - A share does not provide access if there is a relevant Prohibit rule.
  - A share of one object does not provide access to related resources (such as data or linked reports), parent objects (such as folders), or embedded objects (such as images or files).

> **TIP** Because sharing an object does not provide access to the object's parent folders, a share recipient might not be able to navigate to the shared object. For alternatives, see "Inability to Browse to an Object" on page 436.

- Shares cannot provide variations of access that differ from the defined share types.
- Shares cannot be conditional.
- Shares cannot target a media type.
- Shares for a container object always target both the object's URI and the object's container URI.
- Shares are static. For example, a share of reportA by UserA to UserB is unaffected by any subsequent loss of UserA's access to reportA.

Attributes of a Share

resourceUri
: specifies the URI of the content object that is shared. Also referred to as the share *target*.

sharedBy
: specifies the ID of the user that created the share.

sharedWith
: specifies the ID of the user or group that receives the share. Also referred to as the share *recipient*.

sharedWithType
: specifies the type of identity that receives the share.

type
: specifies the type of access that the share provides. Also referred to as the share *level*.

- Users select either **`Can read`** or **`Can read and edit`** when they create a share.
- The service stores one of the following values for each share: **`read`**, **`readEdit`**, **`readShare`**, **`readEditShare`**.

  Note: The appended **`Share`** indicates that a share allows further sharing.

Effects by Share Type

The effects of a share are determined by whether the shared object is a container and what type of access the share provides. The following tables provide details:

Table A.10 Effects of Sharing a Non-Container Object

Share Type	Effects
read	Corresponds to a grant of the Read permission.
readShare	Corresponds to a grant of the Read permission. Also provides the ability to further share Read access.
readEdit	Corresponds to grants of the Read, Update, and Delete permissions.
readEditShare	Corresponds to grants of the Read, Update, and Delete permissions. Also provides the ability to further share some or all of the same access.

Table A.11 Effects of Sharing a Container Object

Share Type	Effects*
read	Corresponds to a grant of the Read permission on the container’s object URI and container URI.
readShare	Corresponds to a grant of the Read permission on the container’s object URI and container URI. Also provides the ability to further share the same access.
readEdit	Corresponds to grants of all permissions except Secure on the container’s object URI and container URI.

Share Type	Effects*
readEditShare	Corresponds to grants of all permissions except Secure on the container's object URI and container URI. Also provides the ability to further share some or all of the same access.

* Sharing a container object automatically provides conveyed access to child members.

Note: The ability to re-share is not inherited. For example, the recipient of a readShare on a folder can re-share the folder, but cannot re-share individual items within that folder.

How to Limit Re-sharing

Implications

Initially, downstream sharing is supported, even for users who lack Secure access to the object that is being shared. Disabling re-sharing has the following effects:

- New shares have values of `read` or `readEdit`, not `readShare` or `readEditShare`.
- Existing shares retain their values of `readShare` or `readEditShare`. However, the `Share` part of those values has no effect. Only those users who have Secure access to a target object can share that object.
- Any access that was already extended through re-sharing remains in effect. Disabling re-sharing does not revoke access that already exists.
- In the Authorization window, no sharing-related information is displayed for the Secure permission.
- There is no **Contributing Shares** tab for the Secure permission.

Instructions

To limit downstream sharing to only those users who have Secure access to the target object:

1. In the vertical navigation bar in SAS Environment Manager, select 🔧.
2. From the **View** drop-down list, select **Definitions**.
3. In the list of definitions, select **sas.authorization**.
4. If no configurations exist, click **New Configuration**. Otherwise, select an existing configuration to edit.
5. In the New sas.authorization Configuration window (or the Edit sas.authorization Configuration window), set the reshareEnabled property to `Off`.

6. Click **Save**. Your change takes effect within 30 seconds. You do not have to restart the authorization service.
7. Verify the result as follows:
   a. Add a folder beneath your My Folder.
   b. Share that folder with another user.
   c. Ask that user to sign in (without assuming any administrative privileges) and make sure they cannot share the folder that you shared with them.

How to Prevent All Sharing

Implications

Initially, sharing is enabled. Disabling sharing has the following effects:

- No new shares can be created.
- All existing share-based rules are disabled.
- In the Authorization window, no sharing-related information is displayed.
- SAS Drive hides all actions and information that are related to sharing.
- SAS Drive provides direct access to the Authorization window for authorized users.

Instructions

To disable sharing:

1 In the vertical navigation bar in SAS Environment Manager, select 🔧.

2 From the **View** drop-down list, select **Definitions**.

3 In the list of definitions, select **sas.authorization**.

4 If no configurations exist, click **New Configuration**. Otherwise, select an existing configuration to edit.

5 In the New sas.authorization Configuration window (or the Edit sas.authorization Configuration window), set the sharingEnabled property to `Off`.

6 Click **Save**. Your change takes effect within 30 seconds. You do not have to restart the authorization service.

7 Verify the result as follows:

   - In SAS Drive, make sure the **Share** action is not available.
   - In SAS Drive, make sure that displays such as **Shared with me** and **Items I've shared** displays are suppressed.
   - In SAS Environment Manager, examine the properties of a share-based rule. Make sure the rule is disabled.

Administrative Oversight on the Rules Page

As an administrator, you can view and delete share-based rules on the Rules page in SAS Environment Manager. When you delete a share-based rule, the corresponding share is automatically deleted along with the rule.

On the **Rules** page, the following constraints apply to managing share-based rules:

- You cannot add share-based rules. Sharing is primarily a user-driven activity. Share-based rules are automatically created when users share content in SAS Drive.
- You cannot edit or copy share-based rules. Any changes to share-based rules must be coordinated with changes to the associated shares, and must conform to requirements that are specific to share-based rules.

All share-based rules are created with generated text in the **Description** field. Here is an example:

The user "userA" shared an object with the specified principal. (This is a share-based rule.)

On the **Rules** page, you can filter for share-based rules as follows:

- To display only share-based rules, enter the following text in the **Description** filter:

  This is a share-based rule

- To display only share-based rules for a particular share recipient, specify the preceding **Description** filter and specify the share recipient in the **Principal** filter.

- To display only share-based rules that were generated by a particular user, include that user's ID in the **Description** filter text as follows:

  The user "userID" shared an object with the specified principal

> **TIP** You cannot instead use the **Modified by** filter, because all share-based rules are generated with a **Modified by** value of `sas.authorization`. Do not assume that all rules that have a **Modified by** value of `sas.authorization` are share-based rules.

Integrated View in the Authorization Window

Effective Access

In the Authorization window, effective access information reflects shares as follows:

- Effective access information reflects any access that is provided by shares.

- In the **Secure** and **Secure (convey)** columns, the Share icon ⋖ indicates that sharing is possible even though Secure access is not granted.

  Note: If Secure access is granted, the Share icon is not displayed. The ability to share is inherent in Secure access. See "Examine Access" on page 413.

Direct Settings

In the Authorization window, a share is not considered a direct setting.

A share does not cause a diamond to be displayed.

You cannot modify shares in the Authorization window.

Contributing Shares

After you click an effective access icon in the Authorization window, a pop-up window that includes a **Contributing Shares** tab is displayed. The tab provides a read-only list of the shares that are relevant to the selected effective access result.

A share is relevant if it meets all of the following criteria:

- The share specifies the current object (or a parent of that object) as the target.

- The share specifies the current principal (or a group to which that principal belongs) as the recipient.

- The share type is relevant for the selected permission. For example, for the Update and Delete permissions, the `readEditShare` and `readEdit` types are relevant, but the `read` and `readShare` types are not relevant.

  Note: For the Secure permission, the `readShare` and `readEditShare` types are treated as relevant only because those types provide the ability to reshare. Sharing never creates a grant of the Secure permission.

  Note: To add the **Share Type** column to the display, click ⋮ and select **Manage columns**.

25

Encryption for Data at Rest

Encryption for Data at Rest: Overview

SAS Viya provides encryption in two contexts:

- Data at rest is data that is stored in databases, file servers, endpoint devices, and various storage networks. This data can be on-premises, virtual, or in the cloud. This data is usually protected in conventional ways by access controls. Numerous layers of defense are needed, and encrypting sensitive data is another layer.

  This document covers administrative tasks for encrypting files at rest in the SASHDAT format and it shows how tables that are imported into caslibs are encrypted. See Concepts for details.

- Data in motion is data that is being transmitted to another location. Data is most vulnerable while in transit. Sensitive data in transit should be encrypted. You can protect all traffic in transit between servers and clients. See “Overview” in *Encryption in SAS Viya: Data in Motion*.

SAS Viya uses Advanced Encryption Standard (AES) algorithms with 256-bit keys to encrypt data at rest.

Refer to Caslibs, Files, and Tables if you need additional background on data and caslib concepts.

Use one of the following interfaces to encrypt files at rest:

- To manage encryption of data files interactively, use SAS Environment Manager.
- To programmatically encrypt data files, use the CASLIB Statement.

How To (SAS Environment Manager)

Introduction

Authorized administrators use SAS Environment Manager to create and manage data security. The Domains area enables you to create a stored credential (an encryption key) that is available to designated identities to facilitate loading of encrypted files. By default, when you are creating a new caslib, enabling that caslib for encryption is disabled. If you choose to enable encryption, this can be done by creating an encryption domain, and then associating that domain with your path-based caslib.

In the Domains area, you can perform the following tasks:

- Create a new encryption domain or use an existing encryption domain.
- Add users or group identities to the encryption domain.
- Create an encryption key (passphrase).

You must be a member of the *SAS Administrators* group and assume groups when you log on to SAS Environment Manager in order to create and manage Domains. For more information, see "Accessing SAS Environment Manager" on page 665.

From the Data area of SAS Environment Manager, you can use SAS Data Explorer to create and manage caslibs. By default, caslibs are not enabled for encryption. You can encrypt the tables in your caslibs by associating the caslibs with an encryption domain.

CAUTION! Be sure to keep a separate record of your encryption key. Once a caslib is created, it is not possible to change the encryption key setting or change the domain. The encryption key value cannot be retrieved through the software. If the caslib is deleted, the tables remain encrypted. To access those encrypted tables, you need to define a new caslib using the same path, the same domain, or a new domain using the same encryption key.

Note: When the encrypted data is loaded into CAS, it is decrypted at load time. Caslib authorizations apply to accessing the loaded data. Identities that are set in the encryption domain provide authorization for who can request that encrypted tables be loaded. Identities can be custom groups or individual user IDs. Custom groups provide access control by simple group membership.

The following instructions explain how to encrypt the tables in your caslibs using SAS Environment Manager.

About the Domains Page

The **Domains** page in SAS Environment Manager enables you to manage the following types of domains and credentials:

Note: The Domains area is available only if you are a member of the SAS Administrators group.

authentication domain	Makes stored credentials (user IDs and passwords) available to designated identities to facilitate connections to servers that require a password. See "External Credentials: Overview" on page 337.
connection domain	Makes stored credentials (user IDs) available to designated identities to facilitate connections to servers that do not require authentication. See "External Credentials: Overview" on page 337.
encryption domain	Makes a stored credential (an encryption key) available to designated identities to facilitate loading of encrypted files. This document describes this process.

Navigation

The Domains area is available if you are a member of the *SAS Administrators* group and you have opted into your assumable groups. For more information, see Accessing SAS Environment Manager on page 665.

1 In the applications menu (☰), under **Administration**, select **Manage Environment**.

2 In the navigation bar, locate the **Security** section and click **Domains** .

3 You can select one of two views from the **Domains** page. The default view is **Domains.** From the **View** drop-down list, select one of the following views:

Domains	Lists all domains that are displayed. Domains is the default view. There are three types of domains: Authentication, Connection, and Encryption. On the Domains page, you can view the information for each domain that is defined, or you can create a new domain. When you create an encryption domain, you cannot delete that domain. Three default Authentication domains appear in the Domains view. These are described in Authentication Domains Defined In the Deployment on page 346. Note: This view is available only to SAS Administrators. Note: This document discusses only the Encryption domain.
Credentials	Enables you to access external data sources and other third-party products requiring authentication. Credentials are associated with a specific domain for use with a specific data source type. For more information, see External Credentials: Overview on page 337.

Manage Encryption Domains

Create a New Encryption Domain

1 In the **Domains** view, click .

2 In the New Domain window, specify general settings as follows:

ID	Create an ID name. Enter a unique ID for your encryption domain.
Type	Select the type of domain. There are three domain types, Encryption, Authentication, and Connection. Select **Encryption** from the list of available domains.
Description	Add a description.
Identities	Select **Identities**. You can select from users, groups, and custom groups. See below for how to add an identity.
Encryption key	Encryption passphrase or key.
Confirm encryption key	Enter the same passphrase as above.

For additional information about identities, from the New Domain window, click ?.

3. Add **Identities**. From the New Domain window, click 👤.

   a. In the left pane of the Choose Identities window, you can filter by Users, Groups, and Custom Groups. From the drop-down menu, select 👤**Users**, 👥**Groups**, or 👥 **Custom Groups**.

      You can also filter using the search 🔍.

      Note: A best practice is to use a custom group. Then you can add additional users to this custom group as needed to grant access to the data. Be sure to assign correct permission for this custom group in the associated caslib Authorization.

   b. Move the selected user, group, or custom group to the **Selected Identities** pane. Click ➡.

   c. Click **OK**.

4. Add an encryption key.

5. Confirm encryption key.

6. After you have entered all of the parameter settings needed, click **Save**.

View Properties of an Encryption Domain

1. In the **Domains** view, select an ID row for Type Encryption.

2. Right-click, and select **Properties**. Or select ☰ from the taskbar. From the **Domain Properties** window, the ID, Type, Description, Date created, Date modified, who created the domain, and who modified the domain is displayed.

3. Click **close**.

Edit an Encryption Domain

From the Edit Domain window, you can add only a description. To modify the identities, go to the credentials view.

1. In the **Domains** view, select an ID.

2. Right-click, and select **Edit**. Or select ✎ from the taskbar.

3. From the Edit Domain window, add a **Description** of the Encryption Domain.

4. Click **Save**.

View and Modify the Credentials of an Encryption Domain

From the Credentials View for a selected encryption domain, you can view existing credentials that are associated with an Encryption domain, add new credentials either by clicking the edit or the new icons, delete identities, or just view the properties.

From the Domains page, you can view, edit, and add credentials to an existing domain from the Credentials view.

View Credentials

1 From the Domains page, select an ID of type Encryption.

2 Right-click, and select **Credentials**. Or select ** from the taskbar. In the Credentials for Domain view, credential properties are displayed.

3 To see all hidden columns, at the right edge of the table, click the **Options** icon ⋮ , and select **Manage Columns**.

4 From the Manage columns window, select items to move from the **Hidden columns** pane to the **Displayed columns** pane. After selecting the **Hidden columns** to display, click the **Add** arrow ➡.

After selecting the **Displayed columns** to hide, click the **Remove** arrow ⬅.

TIP To select more than one item, use the Shift key.

5 Click **OK**.

Edit Credentials

If you are a member of the domain, you can add identities to and remove identities (users, groups, and custom groups) from an existing Encryption Domain. You cannot change the type of domain or change the Encryption key from the credentials view.

1 In the **Credentials for Domain** view, select a credential entry.

2 Right-click, and select **Edit**. Or select ✎ from the taskbar.

3 To add and delete identities, in the Edit Credential window, click 👤.

a In the left pane of the Edit Credential window, you can filter by Users, Groups, and Custom Groups. From the drop-down menu, select 👤**Users**, 👥**Groups**, or 👥 **Custom Groups**.

You can also filter using the search 🔍.

b To add a selected user, group, or custom group to the **Selected Identities** pane, click ➡.

To remove a selected user, group, or custom group from the **Selected Identities** pane, click the **Remove** arrow ⬅.

c Click **OK**.

4 Click **Save**.

Delete Credentials

Perform the following tasks in the Credentials for Domain view.

1 Right-click the credential that you want to delete, and select **Delete**. Or select 🗑 on the taskbar.

2 In the Delete window, this message is displayed: `“Are you sure you want to delete the credential for the identity "name"?”`

3 Click **Yes**.

Delete Encryption Domains

Encryption domains cannot be deleted in the current release. If you try to delete an Encryption domain, you receive the following message: `You cannot delete the encryption domain named 'domain-name'. Libraries associated with this domain will need to be recreated if the domain is deleted.`

Manage Caslibs That Are Enabled for Encryption

For additional information, see “Working with SAS Data Explorer” in *SAS Data Explorer: User’s Guide*.

Add a New Caslib That Connects to a Remote File System

SAS Data Explorer enables you to discover data and copy it to a SAS Cloud Analytic Services (CAS) server. In this section, we use the SAS Data Explorer to add a new caslib that connects to a PATH, HDFS, or DNFS file system. Only these types of file systems can be encrypted. For detailed information about using SAS Data Explorer, see, “Working with SAS Data Explorer” in *SAS Data Explorer: User’s Guide*.

1 In the applications menu (☰), under **Administration**, select **Manage Environment**.

2 In the navigation bar, click **Data** ▦.

3 Click the **Data Sources** tab. The **Data Sources** tab enables you to add a new caslib that connects to a DNFS, HDFS, or PATH-based remote file system. For more information, see “Data Selection Windows and SAS Data Explorer” in *SAS Data Explorer: User’s Guide*.

4 Click 🔌 on the **Data Sources** tab. The Connection Settings dialog box is displayed.

5 Enter a name for the caslib in the **Name** field. If you change the name, follow the conventions for caslib names as described in “Names for Caslibs, Tables, and Columns” in *SAS Data Explorer: User’s Guide*. The target operating system and data source might have additional constraints on these names.

6 Accept the default CAS server or select another CAS server in the **Server** field.

7 Select **File System** in the connection **Type** field.

8 Select the remote file system that you want to access in the **Select source type** field: **`DNFS`**, **`HDFS`**, or **`PATH`**.

9 Select the **Persist this connection beyond the current session** check box to add a global caslib for this connection. Deselect this check box to add a session-based caslib for this connection. For more information about this option, see “Caslibs on the Data Sources Tab and Import Tab” in *SAS Data Explorer: User’s Guide*.

10 Under the **Settings** tab, enter the physical path to the remote file system in the **Path** field.

If the new caslib is global in scope, you can specify one additional directory level to an existing path that does not exist on the target file system, and that directory will be created. If the new caslib is a session-based caslib, you must specify an existing path, or the connection fails. For more information about global and session-based caslibs, see “Caslibs on the Data Sources Tab and Import Tab” in *SAS Data Explorer: User’s Guide*.

11 Enter a **Description** for the connection, if desired.

12 Click **Test Connection** to test your connection. You will see the green check box and the message "The connection was successful."

13 Click the **Advanced** tab to specify an Encryption domain for the current data source. From the drop-down list, select an **Encryption domain**.

Note: If no encryption domains are shown in the drop-down list, you can create one. See "Create a New Encryption Domain" on page 445.

14 Click **Test Connection** to test your connection. You will see the green check box and the message "The connection was successful."

15 Click **Save** to save your connection.

If the connection succeeds, tables that you are authorized to access in the remote directory will be available from the caslib that you specified in Step 4. Information about caslibs and tables on the **Available** and **Data Sources** tabs is stored in the cache for your web browser. If you think this information does not reflect the current state of your system, click ↻ in the nearest toolbar.

If the connection fails, see "General Usage Notes" in *SAS Data Explorer: User's Guide*.

16 If the target caslib is not visible from your current view, click ⮤ to find the caslib and its tables on the specified CAS server.

17 You can see the options that are available to manage this new caslib, which is now available to receive encrypted tables. See "Working with Caslibs" in *SAS Data Explorer: User's Guide*.

View the Properties of the New Caslib

When you select a caslib on the **Data Sources** tab, properties for that caslib are displayed on the right. These properties were specified when the caslib was added, such as the name, the CAS server, the source type, and the domain name.

1 Click the **Data Sources** tab.

2 From the drop-down list, select the CAS server where your new caslib exists.

3 Scroll down to find the caslib that you created and click on the name of the caslib to display the properties of the caslib. The Authentication domain property should show the name of the encryption domain that you are using.

Edit Properties of the Caslib

You can change the caslib path or change the description of a caslib. The new path uses the same encryption domain defined in this caslib. It is not possible to change the encryption domain assigned to a caslib.

Note: Tables created in the previous topics are still encrypted and require the same key. Changing the path for a caslib with encryption enabled is not recommended. It is recommended that you create a new caslib using the same encryption domain and include a new path.

1 Right-click, and select **Edit** from the drop-down list.

2 In the **Edit Properties** window, change the caslib path or description as needed.

3 Click **Save**.

When editing a caslib the following restrictions apply:

- Only path-based libraries can be edited.
- You cannot edit a personal caslib.
- You cannot change the encryption domain assigned to a caslib. If you need to change the encryption domain, you must create a new caslib.

View and Edit Caslib Authorization

To manage tables (promote, create, drop, edit), delete source tables, alter tables and caslibs, and manage access to a selected caslib, you need the appropriate authorization. You can view your permissions and edit your permissions for a specified caslib. For additional information, see Working With Data in CAS and CAS Authorization on page 384.

View the authorization level for your caslib.

1. Click the **Data Sources** tab.
2. Scroll down to find the caslib that you are working with. Right-click, and select **View Authorization** from the drop-down list.

   The levels of authorization are shown for your caslib.
3. Select **Close**.

Change the access level of the caslib.

1. Right-click, and select **Edit Authorization** from the drop-down list.
   a. Change the access level by sliding the control bar to the level of control that you want.
   b. Click **Save**.
   c. Click **Close**.

From both the View Authorization and the Edit Authorization windows, you can add and remove identities from your caslib.

1. Click the **Data Sources** tab.
2. Right-click your caslib and select **Edit Authorization** from the drop-down list.
3. From the Edit Authorization window, click 👤.
   a. In the left pane of the Choose Identities window, you can filter by Users, Groups, and Custom Groups. From the drop-down menu, select 👤**Users**, 👥**Groups**, or 👥 **Custom Groups**.

      You can also filter using the search 🔍.

      Note: A best practice is to use a custom group. Then you can add additional users to this custom group as needed to grant access to the data. Be sure to assign correct permission for this custom group in the associated caslib Authorization.
   b. Move the selected user, group, or custom group to the **Selected Identities** pane. Click ➡.
   c. Click **OK**.
4. Click **Close**.

Manage Tables

A caslib can contain a mix of encrypted and unencrypted tables. However, loading any table still requires the user to be a member of the encryption domain.

Import an Unencrypted Table into a Caslib Enabled for Encryption

To Import tables into your caslib, use the Add to Import feature. The Add to import option enables you to copy a table or a file on the **Available** tab or the **Data Sources** tab to a global caslib. You can copy a table or a file from any caslib to any global caslib to which you have access. For details, see, "Copying Data from the Available Tab or Data Sources Tab" in *SAS Data Explorer: User's Guide*.

Add tables to the queue that you want to import into your caslib.

1 In the Data window, display the window that contains the **Available** tab or the **Data Sources** tab. See "Data Selection Windows and SAS Data Explorer" in *SAS Data Explorer: User's Guide*.

2 Click the **Available** tab or the **Data Sources** tab.

3 Right-click a table or file that you want to copy and select **Add to import**. The selected table or file is added to the queue on the **Import** tab. The table or file that you want to copy appears on the left side of the **Import** tab view. Import properties for the copy appear on the right side.

4 The **Target table name** field is set to the default target location. For our example, accept this name as appropriate for the copy.

5 Make sure that the **Target desintation** field contains the path to the caslib where you want to store the copy of the selected table. You can change the default destination by selecting a caslib either by clicking on **Find** or on the **Select Destination** icon. From the **Data Sources** tab, you can then set this new destination as the default by right-clicking on the caslib and selecting **Set as default target location**.

   - If the source caslib is global in scope, the name of the source caslib is copied into the **Target destination** field.
   - If a caslib has been set as the default target location, that caslib is selected as the default target destination.
   - If a default cas-shared-default server is deployed with the Public caslib, the Public caslib is set as the default target destination.

   For more information, see "Copy Data from One Caslib to Another" in *SAS Data Explorer: User's Guide*.

6 Specify which action the import operation should take if the target filename exists in the caslib that is specified in the **Target destination** field. The options are to cancel the import or to replace the existing item that has the same name.

7 You can right-click the table or file to be copied and select **Import Item**. If the import succeeds, a copy of the table or file is loaded into memory on the CAS server that is specified in the caslib.

   The copy of the table or file can now be selected from the **Available** tab or the **Data Sources** tab. Click the **Find** button next to the **Target table name** field to automatically display the imported table in the **Data Sources** tab.

   If the import fails, see "General Usage Notes" in *SAS Data Explorer: User's Guide*.

   Note: You will see a note that states that the item is being imported. When it is complete, you will see the following message: "The table was successfully imported on *Date/Time* and is ready for use."

8 Select the **Data Sources** tab and scroll to the caslib that you imported tables into.

   A list of the tables that are imported into the caslib are displayed. If the target caslib is not visible from your current view, click to find the caslib and its tables on the specified CAS server. Select the green arrow to see the tables that were imported into the caslib.

9 Click on any of the listed tables to view the details of the imported tables.

How To (Programming Tasks)

SAS Cloud Analytic Services supports encryption of SASHDAT files at the file level. As an administrator, you might want to simplify encryption for data at rest by configuring caslibs with an encryption password so that all files in a directory are encrypted.

For information that describes how to set an encryption password in a program, see CASLIB Statement, DATASOURCE options in *SAS Cloud Analytic Services: User's Guide*.

For an example of using the CASLIB statement to encrypt caslibs, see Encrypt Tables in a Caslib.

Encryption for Data at Rest: Concepts

Overview

When a caslib is created and enabled for encryption (an encrypted domain is associated), a table imported into that caslib is then encrypted in SASHDAT (.sashdat) format. A domain is associated with a caslib to provide access. Domains are used to store both the credentials (passwords and keys) that are required to access external data sources and the identities that are allowed to use those credentials.

When the encrypted tables are loaded into CAS (in-memory tables), these tables are identified as using an AES encrypted source. However, in-memory tables are not encrypted. The encryption applies to source tables, not to tables that are in memory.

You can specify encryption for a caslib only when it is created. To change it you must re-create the caslib.

SAS Viya supports encrypting files at rest in a path location. Only path-based (PATH, DNFS, HDFS) caslibs are supported.

SAS Viya uses Advanced Encryption Standard (AES) algorithms with 256-bit keys to encrypt data at rest.

Refer to *SAS Cloud Analytic Services: Fundamentals* if you need additional background on data and caslib concepts.

What Is a Domain?

Overview

Domains are used to store both the credentials (passwords and keys) that are required to access external data sources (for example, databases like Oracle, Teradata, and other data sources like Facebook and Amazon) and the identities that are allowed to use those credentials. A domain contains one or more references to identities (users or groups) who have access to the credentials in the domain. A user can gain access to the credentials either directly with their user ID or indirectly as a member of a group that is defined as an Identity.

The ID, or name, of a Domain is used in the definition of a non-path-based caslib to access and load tables from external databases. A domain is associated with a caslib to provide access. Examples of external sources include SAS LASR, Oracle, Teradata, Hadoop, Postgres, and Impala. Users of a caslib with an associated domain do not have to know or enter data source credentials to access or load external data.

There are three Domain types: Authentication, Connection, and Encryption.

What Is an Encryption Domain?

An encryption domain is used to store an encryption key that is required to read data at rest in a path assigned to a caslib. The Identities selected in this encryption domain have access to the key. When you create a path-based caslib, you can choose to enable encryption. You then select an encryption domain to assign an encryption key. Tables imported to this caslib are now encrypted. If the path contains preexisting tables, those tables are not encrypted. Users who are not defined in Identities as individuals or as members of a group cannot load data from this caslib.

Note: It is recommended that you start with an empty caslib and import either encrypted or unencrypted tables. Mixing these types is not recommended.

Encryption domains are used to store encryption keys that can then be associated with a caslib type of PATH, HDFS, or DNFS.

What Is a Connection Domain?

A connection domain is used when the external database has been set up to require a user ID but no password. For example, a third-party database like Hadoop might be configured with accounts for authentication that do not require a password. For information about using connection domains, see “External Credentials: Overview” on page 337.

What Is an Authentication Domain?

An authentication domain is a name that facilitates the matching of logons with the servers for which they are valid. Authentication domains are used to store credentials that are used to access an external source (for example, an Oracle database) that can then be associated with a caslib of the appropriate type. They can also be used for batch processing and scheduling where you store your credentials in a domain to run jobs in batch mode.

Each user ID and password is valid within a specific scope. For example, the user ID and password that you use to log on to your computer at work are probably not the same as the user ID and password that you use to log on to a personal computer at home. It is also common for database servers and web servers to have their own authentication mechanisms, which require yet another, different, user ID and password.

The software attempts to use only the credentials that it expects to be valid for a particular resource or system. The software’s knowledge of which credentials are likely to be valid is based entirely on authentication domain assignments. For this reason, you must correctly assign an authentication domain to each set of resources that uses a particular authentication provider, and also assign that same authentication domain to any stored credentials that are valid for that provider.

Authenticating to SAS can be done through SAS logon. For more information, see “Authentication: Overview” on page 289.

For information about using Authentication Domains, see “External Credentials: Overview” on page 337.

Defaults

In a new deployment, encryption for data at rest is not automatically enabled. You can configure encryption of data that is added to PATH, HDFS, and DNFS caslibs by using SAS Environment Manager and the programming interfaces.

Encrypting Caslibs

SAS Viya supports encryption as an option for tables in caslibs. The encryption applies to source tables, not to tables that reside in memory. When you create a caslib, you enable the caslib to receive encrypted tables when you assign an encryption domain. The tables that you import into the caslib become encrypted and all tables use the same encryption key.

When you import a table (SAS table, .csv file, .txt file, Excel file, and so on) into a caslib, a .sashdat file is created in the same path location. If the caslib is enabled for encryption, those .sashdat files at rest are now encrypted. When these tables are loaded into CAS, these in-memory CAS tables are not encrypted. You must have authorization to access the table in the caslib. Table access can be inherited from the caslib, but it can also be granted or denied at the table level.

If you delete the caslib, the tables in the associated path remain encrypted. To access those tables again, you must create a new caslib, enable encryption, and use the same encryption domain with the same key value, or create a new encryption domain with the same key value.

Considerations for Encrypting Tables in Caslibs

Here are a few best practices and considerations when encrypting data at rest:

- Only PATH, HDFS, or DNFS files can be encrypted.
- It is best not to mix encrypted and unencrypted tables in a caslib path. Only the user IDs and groups in the domain identities are able to read the encrypted data.

  When you create a new caslib and enable encryption, only the newly imported tables are encrypted and stored in the path. A best practice is to first make sure that the path is empty before you import the tables that you want encrypted.
- Encryption of data at rest has some performance costs, and user and administrative overhead. You must balance the goals of security and performance at your site when deciding to encrypt data. Users and administrators must keep track of keys (passphrases and passwords) when accessing the data. The system uses additional CPU resources when loading and saving encrypted tables.

Reference

PROC PWENCODE

In the programming environment, the ENCRYPTIONPASSWORD= option in the CASLIB statement specifies a password for encrypting or decrypting tables. For additional password security, you can use the PWENCODE procedure to encode that password. Encoded passwords can be used in place of plaintext passwords in SAS programs that access databases and various servers. For information, see PWENCODE Procedure.

26

General Services

General Servers and Services: Overview

Servers

List of Servers

SAS Viya contains these servers:

- SAS Cloud Analytic Services
- Programming run-time servers:
  - SAS Compute Server
  - SAS Launcher Server
  - SAS Workspace Server and SAS Object Spawner
  - SAS/CONNECT Server and SAS/CONNECT Spawner
  - embedded SAS Web Application Server
- Infrastructure servers:
  - SAS Cache Locator and SAS Cache Server (Geode)
  - SAS Configuration Server (Consul)

- □ SAS Secret Manager (Vault)
- □ SAS HTTP Proxy Server (Apache HTTP Server)
- □ SAS Message Broker (RabbitMQ)
- □ SAS Infrastructure Data Server (PostgreSQL)

The following diagram shows the relationship of the servers to other components in the SAS Viya full deployment:

Figure A.1 SAS Viya Servers (Full Deployment)

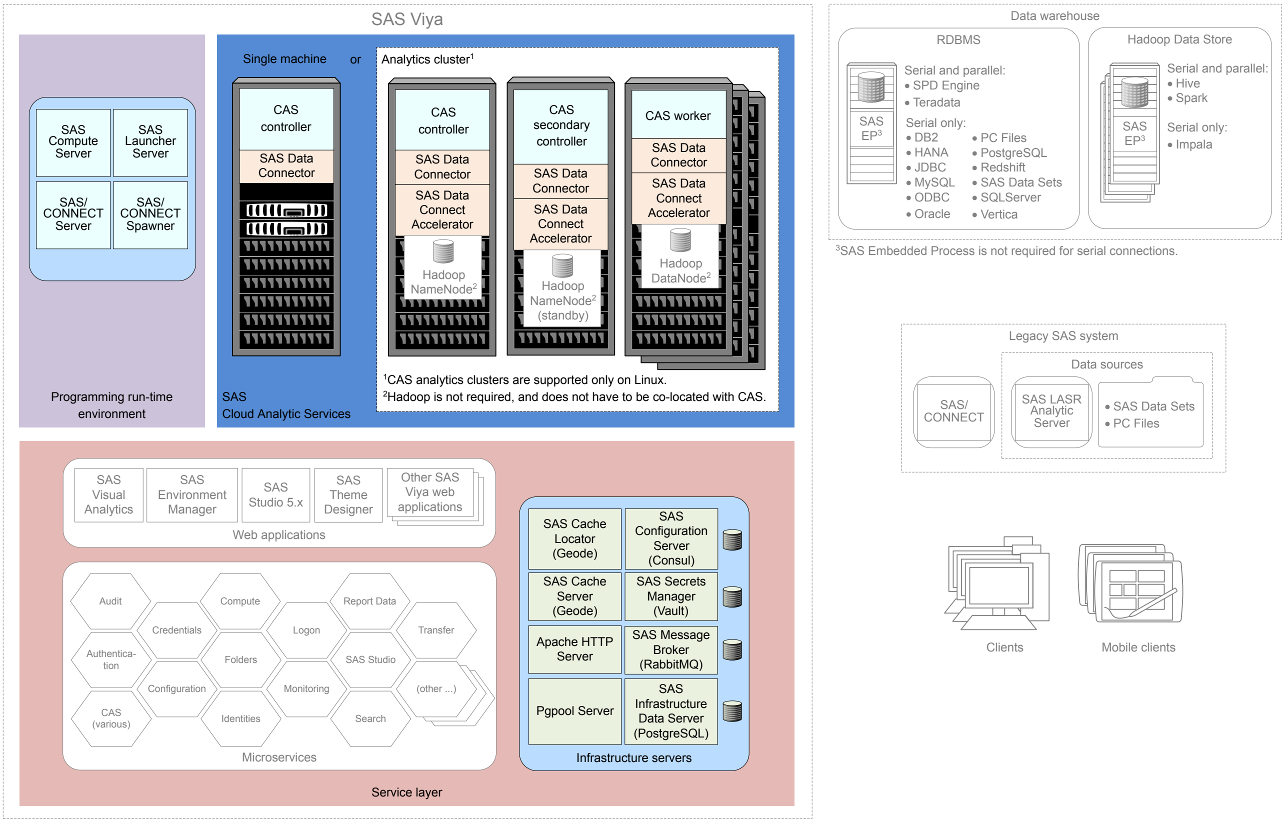

The following diagram shows the relationship of the servers to other components in the SAS Viya programming-only deployment:

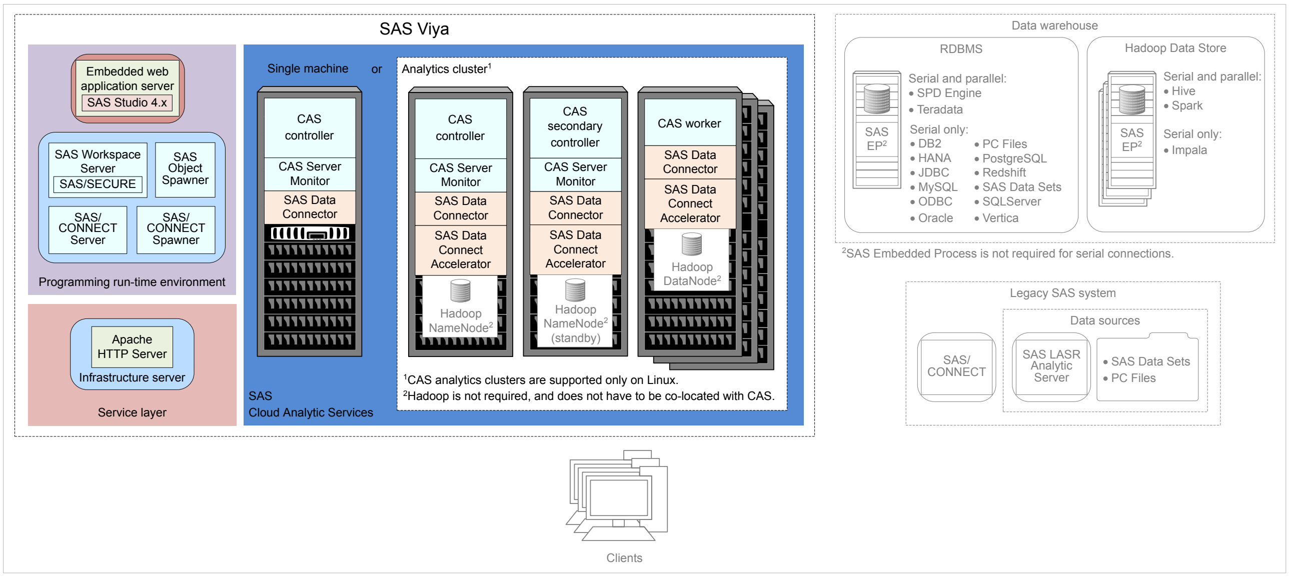

Servers Managed with SAS Environment Manager

Using SAS Environment Manager, you can manage CAS servers and Launcher Servers on page 566:

- CAS Server Tasks on page 481
- Launcher Servers Tasks

Services

SAS Viya contains several services often referred to as *microservices*. A microservice is a discrete service that runs in its own process and that communicates using HTTP.

SAS Viya includes services such as, Audit, Identities, and Monitoring. To see the complete list of SAS Viya services follow the initial steps in “Edit Configuration Instances” on page 75.

Note: A programming-only deployment on page 1does not use most SAS Viya services.

Figure A.2 SAS Viya Services

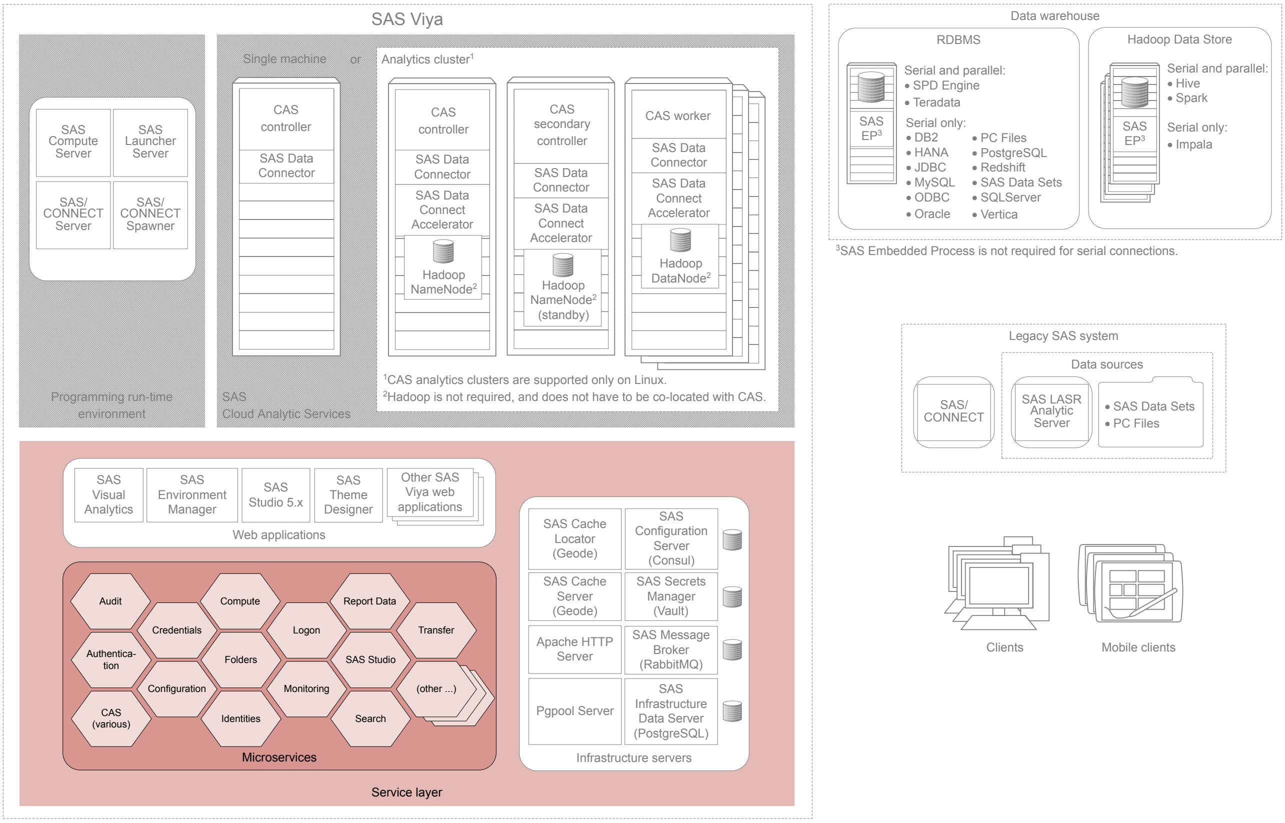

General Servers and Services: Operate (Linux)

Read This First: Start and Stop Servers and Services

CAUTION! There is a sequence for starting and stopping SAS Viya servers and services. You must follow this sequence to avoid operational issues. The SAS Viya start and stop scripts, including the sas-viya-all-services script, do not span multiple machines. You must run the appropriate script, in the correct sequence, on each machine in your SAS Viya topology.

Note: For more information about how to use the individual start and stop scripts, see “Start and Stop a Specific Server or Service”.

Start SAS Viya servers and services in this sequence:

1 Before you start these servers and services, check the system process list and process table, and stop or remove the process for any orphan or zombie service.

For more information, see your Linux documentation.

Note: If this is a multi-tenancy deployment, always start the tenant services last.

2 Start the SAS Configuration Server (Consul) by running the sas-viya-consul-default script.

For highly available deployments that run SAS Configuration Server on multiple machines, start the configuration server on those machines in any order.

To identify which machines contain SAS Configuration Servers, review the `[consul]` host group in your Ansible inventory.ini file.

Refer to information about how to run the sas-viya-consul-default script in *SAS Viya Administration: Infrastructure Servers*.

3 SAS Viya machines that do not host the SAS Configuration Server will host instead the SAS Configuration Server agent. Start the SAS Configuration Server agent on all other machines using the sas-viya-consul-default script.

Note: Most SAS Viya deployments contain either the SAS Configuration server or its agent. However, there are exceptions. To verify that a machine contains a configuration server agent, look for the sas-viya-consul-default script in **`/etc/init.d`**.

Important: On a distributed CAS analytics cluster, it is important to start SAS Configuration Server agents on the CAS worker machines before starting CAS on the CAS controller machine.

4 Start all instances of SAS Secret Manager (Vault).

Note: SAS Secret Manager is deployed wherever the SAS Configuration Server server is deployed. SAS Vault does not reside on machines that host configuration server agents.

Refer to information about how to run the sas-viya-vault-default script on page 600 in *SAS Viya Administration: Infrastructure Servers*.

5 Start SAS Message Broker (RabbitMQ).

If there are multiple instances of SAS Message Broker, start the last instance that went down first. Then, start the other instances.

> **TIP** If you can start all instances of SAS Message Broker within 30 seconds of each other, the order in which you start each instance is unimportant.

Refer to information about how to run the sas-viya-rabbitmq-server-default script on page 636 in *SAS Viya Administration: Infrastructure Servers*.

6 On the pgpool server machine, start the SAS Infrastructure Data Server cluster.

Refer to information about how to run the sas-viya-sasdatasvrc-postgres script on page 603 in *SAS Viya Administration: Infrastructure Servers*.

Note: Check the status of the cluster (sas-viya-sasdatasvrc-postgres status), to make sure that all nodes are running.

7 Start the HTTP proxy server (Apache HTTP Server).

Refer to information about how to run the httpd script on page 642 in *SAS Viya Administration: Infrastructure Servers*.

8 Then, start all remaining services using sas-viya-all-services.

Note: If you have more than one CAS server instance, always start the additional CAS servers after the original CAS server. The order in which you start the additional servers is unimportant. See Step 10.

After `sas-viya-all-services` has finished, run `sas-viya-all-services status`. If any services are reported as down, start it using its script.

9 If this is a multi-tenancy deployment, start the tenant services using the tenant all-services command. Order does not matter.

10 If you have more than one CAS server instance, always start the additional CAS servers after the original CAS server. The order in which you start the additional servers is unimportant.

On the additional CAS server machines, start the SAS Configuration Server agents first, followed by CAS and other SAS Viya services:

a Run the following command, appropriate for your operating system:

- On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-consul-default
```

- On Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-consul-default start
```

b Start all remaining servers and services:

```
sudo /etc/init.d/sas-viya-all-services start
```

Note: After following these steps, if the service still does not start as expected, check the log for the respective service in `/opt/sas/viya/config/var/log/`. For multi-tenant environments, check `/opt/sas/tenant/config/var/log/`.

> **TIP** Because the SAS Viya service start and stop sequence is so important, a best practice is to record the start and stop order of services for your site.

Stop SAS Viya servers and services in this sequence:

1 If this is a multi-tenancy deployment, always stop the tenant services first.

2 If you have more than one CAS server instance, always stop the additional CAS servers before the original CAS server.

Note: If you have more than one additional CAS server, the order in which you stop the CAS servers is unimportant.

On the additional CAS server machines, stop CAS and other SAS Viya services first followed by the SAS Configuration Server agents:

a Run the following command, appropriate for your operating system:

- On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl stop sas-viya-cascontroller-default
```

- On Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-cascontroller-default stop
```

b Stop the other SAS Viya services using their individual service scripts.

3 Stop servers and services on machines that do not contain the following:

- SAS Infrastructure Data Servers and PGPool server
- SAS Configuration Server (Consul)
- SAS Secret Manager (Vault)

Note: If you have machines that contain SAS Configuration Server, SAS Secret Manager, SAS Infrastructure Data Server, and other SAS Viya services, stop the other services first using their individual service scripts. Then, follow the order that is prescribed in Step 4 – Step 6.

4 On the pgpool server machine, stop the SAS Infrastructure Data Server cluster.

Refer to information about how to run the sas-viya-sasdatasvrc-postgres script on page 603 in *SAS Viya Administration: Infrastructure Servers*.

5 Stop all instances of SAS Message Broker (RabbitMQ).

Refer to information about how to run the sas-viya-rabbitmq-server-default script on page 636 in *SAS Viya Administration: Infrastructure Servers*.

6 Stop all instances of SAS Secret Manager and SAS Configuration Server.

Refer to information about how to run the sas-viya-vault-default script on page 636 and the sas-viya-consul-default script on page 597 in *SAS Viya Administration: Infrastructure Servers*.

Note: After following these steps, if the service still does not stop as expected, check the log for the respective service in `/opt/sas/viya/config/var/log/`. For multi-tenant environments, check `/opt/sas/`***`tenant`***`/config/var/log/`.

Start and Stop a Specific Server or Service

SAS Viya provides scripts in `/etc/init.d` that you use to stop, start, restart, and check the status of an individual SAS Viya server and service.

Syntax

How you run the individual server and service scripts depends on your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl status | stop | start | restart sas-viya-*server-or-service*-default

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-*server-or-service*-default status | stop | start | restart

Usage Notes and Tips

- You must be logged on to the machine where the particular service resides that you want to start or stop. Also, you must have root-level privileges to run these scripts.

CAUTION! There is a sequence for starting and stopping SAS Viya servers and services. You must follow this sequence to avoid operational issues. The SAS Viya start and stop scripts, including the sas-viya-all-services script, do not span multiple machines. You must run the appropriate script, in the correct sequence, on each machine in your SAS Viya topology. For more information, see “Read This First: Start and Stop Servers and Services”.

- On multi-tenant SAS Viya systems, individual server and service scripts are named `sas-tenant-ID-server-or-service-default`.

- To see the complete list of SAS Viya server and service scripts, run the following command: **`ls /etc/init.d/sas-viya-*`**. To operate Apache HTTP Server, see “Operate (Linux)”.

- On Linux systems that support systemd, use the `systemctl` command when running the individual service and server scripts. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-all-services` on Red Hat Enterprise Linux 7.*x* with the `service` command, and later attempt to shut down all services using the `systemctl` command, the service stops responding and does not shut down.

Examples

- To check status of SAS Logon Manager on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-saslogon-default
```

- To stop the Comments service on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-comments-default stop
```

- To start SAS Configuration Server on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-consul-default
```

Start and Stop All Servers and Services

SAS Viya provides a script, `sas-viya-all-services`, in **`/etc/init.d`** that you use to stop, start, and check the status of all SAS Viya servers and services.

Syntax

Unlike the individual server and services start scripts, there is only one method for running `sas-viya-all-services`:

`sas-viya-all-services status | stop | start`

Usage Notes and Tips

- You must be logged on to the machine where the SAS Viya servers and services reside, and you must have root-level privileges to run `sas-viya-all-services`.

CAUTION! There is a sequence for starting and stopping SAS Viya servers and services. You must follow this sequence to avoid operational issues. The SAS Viya start and stop scripts, including the sas-viya-all-services script, do not span multiple machines. You must run the appropriate script, in the correct sequence, on each machine in your SAS Viya topology. For more information, see "Read This First: Start and Stop Servers and Services".

- On multi-tenant SAS Viya systems, `sas-viya-all-services` is named `sas-tenant-ID-all-services`.
- `sas-viya-all-services` does not control Apache HTTP Server.
- When checking status, it is normal for certain servers and services to not display host, port, and PID information. The reason is that these servers and services are not registered with the SAS Configuration Server, including the configuration server itself.

Examples

- To check status of all servers and services:

```
sudo /etc/init.d/sas-viya-all-services status
```

- To stop all servers and services:

```
sudo /etc/init.d/sas-viya-all-services stop
```

- To start all servers and services:

```
sudo /etc/init.d/sas-viya-all-services start
```

General Servers and Services: Operate (Windows)

Start and Stop All Servers and Services

Using the Microsoft Management Console (MMC) Services snap-in, you can use the **SAS Services Manager** service to start and stop all SAS Viya servers and services.

Figure A.3 SAS Viya Servers and Services in the Services Snap-In

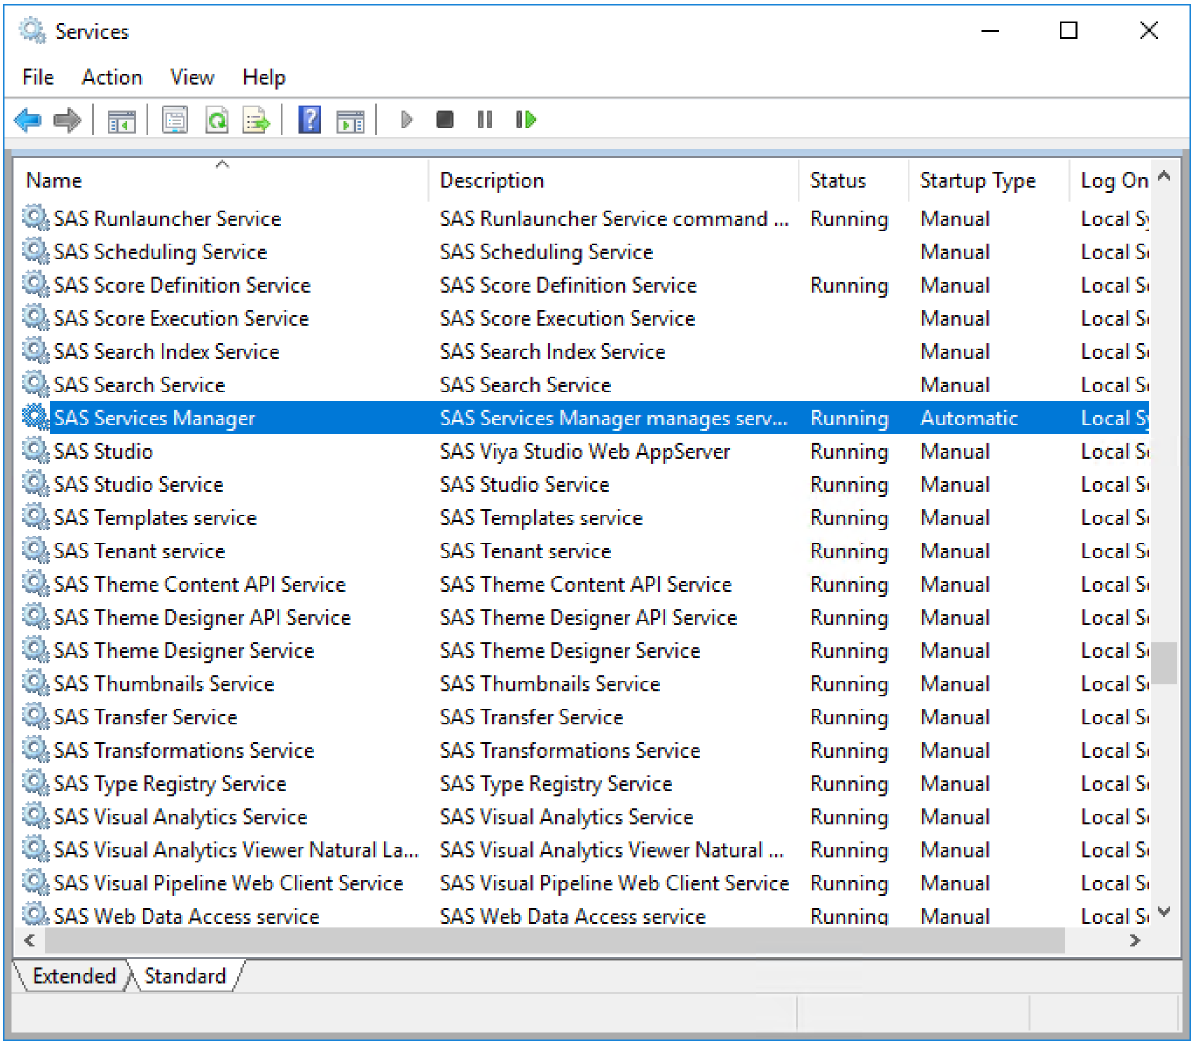

Name	Description	Status	Startup Type	Log On
SAS Runlauncher Service	SAS Runlauncher Service command ...	Running	Manual	Local Sy
SAS Scheduling Service	SAS Scheduling Service		Manual	Local S
SAS Score Definition Service	SAS Score Definition Service	Running	Manual	Local S
SAS Score Execution Service	SAS Score Execution Service		Manual	Local S
SAS Search Index Service	SAS Search Index Service		Manual	Local S
SAS Search Service	SAS Search Service		Manual	Local S
SAS Services Manager	SAS Services Manager manages serv...	Running	Automatic	Local Sy
SAS Studio	SAS Viya Studio Web AppServer	Running	Manual	Local S
SAS Studio Service	SAS Studio Service	Running	Manual	Local S
SAS Templates service	SAS Templates service	Running	Manual	Local S
SAS Tenant service	SAS Tenant service	Running	Manual	Local S
SAS Theme Content API Service	SAS Theme Content API Service	Running	Manual	Local S
SAS Theme Designer API Service	SAS Theme Designer API Service	Running	Manual	Local S
SAS Theme Designer Service	SAS Theme Designer Service	Running	Manual	Local S
SAS Thumbnails Service	SAS Thumbnails Service	Running	Manual	Local S
SAS Transfer Service	SAS Transfer Service	Running	Manual	Local S
SAS Transformations Service	SAS Transformations Service	Running	Manual	Local S
SAS Type Registry Service	SAS Type Registry Service	Running	Manual	Local S
SAS Visual Analytics Service	SAS Visual Analytics Service	Running	Manual	Local S
SAS Visual Analytics Viewer Natural La...	SAS Visual Analytics Viewer Natural ...	Running	Manual	Local S
SAS Visual Pipeline Web Client Service	SAS Visual Pipeline Web Client Service	Running	Manual	Local S
SAS Web Data Access service	SAS Web Data Access service	Running	Manual	Local S

SAS Services Manager is a registered service that is configured to start automatically. **SAS Services Manager** starts and stops SAS Viya servers and services in the proper sequence.

Important:

Do not attempt to restart SAS Services Manager, using the **restart** command in the Services snap-in.

In the Services snap-in, use the **stop** command to stop SAS Services Manager, and wait until all the SAS Viya services have stopped (are down). After all the services have stopped, use the **start** command in the Services snap-in to start SAS Services Manager. Failure to wait until all SAS Viya services have stopped before you start SAS Services Manager causes errors.

Start and Stop a Specific Server or Service

Use the Microsoft Management Console (MMC) Services snap-in to start, stop, and restart individual SAS Viya servers and services.

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Remove Erroneous GPU Reservations

The GPU Reservation service is a SAS Viya service that aids SAS processes in resource sharing and utilization of graphics processing units (GPUs) that are available on a system. A companion tool to the GPU Reservation service is gputool.

Communication with the GPU Reservation service to remove an erroneous reservation is the primary use for gputool. An erroneous reservation can occur when a SAS Viya server that has made the reservation abnormally exits without releasing the reservation.

1 Log on to the machine that hosts the GPUs.

2 Change to the **`/opt/sas/viya/home/SASFoundation/utilities/bin/`** directory.

3 If you know the PCI bus ID for the GPU, skip this step and go to Step 4.

   Otherwise, run the following command to obtain the bus ID for the GPU:

   `./gputool -l`

   You receive output similar to the following:

```
----------------------------
PCIBusId 00000000:03:00.0
Device 0 is reserved
----------------------------
PCIBusId 00000000:82:00.0
Device 1 is available
```

4 To remove a GPU reservation, enter the following command:

   `./gputool -release PCI-bus-ID`

   where *PCI-bus-ID* is the bus ID for the GPU (obtained in Step 3).

   Here is an example of the command and output:

   `./gputool -release 00000000:03:00.0`

```
NOTE: GPU released
```

5 Gputool also provides the following command-line options:

- **-d**

  Causes the GPU Reservation service to dump its reservation statistics to `/var/tmp/sasgpud_resdump`.

- **-g**

  Prints the utilization for each GPU to the screen.

- **-h**

  Prints help (the gputool usage commands) to the screen.

Fault Tolerance in SAS Viya (Linux)

The following figure shows the recommended fault-tolerant topology for SAS Viya on Linux.

Figure A.4 Recommended Fault-Tolerant Topology for SAS Viya on Linux

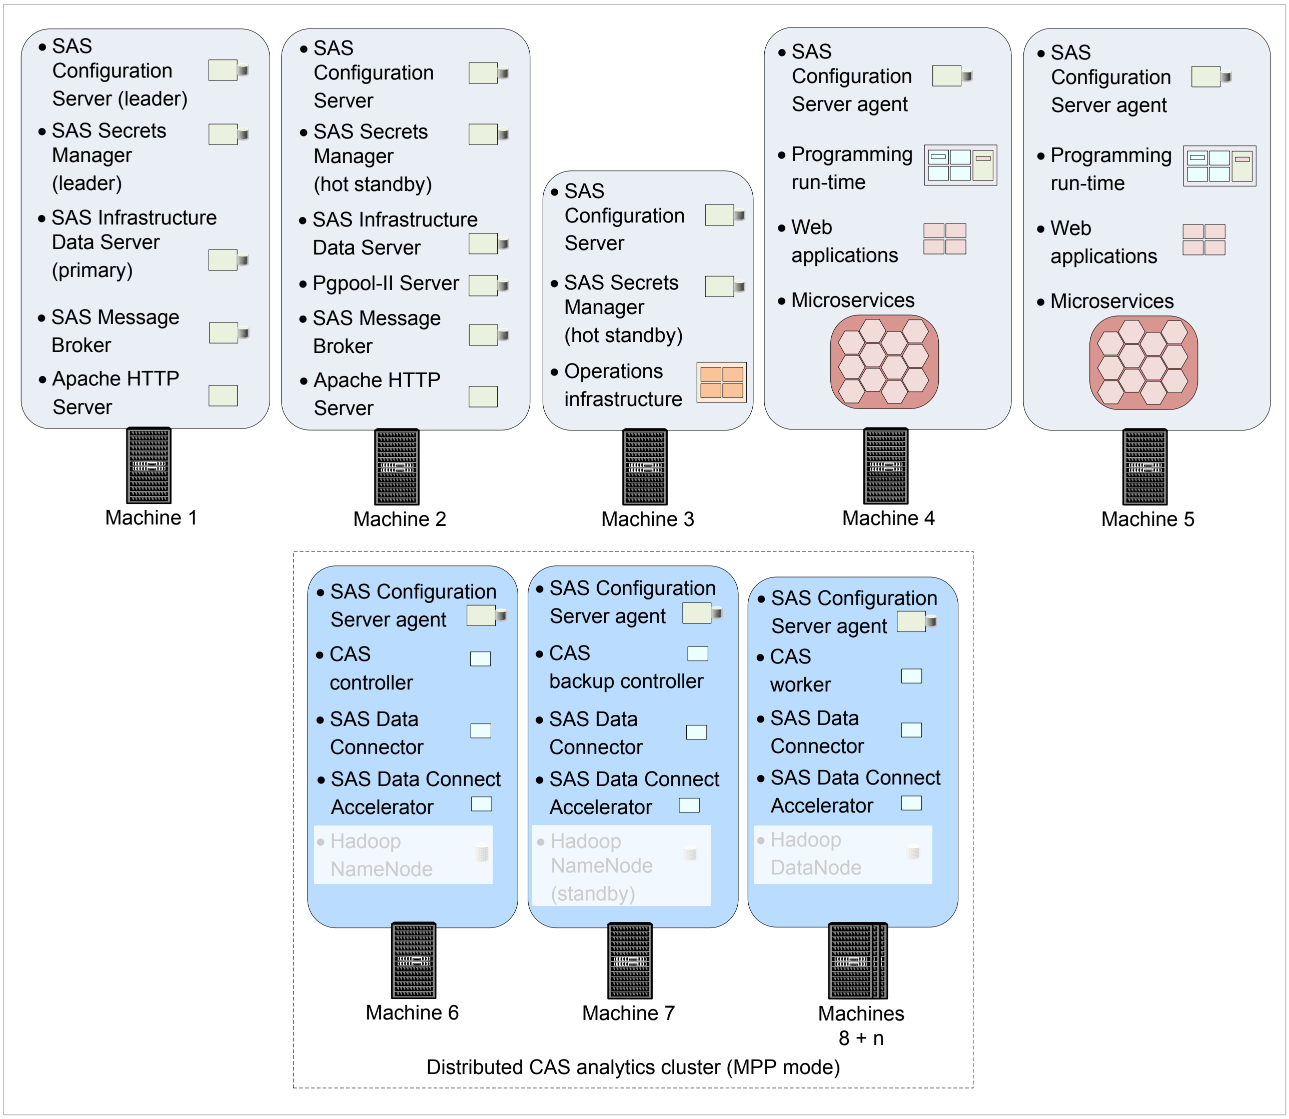
SAS Configuration Server (leader)
SAS Secrets Manager (leader)
SAS Infrastructure Data Server (primary)
SAS Message Broker
Apache HTTP Server
Machine 1
SAS Configuration Server
SAS Secrets Manager (hot standby)
SAS Infrastructure Data Server
Pgpool-II Server
SAS Message Broker
Apache HTTP Server
Machine 2
SAS Configuration Server
SAS Secrets Manager (hot standby)
Operations infrastructure
Machine 3
SAS Configuration Server agent
Programming run-time
Web applications
Microservices
Machine 4
SAS Configuration Server agent
Programming run-time
Web applications
Microservices
Machine 5
SAS Configuration Server agent
CAS controller
SAS Data Connector
SAS Data Connect Accelerator
Hadoop NameNode
Machine 6
SAS Configuration Server agent
CAS backup controller
SAS Data Connector
SAS Data Connect Accelerator
Hadoop NameNode (standby)
Machine 7
SAS Configuration Server agent
CAS worker
SAS Data Connector
SAS Data Connect Accelerator
Hadoop DataNode
Machines 8 + n
Distributed CAS analytics cluster (MPP mode)

Most of the SAS Viya infrastructure servers reside on machines 1 and 2.SAS Configuration Server (Consul) is unique because it uses a consensus algorithm that requires a leader and a quorum to maintain and replicate logs. For more information, see www.consul.io/docs/internals/consensus.html.

Three configuration server instances (machines 1–3) is the minimum number of instances that is required to provide fault tolerance. SAS Secrets Manager (Vault) is always deployed on the same machine as the configuration server. (Machines that contain configuration agents do not also contain SAS Secrets Manager.)

Each machine in a SAS Viya deployment runs a SAS Configuration Server or an agent process that performs health checks on the SAS Viya services that are running and on the machine itself. Each configuration server agent provides health information to one or more configuration servers. The configuration servers store and replicate service information.

SAS Viya services send queries to configuration servers or configuration server agents in order to discover other services.

The SAS Infrastructure Data Server on machine 1 is the primary data server. If machine 1 goes down, the data server should fail over to machine 2. However, the data server is not truly fault tolerant because SAS Viya does not support multiple pgpool servers. For example, if the pgpool server goes down (machine 2), the data server will go offline.

The SAS Viya operations infrastructure gathers metrics and log data from SAS Viya processes and stores the data in CAS tables. In turn, SAS Viya applications use this data to produce reports and views, such as the Logs window in SAS Environment Manager. Because this infrastructure is not fault tolerant, it resides only on one machine (machine 3).

The SAS Viya programming run-time environment, the web applications, and the microservices (machine 4 and 5) are separated from the SAS Viya infrastructure servers (machines 1 and 2) to enhance performance. (Machine 5 provides fault tolerance for machine 4.)

For Cloud Analytic Services (CAS) that are running in MPP mode and are distributed across an analytics cluster, unique forms of fault tolerance are available for each machine type: the CAS worker and the CAS controller.

Because CAS worker machines outnumber CAS controller machines, worker machines are more likely to experience failure. Fault tolerance is provided automatically for worker machines that contain CAS tables, which are created with redundant copies of blocks.

The less common CAS controller failure problem is addressed with an optional CAS backup controller (also referred to as the *secondary controller*). For more information, see "Fault Tolerance" on page 504.

In the recommended fault-tolerant topology, the CAS controller is deployed on Machine 6, its backup controller is deployed on Machine 7, and CAS workers are on Machines 8 + n.

General Servers and Services: Troubleshooting

Starting sas-viya-consul-default
Timed out waiting for the consul service to start Exiting

Explanation:

The SAS Configuration Server (Consul) does not start. One cause might be that certain configuration files are corrupted. For example, a nonzero length configuration file might contain a definition that is not supported by the current version of Consul. This can cause a failure for Consul to start if the definition is no longer valid, and the affected SAS Viya product did not fix its definition during a software upgrade.

Resolution:

Determine what is causing the start-up failure. Then use the command-line interface, sas-bootstrap-config, to remove invalid checks and services using its `agent check deregister` and `agent service deregister` commands.

One or more SAS Viya microservices fail to start up UnknownHostException: rabbitmq.service.consul

Explanation:

All microservices use the same API to publish and receive events from SAS Message Broker (RabbitMQ). The microservices are attempting to fetch one or more message broker server host names from SAS Configuration Server (Consul) but that information is not registered correctly because of missing information in `/etc/hosts`.

Resolution:

Make sure that `/etc/hosts` contains every machine name in your SAS Viya deployment, and that `/etc/hosts` has been copied to every machine in your SAS Viya system.

sas-viya-all-services status command returns ‘not ready’

Explanation:

Machines in your SAS Viya deployment might be defined in `/etc/hostname` with a short host name.

Resolution:

Using the operating system commands that are appropriate for your Linux distribution (for example, `hostname` and `hostnamectl`), determine how the operating system identifies each machine in your SAS Viya deployment. Then, redefine the machines in your SAS Viya deployment using their fully qualified domain names (for example, my-machine.example.com). For more information, see your Linux documentation.

How do I verify that SAS Viya servers and services are healthy after an outage?

Explanation:

After an outage, you want to know that the SAS Viya servers and services are running normally.

Resolution:

1. Run the following command:

   ```
   sudo /etc/init.d/sas-viya-all-services status
   ```

2. For any server or service that is listed as `down`, check its log.

   The logs reside in `/opt/sas/viya/config/var/log/`***`server-service-name/deployment-instance`***.

3. For any server or service that is listed as `down`, run the following command:

   ```
   /opt/sas/viya/home/bin/sas-csq service-health --service server-service-name
   ```

4. Open a SAS Technical Support track and attach the output from steps 2 and 3.

Job cannot run because no GPU can be reserved.

Explanation:

A SAS Viya server runs a job using a thread that makes a GPU reservation. The thread that made the reservation exits (normally or abnormally) without removing the reservation. Because the GPU Reservation service tracks GPUs using only process IDs, the service does not know to release the GPU reservation. In this situation, the GPU is reserved indefinitely and prevents future SAS jobs from being able to make a reservation.

Resolution:

If no one is using the GPU reservation, remove the reservation using gputool.

27

SAS Cloud Analytic Services

SAS Cloud Analytic Services: Overview

SAS Cloud Analytic Services (CAS) is a server that provides the cloud-based, run-time environment for data management and analytics with SAS. Suitable for both on-premises and Cloud deployments, CAS uses a combination of hardware and software where data management and analytics take place on either a single machine or as a distributed server across multiple machines.

The following diagram shows the relationship of CAS to other components in the SAS Viya full deployment:

Figure A.1 SAS Cloud Analytic Services (Full Deployment)

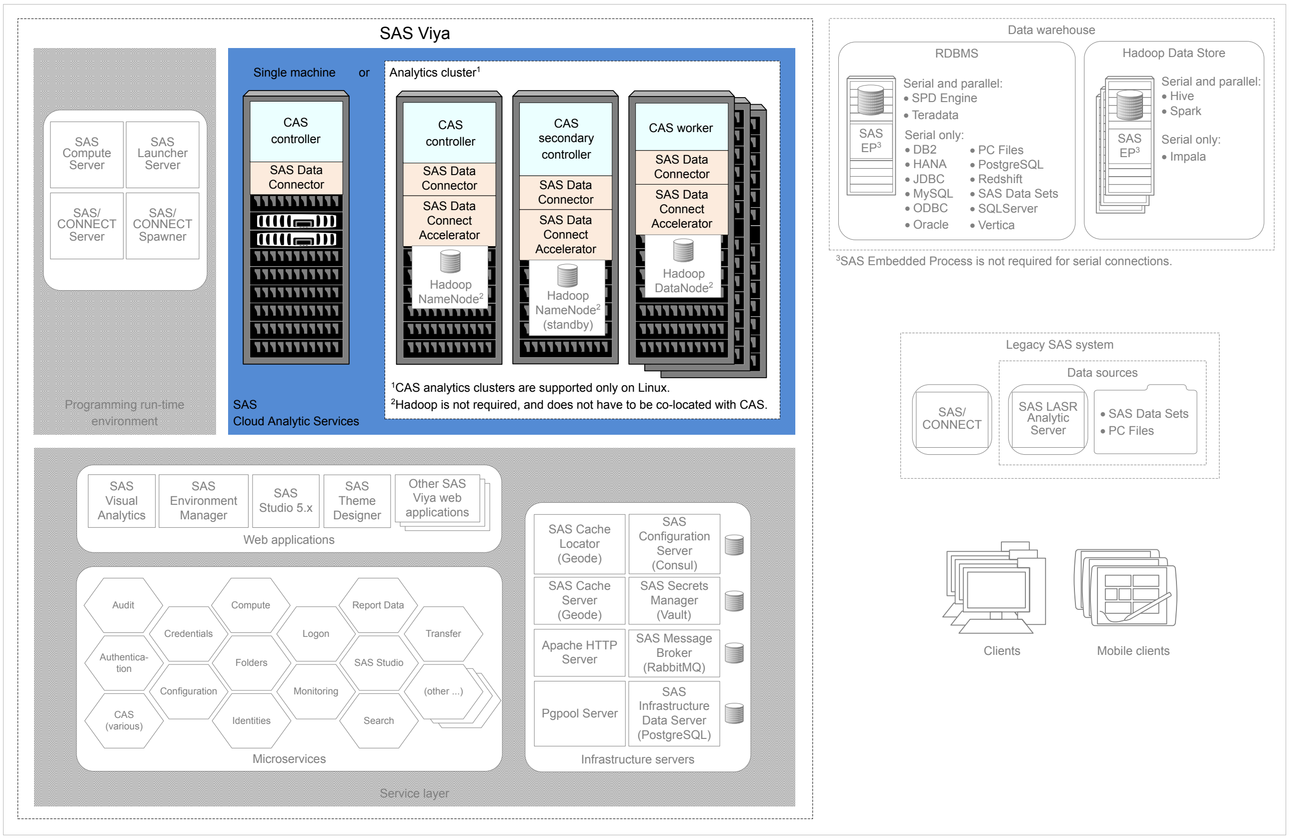

Note: CAS analytics clusters are supported only on Linux.

The following diagram shows the relationship of CAS to other components in the SAS Viya programming-only deployment:

Figure A.2 SAS Cloud Analytic Services (Programming-only Deployment)

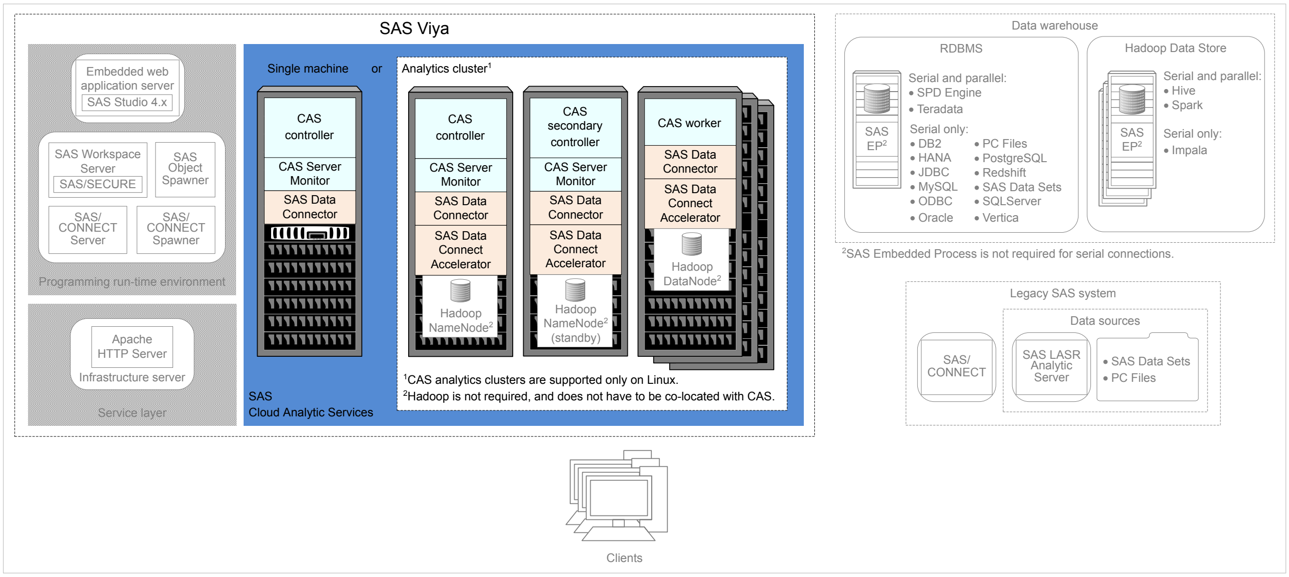

SAS Cloud Analytic Services: How To (Scripts)

Operate (Linux)

SAS Viya provides a script in **/etc/init.d** that you use to stop, start, restart, and check the status of a CAS server. The script is named, `sas-viya-cascontroller-default`.

Syntax

How you run `sas-viya-cascontroller-default` depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

  sudo systemctl status | stop | start | restart sas-viya-cascontroller-default

- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

  sudo service sas-viya-cascontroller-default status | stop | start | restart

Usage Notes and Tips

- You must be signed in to the machine where the CAS controller resides and you must have root-level privileges to run this script. Running sas-viya-cascontroller-default affects all worker nodes.
- `sas-viya-cascontroller-default` checks the status of the CAS controller only. To check the status of an individual worker node, use SAS Environment Manager or CAS Server Monitor.
- On multi-tenant SAS Viya systems, the script is named `sas-tenant-ID-sas-viya-cascontroller`.
- Your site's Linux administrator might want to create a regular account (for example, sas-service-admin) and give that account the sudo permissions to manage the SAS services. Make sure that the services are defined as “start on reboot” so that the CAS server automatically starts when the machine is rebooted.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see “Start and Stop All Servers and Services” on page 462.
- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-cascontroller-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-cascontroller-default` on RHEL 7.x with the `service` command, and later attempt to shut down the CAS server using the `systemctl` command, the CAS server stops responding and does not shut down.

Examples

- To check status of the CAS controller on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-cascontroller-default
```

- To stop the CAS controller and its worker nodes on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-cascontroller-default stop
```

- To start the CAS controller and its worker nodes on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-cascontroller-default
```

- To restart the CAS controller and its worker nodes on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-cascontroller-default restart
```

Operate (Windows)

Using the Services snap-in in the Microsoft Management Console, you can start, stop, and restart SAS Cloud Analytic Services.

Figure A.3 *SAS Cloud Analytic Services in the Services snap-in*

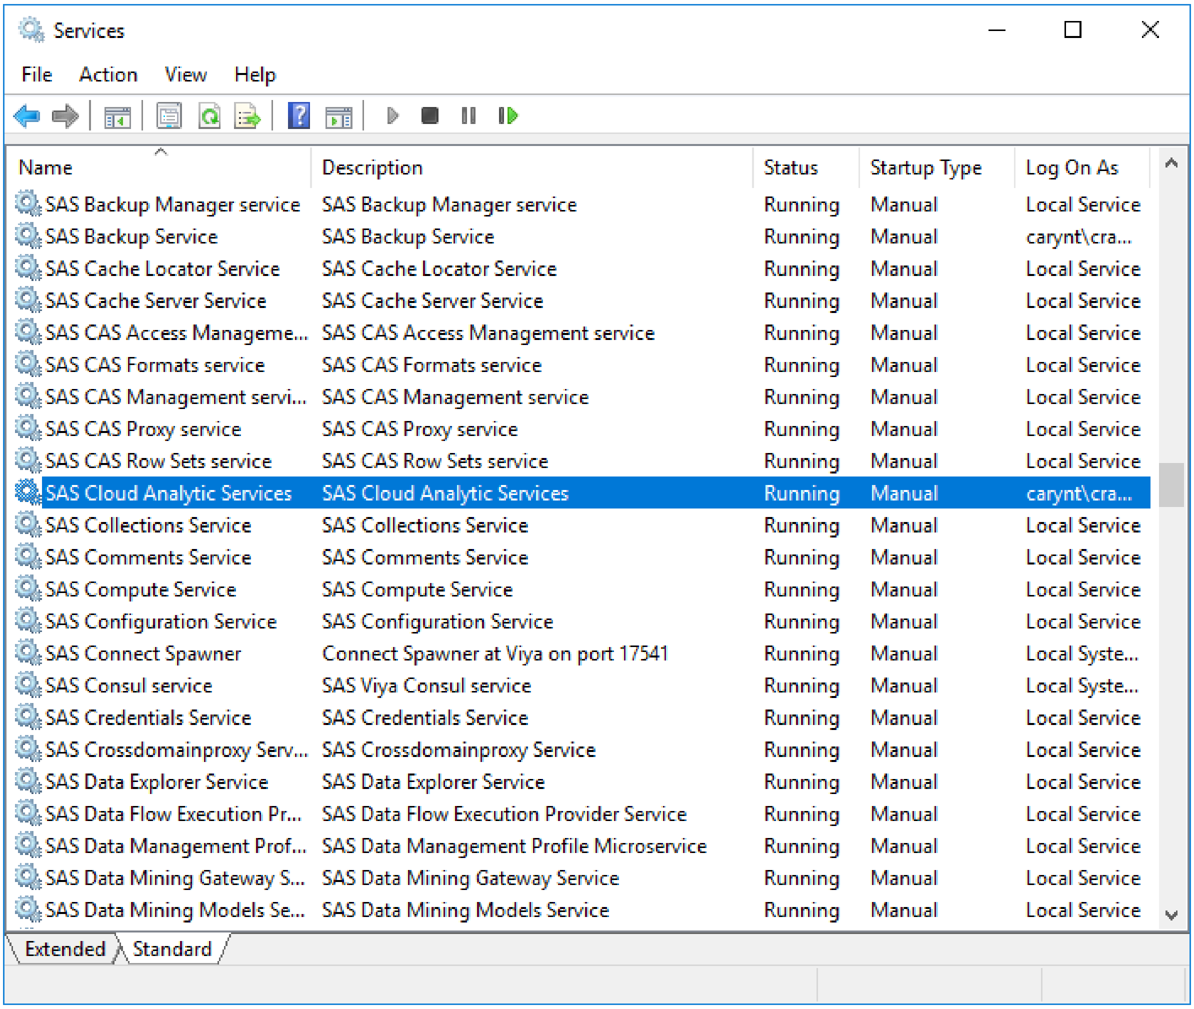

Name	Description	Status	Startup Type	Log On As
SAS Backup Manager service	SAS Backup Manager service	Running	Manual	Local Service
SAS Backup Service	SAS Backup Service	Running	Manual	carynt\cra...
SAS Cache Locator Service	SAS Cache Locator Service	Running	Manual	Local Service
SAS Cache Server Service	SAS Cache Server Service	Running	Manual	Local Service
SAS CAS Access Manageme...	SAS CAS Access Management service	Running	Manual	Local Service
SAS CAS Formats service	SAS CAS Formats service	Running	Manual	Local Service
SAS CAS Management servi...	SAS CAS Management service	Running	Manual	Local Service
SAS CAS Proxy service	SAS CAS Proxy service	Running	Manual	Local Service
SAS CAS Row Sets service	SAS CAS Row Sets service	Running	Manual	Local Service
SAS Cloud Analytic Services	SAS Cloud Analytic Services	Running	Manual	carynt\cra...
SAS Collections Service	SAS Collections Service	Running	Manual	Local Service
SAS Comments Service	SAS Comments Service	Running	Manual	Local Service
SAS Compute Service	SAS Compute Service	Running	Manual	Local Service
SAS Configuration Service	SAS Configuration Service	Running	Manual	Local Service
SAS Connect Spawner	Connect Spawner at Viya on port 17541	Running	Manual	Local Syste...
SAS Consul service	SAS Viya Consul service	Running	Manual	Local Syste...
SAS Credentials Service	SAS Credentials Service	Running	Manual	Local Service
SAS Crossdomainproxy Serv...	SAS Crossdomainproxy Service	Running	Manual	Local Service
SAS Data Explorer Service	SAS Data Explorer Service	Running	Manual	Local Service
SAS Data Flow Execution Pr...	SAS Data Flow Execution Provider Service	Running	Manual	Local Service
SAS Data Management Prof...	SAS Data Management Profile Microservice	Running	Manual	Local Service
SAS Data Mining Gateway S...	SAS Data Mining Gateway Service	Running	Manual	Local Service
SAS Data Mining Models Se...	SAS Data Mining Models Service	Running	Manual	Local Service

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see “Start and Stop All Servers and Services”.

Change the Process Owner Account

1 Log on to the CAS controller machine as the SAS install user (sas) or with root-level privileges.

2 Using a text editor, open **/opt/sas/viya/config/etc/sysconfig/cas/default/sas-cas-usermods**.

3 Locate or add the following lines:

```
SASUSER="user-account"
SASGROUP="primary-group"
```

Note: The default process owner account is cas. The default primary group for cas is sas.

Enter the new CAS process owner account. If needed, enter a new primary group for the CAS process owner, and save the file.

4 Open **/opt/sas/viya/home/SASFoundation/utilities/bin/launchconfig-viya-default**.

5 Locate the following line:

```
restrictServerLaunch=user-account
```

Enter the new CAS process owner account, and save the file.

CAS uses the new process owner account the next time it is run.

Add New Worker Nodes or a Backup Controller

The processes for adding new worker nodes or a backup controller (also known as a secondary controller) to your CAS server are very similar.

1 Make sure that you are licensed for the additional nodes or a backup controller that you are planning to add to your analytic cluster.

2 If you are adding a backup controller, make sure that the backup controller and the CAS controller (the primary controller) both use the same shared file system.

When starting CAS sessions from SAS Studio or from any interface by users in the `CASHostAccountRequired` group, the users' home directories (`$HOME`) must be shared so that they can be accessed on both the controller machine and the backup controller machine. Sharing users' home directories ensures the path for the `CASUSER` library is available during CAS session start-up.

For most other CAS session scenarios, the `CASUSER` library is set to a path in the shared file system described in "Set Up a Shared File System for CAS Controllers (Post-Deployment)".

3 When adding to an existing SAS Viya deployment, SAS downloads and installs the latest software available from the software repositories. Therefore, make sure that you are using a mirror repository.

For more information, see "Create a Mirror Repository" in *SAS Viya for Linux: Deployment Guide*.

4 Every machine on which you are installing CAS worker nodes or a backup controller must have the CAS user account (cas) and group (sas) set up.

For more information, see "Set Up the cas Account" in *SAS Viya for Linux: Deployment Guide*.

5 Sign in to the Ansible controller as the user account that deploys the software.

For more information, see "Set Up the User Account that Deploys the Software" in *SAS Viya for Linux: Deployment Guide*.

6 Choose which playbook to use. If you are adding a new:

- backup controller

  Use the site playbook (site.yml).

- worker node

  Decide which playbook to use.

  Note: When adding worker nodes to a CAS controller running in SMP mode, use the site playbook.

CAUTION! Use the deploy-casworker playbook for adding worker nodes only. Do not change other CAS server configuration settings using the deploy-casworker playbook. Doing so can cause a mismatch between configuration in memory versus configuration on disk.

7 In the inventory file, define the machines on which you are adding the worker nodes or a backup controller.

> **TIP** If you used the recommended location for uncompressing your playbook, the file is located at **/sas/install/sas_viya_playbook/inventory.ini**.

Here is an example of adding a backup controller:

```
controller_02 ansible_host=controller_02.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
```

Here is an example of adding a worker node:

```
worker_023 ansible_host=worker_23.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
```

For more information, see "Specify the Machines in the Deployment" in *SAS Viya for Linux: Deployment Guide*.

8 Also, in the inventory file, add the machines that you are adding to the appropriate group:

- backup controller

  Use the `sas-casserver-secondary` group.

  Note: If your inventory file does not contain a `sas-casserver-secondary` group, then create one, using the example that follows as a guide.

  In this example, a backup controller is being added to the controller_02 machine:

  ```
  [sas-casserver-secondary]
   controller_02
  ```

- worker nodes

  Use the `sas-casserver-worker` group.

  Note: If your inventory file does not contain a `sas-casserver-worker` group, then create one, using the example that follows as a guide.

  In this example, a worker node is being added to the worker_23 machine:

  ```
  [sas-casserver-worker]
   worker_019
   worker_020
   worker_021
   worker_022
   worker_023
  ```

9 Run Ansible, using the playbook that you chose in Step 6:

```
ansible-playbook -i inventory.ini playbook.yml
```

10 If you ran the deploy-casworker playbook and want to immediately start the worker nodes that you have added, sign in to CAS Server Monitor to add (join) on page 488 the new worker nodes.

Set Up a Shared File System for CAS Controllers (Post-Deployment)

If you want to make your CAS controller fault tolerant, during installation you can choose to deploy a CAS backup controller (also referred to as a secondary controller). A requirement for operating CAS with a backup controller is that it and the CAS controller (the primary controller) must both use the same shared file system.

1 Shut down the CAS controller.

2 Copy all of the data from **/opt/sas/viya/config/data/cas** to a network share that both the CAS controller and its backup controller can access.

   Note: Make sure that the copy preserves directory ownership and permissions.

   Here is an example:

   ```
   cp -Rp /opt/sas/viya/config/data/cas /share
   ```

   > **TIP** After you copy **/opt/sas/viya/config/data/cas**, remember that you will have a **cas** directory under your target. For example, if you copy **/opt/sas/viya/config/data/cas** to **/share**, you will have a cas directory under **/share** (**/share/cas/**).

3 Verify that all files and directories have been copied.

4 On the CAS controller, delete the old data directory.

   Here is an example:

   ```
   rm -r /opt/sas/viya/config/data/cas
   ```

5 On both the CAS controller and the backup controller machines, create a Linux symbolic link in **/opt/sas/viya/config/data/cas** that points to the new shared file system.

   Here is an example:

   ```
   ln -sf /share/cas /opt/sas/viya/config/data/cas
   ```

6 Start the CAS controller.

Recover a Failed Controller

Note: While CAS is operating in the failed-over state, do not restart the primary (failed) controller service.

1 Shut down the CAS controller using SAS Environment Manager. (During a failover, the backup controller becomes the primary controller.)

   For more information, see “Stop a CAS Server”.

2 Perform whatever steps necessary to either repair or replace the failed primary controller.

3 Restore the permstore directory from the backup controller to the primary controller.

   For more information, see “Restore the Most Recent Permstore on Linux in the Event of a Failover” on page 68.

4 Do the following:

   a Restart all CAS worker nodes in your deployment. (Do not start your CAS controller.)

b Start the CAS controller.

For more information, see “Operate (Linux)” on page 472.

c At the Linux command prompt, enter the **sas-viya-controller-*deployment-instance*** command and verify that the CAS controller is indeed running.

Here is an example:

```
./sas-viya-cascontroller-default status
 sas-viya-cascontroller-default is running

Host role in cluster:

Cluster Information:
    Tenant                = shared
    Instance              = default
    Primary Controller    =
```

Add a CAS Server

To add a CAS server, follow these steps:

Note: If you want to only add a CAS worker node, or a CAS backup controller, refer to “Add New Worker Nodes or a Backup Controller”.

1 Before you proceed, make sure that your SAS Viya system meets the requirements.

2 When adding to an existing SAS Viya deployment, SAS downloads and installs the latest software available from the software repositories. Therefore, make sure that you are using an existing mirror repository.

For more information, see “Create a Mirror Repository” in *SAS Viya for Linux: Deployment Guide*.

Important: If you did not use a mirror repository to deploy SAS Viya, contact SAS Technical Support before proceeding.

3 Sign in to the Ansible controller as the user account that deploys the software.

For more information, see “Set Up the User Account that Deploys the Software” in *SAS Viya for Linux: Deployment Guide*.

4 Each CAS server that you are adding must have its own vars.yml and inventory.ini files.

Important: The recommended method for managing multiple instances of vars.yml and inventory.ini is to create a copy of each. Include the name of the CAS server in the filenames of these copies. Store these new copies in the playbook directory. The examples used in this document are casserv02.vars.yml and casserv02.inventory.ini.

Change to your Ansible playbook directory and make a copy of vars.yml.

> **TIP** If you used the recommended location for uncompressing your playbook, vars.yml is located in **/sas/install/sas_viya_playbook/**.

Here is an example:

```
cp vars.yml casserv02.vars.yml
```

5 Create a unique instance name for the CAS server that you are adding.

Open the copy of vars.yml (in this document, casserv02.vars.yml) and locate the `CAS CONFIGURATION` section. Under `casenv_user`, add the following line:

casenv_instance: *new-CAS-server-name*

Important: CAS server instance names must contain only alphanumeric characters. Any case is allowed. Instance names must not contain any special characters.

Here is an example:

```
############################################################################
## CAS Configuration
############################################################################

# The user that the CAS process will run under

casenv_user: cas

casenv_instance: server02

# The group that the CAS user belongs to

casenv_group: sas
```

> **TIP** The value of `casenv_instance` is appended to the base name, `cas-shared`, to form the deployment instance name for the new CAS server. In this example, the full deployment instance name is: `cas-shared-server02`.

6 Evaluate each CAS server configuration option in the `CAS CONFIGURATION` section, and update each as appropriate for the new CAS server instance.

For more information, see "Configuration File Options" on page 512.

7 Also, each CAS server that you are adding must have its own inventory.ini file. Make a copy of your inventory.ini file.

Here is an example:

```
cp inventory.ini casserv02.inventory.ini
```

8 In a text editor, open the copy of the new inventory.ini file (casserv02.inventory.ini), and do the following:

a At the very top of the file, there is a list of deployment targets. Do not remove this list. (You add to this list later in this document.)

```
controller_01 ansible_host=controller_101.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
controller_02 ansible_host=controller_102.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_01 ansible_host=worker_101.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_02 ansible_host=worker_102.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_03 ansible_host=worker_103.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_04 ansible_host=worker_104.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_05 ansible_host=worker_105.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
```

b Remove the original deployment targets from all host groups in the file, except for `[consul]`, `[httpproxy]` and `[sas-all:children]`. These three host groups must contain their original entries.

Important: A host group can have no entries under it. But, a host group should not be removed, even if it is empty. Do not modify `[consul]`, `[httpproxy]` and `[sas-all:children]` or deployment failures are likely to occur.

9 If you are adding a distributed CAS server across multiple machines, skip to Step 12.

10 If you adding a CAS non-distributed server to a single machine, perform these steps:

a Using a text editor open the copy of your inventory.ini (in this document, casserv02.inventory.ini). At the very top of the file, map a deployment target to a machine, an Ansible user, and a private key file.

Here is an example:

```
new_controller ansible_host=controller_201.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
```

b Assign the deployment targets that you mapped in Step 10a to the `[sas-casserver-primary]` host group.

Here is an example:

```
[sas-casserver-primary]
new_controller
```

c Verify that the `[consul]`, `[httpproxy]` and `[sas-all:children]` host groups are present, and that they contain the entries from your original SAS Viya deployment.

Here is an example:

```
[consul]
original_deployment_target

[httpproxy]
original_deployment_target

[sas-all:children]
hostgroup1
hostgroup2
hostgroup3
hostgroup4
.
.
.
```

11 Skip to Step 13.

12 If you are adding a distributed CAS server across multiple machines, perform these steps:

a Using a text editor open the copy of your inventory.ini (in this document, casserv02.inventory.ini). At the very top of the file, map a deployment target to a machine, an Ansible user, and a private key file.

Here is an example for mapping five deployment targets: one controller, one backup controller (optional), and three workers (highlighted text that follows):

```
controller_01 ansible_host=controller_101.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
controller_02 ansible_host=controller_102.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_01 ansible_host=worker_101.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_02 ansible_host=worker_102.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_03 ansible_host=worker_103.example.com ansible_user=user1
```

```
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_04 ansible_host=worker_104.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
worker_05 ansible_host=worker_105.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
new_controller_01 ansible_host=controller_201.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
new_controller_02 ansible_host=controller_202.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
new_worker_03 ansible_host=worker_203.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
new_worker_04 ansible_host=worker_204.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
new_worker_05 ansible_host=worker_205.example.com ansible_user=user1
ansible_ssh_private_key_file= ~/.ssh/id_rsa
```

b Assign the deployment targets that you mapped in Step 12a to the required host groups.

Here is an example:

```
[sas-casserver-primary]
new_controller_01

[sas-casserver-secondary]
new_controller_02

[sas-casserver-worker]
new_worker_03
new_worker_04
new_worker_05
```

c Verify that the `[consul]`, `[httpproxy]` and `[sas-all:children]` host groups are present, and that they contain the entries from your original SAS Viya deployment.

Here is an example:

```
[consul]
original_deployment_target

[httpproxy]
original_deployment_target

[sas-all:children]
hostgroup1
hostgroup2
hostgroup3
hostgroup4
.
.
.
```

13 Run Ansible:

```
ansible-playbook -i CAS-server-inventory-file-name site.yml -e "@CAS-server-vars-file-name"
```

Here is an example:

```
ansible-playbook -i casserv02.inventory.ini site.yml -e "@casserv02.vars.yml"
```

For more information, see “Adding SAS Viya Software to a Deployment and Upgrading Products in SAS Viya 3.4” in *SAS Viya for Linux: Deployment Guide*.

14 Verify that the CAS server that you added is running by launching gridmon.sh.

a Log on to the CAS controller machine as a user with passwordless SSH access to all CAS nodes.

b Run the following command to start gridmon.sh:

```
/opt/sas/viya/home/SASFoundation/utilities/bin/gridmon.sh
```

c You should see one or more jobs running.

Figure A.4 gridmon.sh Running in Job Mode

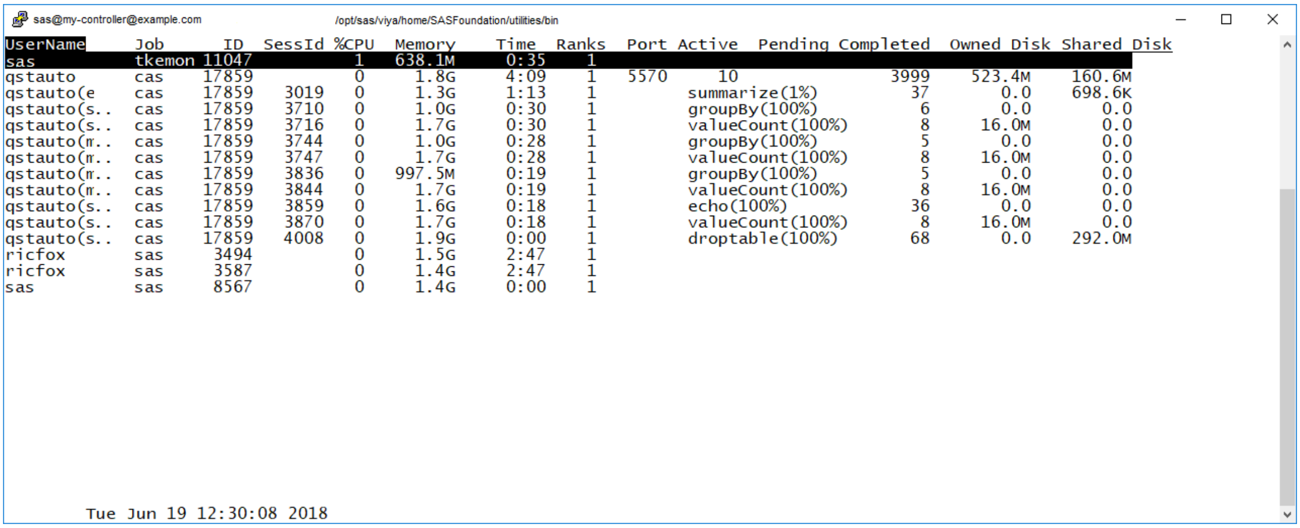

sas@my-controller@example.com /opt/sas/viya/home/SASFoundation/utilities/bin

```
UserName     Job      ID  SessId %CPU  Memory    Time  Ranks  Port Active  Pending Completed  Owned Disk Shared Disk
sas          tkemon 11047          1   638.1M     0:35    1
qstauto      cas    17859          0     1.8G     4:09    1   5570     10              3999  523.4M    160.6M
qstauto(e    cas    17859   3019   0     1.3G     1:13    1         summarize(1%)        37     0.0    698.6K
qstauto(s..  cas    17859   3710   0     1.0G     0:30    1         groupBy(100%)         6     0.0       0.0
qstauto(s..  cas    17859   3716   0     1.7G     0:30    1         valueCount(100%)      8   16.0M       0.0
qstauto(m..  cas    17859   3744   0     1.0G     0:28    1         groupBy(100%)         5     0.0       0.0
qstauto(m..  cas    17859   3747   0     1.7G     0:28    1         valueCount(100%)      8   16.0M       0.0
qstauto(m..  cas    17859   3836   0   997.5M     0:19    1         groupBy(100%)         5     0.0       0.0
qstauto(m..  cas    17859   3844   0     1.7G     0:19    1         valueCount(100%)      8   16.0M       0.0
qstauto(s..  cas    17859   3859   0     1.6G     0:18    1         echo(100%)           36     0.0       0.0
qstauto(s..  cas    17859   3870   0     1.7G     0:18    1         valueCount(100%)      8   16.0M       0.0
qstauto(s..  cas    17859   4008   0     1.9G     0:00    1         droptable(100%)      68     0.0    292.0M
ricfox       sas     3494          0     1.5G     2:47    1
ricfox       sas     3587          0     1.4G     2:47    1
sas          sas     8567          0     1.4G     0:00    1






        Tue Jun 19 12:30:08 2018
```

SAS Cloud Analytic Services: How To (SAS Environment Manager)

Introduction

These instructions explain how to view and modify SAS Cloud Analytic Services (CAS) settings using SAS Environment Manager.

Navigation

To access the Servers page from SAS Environment Manager:

1 In the applications menu (☰), under **Administration**, select **Manage Environment**.

2 In the vertical navigation bar, click 🖥.

The tasks described in this section are performed from the Servers page and most can be performed only by SAS Administrators.

View CAS Server Properties and System Information

You can view CAS server properties (such as machine name and port) and system information (such as CAS version and build date).

Note: You can also view CAS server metrics.

1 Select the CAS server whose properties and system information that you want to view.

2 On the right side of the list, click ☷.

3 To close the pane, in the top left corner of the pane, click ».

View CAS Server Configuration

You can view CAS server configuration values and identify how they were set (for example, the maximum CAS table size, the location of the permstore, the HTTP port being used, and so on).

1 Select the CAS server whose configuration you want to view.

2 Click 🔧.

3 Make sure that the CAS Configuration tab is selected.

4 To return to the Servers page, in the top left corner of the window, click ☰.

View CAS Start-up Options and Environment Variables

You can view environment variable and command-line option values used to run a CAS server.

1 Select the CAS server whose run-time environment you want to view.

2 Click 🔧.

3 Make sure that the **Nodes** tab is selected.

> **TIP** If the list of server nodes is not displayed, immediately above the list, click ☰.

4 In the list of server nodes, select the CAS controller, backup controller, or worker whose run-time environment you want to view.

5 Click ▦.

6 To return to:

- the list of CAS server nodes.

  Immediately above the list, click ☰.

- the Servers page.

  In the top left corner of the window, click ☰.

View CAS User Session Information

You can view information about a CAS server session such as the session name, session ID, connection state, and so on. If you are a member of the Superuser on page 363role and assume that role when prompted during sign-in to SAS Environment Manager, you can terminate sessions.

1 Select the CAS server whose sessions you want to view.

2 Click .

3 Make sure that the Sessions tab is selected.

> **TIP** To view additional details, add columns to the table. Click and select **Manage columns**.

4 To return to the Servers page, in the top left corner of the window, click ☰.

Terminate a CAS User Session

You can terminate all CAS sessions that you started. To terminate other users' CAS sessions, you must be a member of the Superuser on page 363role and assume that role when prompted during sign-in to SAS Environment Manager.

1 Select the CAS server whose session you want to terminate.

2 Click .

3 With the specified server highlighted, click .

4 If it is not already selected, select **Sessions**.

5 On the Sessions tab, select the check boxes for the sessions that you want to terminate, and click .

6 In the alert box that is displayed, confirm your selection by clicking **Terminate**.

7 In the top right of the window, click **Relinquish** to surrender the Superuser role for the specified server.

8 To return to the Servers page, in the top left corner of the window, click ☰.

View CAS Server Nodes

You can view basic information such as machine name, connection state, and role for all of the nodes that belong to a CAS server.

1 Select the CAS server whose nodes you want to view.

2 Click .

3 Make sure that the Nodes tab is selected.

4 To return to the Servers page, in the top left corner of the window, click ☰.

Manage CAS Server Nodes

To manage CAS server worker nodes, you must be a member of the Superuser on page 363role and assume that role when prompted during sign-in to SAS Environment Manager.

1 Select the distributed CAS server whose worker nodes you want to manage.

Important: CAS servers running in SMP mode are non-distributed and do not have worker nodes.

2 Click 👤☆.

3 With the specified server highlighted, click ✎.

4 If it is not already selected, select **Nodes** to see the worker nodes for the specified CAS server.

> **TIP** If the list of server nodes does not display, immediately above the list, click ☰.

5 To add or to remove a worker node, click ✎.

6 In the Edit Nodes window, perform one of the following actions:

- Add a worker node

  Click + and enter the fully qualified domain name for the machine of the CAS worker node that you are adding.

  Important: Adding a node is best suited for times when the system is mainly processing batch jobs, where a delay is not a concern. CAS waits for all sessions to complete any current actions before it runs the AddNode action. CAS prevents any new actions from running until the AddNode action finishes. After the node is added, existing sessions are updated to include the added node, except subset sessions. Subset sessions are not modified to use the added node.

- Remove a worker node

  Select the CAS worker node in the table, click 🗑, and, in the alert box, confirm the removal by clicking **Yes**.

  Important: Removing a node is best suited for times when the system is mainly processing batch jobs, where a delay is not a concern. The process of dropping (removing) a worker node ensures that active and backup copies of table blocks are preserved. This requires that all sessions complete their actions and pause while the blocks are moved. Long-running actions are canceled based on the values of cas.REMOVENODECANCELTIMEOUT and cas.REMOVENODEKILLTIMEOUT. Occasionally, a session might need to be killed so that the operation can proceed. The data movement often takes minutes.

7 To save any changes, click **Save**. Otherwise, click **Cancel**.

8 In the top right of the window, click **Relinquish** to surrender the Superuser role for the specified server.

9 To return to the Servers page, in the top left corner of the window, click ☰.

Manage Path Lists (Whitelists and Blacklists)

To change CAS server path list (whitelist and blacklist) settings, you must be a member of the Superuser on page 363role and assume that role when prompted during sign-in to SAS Environment Manager. For more information about how CAS manages whitelists and blacklists, see “Paths List”.

Important: CAS does not support blacklists for caslibs on Amazon S3.

1 Select the CAS server whose whitelist or blacklist you want to access.

2 Click 👤☆.

3 With the specified server highlighted, click ⚙.

4 Make sure that the **Paths List** tab is selected.

5 To modify the active list, or to switch between a whitelist, blacklist, or no list, on the right side of the Server Settings window, click ✎.

If you select the blacklist or whitelist, you can add or remove paths to the list.

Note: By default, the SAS Viya install and various configuration directories are on the blacklist.

Important: On Linux, if a blacklist path is changed to a symbolic link, then the blacklist should be updated using the fully resolved path.

6 To save any changes, click **Save**. Otherwise, click **Cancel**.

7 When you are finished, click **Close**.

8 In the top right of the window, click **Relinquish** to surrender the Superuser role for the specified server.

Adjust Caslib Management Privileges

To adjust caslib management privileges for a particular CAS server in SAS Environment Manager, you must be a member of the Superuser on page 363role and assume that role when prompted during sign-in to SAS Environment Manager.

1 Select the CAS server whose caslib management privileges you want to adjust.

2 Click 👤.

3 With the specified server highlighted, click ⚙.

4 Select **Caslib Management Privileges** to view identities and their caslib privileges.

5 To modify privileges, click ✎.

6 For the identities listed, choose to enable (or disable) the ability to add and delete session and global caslibs.

Regardless of access controls, the Superuser can add and manage all caslibs.

Note: This display shows directly granted privileges. Indirectly granted privileges and denials of privileges are not reflected in this display.

7 To save any changes, click **Save**. Otherwise, click **Cancel**.

8 When you are finished, click **Close**.

9 In the top right of the window, click **Relinquish** to surrender the Superuser role for the specified server.

Stop a CAS Server

To shut down a CAS server, you must be a member of the Superuser on page 363role and assume that role when prompted during sign-in to SAS Environment Manager.

1 Select the CAS server that you want to stop.

2 Click 👤.

3 Right-click the CAS server, and select **Stop server**.

4 In the alert box that is displayed, confirm your selection by clicking **Stop the Server**.

5 In the top right of the window, click **Relinquish** to surrender the Superuser role for the specified server.

Remove CAS Worker Software

1 To remove a CAS worker, you must be a member of the Superuser on page 363role and assume that role when prompted during sign-in to SAS Environment Manager.

2 Drop the CAS worker node whose software you want to remove.

For more information, see “Manage CAS Server Nodes”.

3 Sign in to the CAS controller machine with root-level privileges.

4 On the CAS controller machine, remove the machine name of the worker node from `/opt/sas/viya/config/etc/cas/default/cas.hosts`.

5 Edit the inventory file on the Ansible controller machine to remove the deploy target definition at the top of the file. Also remove the deploy target name from [sas-casserver-worker].

6 Sign in to the CAS worker machine with root-level privileges.

7 Run the following commands from an operating system prompt:

```
sudo yum erase "@SAS*" "@CAS*"

sudo /opt/sas/viya/home/utils/uninstall_viya.sh

sudo mv /opt/sas/viya/ /opt/sas/viya_$(date +"%m_%d_%Y")
```

SAS Cloud Analytic Services: How To (CAS Server Monitor)

View CAS Server Properties and System Information

1 Sign in to CAS Server Monitor with a valid user ID and password.

2 In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click 🖥.

3 On the System State page, make sure that **Controller** is selected.

Important: After a CAS license is renewed, the **License File** field is not updated until the CAS server is restarted.

View CAS Server Configuration

To use CAS Server Monitor to view the current list of CAS Server options and their values, follow these steps:

1 Sign in to CAS Server Monitor with a valid user ID and password.

2 In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click 🔧.

3 On the Configuration page, make sure that **CAS Configuration** is selected.

View CAS Start-up Options and Environment Variables

You can use CAS Server Monitor to view the option used when a CAS server was started and to see the current list of CAS environment variables and their values.

To view CAS start-up options and environment variable values, follow these steps:

1. Sign in to CAS Server Monitor with a valid user ID and password.
2. In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click 🖥.
3. On the System State page, select **Runtime Environment**.

View CAS User Session Information

You can use CAS Server Monitor to view information about a user's session, such as connection port, length of connection time, and so on.

To view user session information, follow these steps:

1. Sign in to CAS Server Monitor with a valid user ID and password.
2. In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click 🖥.
3. On the System State page, select **User Sessions**.

Cancel a CAS User Session

To cancel a CAS server session, follow these steps:

1. Sign in to CAS Server Monitor with a valid user ID and password.

   Note: If you are canceling another user's session, you must sign in to CAS Server Monitor with a user ID that has CAS Administrator privileges on page 362.

2. In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click 🖥.
3. On the System State page, select **User Sessions**.
4. At the end of the row for the session that you want to cancel, click ⋮ and select **Cancel Session**.

> **TIP** You can also use ⋮ to launch the Resource Monitor, cancel the CAS action, and terminate the session.

Terminate a CAS User Session

> **TIP** Terminate a session only after having tried canceling a session. Using terminate might not release resources (for example, mapped memory and memory involving database connections, and so on).

To terminate your CAS server session, follow these steps:

1. Sign in to CAS Server Monitor with a valid user ID and password.

2. In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click 🖥.

3. On the System State page, select **User Sessions**.

4. At the end of the row for the session that you want to terminate, click ⋮ and select **Terminate Session**.

   > **TIP** You can also use ⋮ to launch the Resource Monitor, cancel the CAS action, and terminate the session.

View and Manage CAS Nodes

You can use CAS Server Monitor to view, add, and remove CAS nodes in your analytics cluster.

Note: In order to add CAS nodes, the requisite software must already have been deployed on the machines that you are adding. To add new machines, deploy SAS on them first, and then you can add them using the CAS Server Monitor.

To view or manage a node, follow these steps:

1. Sign in to CAS Server Monitor with a user ID that has CAS Administrator privileges on page 362.

   Note: If you want only to view CAS server nodes, CAS Administrator privileges are not required.

2. In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click 🖥.

3. On the System State page, select **Nodes**.

4. From the **Nodes** table, you can:

   - View information about all the nodes in your analytics cluster.

   - Add nodes to your analytics cluster:

     Note: Before you can add *new* nodes to your cluster, you must have already added the CAS worker node software to the machine. For more information, see "Add New Worker Nodes or a Backup Controller".

     - Click **Add Nodes**.

     - On the Add Nodes dialog box, in **Hostname**, enter a simple host name, such as mygrid011, and click **OK**. The server monitor runs the CAS addNode action, which starts (or restarts) the node and joins it to the cluster.

       Separate multiple host names with a comma.

       If your hosts are named in numeric order (for example, host002, host003, ...) you can enter a range of host names. Use the form, **`host[`*`start-number-end-number`*`]`** (for example, `mygrid[002-030]`).

       Important: Adding a node is best suited for times when the system is mainly processing batch jobs, where a delay is not a concern. CAS waits for all sessions to complete any current actions before it runs the AddNode action. CAS prevents any new actions from running until the AddNode action finishes. After the node is added, existing sessions are updated to include the added node, except subset sessions. Subset sessions are not modified to use the added node.

   - Drop worker nodes from your analytics cluster:

     Next to the node that you want to drop, click ⋮ and select **Remove Nodes**. The server monitor runs the CAS removeNode action, which stops the node and redistributes its data to other nodes in the cluster.

     Important: Removing a node is best suited for times when the system is mainly processing batch jobs, where a delay is not a concern. The process of dropping (removing) a worker node ensures that active and backup copies of table blocks are preserved. This requires that all sessions complete their actions and pause while the blocks are moved. Long-running actions are canceled based on the values of cas.REMOVENODECANCELTIMEOUT and cas.REMOVENODEKILLTIMEOUT. Occasionally, a session might need to be killed so that the operation can proceed. The data movement often takes minutes.

- View information about processes running on a particular node:

  Next to the node that you want to view process information about, click ⋮ and select **Show Processes**.

- Stop the server immediately by sending a kill signal to the server process:

  Next to the node that you want to stop, click ⋮ and select **Terminate Server Instance**. The server monitor issues a command to kill the node process.

  Note: Terminate a server instance only after having tried removing a node. Using terminate might not release resources (for example, mapped memory and memory involving database connections, and so on).

Adjust Caslib Management Privileges

To enable non-administrators to add global caslibs:

1. Sign in to CAS Server Monitor with a valid user ID and password that has administrator privileges.
2. In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click 🔧.
3. On the Configuration page, select **Access Controls**.
4. In the **Caslibs** list, select **Global Caslib Creation**.

   > **TIP** If the **Global Caslib Creation** caslib is not listed, you are not signed in as an administrator.

5. In the upper right, click **Edit**.
6. In the **Edit Access Controls** window, adjust values as needed.

Intent	Instructions
Enable all users to add global caslibs.	In the existing row for **Authenticated Users**, select the **Grant** radio button.
Enable a group to add global caslibs.	Click **Add Row**. Select **Group**, enter the group name, and select the **Grant** radio button.
Enable an individual user to add global caslibs.	Click **Add Row**. Select **User**, enter the user name, and select the **Grant** radio button.

7. Click **OK** to save your changes.
8. Under **Access Controls**, review the results of your changes.
9. Verify that users who should be able to add global caslibs can do so.

Here are details:

- User and group names that you enter are not validated.
- Regardless of access controls, administrators can add and manage all caslibs.

- For the special caslibs (**Global Caslib Creation** and **Session Caslib Creation**), the only available value in the **Activity** column is **Manage Access**. The special caslibs are protected by role requirements, not by the ManageAccess permission. Granting or denying the ManageAccess permission on the special caslibs affects only the ability of non-administrators to manage other caslibs.
- If you want to restrict the ability to manage session caslibs, select **Session Caslib Creation** in the **Caslibs** list. Add direct denials as needed.

Stop a CAS Server

You can use CAS Server Monitor to shut down the CAS server. For a distributed CAS server, in addition to the controller, clicking the **shutdown** button stops all the CAS worker nodes and the backup controller (if present).

1. Sign in to CAS Server Monitor with a user ID that has CAS Administrator privileges on page 362.
2. In CAS Server Monitor, beneath the **Cloud Analytic Services** banner, click .
3. On the System State page, make sure that **Controller** is selected.
4. At the top, on the right side of the page, click **Shutdown**.

Remove CAS Worker Software

1. Sign in to CAS Server Monitor with a user ID that has CAS Administrator privileges on page 362.
2. Drop the CAS worker node whose software you want to remove.

   For more information, see “View and Manage CAS Nodes” on page 488.
3. Sign in to the CAS controller machine with root-level privileges.
4. On the CAS controller machine, remove the machine name of the worker node from `/opt/sas/viya/config/etc/cas/default/cas.hosts`.
5. Edit the inventory file on the Ansible controller machine to remove the deploy target definition at the top of the file. Also remove the deploy target name from [sas-casserver-worker].
6. Sign in to the CAS worker machine with root-level privileges.
7. Run the following commands from an operating system prompt:

```
sudo yum erase "@SAS*" "@CAS*"
```

```
sudo /opt/sas/viya/home/utils/uninstall_viya.sh
```

```
sudo mv /opt/sas/viya/ /opt/sas/viya_$(date +"%m_%d_%Y")
```

SAS Cloud Analytic Services: How To (gridmon.sh)

Overview

Note: gridmon.sh is supported only on Linux platforms.

gridmon.sh is a console or terminal application that can be run from a Linux terminal or a terminal emulator like PuTTY. gridmon.sh displays data streamed from all the machines on your CAS server showing information about jobs, individual machines on the server, and attached disks.

gridmon.sh enables you to perform several limited actions, such as killing a job, killing a rank, or running gstack. (For a complete list of functionality, see "gridmon.sh Commands".) If an X Server resides on the CAS controller, then you can launch an Xterm, a Perf Top, or an Attach Debugger session directly from gridmon.sh.

Note: Attach Debugger is for use only when directed by SAS Technical Support or by SAS R&D.

If you run gridmon in record mode, gridmon captures this streamed data. Using the playback feature, you can investigate the state of your CAS server while it was recorded.

Use gridmon.sh (Linux)

1 Log on to the CAS controller machine as a user with passwordless ssh access to all CAS nodes. The user also needs sudo privileges on all CAS nodes to run Grid Monitor commands that require root access, such as viewing process limits and killing jobs.

2 To start gridmon.sh, run the following command:

```
/opt/sas/viya/home/SASFoundation/utilities/bin/gridmon.sh
```

3 By default, gridmon.sh runs in job mode.

Figure A.5 *gridmon.sh Running in Job Mode*

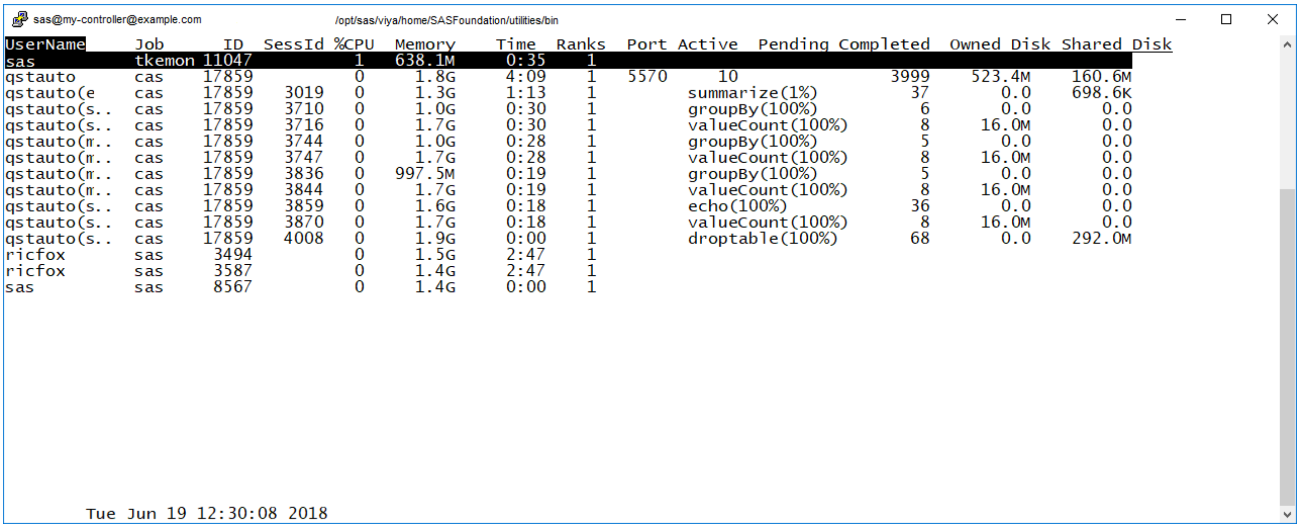

sas@my-controller@example.com
/opt/sas/viya/home/SASFoundation/utilities/bin

```
UserName     Job      ID SessId %CPU  Memory    Time  Ranks  Port Active  Pending Completed  Owned Disk Shared Disk
sas          tkemon 11047          1   638.1M     0:35    1
qstauto      cas    17859          0     1.8G     4:09    1  5570     10                3999    523.4M    160.6M
qstauto(e    cas    17859   3019   0     1.3G     1:13    1        summarize(1%)          37       0.0    698.6K
qstauto(s..  cas    17859   3710   0     1.0G     0:30    1        groupBy(100%)           6       0.0       0.0
qstauto(s..  cas    17859   3716   0     1.7G     0:30    1        valueCount(100%)        8     16.0M       0.0
qstauto(m..  cas    17859   3744   0     1.0G     0:28    1        groupBy(100%)           5       0.0       0.0
qstauto(m..  cas    17859   3747   0     1.7G     0:28    1        valueCount(100%)        8     16.0M       0.0
qstauto(m..  cas    17859   3836   0   997.5M     0:19    1        groupBy(100%)           5       0.0       0.0
qstauto(m..  cas    17859   3844   0     1.7G     0:19    1        valueCount(100%)        8     16.0M       0.0
qstauto(s..  cas    17859   3859   0     1.6G     0:18    1        echo(100%)             36       0.0       0.0
qstauto(s..  cas    17859   3870   0     1.7G     0:18    1        valueCount(100%)        8     16.0M       0.0
qstauto(s..  cas    17859   4008   0     1.9G     0:00    1        droptable(100%)        68       0.0    292.0M
ricfox       sas     3494          0     1.5G     2:47    1
ricfox       sas     3587          0     1.4G     2:47    1
sas          sas     8567          0     1.4G     0:00    1




        Tue Jun 19 12:30:08 2018
```

Note: **Shared Disk** consists of the sum of HDFSSize, DNFSSize, and Global FSSize across all ranks. **Owned Disk** is the sum of disk space in CAS_DISK_CACHE across all CAS worker nodes. For more information, see Table 27.89.

4 To run in machine mode, enter `m`.

Figure A.6 gridmon.sh Running in Machine Mode

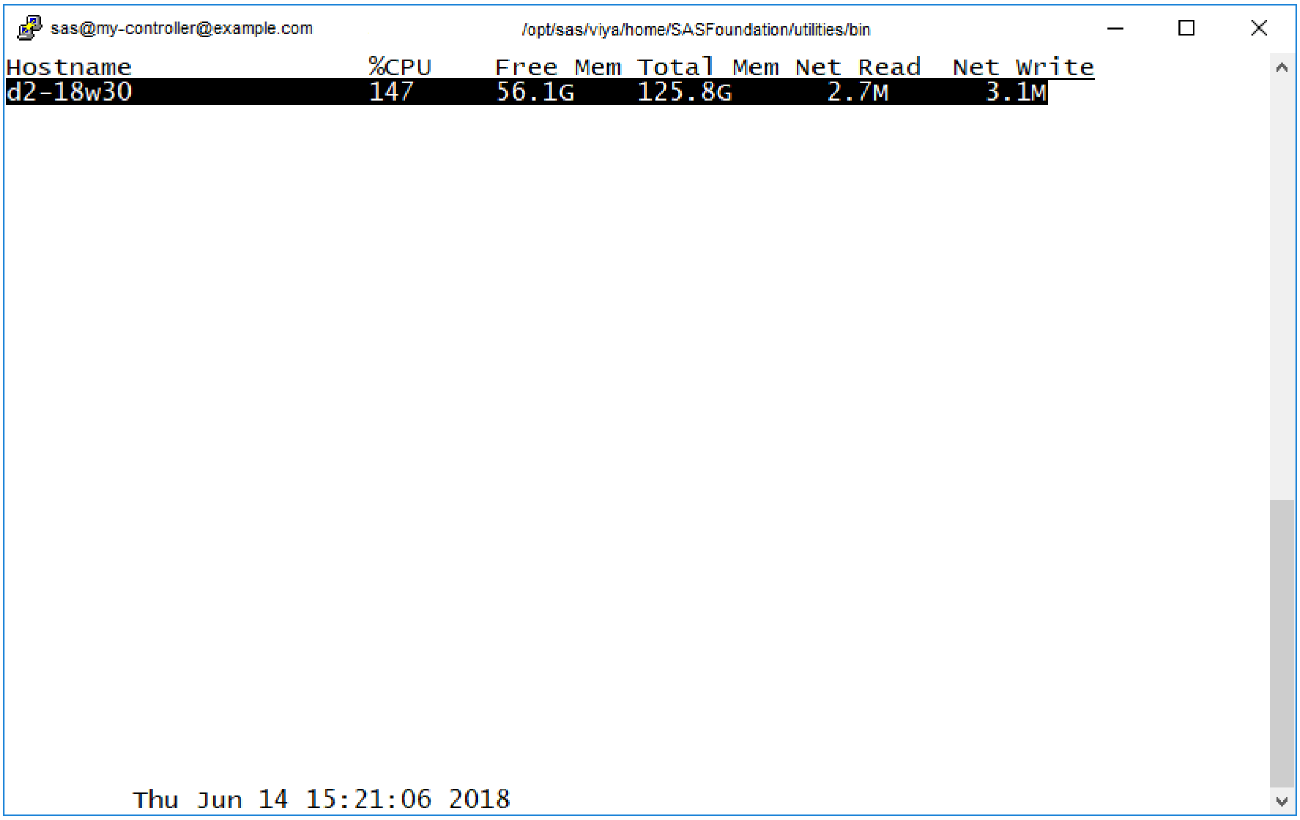
sas@my-controller@example.com
/opt/sas/viya/home/SASFoundation/utilities/bin

```
Hostname                   %CPU     Free Mem Total Mem Net Read  Net Write
d2-18w30                   147      56.1G    125.8G        2.7M        3.1M









         Thu Jun 14 15:21:06 2018
```

5 To run in disk mode, enter d.

Figure A.7 gridmon.sh Running in Disk Mode

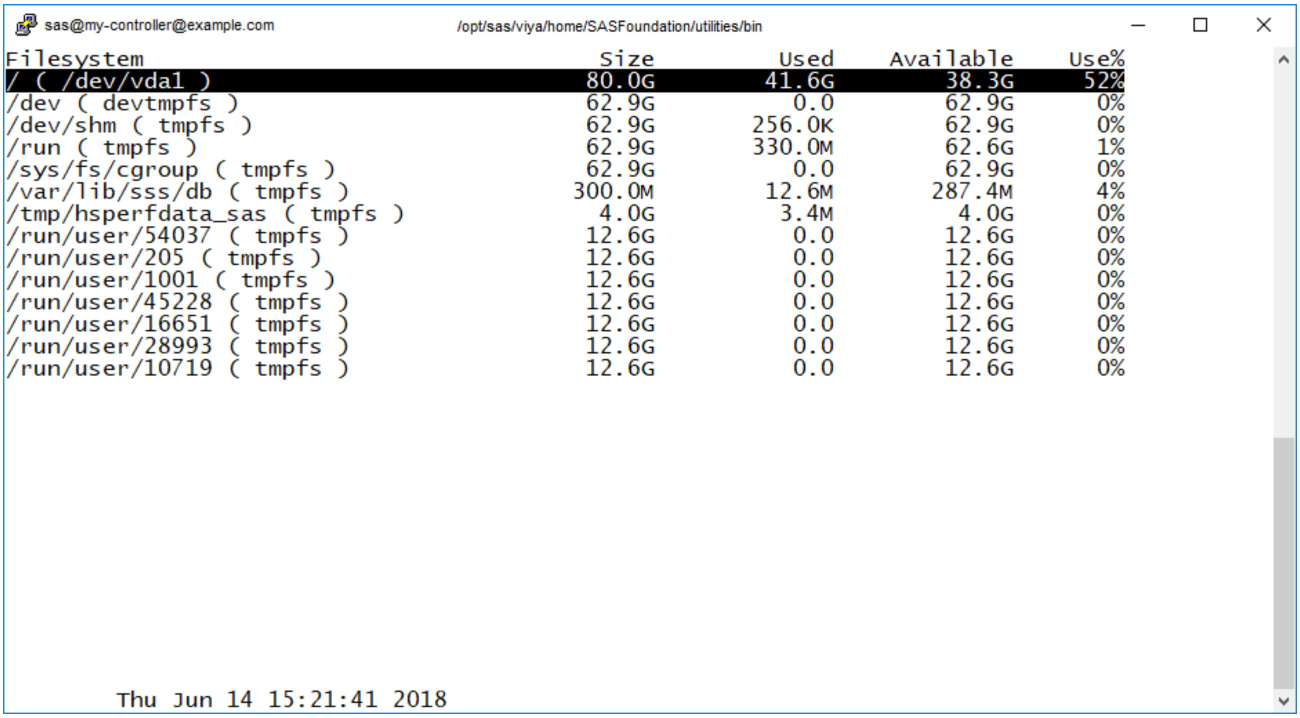
sas@my-controller@example.com
/opt/sas/viya/home/SASFoundation/utilities/bin

```
Filesystem                                  Size        Used    Available   Use%
/ ( /dev/vda1 )                            80.0G       41.6G        38.3G    52%
/dev ( devtmpfs )                          62.9G         0.0        62.9G     0%
/dev/shm ( tmpfs )                         62.9G      256.0K        62.9G     0%
/run ( tmpfs )                             62.9G      330.0M        62.6G     1%
/sys/fs/cgroup ( tmpfs )                   62.9G         0.0        62.9G     0%
/var/lib/sss/db ( tmpfs )                 300.0M       12.6M       287.4M     4%
/tmp/hsperfdata_sas ( tmpfs )               4.0G        3.4M         4.0G     0%
/run/user/54037 ( tmpfs )                  12.6G         0.0        12.6G     0%
/run/user/205 ( tmpfs )                    12.6G         0.0        12.6G     0%
/run/user/1001 ( tmpfs )                   12.6G         0.0        12.6G     0%
/run/user/45228 ( tmpfs )                  12.6G         0.0        12.6G     0%
/run/user/16651 ( tmpfs )                  12.6G         0.0        12.6G     0%
/run/user/28993 ( tmpfs )                  12.6G         0.0        12.6G     0%
/run/user/10719 ( tmpfs )                  12.6G         0.0        12.6G     0%





        Thu Jun 14 15:21:41 2018
```

6 Refer to “gridmon.sh Commands” for the commands that you can use in each mode.

7 In job mode there are two menus.

Run in job mode (enter `j`), select a job, and press **`Enter`**.

The **Show Ranks** menu is displayed:

Figure A.8 Show Ranks Menu

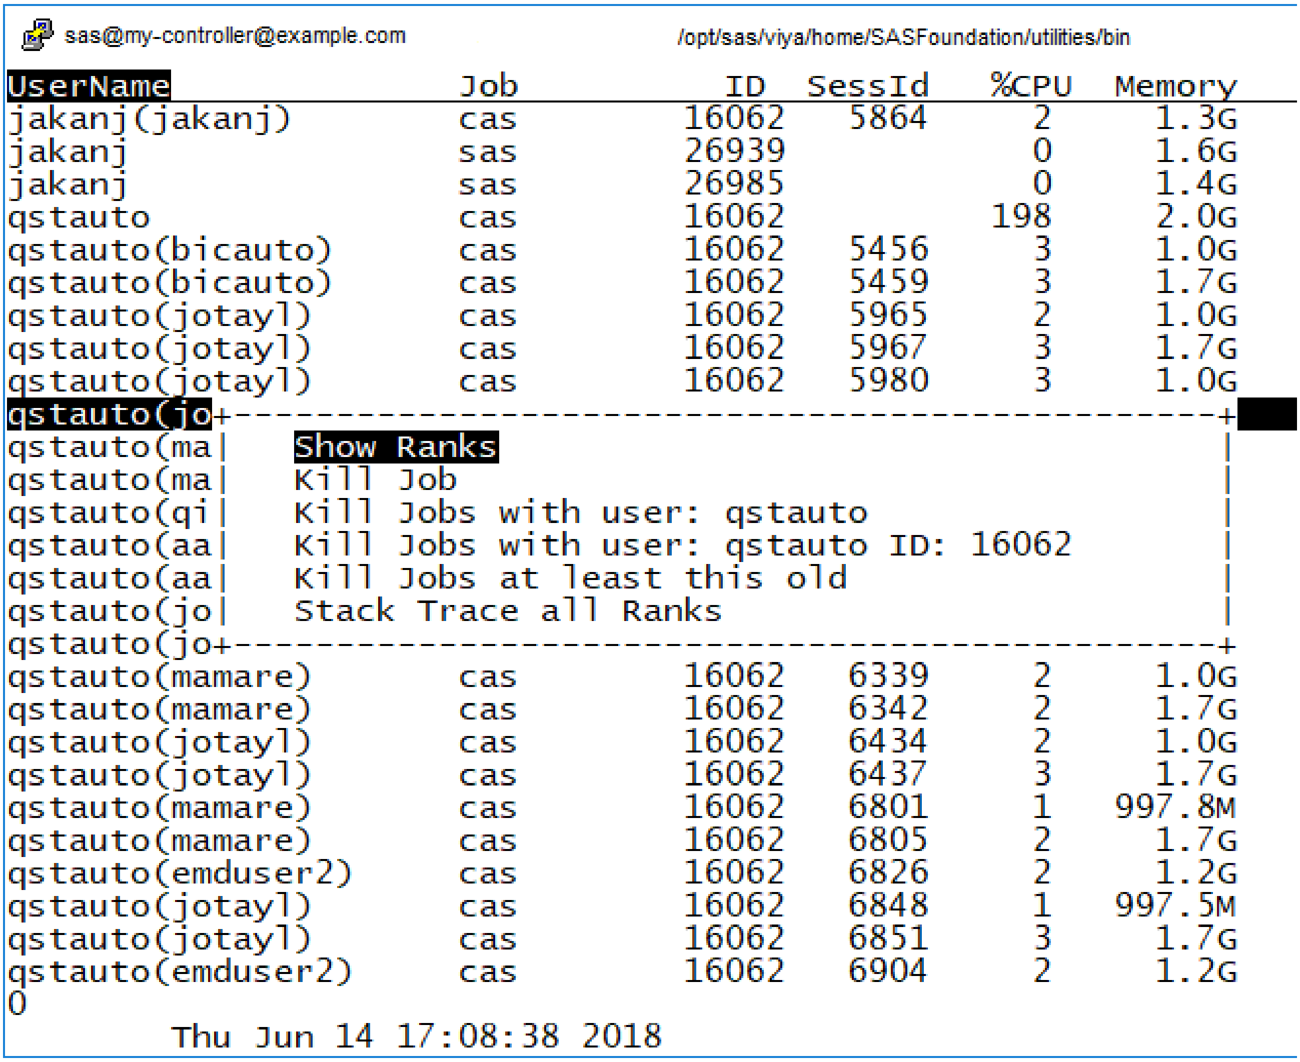

```
sas@my-controller@example.com                /opt/sas/viya/home/SASFoundation/utilities/bin

UserName              Job          ID  SessId   %CPU  Memory
jakanj(jakanj)        cas        16062   5864     2     1.3G
jakanj                sas        26939            0     1.6G
jakanj                sas        26985            0     1.4G
qstauto               cas        16062          198     2.0G
qstauto(bicauto)      cas        16062   5456     3     1.0G
qstauto(bicauto)      cas        16062   5459     3     1.7G
qstauto(jotayl)       cas        16062   5965     2     1.0G
qstauto(jotayl)       cas        16062   5967     3     1.7G
qstauto(jotayl)       cas        16062   5980     3     1.0G
qstauto(jo+------------------------------------------------+
qstauto(ma|   Show Ranks                                   |
qstauto(ma|   Kill Job                                     |
qstauto(qi|   Kill Jobs with user: qstauto                 |
qstauto(aa|   Kill Jobs with user: qstauto ID: 16062       |
qstauto(aa|   Kill Jobs at least this old                  |
qstauto(jo|   Stack Trace all Ranks                        |
qstauto(jo+------------------------------------------------+
qstauto(mamare)       cas        16062   6339     2     1.0G
qstauto(mamare)       cas        16062   6342     2     1.7G
qstauto(jotayl)       cas        16062   6434     2     1.0G
qstauto(jotayl)       cas        16062   6437     3     1.7G
qstauto(mamare)       cas        16062   6801     1   997.8M
qstauto(mamare)       cas        16062   6805     2     1.7G
qstauto(emduser2)     cas        16062   6826     2     1.2G
qstauto(jotayl)       cas        16062   6848     1   997.5M
qstauto(jotayl)       cas        16062   6851     3     1.7G
qstauto(emduser2)     cas        16062   6904     2     1.2G
0
        Thu Jun 14 17:08:38 2018
```

For specific information about each **Show Ranks** menu command, see“Show Ranks Menu Commands ”.

8 From the **Show Ranks** menu, select **Show Ranks** and the Ranks window is displayed. Press **Enter** and the **Show Details** menu is displayed.

Figure A.9 Show Details Menu

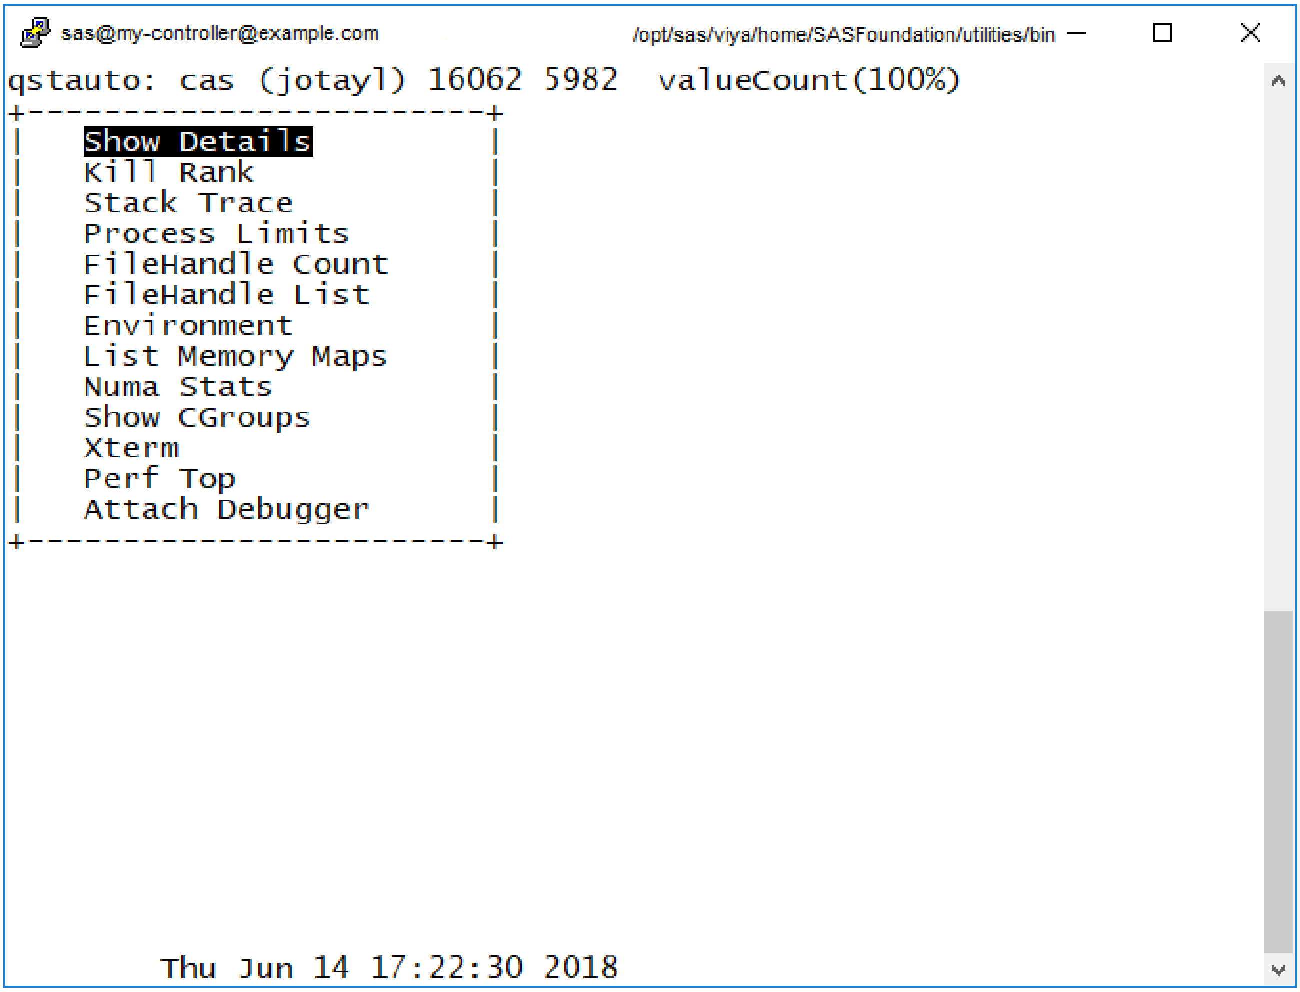

sas@my-controller@example.com
/opt/sas/viya/home/SASFoundation/utilities/bin

```
qstauto: cas (jotayl) 16062 5982  valueCount(100%)
+-------------------------+
|    Show Details         |
|    Kill Rank            |
|    Stack Trace          |
|    Process Limits       |
|    FileHandle Count     |
|    FileHandle List      |
|    Environment          |
|    List Memory Maps     |
|    Numa Stats           |
|    Show CGroups         |
|    Xterm                |
|    Perf Top             |
|    Attach Debugger      |
+-------------------------+









        Thu Jun 14 17:22:30 2018
```

For specific information about each **Show Details** menu command, see“Show Details Menu Commands ”.

9 Press `Esc` to leave the **Show Details** menu.

10 In machine mode there is one menu.

Run in machine mode (enter `m`), select a machine and press `Enter`.

The **Details** menu is displayed:

Figure A.10 Details Menu

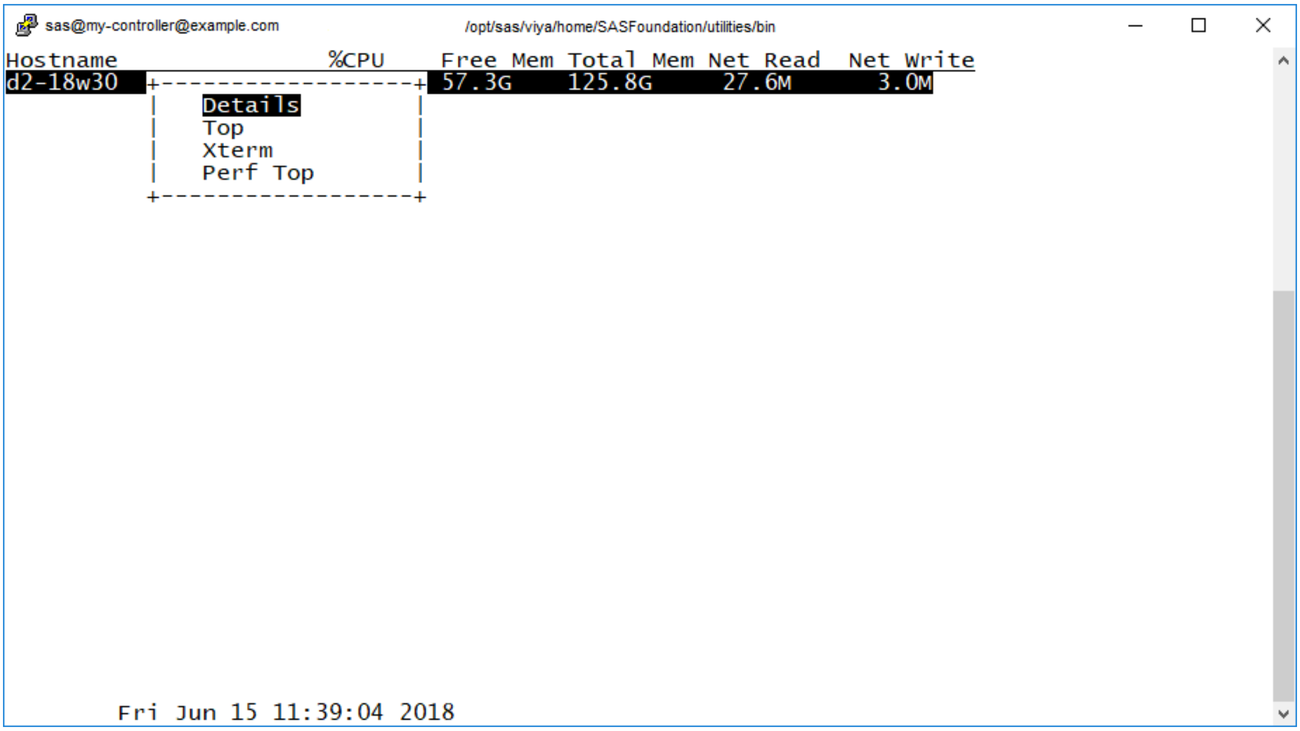
sas@my-controller@example.com
/opt/sas/viya/home/SASFoundation/utilities/bin

```
Hostname                      %CPU     Free Mem Total Mem Net Read  Net Write
d2-18w30  +------------------+ 57.3G    125.8G       27.6M      3.0M
           |     Details      |
           |     Top          |
           |     Xterm        |
           |     Perf Top     |
          +------------------+

          Fri Jun 15 11:39:04 2018
```

For specific information about each **Details** menu command, see “Details Menu Commands ”.

11 Enter `q` to exit gridmon.sh.

Run gridmon.sh in Record Mode (Linux)

You can run gridmon.sh in record mode in order to capture data that is streamed from each machine on your CAS server at approximately one second intervals. You can review this captured data later by running gridmon.sh in playback mode.

1 Log on to the CAS controller machine as a user with passwordless ssh access to all CAS nodes. The user also needs sudo privileges on all CAS nodes to run Grid Monitor commands that require root access, such as viewing process limits and killing jobs.

2 Change to the following directory:

```
cd /opt/sas/viya/home/SASFoundation/utilities/bin/
```

3 To start gridmon.sh in record mode, run the following command:

```
./gridmon.sh -record path/output-filename
```

where ***path/output-filename*** is the absolute path and filename for where gridmon.sh writes its output.

Here is an example:

```
./gridmon.sh -record /my_data/tkgridmon_output
```

Run gridmon.sh in Playback Mode (Linux)

You can run gridmon.sh in playback mode to review data streamed from all the machines on your CAS server that you captured earlier while running gridmon.sh in record mode.

1 Log on to the CAS controller machine as a user with passwordless ssh access to all CAS nodes. The user also needs sudo privileges on all CAS nodes to run Grid Monitor commands that require root access, such as viewing process limits and killing jobs.

2 To start gridmon.sh in playback mode, run the following command:

```
./gridmon.sh -playback path/output-filename
```

Here is an example:

```
./gridmon.sh -gridhost -playback /my_data/tkgridmon_output
```

SAS Cloud Analytic Services: Concepts

CAS Controller

Controller is one of three roles that can be assigned to a machine for SAS Cloud Analytic Services (CAS): controller, backup controller, and worker. For both server architectures—distributed and single-machine—one machine is assigned the controller role. When the server starts, the controller process is started. This process is sometimes referred to as the server controller. The controller accepts connections from clients.

Single-machine CAS Server

The single-machine architecture uses symmetric multiprocessing (SMP). The functionality for a single-machine server is nearly identical to MPP, except that there is no cluster communication. In this architecture, the server acts as a controller. Before a client connects, the server listens on a port for connections.

After a client connects, a session is created and the session connects back to the client. (This is identical to the method that is performed by a CAS server that uses MPP. Compare with “Distributed CAS Server”.)

Figure A.11 Single-machine CAS Server

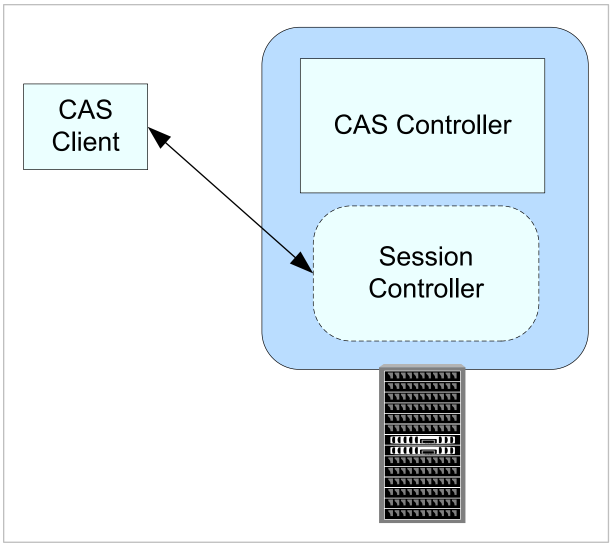

CAS Backup Controller

A SAS Cloud Analytic Services (CAS) backup controller (sometimes referred to as *secondary controller*) provides fault tolerance for the CAS controller. A backup controller is used only in a distributed server architecture. Deploying a backup controller is optional. CAS supports one backup controller only.

Note: A requirement for operating CAS with a backup controller is that it and the CAS controller (the primary controller) must both use the same shared file system. For more information, see “Set Up a Shared File System for CAS Controllers (Post-Deployment)”.

When CAS starts, the backup controller process is also started. In the event that the controller experiences a disruption (such as a loss of network connectivity, disk full scenarios, and so on) the backup controller enables the CAS server to continue running. When the backup controller takes control of client communication, the transfer is seamless. For more information, see Architecture in *SAS Cloud Analytic Services: Fundamentals*.

CAS Workers

When a server is running in massively parallel processing (MPP) mode, in addition to a controller, the server also has multiple machines that are assigned the worker role.

The controller parses out work to each worker node. Each worker node sends the results of its computations back to the controller.

Distributed CAS Server

CAS can be co-located with Hadoop on a cluster of machines. This massively parallel processing (MPP) architecture is appropriate for analyzing large data sets. Analysis proceeds on tables that are already made available to the server (loaded) or on tables that are gathered or created by the server on demand. A distributed CAS server consists of a controller, at least one worker, and a backup controller running on a minimum of two machines. (If a backup controller is deployed, then the minimum number of machines is three.)

Figure A.12 Distributed CAS Server

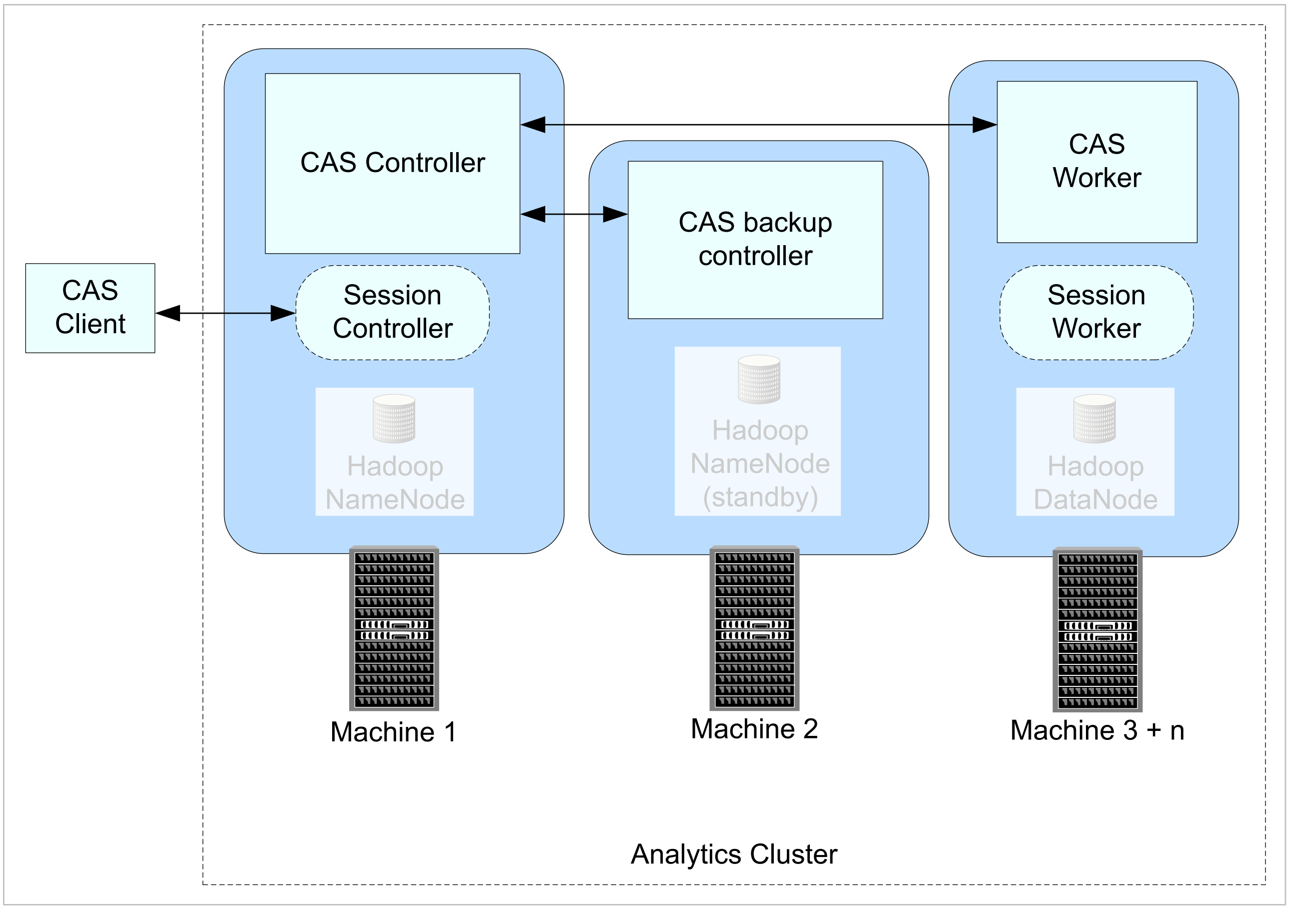

Multiple CAS Servers

Overview

As of SAS Viya 3.4, it is now possible to have multiple instances of CAS servers within a single instance of SAS Viya.

What constitutes a CAS server depends on the type of CAS environment that you are running:

- In a symmetric multiprocessing (SMP) environment, a CAS server consists of a controller and runs on a single machine.

Figure A.13 *Multiple Single-Machine CAS Servers (SMP Mode)*

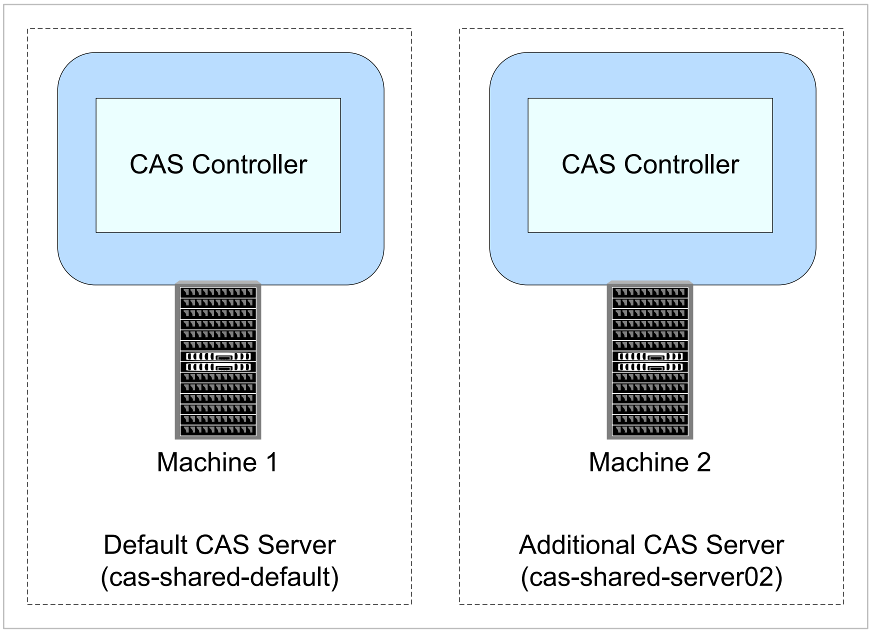

- In a massively parallel processing (MPP) environment, a distributed CAS server consists of one controller, one or more workers, and one backup controller (optional) each running on a separate machine.

Figure A.14 *Multiple Distributed CAS Servers (MPP Mode)*

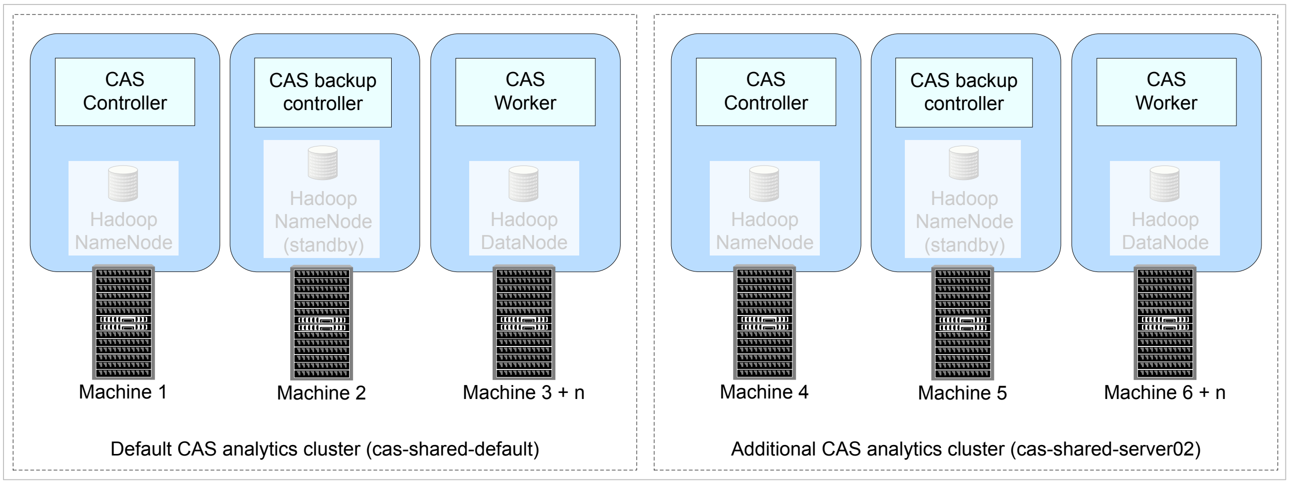

Requirements

- You can add a CAS server to a machine that does not already host existing SAS Viya software.
- CAS servers that are added to a SAS Viya deployment cannot be removed, without removing your entire SAS Viya deployment.
- A multitenant environment does not support multiple CAS servers per tenant. (Each tenant has exclusive access to a single CAS Server.)
- In SAS Viya environments that have more than one CAS server, the default CAS server (typically, cas-shared-default) must be running. The default CAS server needs to be running even if a customer is using another CAS server, so that certain caslibs such as AppData, ReferenceData, and SystemData can be accessed from the default CAS server. (These caslibs contain data needed by applications such as SAS

Visual Analytics, which depends on AppData for map data.) The default CAS server is defined in the `sas.casmanagement.global.casServer` configuration property.

- Make sure that you are licensed for the additional CAS servers that you are planning to add.

  When properly set, the CAS server option cas.MAXCORES specifies the limit for the total number of physical cores that are available to a CAS server. For more information, see cas.MAXCORES.

- Your SAS Viya deployment must be a Linux deployment.

  Multiple CAS servers are not supported on Windows environments.

- You are limited to one CAS controller or one CAS backup controller per machine.

- If you are adding a distributed CAS server that contains a backup controller, make sure that the backup controller and the CAS controller (the primary controller) both use the same shared file system.

  For more information, see "Set Up a Shared File System for CAS Controllers (Post-Deployment)" on page 476.

- Every machine on which you are installing a CAS server must contain the CAS user account (cas) and group (sas).

  For more information, see "Set Up the cas Account" in *SAS Viya for Linux: Deployment Guide*.

- For programming-only deployments and visual deployments that use the `CASHostAccountRequired` custom group, there is an additional requirement for users' home directories. In these two cases, a user's Casuser caslib is mapped to **`~/casuser`**. Therefore, the home directories (`$HOME`) for all CAS users must be shared so that they can be accessed from both the controller and the backup controller machines. Sharing users' home directories ensures that the path for the `CASUSER` library is available during CAS session start-up.

  For most other CAS session scenarios, the `CASUSER` library is set to a path in the shared file system described in "Set Up a Shared File System for CAS Controllers (Post-Deployment)".

See Also

"Add a CAS Server" on page 477

Session Processes

When a user connects to the server with a client, the server starts a session process for the user. Afterward, the client communicates with the session process.

A server running in symmetric multiprocessing mode (SMP mode) consists of a controller only, and the server starts a session controller process only. It is the session controller process that operates on rows of data.

In a distributed server (MPP mode), a session process is created on each machine in the cluster. These processes are sometimes referred to as the session controller and session worker processes.

Even though the sessions have their own operating system processes, the server processes must continue to run. When the server process terminates, the session processes also terminate.

Paths List

From a CAS server, all access to file system paths (host and HDFS directories) is through caslibs. To limit the paths that are available to non-administrators when they create or edit a caslib, use one of the following approaches:

- Create a blacklist of paths that should not be available.
- Create a whitelist of paths that should be available.

You can view and modify the lists for CAS using:

- SAS Environment Manager.

  See, “Manage Path Lists (Whitelists and Blacklists)”.

- or, the programming interfaces.

  See, Access Control Action Set: Syntax in *SAS Viya System Programming Guide*.

Here are key points:

- Paths must be absolute.
- Paths must be unique.

  CAS automatically removes any duplicate paths.

- On Linux, if a blacklist path is changed to a symbolic link, then the blacklist should be updated using the fully resolved path.
- All subdirectories of each specified path are affected.
- Paths list constraints do not affect access to existing caslibs.
- If you do not define a blacklist or whitelist, no paths list constraints are in effect.
- Paths list constraints do not apply to users who assume the Superuser role or the Data role.
- Only users who assume the Superuser role for a server can see and manage that server’s paths list.

Note: Access to third-party databases is not affected by a server’s blacklist or whitelist.

See Also

“Lock Down SAS Workspace Servers” on page 573

Caslib Management Privileges

Table A.1 Caslib Management Privileges

Task	Who Can Perform the Task[*]
Add global caslibs.	Superusers and Data administrators. Users who have global caslib management privileges.
Add session caslibs.	Superusers and Data administrators. Users who have session caslib management privileges.
Delete global caslibs.	Superusers and Data administrators. Users who have global caslib management privileges can delete any global caslib for which they have the ReadInfo and ManageAccess permissions.
Delete session caslibs.	Superusers and Data administrators. Users who have session caslib management privileges can delete any session caslib for which they have the ReadInfo and ManageAccess permissions.
Adjust caslib management privileges.	Superusers and Data administrators.

[*] Global caslib management privileges correspond to the ManageAccess permission on the _GLOBAL caslib. Session caslib management privileges correspond to the ManageAccess permission on the _SESSION caslib.

Note: Data administrators are displayed in CAS Server Monitor only.

See Also

- Adjust Caslib Management Privileges using SAS Environment Manager
- Adjust Caslib Management Privileges using CAS Server Monitor on page 489

Two Playbooks for Adding Worker Nodes

Adding worker nodes to your CAS server is accomplished using the third-party orchestration tool, Ansible. When run on a CAS server that already has workers (MPP mode), or on a CAS server that does not have workers (SMP mode), Ansible performs these steps:

- configures SSH for the new machines and all existing machines that comprise the CAS server.
- installs software on all machines listed in the `sas-casserver-worker` group contained in the Ansible inventory file.

There are two playbooks to add nodes. Each playbook offers a different set of CAS usage characteristics:

- site playbook (site.yml)

  Designed for when you want to:

  - permanently add workers and have a maintenance window.

    Added workers are automatically joined to the CAS server.

  - change your CAS server configuration (for settings other than adding workers).

  Note: Ansible restarts the CAS server when site.yml is used.

- deploy-casworker playbook (deploy-casworker.yml)

  Designed for when you want to:

  - temporarily add workers that persist only until the CAS server restarts or until the worker is dropped manually using the CAS Server Monitor **Remove Node** command.

    Added workers must be joined manually to the CAS server using the CAS Server Monitor.

  - permanently add workers, but do not have a maintenance window when the CAS server can be restarted.

    On first use, added workers must be joined manually to the CAS server using the CAS Server Monitor. On subsequent invocations, the added workers are automatically joined to the CAS server.

  Note: Ansible does not restart the CAS server when deploy-casworker.yml is used.

  Note: When you add worker nodes to a CAS controller running in SMP mode, use the site playbook. The CAS server requires a restart when moving from SMP to MPP mode. The site playbook restarts the server. The deploy-casworker playbook does not restart the CAS server.

  CAUTION! Use the deploy-casworker playbook for adding worker nodes only. Do not change other CAS server configuration settings using the deploy-casworker playbook. Doing so can cause a mismatch between configuration in memory versus configuration on disk.

Understanding Configuration Files and Start-up Files

Several SAS applications require application-specific configuration and start-up before SAS Cloud Analytic Services begins processing client requests. It is worthwhile to understand how the files are processed so that you can use the same technique to customize your server deployment.

Configuration Home Directory

The installation and deployment software creates a configuration home directory for each server instance.

Here is an example:

```
/opt/sas/viya/config/etc/cas/default
```

The final directory in the path, *default*, is the deployment instance for the server.

Standard Configuration Files

The configuration home directory includes several files with standard names. The server automatically processes these files when the standard names are used.

The following table describes the purpose and use for each of the standard files.

Table A.2 *Server Configuration Files*

Standard Filename	Description
casconfig.lua	This file contains most of the configuration settings for the server instance, such as the network port that the server listens on. During deployment, RPM owns the casconfig.lua file and during updates can override any user configuration. For more information about the settings in the file, see "Configuration File Options" on page 512.
casconfig_deployment.lua	This file contains CAS configuration settings that are created during deployment by Ansible from vars.yml. During updates, user configuration settings are overwritten.
casconfig_usermods.lua	This file contains modifications made by the SAS administrator. Using casconfig_usermods.lua ensures that your modifications are not overwritten when you upgrade CAS.
conf.d/	This directory can contain one or more configuration files that are similar to the casconfig.lua file. The files are processed in alphabetical order. The files in this directory are processed before the casconfig.lua file.
node.lua	This file contains host-specific configuration. One possible use is for security setup that relies on the host name.
node_usermods.lua	This file contains modifications that are made by the SAS administrator. Using node_usermods.lua ensures that your modifications are not overwritten when you upgrade CAS.
logconfig.xml	This file contains the SAS logging facility instructions that control server logging.
perms.xml	This file contains the initial permission settings. This file is not used after the first time the server is started and the permstore is populated.
cas.hosts	This file contains the initial set of host names and roles (controller or worker) for the server.
cas.settings*	This file contains CAS and system environment variables that are created during deployment by Ansible from vars.yml. During updates, user configuration settings can be overwritten.

Standard Filename	Description
cas_usermods.settings	This file contains modifications that are made by the SAS administrator. Using cas_usermods.settings ensures that your modifications are not overwritten when you upgrade CAS.

* There is a global version of cas.settings that resides in /opt/sas/viya/home/SASFoundation. CAS processes the global version of cas.settings **before** processing the configuration-specific version of cas.settings.

When the server starts, the configuration files that are described in the preceding table are processed. After the configuration is complete, the server runs start-up scripts.

The following table describes the standard names for the start-up files in the configuration home directory. The start-up scripts run before the server accepts any client connections. This is also referred to as *session-zero processing*.

Table A.3 *Server Start-up Files (Session 0 Processing)*

Standard Filename	Description
casstartup.lua	This file contains the actions to run as the CAS server starts, such as some addFmtLib actions and a setServOpt action, that are created during deployment by Ansible from vars.yml. CAS processes casstartup.lua before any of the other start-up files residing in the **`start.d/`** directory.

casstartup_usermods.lua	This file contains modifications that are made by the SAS administrator to casstartup.lua, such as adding global-scope caslibs and loading global-scope tables. Using casstartup_usermods.lua ensures that your modifications are not overwritten when you upgrade CAS.
	CAS processes casstartup_usermods.lua before any of the other start-up files residing in the **`start.d/`** directory.
	Use Lua syntax such as the following. Do not forget to use global scope.
	`s:table_addCaslib{caslib="worldbank",`
	`     dataSource={srcType="path"},`
	`     path="/rdstore/data/smp/world_bank",`
	`     session=false}`

start.d/	This directory contains one or more start-up scripts that are similar to the casstartup.lua file. The files are processed in alphabetical order. The files in this directory are processed after the casstartup.lua file **Important:** CAS processes only files with a .lua file extension as start-up files.

Fault Tolerance

The distributed CAS server has a communication layer that supports fault tolerance. A distributed server can continue processing requests even after losing connectivity to some nodes. The communication layer also enables you to remove or add worker nodes from a server while it is running.

If a deployment uses a backup controller, the backup controller is started along with the rest of the nodes in the CAS server. The backup controller continuously tracks the state of the server, and is current if it has to take over control. The binary protocol is the socket communication between CAS client and corresponding server session that is typically provided using port 5570. Clients like SAS, Python, SWAT, and the CASClient Java library (a JAR) normally communicate using this interface.

If the primary controller fails, the site operates without fault tolerance for the controller until a planned outage. During the planned outage, the site should recover the failed controller and return to a redundant state.

After the outage, the primary controller accepts connections from clients and the backup (or secondary) controller resumes its role of providing fault tolerance.

If the backup controller fails or is shut down while the primary controller continues to operate, the site can continue to operate without fault tolerance.

Note: For information about fault tolerance in other parts of SAS Viya, see "Fault Tolerance in SAS Viya (Linux)" on page 465.

Operating system tuning such as ulimits should be identical on both controller hosts.

For sites that co-locate the server with Hadoop:

- You can set the primary controller and backup controller to use the same hosts as the active NameNodes and the standby NameNodes. This is not a requirement.
- The HADOOP_NAMENODE environment variable can include the two host names. You can specify the active and the standby NameNodes hosts, separated by a colon.

Access controls and caslib information are stored in a directory that is known as a permstore. While the primary controller and the backup controller are running, the permstore for the backup controller is kept in sync with the primary controller. After a failover, the permstore for the backup controller becomes the most current. Part of the task of restoring the primary controller is to copy the files from the permstore on the backup controller to the permstore on the primary controller. For more information, see *SAS Viya Administration: Backup and Restore*.

See Also

- "Set Up a Shared File System for CAS Controllers (Post-Deployment)"
- "Recover a Failed Controller"

CAS Resource Management

Overview

Important: CAS resource management use of CPU shares from SAS Configuration Server (Consul) applies only on Linux. But CAS resource management use of quotas is applicable on both Linux and Windows.

On Linux platforms, CAS has resource management capabilities that enable administrators to control CAS table size and CPU consumption. CAS relies on the Linux kernel feature, cgroups, to provide the CPU and memory consumption control. You must implement cgroups before you can fully use CAS resource management. For more information, see "(Optional) Additional Requirements for CAS Resource Management" in *SAS Viya for Linux: Deployment Guide*.

You can implement resource management broadly through the use of the cas.MEMORYSIZE and cas.CPUSHARES server configuration options, or more specifically through policies that you create with the SAS Viya command line interface. CAS resource management is achieved on the CAS server, with a policy or server option that applies to a specific server. If you have additional CAS servers on which you want to manage resources, then you must create a unique policy or server option definition for each. No steps need to be performed on the client side.

TIP CAS Java clients already handle all errors that are returned by the server—including those due to a lack of resources. However, Java clients can test specifically for CAS table quota failures by checking these two status code values: `SESSION_TABLE_QUOTA_EXCEEDED` and `GLOBAL_CASLIB_QUOTA_EXCEEDED`. Both of these Java constants reside in the `com.sas.actions.StatusCodes` class. Compare these constants to the value returned by `getStatusCode()` in `CASException`.

Policy Details

Here are some key details about CAS resource management policies:

- You must have CAS Superuser privileges to create and manage resource management policies.
- Supported only in full deployments.

  Important: For CAS to use resource management policies, the environment variable, CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT must be turned on. For more information, see CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT.
- CAS resource management use of CPU shares from SAS Configuration Server (Consul) applies only on Linux. But CAS resource management use of quotas is applicable on both Linux and Windows.
- Supported on both single-machine (SMP) and distributed (MPP) CAS servers.
- Policies are specific to a CAS server.
- Policy limits apply per machine for distributed CAS servers.
- Policies are managed through the SAS Viya command line interface and are turned off by default.

  TIP With the CAS command line interface, you can create a policy template that can serve as a starting point for creating your own CAS resource management policies. For more information, see “Create Policies from JSON Templates” on page 694.
- CAS stores policies as key-value pairs in the SAS Configuration Server (Consul).

How Policies Work

Administrators create policies using the SAS Viya command line interface, which stores the policies as key-value pairs in the SAS Configuration Server (Consul).

There are two types of policies:

- global caslibs (`globalCaslibs`)

  One policy per CAS server that places a disk cache space quota on global caslibs.
- priority-level (`CAS-server-name-priority-n`)

  A maximum of five policies per CAS server that places a disk cache space quota on personal caslibs and a CPU consumption limit on CAS sessions.

If policies are defined, a CAS server reads its policy information directly from the configuration server. For each priority-level policy it finds that contains a CPU consumption value, CAS creates a corresponding Linux cgroup. These cgroups are created at CAS server start-up, and destroyed when the CAS server is shut down. CAS provides one policy for global caslibs and up to five priority level policies to which sites can assign resources. The administrator can assign users to these priority-level policies based on their identity group memberships, or they can explicitly assign priority-level policies to individual users.

Figure A.15 How CAS policies work

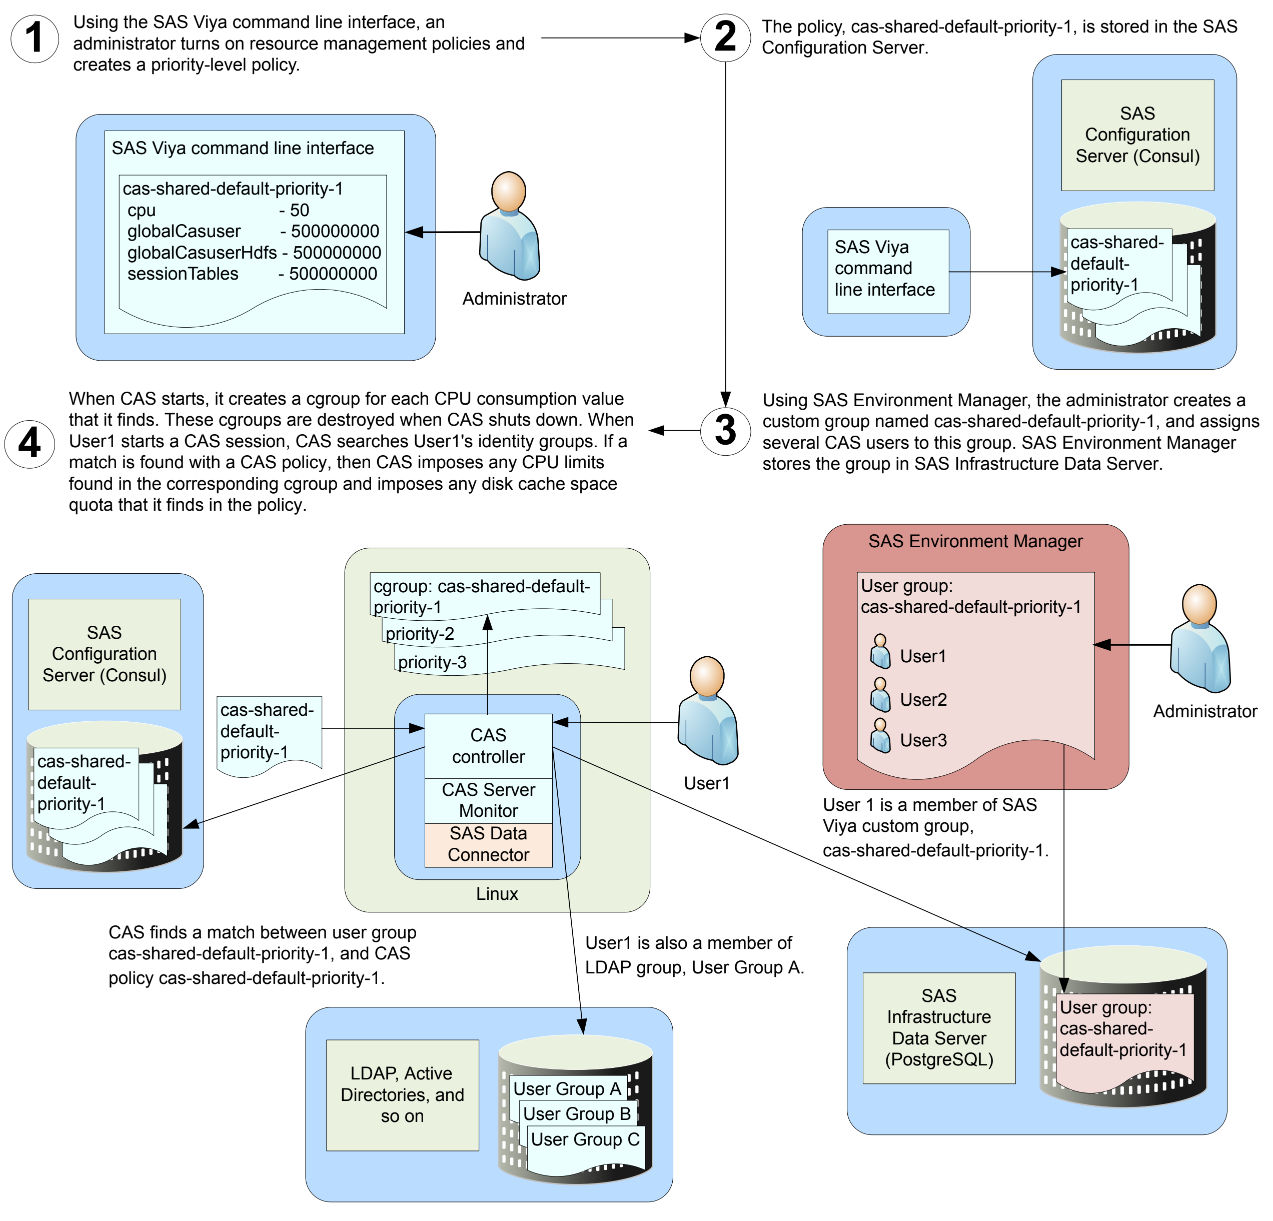
1
Using the SAS Viya command line interface, an administrator turns on resource management policies and creates a priority-level policy.
2
The policy, cas-shared-default-priority-1, is stored in the SAS Configuration Server.
SAS Viya command line interface
cas-shared-default-priority-1
cpu - 50
globalCasuser - 500000000
globalCasuserHdfs - 500000000
sessionTables - 500000000
Administrator
SAS Configuration Server (Consul)
SAS Viya command line interface
cas-shared-default-priority-1
4
When CAS starts, it creates a cgroup for each CPU consumption value that it finds. These cgroups are destroyed when CAS shuts down. When User1 starts a CAS session, CAS searches User1's identity groups. If a match is found with a CAS policy, then CAS imposes any CPU limits found in the corresponding cgroup and imposes any disk cache space quota that it finds in the policy.
3
Using SAS Environment Manager, the administrator creates a custom group named cas-shared-default-priority-1, and assigns several CAS users to this group. SAS Environment Manager stores the group in SAS Infrastructure Data Server.
SAS Configuration Server (Consul)
cas-shared-default-priority-1
cas-shared-default-priority-1
cgroup: cas-shared-default-priority-1
priority-2
priority-3
CAS controller
CAS Server Monitor
SAS Data Connector
Linux
User1
SAS Environment Manager
User group: cas-shared-default-priority-1
User1
User2
User3
Administrator
User 1 is a member of SAS Viya custom group, cas-shared-default-priority-1.
CAS finds a match between user group cas-shared-default-priority-1, and CAS policy cas-shared-default-priority-1.
User1 is also a member of LDAP group, User Group A.
SAS Infrastructure Data Server (PostgreSQL)
User group: cas-shared-default-priority-1
LDAP, Active Directories, and so on
User Group A
User Group B
User Group C

To help explain how policies work, here is a sample policy definition for a CAS server. In this definition, there is a policy for global caslibs, three priority-level policies (out of a possible of five), and priority-level policy assignments for three individual users:

```
resourceManagement
   globalCaslibs
      _ALL_      400000000
      HPS        200000000
      MyGlobal   100000000

   priorityLevels
     cas-shared-default-priority-1
        cpu                 - 50
        globalCasuser       - 500000000
        globalCasuserHdfs - 500000000
        sessionTables       - 500000000
     cas-shared-default-priority-2
        cpu                - 20
        globalCasuser      - 50000000
        globalCasuserHdfs - 50000000
        sessionTables     - 50000000
     cas-shared-default-priority-3
        cpu                - 30
        globalCasuser      - 10000000
        globalCasuserHdfs - 10000000
        sessionTables     - 10000000

   priorityAssignments
      userA    1
      userB    3
      userC    2
```

CAS has built-in functionality to recognize user group names that start with the name of the CAS server as a group related to CAS resource management. When a user authenticates to create a new session, CAS always scans the user's identity groups searching for resource management user group names. If CAS finds that the user belongs to a resource management user group—a group whose name starts with the CAS server name—then CAS attempts to match the group name with the list of resource management policies.

Relying on this behavior, the administrator can create corresponding custom groups using identical names in SAS Environment Manager and add the appropriate user names to these groups.

In our example, if the user that is authenticating to CAS is a member of the group, **cas-shared-default-priority-2**, then CAS enforces the policy of the same name. Therefore, this user has a 50GB quota for all loaded tables and has a CPU share of 20. (CPU shares are discussed later in this document.)

If a user is a member of multiple resource management groups, then CAS assigns the lowest priority number. If the user is not a member of a resource management group, then CAS searches for an explicit assignment for the user under the policy definition `priorityAssignments` section. If the user is not a member of a resource management group, and the user has not been assigned any priority level, then the user has no limits.

CPU Shares (Linux)

Important: CAS resource management use of CPU shares from SAS Configuration Server (Consul) applies only on Linux. But CAS resource management use of quotas is applicable on both Linux and Windows.

On Linux, CAS servers support session CPU limits. Because there is a unique user for each CAS server session, CPU consumption can be controlled on a per user basis.

Note: You can specify cgroup shares for the CAS server as a whole with the cas.CPUSHARES server configuration option.

Every priority-level policy can be assigned a CPU share. A share is a value relative to the sum of all the shares used in a CAS server's policies. The recommended practice is to define shares for all policies to total 100—thus defining percentages.

Here is our example repeated again:

```
resourceManagement
   globalCaslibs
      _ALL_      400000000
      HPS        200000000
      MyGlobal   100000000

   priorityLevels
     cas-shared-default-priority-1
        cpu                - 50
        globalCasuser      - 500000000
```

```
          globalCasuserHdfs - 500000000
          sessionTables    - 500000000
       cas-shared-default-priority-2
          cpu               - 20
          globalCasuser     - 50000000
          globalCasuserHdfs - 50000000
          sessionTables    - 50000000
       cas-shared-default-priority-3
          cpu               - 30
          globalCasuser     - 10000000
          globalCasuserHdfs - 10000000
          sessionTables    - 10000000

     priorityAssignments
        userA    1
        userB    3
        userC    2
```

In this example, when a session is created by a user who is a member of the `cas-shared-default-priority-2` user group (and is not a member of `cas-shared-default-priority-1`), then CAS selects the priority 2 resource management group for the session. And, the session is placed in the `cas-shared-default-priority-2` cgroup with a 20% share.

On a completely busy system, that user (together with all other users assigned to the priority 2 cgroup) are limited to using 20% of the CPU capacity. But, during off-hours, when there is excess CPU capacity, that same user might be able to consume all available capacity.

Note that priority 3 was assigned a higher percentage than priority 2. Does that mean it would be more advantageous to be assigned to priority 3? Probably not. Because there are likely many more users assigned to the priority 3 policy sharing 30% of the CPU capacity. When making share assignments for policies, be sure to consider the potential number of users and their load that might be members of the corresponding resource management user group.

See Also

"CAS Resource Management Policies"

Using CAS Server Monitor

What Is CAS Server Monitor?

CAS Server Monitor is a web application that you use to monitor your CAS Server and to perform some administration tasks.

Important: By default, CAS Server Monitor is turned off in full deployments of SAS Viya.

Monitor Cheat Sheet

In addition to providing information, CAS Server Monitor supports the following tasks:

Task	Required Role
Stop the server	CAS (Superuser)

Task	Required Role
Add or remove nodes**	CAS (Superuser)
Terminate your sessions	(none)
Terminate other sessions	CAS (Superuser) or Data
Designate administrators	CAS (Superuser)
Set caslib permissions on page 390	CAS (Superuser) or Data*

* In addition, any user who has the ManageAccess permission for a global caslib can set permissions on that caslib.

** On Cloud Foundry, do not attempt to add nodes, remove nodes, or terminate a server instance from the CAS Server Monitor (or with the addNode and removeNode CAS actions). Instead, use the appropriate BOSH command.

Access the Monitor

Important: By default, CAS Server Monitor is turned off in full deployments of SAS Viya. Instead, use SAS Environment Manager.

1 You can access the monitor without first starting a CAS server session. However, if there are no sessions, the session list in the monitor is empty.

If you already have a CAS server session, skip to Step 2. Otherwise, to start a session, perform the following steps:

a Open a web browser and sign in to SAS Studio with administrator privileges:

https://*reverse-proxy-server*/SASStudio

b In the **Code** tab, start a CAS server session by entering the following:

```
cas my_session;
```

c Click 🏃.

You should see output similar to the following:

```
56         cas my_session;
 NOTE: The session MY_SESSION connected successfully to Cloud Analytic Services 10.120.9.159 using port 5570.
The UUID is
       5120eb8f-ca06-8c44-9f93-2dd2e557b1cb. The user is myacct and the active caslib is CASUSER(myacct).
 NOTE: The SAS option SESSREF was updated with the value MY_SESSION.
 NOTE: The SAS macro _SESSREF_ was updated with the value MY_SESSION.
 NOTE: The session is using 0 workers.
```

2 Open a web browser and enter the following URL in the address field:

https://*http-reverse-proxy-machine-name*/cas-shared-*deployment-instance-name*-http

Here is an example:

https://myproxy.example.com/cas-shared-default-http

TIP To designate administrators, select **Configuration** ⇨ **Administrators**, click **Add**, and select **CAS (Superuser)**.

Usage notes:

- When the CAS server terminates, the CAS Server Monitor also terminates.
- The session view remains displayed even if the session is terminated. A session view is displayed until you close it.
- Sessions are removed from the list if the session is terminated. You can click the refresh button to get the current list of sessions.
- You can also access CAS Server Monitor from SAS Studio.
- CAS Server Monitor does not use a session time-out. You must click ▼ in the top right corner of the window and select **Sign Out** to exit.

Set Monitor Preferences

To control the view that you see by default, click ▼ in the top right corner of the window and select **Settings**.

SAS Cloud Analytic Services: Troubleshooting

Failed to open temporary file for upload (80BFE801): /tmp/cascache1/_f_43d6c87c_7f5d854996e8.sas7bdat

Explanation:

Insufficient disk space in CAS_DISK_CACHE on the CAS controller.

Resolution:

Add more disk space.

The first data quality operation that is performed after CAS starts takes longer than normal to execute.

Explanation:

Data quality operations require access to the CAS table containing the QKB on all workers in the analytics cluster. Loading this table takes longer for the first data quality operation after CAS starts because the table has not been loaded before. (Subsequent loads take less time.) CAS automatically loads this table when it starts. However, CAS tables for all non-default QKBs are not loaded automatically.

Resolution:

If CAS sessions use non-default QKBs, load these non-default QKBs as part of starting CAS. In the CAS configuration file, casstartup_usermods.lua, add a qkb.loadQKB() action for each non-default QKB. For more information, see “Understanding Configuration Files and Start-up Files”.

From SAS Studio 4.4, CAS Server Monitor is inaccessible

Explanation:

In full SAS Viya deployments, CAS Server Monitor is turned off by default and cannot be launched from the SAS Studio 4.4 banner.

Resolution:

Use another interface to perform CAS administration such as SAS Environment Manager or the SAS Viya command line interface.

A CAS server with process id <*process-id*> is currently running and has exclusive access to the access control storage location <*storage-location*>

Explanation:

This error message is displayed when a CAS server fails to access the permstore and terminates because the permstore is already being accessed by another CAS server.

On Windows, the error message is as follows: `A CAS server is currently running and has exclusive access to the access control storage location <storage-location>.`

SAS Cloud Analytic Services: Reference

Configuration File Options

How Do I Use CAS Configuration File Options?

You set SAS Cloud Analytic Services options in the CAS controller's configuration file, casconfig_usermods.lua. During CAS server start-up, the controller shares the configuration as each worker and the backup controller (if present) connects. By default, casconfig_usermods.lua is located in **`/opt/sas/viya/config/etc/cas/default`** on Linux and in **`\ProgramData\SAS\Viya\etc\cas\default`** on Windows.

If you want to isolate a configuration option change to a particular CAS node, then make your change in node_usermods.lua residing on the particular node machine.

TIP For sites that use Ansible, it is recommended that you make your CAS server configuration changes to vars.yml and rerun Ansible to apply these changes. For more information, see "Modify the vars.yml File" in *SAS Viya for Linux: Deployment Guide*.

There are additional CAS configuration files and directories. For more information, see "Understanding Configuration Files and Start-up Files".

For the order of precedence for server configuration options, see How the Session Option Values Are Determined.

When a session starts, session options specified in the casconfig files are set for the session. For the order of precedence for session options specified in the casconfig files, see Table 27.81 on page 503.

TIP Remember that you can also set operating system environment variables in casconfig_usermods.lua. For example, `env.HADOOP_HOME='/hadoop/hadooop-someversion' env.HADOOP_NAMENODE='name_node.example.com'`.

You can override CAS configuration file settings for a session by changing the equivalent session option. For more information, see Setting Session Options.

Configuration File Options Reference

Note:

Other security-related configuration file options can be found in "Configuration File Options for Data Transfer" in *Encryption in SAS Viya: Data in Motion*.

cas.APPTAG='*tag-string*'

specifies an arbitrary string to prefix to log messages.

Using `apptag` helps determine which log messages are associated with an application.

Valid in	CAS statement SESSOPTS option casconfig_usermods.lua file

Category	Session
Default	No *tag-string*
Note	The CAS server uses `apptag` when writing to its log.
See	cas.LOGCFGLOC
Example	`apptag='my_app'`

cas.CMPOPT='*optimization-value* <*optimization-value* <…>>' | 'all' | 'none'
specifies the type of code generation optimizations to use in the SAS language compiler.

- *optimization-value*

  specifies the type of optimization that the SAS compiler is to use. Specify one or more of the following as a space-delimited list enclosed in quotation marks:

  - 'dumptkgcode' | 'nodumptkgcode'

    specifies whether all CAS server nodes create an output file with the generated program. The CAS log lists where CAS writes the output file.

  - 'extramath' | 'noextramath'

    specifies whether the compiler is to retain or remove the extra mathematical operations that do not affect the outcome of a statement.

  - 'funcdifferencing' | 'nofuncdifferencing'

    specify **funcdifferencing** to calculate numeric-differencing derivatives for user-defined functions. Specify **nofuncdifferencing** to calculate analytic derivatives for user-defined functions.

  - 'guardcheck' | 'noguardcheck'

    specifies whether the compiler checks for array boundary problems.

  Note: **noguardcheck** is set when cmpopt is set to 'all' or 'none'.

  - 'misscheck' | 'nomisscheck'

    specifies whether to check for missing values in the data.

  - 'precise' | 'noprecise'

    specify precise to handle exceptions at the operation boundary. Specify **noprecise** to handle exceptions at the statement boundary.

- 'all'

  specifies that the compiler is to optimize the machine language code by using the **noextramath**, **nomisscheck**, **noprecise**, **noguardcheck**, and **nofuncdifferencing** optimization values.

  Note: 'all' cannot be specified with other values.

- 'none'

  specifies that the compiler is not set to optimize the machine language code by using the **extramath**, **misscheck**, **precise**, **noguardcheck**, and **funcdifferencing** optimization values.

  Note: 'none' cannot be specified with other values.

Valid in	CAS statement SESSOPTS option casconfig_usermods.lua file
Category	Action

Default	**noextramath**, **nofuncdifferencing**, **noguardcheck**, **nomisscheck**, and **noprecise**
Note	If the data contains a significant amount of missing data, specify **misscheck** to optimize the compilation. Otherwise, specify **nomisscheck**.
Example	In this example, the SAS compiler is set to retain the extra mathematical operations, check for missing values, and handle exceptions at an operation boundary: `cas.cmpopt='extramath misscheck precise'`

cas.COLLATE='mva' | 'uca'

specifies the collating sequence for sorting.

`mva` specifies SAS client collating. `uca` specifies a locale-appropriate collating sequence.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	Sort
Default	**'uca'**
Example	`cas.collate='mva'`

cas.COLOCATION='none' | 'hdfs'

specifies whether to create a personal caslib (`hdfs`) at CAS server start-up.

A server started in MPP mode defaults to `hdfs` because it assumes it is co-located with Hadoop. Specify `none` for the server running in MPP mode not to create a personal caslib at start-up.

Valid in	casconfig_usermods.lua file
Category	Caslib
Default	**`cas.colocation='hdfs'`**
Restriction	Applies to Linux only.
Requirement	Used with `cas.mode='mpp'` and `cas.hdfsuserloc`.
Example	In this example, the CAS server is running in MPP mode and is not co-located with Hadoop. At start-up, the CAS server does not create a personal caslib for the user ID under which the server is run. `cas.colocation='none'`

cas.CPUSHARES='*number*'

on Linux operating systems, specifies cgroup shares for the CAS server as a whole. The higher the value, the more priority is given to the CAS server when Linux allocates CPU resources.

on Windows operating systems, specifies the percentage of CPU time that should be allocated to the CAS server. The range on Windows is 0–100. The default is 95. A zero indicates that the CAS server should choose a value between 1–100 that is appropriate for most environments.

Note: A recommended value for `cas.CPUSHARES` for a SAS Viya full deployment on Windows is 70.

Valid in	casconfig_usermods.lua file
Category	Administration

Operating environments	Use cas.CPUSHARES to balance a CAS server against other processes in other top-level cgroups, including other CAS servers.
Windows specifics	The exact division of CPU time between CAS and the other processes is performed by the Windows operating system. CAS provides guidance only to Windows as to how this division should be performed. A value above 95 is not recommended. The maximum value (100) results in CAS not requesting that the operating system to limit the CPU cycles provided to the CAS server. A value of 100 can result in other critical processes, possibly including other processes that are essential parts of a SAS Viya installation, becoming unresponsive, and potentially failing.
Note	When other processes are idle, CAS can still consume up to 100% of available CPU cycles. CAS resource management use of CPU shares from SAS Configuration Server (Consul) applies only on Linux. But, CAS resource management use of quotas is applicable on both Linux and Windows.
See	CPU Shares (Linux)
Example	In this example, on Windows, at least 10% of the available CPU time is reserved for other processes, with up to 90% of the available CPU time being allocated to CAS processes: `cas.cpushares='90'`

cas.DATASTEPFMTERR=true | false

corresponds to the FMTERR in SAS. Specifies how the DATA step reacts when a format is not available. When `true`, the DATA step writes an error and stops. When `false`, the DATA step uses `$w.` or `BEST12.` instead of the unavailable format. (The unavailable format is still associated with variables in the output table.)

Valid in	CAS statement SESSOPTS option casconfig_usermods.lua file
Category	DATA Step
Alias	FMTERR
Default	True
Note	The values `true` and `false` are case sensitive.
See	FMTERR System Option
Example	In this example, the DATA step uses $w. or BEST12. instead of the unavailable format. `cas.datastepfmterr=false`

cas.DATASTEPMSGSUMLEVEL='all' | 'none' | 'put'

specifies the DATA step message summary level. When the DATA step runs on multiple threads, the same message can be generated on each thread. This option controls the summary level of duplicate messages.

- 'all'

  The first occurrence of all message and put statements are sent to the client when they occur. Duplicate occurrences of all message and put statements are summarized and sent to the client when the DATA step exits. This is the default.

- 'none'

  All message and put statements from every thread are written to the client log. No summarization occurs.

- 'put'

The first occurrence of all message and put statements are sent to the client. Duplicate occurrences of messages are summarized and sent to the client when the DATA step exits. Put statements are not summarized; rather, they are sent to the client when they occur.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	DATA Step
Default	All
Example	In this example, all message and put statements from every thread are written to the client log. No summarization occurs. `cas.datastepmsgsumlevel='none'`

cas.DATASTEPREPLACETABLE=true | false
specifies whether a DATA step can replace an existing table.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	DATA Step
Default	true
Note	The values `true` and `false` are case sensitive.
Example	`cas.datastepreplacetable=true`

cas.DCHOSTNAMERESOLUTION= '[ep | ep_fqdn | cas | cas_ipv6]'
specifies how CAS sends the CAS node machine name to SAS Embedded Process. cas.DCHOSTNAMERESOLUTION is valid for any data store that supports SAS Embedded Process.

- 'ep'

  Send CAS node machine names to SAS Embedded Process exactly how they are defined in cas.hosts. SAS Embedded Process resolves the names using either IPv4 or IPv6.

- 'ep_fqdn'

  Send CAS node machine names to SAS Embedded Process as fully qualified domain names. SAS Embedded Process resolves the names.

- 'cas'

  (Default) send CAS node machine names to SAS Embedded Process as IPv4 addresses.

- 'cas_ipv6'

  Send CAS node machine names to SAS Embedded Process as either IPv4 or IPv6 addresses.

Valid in	casconfig_usermods.lua file
Category	Network
Default	'cas'
Restriction	Applies to Linux only.
Note	

Example	`cas.dchostnameresolution='ep'`

cas.DQLOCALE='*locale-code*'

specifies the default locale to use for data quality (DQ) operations, using the five-letter SAS Quality Knowledge Base (QKB) ISO locale code.

For more information, see QKB Locale ISO Codes.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	Data Quality
Example	In this example, the default locale for DQ operations is French Canadian: `cas.dqlocale='fr_CA'`

cas.DQSETUPLOC='*QKB-name*'

specifies the name of the default SAS Quality Knowledge Base (QKB) to use for data quality (DQ) operations.

QKB-name is the absolute path to a SAS Quality Knowledge Base.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	Data Quality
Example	`dqSetupLoc='/opt/sashome/SASQualityKnowledgeBases/en/my_qkb'`

cas.ELASTIC=true | false

indicates that new machines are allowed to join the analytics cluster.

Important: Adding a node is best suited for times when the system is mainly processing batch jobs, where a delay is not a concern. CAS waits for all sessions to complete any current actions before it adds the new worker node. CAS prevents any new actions from running until the node is added. After the node is added, existing sessions are updated to include the added node, except subset sessions. Subset sessions are not modified to use the added node.

Valid in	casconfig_usermods.lua file
Category	Administration
Default	false
Restriction	Applies to Linux only.
Requirement	Used with `cas.gcport`.
Supports	CAS servers running MPP.
Note	The values `true` and `false` are case sensitive.
See	cas.GCPORT
Example	In this example, the CAS controller allows new worker nodes to join the analytic cluster: `cas.elastic=true`

cas.EVENTDS='*event-data-set*'

specifies one or more event objects that define custom date events.

event-data-set specifies the name of a data set that contains event definitions. You can use a one-level name or a two-level name, such as `libref.dataset`. When specifying multiple names, separate each name with a space.

Enclose *event-data-set* in single quotation marks.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	Input Control
Example	`cas.eventds='mydataset'`

cas.FMTSEARCH='*logicalformatlibname logicalformatlibname2 …* '

specifies the format search list to be automatically set at session start-up.

Server configuration files can be used to add and promote common format libraries when a server starts. Promotion makes the format libraries available to all sessions.

During session start-up, the cas.FMTSEARCH value is used to generate and execute the setFmtsearch action. The setFmtsearch action specifies the order in which format libraries are searched. Table variables can have a user-defined format association. During table processing, when a user-defined format needs to be applied, CAS searches the list of format libraries established by the setFmtsearch action.

Category	Formats
Default	blank
Interaction	Specifying the setServOpt action for format search in startup.lua takes precedence over cas.FMTSEARCH used in casconfig.lua.
Note	The format library names are logical names known to a session. The names are case insensitive.
Example	`cas.fmtsearch='myformatsA myformatsB'`

cas.GCPORT=*port*

specifies the network port that is used on a distributed server for communication between the controller and its worker nodes.

The commonly configured port is 5580.

Valid in	casconfig_usermods.lua file
Category	Network
Default	0 (random port in the range 32678–61000)
Restriction	Applies to Linux only.
Supports	CAS servers running MPP.
See	cas.HTTPPORT and cas.PORT
Example	`cas.gcport=5580`

cas.HDFSUSERLOC='/*hdfs-path*/%USER'

for CAS servers running in MPP mode, specifies that the server create a personal caslib for each user at session start-up time in the specified HDFS path.

`'hdfs-path/%USER'` refers to a directory named for the user's user ID under the specified HDFS path.

Enclose *hdfs-path* in single quotation marks.

Valid in	casconfig_usermods.lua file
Category	Data
Restriction	Applies to Linux only.
Requirement	cas.MODE='mpp' on page 526 and cas.COLOCATION='hdfs' on page 514
Example	In this example, the user's caslib directory is a subdirectory named for the user ID under **`/user`**: `cas.hdfsuserloc='/user/%USER'`

cas.HOSTKNOWNBY='*machine-name* | *IP-address*'

specifies the preferred network interface for CAS to use on machines that contain multiple network interface cards.

machine-name is the fully qualified host name that is associated with the preferred network interface. *IP-address* is the IP address of the preferred network interface.

On the controller machine, the host names in the machine list file should be either the IP addresses of the preferred network interface to use on each peer node or the fully qualified host name that is associated with the preferred network interfaces.

When `cas.ELASTIC=true` every machine in the CAS analytic cluster should specify `cas.HOSTKNOWNBY` and the `cas.ELASTIC=true` options. All worker nodes should be started with the `join` command. The machine name that is specified after the join command should be either the IP address of the preferred network interface on the controller or the fully qualified host name that is associated with that IP address.

Valid in	casconfig_usermods.lua file
Category	Network
Restriction	Applies to Linux only.
See	cas.ELASTIC and cas.MACHINELIST
Example	`cas.hostknownby='primary.private.example.com'`

cas.HTTPPORT=*port* | *port-range*

The port (or range of ports) that SAS Cloud Analytic Services listens to for HTTP communication.

The commonly configured port is 8777.

Valid in	casconfig_usermods.lua file
Category	Network
Default	0 (random port)
Note	If the first port in the range is already taken, CAS tries the next port until it finds a port that is free.
See	cas.HTTPPORTMAX , cas.GCPORT , and cas.PORT

Examples	`cas.httpport=8777`
	`cas.httpport=8777-9000`

cas.HTTPPORTMAX=*maximum-port-range*

specifies the maximum port range that SAS Cloud Analytic Services listens to for HTTP communication.

Valid in	casconfig_usermods.lua file
Category	Network
Default	0
Range	0–65535
See	cas.HTTPPORT
Example	`cas.httpport=8777-9000`

cas.INITIALBACKUPS= -1 | 0 | *positive-number*

specifies whether SAS Cloud Analytic Services (CAS) waits for backup controllers to connect to the CAS analytics cluster before CAS begins to accept client connections.

Valid values are:

- -1

  Use the value specified in cas.hosts for the number of backup controllers to connect to the analytics cluster before CAS begins accepting connections from clients.

- 0 (zero)

  Do not wait for any backup controllers to connect to the analytics cluster before CAS begins accepting connections from clients.

- 1

  Wait for the backup controller to connect to the analytics cluster before CAS begins accepting connections from clients.

Valid in	casconfig_usermods.lua file
Category	Server
Default	-1
Range	-1–1
Restriction	Applies to Linux only.
Example	`cas.INITIALBACKUPS=-1`

cas.INITIALWORKERS='*n*'

specifies the number of CAS worker nodes that must join the analytic cluster before CAS begins processing user connections.

cas.initialworkers enables administrators to establish an expected cluster size for configurations, where it is typical for all or most worker nodes to join elastically.

Valid in	casconfig_usermods.lua file
Category	Server

Default	-1
Range	-1 to 32767
Restriction	Applies to Linux only.
Requirement	cas.elastic must be set to true.
Notes	A value of zero indicates that the controller does not wait for any worker nodes to join the cluster before it begins to establish user connections.
	A value of -1 indicates that the controller waits for the number of workers that is specified in the machine list file.
See	cas.ELASTIC and cas.MACHINELIST
Example	In this example, the CAS controller waits for 16 workers to join the analytic cluster before it begins processing user connections. `cas.initialworkers='16'`

cas.INTERVALDS='*interval-1=libref.dataset-name-1 <interval-2=libref.dataset-name-2 …>*'

specifies one or more interval-name=value pairs, where the value is the name of a data set that contains user-defined intervals.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	Input Control
See	INTERVALDS= System Option
Example	`cas.intervalds='subsid1=subsid.storeHours'`

cas.JREOPTIONS='(*JRE-option <JRE-option> <…>>*)'

specifies the Java Virtual Machine (JVM) options that SAS Cloud Analytic Services uses at start-up. Separate JRE options with a whitespace character. Enclose any paths in quotation marks.

For the list of the valid Java options, and what they do, see http://docs.oracle.com/javase/6/docs/technotes/tools/windows/java.html

Valid in	casconfig_usermods.lua file
Category	Java
Default	(null)

Example	In the following example, the initial and maximum sizes of the memory allocation pool are set to 256 and 1024MB, respectively. Also, the log4j configuration file path and Java classpath are set:
	`cas.jreoptions = '(-Xms256m -Xmx1024m`
	`  -Dlog4j.configuration=' .. ' -Djava.class.path=' .. env.CAS_HOME .. '/lib/base/base-tkjni.jar')'`

cas.KEYFILE='*pathname*'

identifies to the CAS controller the path and filename to the X.509 digital certificate file that is used to start the server. The certificate must be signed by a CA that is trusted by the CAS server.

Enclose *pathname* in single quotation marks.

Valid in	casconfig_usermods.lua file

Category	Security
Restriction	Applies to Linux only.
Requirement	Used with `cas.elasticssl` and `cas.mode='mpp'`.
Supports	CAS servers running MPP.
See	cas.GCPORT and cas.ELASTIC
Example	`cas.keyfile='/opt/TKGrid/certs/controller.pem'`

cas.LIFETIME=*minutes*

indicates the duration, in minutes, that a server remains running.

Valid in	casconfig_usermods.lua file
Category	Administration
Default	0
Example	In the following example, the server shuts itself down in 120 minutes: `cas.lifetime=120`

cas.LOCALE='*POSIX-locale-string*'

specifies the locale to use for sorting and formatting. For a list of valid POSIX locale strings, see SAS National Language Support (NLS): Reference Guide

Valid in	CAS statement SESSOPTS option casconfig_usermods.lua file
Category	Localization
Default	'en_US'
Example	`cas.locale='fr_FR'`

cas.LOGCFGLOC='*pathname*'

specifies the path to the SAS logging facility logging configuration file.

Enclose *pathname* in single quotation marks.

Valid in	casconfig_usermods.lua file
Category	Log
See	cas.APPTAG on page 512

Examples	`cas.logcfgloc='/opt/sas/cas1/etc/logconfig.xml'`
	Here is an example on Windows:
	`cas.logcfgloc='C:\\ProgramData\\SAS\\Viya\\etc\\cas\\default\\logconfig.xml'`

cas.LOGFLUSHTIME=-1 | 0 | *number*

specifies the log flush time, in milliseconds.

- -1

  flushes logs after each action completes.

- 0

  flushes logs as they are produced.
- *number*

  flushes logs in *number* milliseconds.

Valid in	CAS statement SESSOPTS option casconfig_usermods.lua file
Category	Log
Default	100
Range	-1–86400
Example	In the following example, the CAS server writes buffered lines to the log every 500 milliseconds: `cas.logflushtime=500`

cas.MACHINELIST='*path/machine-list-file*'
identifies the path and filename on the controller machine that contains the list of machines in the CAS analytics cluster.

Enclose *path* in single quotation marks.

machine-list-file contains all of the machines in the analytics cluster in the form:

```
<fully-qualified-domain-name controller | worker>
```

Place each machine on a separate line. For example:

```
my_machine01.example.com controller
my_machine02.example.com worker
my_machine03.example.com worker
my_machine04.example.com worker
my_machine05.example.com worker
```

Valid in	casconfig_usermods.lua file
Category	Administration
Restriction	Applies to Linux only.
Requirement	Used with `cas.mode='mpp'`.
Interaction	Do not specify `cas.mode='smp'` when a valid machine list file is used.
See	cas.MODE
Example	`cas.machinelist='/etc/my_machine_list'`

cas.MAXCORES='*number-of-cores*'
specifies the limit for the total number of physical cores that are available to a CAS server. For distributed CAS servers, cas.MAXCORES refers to the total number of cores for the controller and workers. (Cores for the backup controller are not counted.)

When specified, the lesser of the specified `cas.MAXCORES` value and the product licensed cores is used during CAS action invocation.

If hyperthreading is enabled on a worker, CAS uses two virtual cores, both on the same physical core. The controller also allocates threads. (See the examples that follow for more information.)

Valid in	casconfig_usermods.lua file
Category	Administration
Interaction	If cas.MAXCORES is specified too low, CAS is not able to establish a session (TKCASA_LICENSE_TOO_SMALL).
Note	The core count limit is server-wide, and for distributed CAS servers the value should be at least the same as the total number of machines. The total number of machines is the primary controller and workers. (The backup controller is not included in this total.) For example, if a distributed CAS server has one controller and one worker, and `cas.MAXCORES=4`, the maximum number of cores that the worker can use is two. If you set `cas.MAXCORES` too low, CAS writes a licensing error.
Examples	In this example, the total number of cores available to the CAS server is 17: `cas.maxcores='17'`

	In another example, we want to ensure that exactly 128 hyperthreads per worker are run. (Hyperthreads equal two times the number of cores.) For a single-machine CAS server:
	`cas.MAXCORES='64'`
	For a distributed CAS server use the formula, (Nworkers+1) * 64. To ensure that 128 hyperthreads per worker are run for a distributed CAS server that has a controller plus eight workers:
	`cas.MAXCORES='576'`

cas.MAXSESSIONS='*n*'

specifies the maximum number of concurrent sessions. Users who can assume an administrative role are not subject to the limit.

Valid in	casconfig_usermods.lua file
Category	Server
Default	5000
Range	0–100000
Notes	Specifying zero (0) indicates that there is no session limit.
	This option cannot be changed after the system initializes.
Example	In this example, the maximum number of concurrent CAS sessions is 1,000: `cas.maxsessions='1000'`

cas.MAXTABLEMEM=*number* | '[*number* k | m | g | t] '

specifies the maximum amount of physical memory to allocate for a table.

> **TIP** The intent of cas.MAXTABLEMEM is to manage the efficiency of accessing CAS tables on disk, not to control the amount of data that CAS keeps resident in RAM. When you need to manage CAS memory, consider modifying cas.MEMORYSIZE.

- *number*

  specifies the maximum amount of physical memory, in bytes, to allocate for a table.
- '[*number* k | m | g | t]'

specifies the maximum amount of physical memory to allocate for a table in a unit other than bytes: **k** (kilobytes), **m** (megabytes), **g** (gigabytes), and **t** (terabytes).

Valid in	CAS statement SESSOPTS option casconfig_usermods.lua file
Category	Caslib
Default	16M
Note	After this threshold is reached, the server uses temporary files and operating system facilities for memory management.
See	cas.MEMORYSIZE
Example	In this example, the CAS server can allocate up to 32MB of physical memory for a table: `cas.maxtablemem='32m'`

cas.MEMORYSIZE=*number* | '[*number* k | m | g | t] '

specifies the maximum amount of physical memory to allocate for the CAS cgroup for each machine. This limit also applies to the YARN request, when cas.USEYARN is specified.

- *number*

  specifies the maximum amount of physical memory, in bytes, to allocate for the CAS cgroup and the YARN request.

- '[*number* k | m | g | t]'

  specifies the maximum amount of physical memory to allocate for the CAS CGroup and the YARN request in a unit other than bytes: **k** (kilobytes), **m** (megabytes), **g** (gigabytes), and **t** (terabytes).

Valid in	casconfig_usermods.lua file
Category	Administration
Default	0
Restriction	Applies to Linux only.
See	cas.USEYARN on page 532 How to Limit Memory Use in *SAS Cloud Analytic Services: Fundamentals*
Example	In the following example, the maximum amount of physical memory allocated for the CAS cgroup and the YARN request is 256GB: `cas.memorysize='256g'`

cas.MESSAGELEVEL='all' | 'default' | 'error' | 'none' | 'note' | 'warning'

specifies the log message level.

Valid in	CAS statement SESSOPTS option casconfig_usermods.lua file
Category	Log
Default	'all'

Example `cas.messagelevel='default'`

cas.METRICS=true | false

causes CAS server metrics information to be displayed (`true`) or not displayed (`false`) in the SAS log.

When `cas.metrics=true`, you see information similar to the following displayed in the SAS log:

```
NOTE: Action 'nobs' used (Total process time):
NOTE:       real time              2.100185 seconds
NOTE:       cpu time               0.010999 seconds (0.52%)
NOTE:       total nodes            6 (192 cores)
NOTE:       total memory           1.11T
NOTE:       memory                 7.00K (0.00%)
The analytic server processed the request in 2.100185 seconds.
```

Valid in CAS statement SESSOPTS option

casconfig_usermods.lua file

Category Log

Default false

Note The values `true` and `false` are case sensitive.

See CASLIB Statement

Example In the following example, CAS server metrics information is not displayed in the SAS log:

```
cas.metrics=false
```

cas.MODE='smp' | 'mpp'

forces a server to be started in symmetric multiprocessing mode (`smp`) or in massively parallel processing mode (`mpp`).

Valid in casconfig_usermods.lua file

Category Administration

Restriction Applies to Linux only.

Interaction cas.MODE is implicitly set to 'mpp' when cas.ELASTIC is set or cas.MACHINELIST is set and contains at least one CAS worker node or backup controller. Otherwise, cas.MODE is implicitly set to 'smp'.

Note The server returns an error when `cas.mode='smp'` is specified for a server with a valid machine list.

Example In the following example, the CAS server is forced to start in massively parallel processing mode (MPP).

```
cas.mode='mpp'
```

cas.NODE='*filename*'

specifies the configuration file that is run on all CAS worker nodes.

Valid in casconfig_usermods.lua file

Category Server

Default node.lua

Restrictions	Applies to Linux only.
	Any changes to node.lua should be made to node_usermods.lua.
See	"Understanding Configuration Files and Start-up Files"
Example	`cas.node='node.lua'`

cas.NWORKERS=*number*

specifies the number of worker nodes associated with this session.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	Administration
Default	0
Range	0–5000
Restriction	Applies to Linux only.
Example	`cas.nworkers=8`

cas.PERMSTORE='*path*'

specifies the path to a directory where the CAS server stores permissions.

Enclose *path* in single quotation marks.

The server saves its caslib and access control information to the `cas.PERMSTORE` directory periodically and when it shuts down.

Each subsequent time that the server starts, caslib and access control information is initialized from the server's `cas.permstore` location.

Important: When you update `cas.PERMSTORE` in casconfig_usermods.lua, you must also update SASPERMSTORE in **`/opt/sas/viya/config/etc/sysconfig/cas/default/sas-cas-usermods`**. Add the line: **`export SASPERMSTORE=`*`path`***.

CAUTION! Backups of access controls are not automatically performed. It is strongly recommended that you periodically back up each CAS server's stored access control and caslib information. In particular, it is important to create a backup after you modify access controls or add, delete, or modify global caslibs. See *SAS Viya Administration: Backup and Restore*.

Valid in	casconfig_usermods.lua file
Category	Access Control
Restriction	Do not directly edit the files in a `cas.PERMSTORE` location.
Note	Each CAS server should have its own `cas.PERMSTORE` location. To minimize the potential for network timing issues, it is recommended that each `cas.PERMSTORE` location be on the controller machine and not on a network file system. The server creates a directory with the name of the fully qualified DNS name of the machine that the main controller is running on in the specified permstore directory. For example, if you specify `cas.PERMSTORE='/var/my_permstore'`, CAS writes its permstore to `/var/my_permstore/controller.machine.example.com`.
Examples	Here is an example defining the CAS permstore location on Linux:

```
cas.permstore='/opt/sas/viya/config/etc/cas/default/permstore'
```

Here is an example defining the CAS permstore location on Windows:

```
cas.permstore='C:\\ProgramData\\SAS\\Viya\\etc\\cas\\default\\permstore'
```

cas.PORT=*port*

specifies the port to which the CAS server listens.

The maximum allowable port number is 65535. If a valid port is not specified, the server listens on a port selected by the operating system through the TCP/IP ephemeral port range. A common range is 32768-61000.

The commonly configured port is 5570.

Valid in	casconfig_usermods.lua file
Category	Network
See	cas.GCPORT and cas.HTTPPORT
Example	`cas.port=5570`

cas.PROVLIST='ext' | 'kerb' | 'oauth'

specifies the authentication providers that the CAS server uses to authenticate incoming user connections.

- 'ext'

  The external provider provides support for an external PAM authentication method when root access is required for authentication.

  Note: The 'ext' option applies to Linux only.

- 'kerb'

  The Kerberos provider is used only when a Kerberos ticket is provided for authentication. For more information, see "Kerberos Security" in *SAS Viya for Linux: Deployment Guide*.

- 'oauth'

  OAuth provider is always loaded (even when not listed) to support REST endpoints and communications between CAS worker nodes and the controller.

Valid in	casconfig_usermods.lua file
Category	Security
Default	oauth
Note	The CAS server configures the specified providers and uses each in order until an authenticated connection is successful.
Example	In this example, an external provider provides support for an external PAM authentication method. Although not specified, OAuth is always loaded to support REST endpoints and communications between CAS worker nodes and the controller. `cas.provlist='ext'`

cas.REMOVENODECANCELTIMEOUT='*interval*'

when quiescing sessions in preparation for moving data from nodes that are being removed, the amount of time that CAS waits for long-running actions to complete before canceling them.

Valid in	casconfig_usermods.lua file

Category	Server
Default	120 seconds
Restriction	Applies to Linux only.
Note	A value of zero indicates that a cancel request should never be sent.
Example	`cas.removeNodeCancelTimeout='600'`

cas.REMOVENODEKILLTIMEOUT='*interval*'

when quiescing sessions in preparation for moving data from nodes that are being removed, the amount of time that CAS waits for a canceled action to stop before killing its session.

Valid in	casconfig_usermods.lua file
Category	Server
Default	15 seconds
Restriction	Applies to Linux only.
Note	A value of zero indicates that the sessions should never be killed.
See	cas.REMOVENODECANCELTIMEOUT on page 528
Example	`cas.removeNodeKillTimeout='300'`

cas.RESOLVEWORKERADDRESS=true | false

specifies how CAS list node actions return CAS worker node host names.

When `true`, CAS list node actions attempt to return the list of worker node host names. If the directory name service (DNS) lookup is unresponsive, CAS cancels the lookup and resolveworkeraddress is automatically set to `false`.

When `false`, list node actions return only the IP address of elastically added nodes.

> **TIP** Setting cas.RESOLVEWORKERADDRESS to `false` ensures that the analytic cluster is less impacted by an unresponsive DNS configuration. However, some output is displayed as IP addresses instead of host names.

Valid in	casconfig_usermods.lua file
Category	Server
Default	True
Restriction	Applies to Linux only.
See	cas.ELASTIC
Example	`resolveworkeraddress=false`

cas.SERVICESBASEURL='*URL*'

specifies the URL that enables CAS server to authenticate and to use SAS Viya services. *URL* points to the deployed SAS Viya web services.

When set, cas.SERVICESBASEURL creates a hybrid authentication environment where username-password authentication is converted to an OAuth token and CAS can fetch groups from SAS Viya services and use features such as the credentials vault.

Valid in	casconfig_usermods.lua file
Category	Security
Notes	*URL* must match the reverse proxy server host name and port to enable CAS to communicate with SAS Viya web services.
	Enclose *URL* in single quotation marks.
Example	`cas.servicesbaseurl='http://company.example.com'`

cas.STARTUP='*filename*'

specifies the configuration file that the CAS server runs before the server accepts any client connections. This start-up file contains CAS actions that the server runs as it starts up.

Valid in	casconfig_usermods.lua file
Category	Server
Default	casstartup.lua
Restriction	Any changes to casstartup.lua should be made to casstartup_usermods.lua.
See	"Understanding Configuration Files and Start-up Files"
Example	`cas.startup='casstartup.lua'`

cas.STARTUPDIR='*path*'

specifies the location for the SAS Cloud Analytic Services start-up directory.

Valid in	casconfig_usermods.lua file
Category	Server
Default	**`/opt/sas/viya/config/default`**
See	"Understanding Configuration Files and Start-up Files"
Examples	Here is an example on Linux: `cas.startupdir='/opt/sas/viya/config/default/my_cas_startup'`
	Here is an example on Windows: `cas.startupdir='C:\ProgramData\SAS\Viya\etc\cas\default\MyCASStartup'`

cas.SUBSETSESSIONCOPIES=*number-of-blocks*

specifies the number of extra block copies made for failover in either of the following scenarios:

- a session is smaller than the full server.
- CAS reads blocks of an HDFS remotely.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	Administration

Default	0
Restriction	Applies to Linux only.
Example	`cas.subsetsessioncopies=3`

cas.TAG='*string*'

specifies a string to name the CAS server instance that is visible in the operating system, such as in the process list.

The cas.TAG option can be useful when debugging CAS.

Valid in	casconfig_usermods.lua file
Category	Server

Example	In this example, the string 'cas-my_tag' is used to name the CAS server instance:
	`cas.tag='cas-my_tag'`
	When you view the process list from a Linux command prompt, you see something similar to the following:
	`27019 ?        00:00:01 cas-my_tag`

cas.TENANTID='*string*'

specifies the user ID for the CAS tenant. cas.TENANTID is used to validate that the authenticating user belongs to the correct CAS tenant.

Valid in	casconfig_usermods.lua file
Category	Administration
Restrictions	Used only in multi-tenant CAS deployments.
	Applies to Linux only.
Example	`cas.tenantid='tenant1'`

cas.TIMEOUT=*seconds*

specifies the SAS Cloud Analytic Services session time-out in seconds for a new or existing session.

Valid in	CAS statement SESSOPTS option
	casconfig_usermods.lua file
Category	Session
Default	In order of descending precedence: 1. CAS statement TIMEOUT= option value, if specified 2. SAS system option CASTIMEOUT=, if you explicitly set it in SAS to a value greater than 0 3. 60
Range	0–31536000
Notes	The session time-out starts when the number of connections to the session becomes zero and no actions are executing.
	If a connection is established before the time-out expires, the time-out is canceled. Otherwise, the session is automatically terminated when the time-out expires.

When set to 0, the session is terminated immediately when the connection count becomes zero.

See	CASTIMEOUT= System Option
Example	`cas.timeout=100`

cas.USERLOC='%HOME' | '*pathname*/%USER'

specifies that the CAS server create a personal caslib for each user at session start-up time in the specified location.

`'%HOME'` equates to the user's operating system `$HOME` directory.

`'pathname/%USER'` refers to a directory named for the user's user ID under the specified file system path.

Enclose *pathname* in single quotation marks.

Important: When you update `cas.USERLOC` in casconfig_usermods.lua, you must also update `SASUSERLOCDIR` in **`/opt/sas/viya/config/etc/sysconfig/cas/default/sas-cas-usermods`**. Add the line: **`export SASUSERLOCDIR=path/%USER`**.

Valid in	casconfig_usermods.lua file
Category	Caslib
Restriction	Applies to Linux only.
Examples	In this example, the personal caslib directory is the user's operating system $HOME directory: `cas.userloc='%HOME'`
	In this example, the user's personal caslib directory is named for his or her user ID and is located under **`/local`**: `cas.userloc='/local/%USER'`

cas.USEYARN=true | false

adds a reservation request to YARN for CAS memory size.

The memory limit for the YARN request is set with cas.memorysize.

Valid in	casconfig_usermods.lua file
Category	Administration
Default	false
Restriction	Applies to Linux only.
See	cas.MEMORYSIZE on page 525
	How to Limit Memory Use in *SAS Cloud Analytic Services: Fundamentals*
Example	`cas.useyarn=true`

Grouped by Categories

Access Control Options

- cas.PERMSTORE on page 527

Action Options

- cas.CMPOPT on page 513

Administration Options

- cas.CPUSHARES on page 514
- cas.ELASTIC on page 517
- cas.LIFETIME on page 522
- cas.MACHINELIST on page 523
- cas.MAXCORES on page 523
- cas.MEMORYSIZE on page 525
- cas.MODE on page 526
- cas.NWORKERS on page 527
- cas.SUBSETSESSIONCOPIES on page 530
- cas.TENANTID on page 531
- cas.USEYARN on page 532

Caslib Options

- cas.COLOCATION on page 514
- cas.MAXTABLEMEM on page 524
- cas.USERLOC on page 532

Data Options

- cas.HDFSUSERLOC on page 519

Data Quality Options

- cas.DQLOCALE on page 517
- cas.DQSETUPLOC on page 517

DATA Step Options

- cas.DATASTEPFMTERR on page 515
- cas.DATASTEPMSGSUMLEVEL on page 515
- cas.DATASTEPREPLACETABLE on page 516

Formats Options

- cas.FMTSEARCH on page 518

Input Control Options

- cas.EVENTDS on page 518
- cas.INTERVALDS on page 521

Java

- cas.JREOPTIONS on page 521

Localization

- cas.LOCALE on page 522

Log Options

- cas.LOGCFGLOC on page 522
- cas.LOGFLUSHTIME on page 522
- cas.MESSAGELEVEL on page 525
- cas.METRICS on page 526

Network Options

- cas.DCHOSTNAMERESOLUTION on page 516
- cas.GCPORT on page 518
- cas.HOSTKNOWNBY on page 519
- cas.HTTPPORT on page 519
- cas.HTTPPORTMAX on page 520
- cas.PORT on page 528

Security Options

- cas.KEYFILE on page 521
- cas.PROVLIST on page 528
- cas.SERVICESBASEURL on page 529

Note:

Other security-related configuration file options can be found in "Configuration File Options for Data Transfer" in *Encryption in SAS Viya: Data in Motion*.

Server Options

- cas.INITIALBACKUPS on page 520
- cas.INITIALWORKERS on page 520
- cas.MAXSESSIONS on page 524
- cas.NODE on page 526
- cas.REMOVENODECANCELTIMEOUT on page 528
- cas.REMOVENODEKILLTIMEOUT on page 529
- cas.RESOLVEWORKERADDRESS on page 529
- cas.STARTUP on page 530
- cas.STARTUPDIR on page 530
- cas.TAG on page 531

Session Options

- cas.APPTAG on page 512
- cas.TIMEOUT on page 531

Sort Options

- cas.COLLATE on page 514

CAS Environment Variables

Where Do I Set CAS Environment Variables?

With one exception (see CAUTION later), you set SAS Cloud Analytic Services environment variables in the CAS controller's configuration file, casconfig_usermods.lua. During CAS server start-up, the controller shares the configuration as each worker and the backup controller (if present) connect. By default, casconfig_usermods.lua is located in `/opt/sas/viya/config/etc/cas/default` on Linux and in `\ProgramData\SAS\Viya\etc\cas\default` on Windows.

If you want to isolate an environment variable change to a particular CAS node, then make your change in node_usermods.lua residing on the particular node machine.

> **TIP** For sites that use Ansible, it is recommended that you make your CAS server environment variable changes to vars.yml and rerun Ansible to apply these changes. For more information, see "Modify the vars.yml File" in *SAS Viya for Linux: Deployment Guide*.

On Linux, one CAS environment variable, LD_LIBRARY_PATH, must be specified before CAS starts. Therefore, you must add this variable to the cas_usermods.settings file. If you edit cas_usermods.settings on a single node, only that node is affected. If you want to set an environment variable on every node, you must edit cas_usermods.settings on every node. By default, cas_usermods.settings is located in `/opt/sas/viya/home/SASFoundation` on Linux.

CAUTION! SAS Cloud Analytic Services ignores any instance of LD_LIBRARY_PATH found in the server configuration file. Specify LD_LIBRARY_PATH in the cas_usermods.settings file only.

There are additional CAS configuration files and directories. For more information, see "Understanding Configuration Files and Start-up Files".

CAS Environment Variables Reference

Note: For information about SAS Cloud Analytic Services TLS environment variables, see "CAS TLS Environment Variables" in *Encryption in SAS Viya: Data in Motion*.

env.CAS_ACTION_THREAD_NICE='*niceness-priority*'

specifies the niceness priority for the CPU intensive threads that do CAS action processing.

Increasing the niceness priority value—especially during high CPU utilization—provides an opportunity for the CAS heartbeat thread to run, and prevents workers from being disconnected prematurely.

Valid in	casconfig_usermods.lua file
Category	Environment
Default	1
Range	0–19

Restrictions	`env.CAS_ACTION_THREAD_NICE` is case sensitive.
	Applies to Linux only.
See	The man page for the Linux nice command.
Example	`env.CAS_ACTION_THREAD_NICE='1'`

env.CAS_ADDITIONAL_YARN_OPTIONS='*yarn-option…*'

specifies a yarn queue or other values to provide similar flexibility to the CAS server.

Valid in	casconfig_usermods.lua file
Category	Server
Restrictions	`env.CAS_ADDITIONAL_YARN_OPTIONS` is case sensitive.
	Applies to Linux only.
Example	In this example, the YARN queue and priority are specified: `env.CAS_ADDITIONAL_YARN_OPTIONS='--queue sas_cas --priority 3'`

env.CAS_CFG_CONSULPATH='*path-to-kv-pairs*'

specifies the SAS Configuration Server (Consul) path to key-value pairs loaded into the CAS Options namespace and OSENV namespace, respectively. If a key does not map to a CAS server option, CAS ignores it.

Valid in	casconfig_usermods.lua file
Category	Server
Restrictions	Applies to Linux only.
	`env.CAS_CFG_CONSULPATH` is case sensitive.
Example	In this example, the location to the SAS Configuration Server key-value pairs for CAS configuration options and for CAS environment variables is specified: `env.CAS_CFG_CONSULPATH='/opt/sas/viya/config/data/cas/'`

env.CAS_CONTROLLER_TEMP='*path* [[:*path*] …]'

specifies the disk paths to temporarily store data on the CAS controller machine for the CAS upload action and when saving CSV files to certain cloud platforms such as Amazon S3.

Delimit multiple paths with a colon (:).

Valid in	casconfig_usermods.lua file
Category	Data
Default	When env.CAS_CONTROLLER_TEMP is not specified, CAS defaults to the locations pointed to by env.CAS_DISK_CACHE.
Restrictions	Applies to Linux only.
	`env.CAS_CONTROLLER_TEMP` is case sensitive.
	env.CAS_CONTROLLER_TEMP should point to a file system that uses only a local disk.
Example	`env.CAS_CONTROLLER_TEMP = '/upload_data/disk1:/upload_data/disk2'`

env.CAS_DISK_CACHE='*path* [[:*path*] ...]'

specifies the disk paths to cache data on CAS worker machines.

On Linux, delimit multiple paths with a colon (:). On Windows, delimit multiple paths with a semicolon (;).

Important: It is a best practice to always specify `env.CAS_DISK_CACHE`. When not specified, `env.CAS_DISK_CACHE` defaults to **`/tmp`** on Linux and to the `TEMP` environment variable on Windows.

Valid in	casconfig_usermods.lua file
Category	Data
Restrictions	`env.CAS_DISK_CACHE` is case sensitive.
	`env.CAS_DISK_CACHE` should point to a file system that uses only a local disk.
Tip	There is an advantage to using multiple physical disks. When using multiple threads, mapping files can occur concurrently if multiple disks are used. Hadoop also uses this method. Therefore, there is an advantage to using a set of disks that map to both hadoop_data and CAS_DISK_CACHE directories.
Examples	Here is an example on Linux: `env.CAS_DISK_CACHE = '/data/disk1:/data/disk2'`
	Here is an example on Windows: `env.CAS_DISK_CACHE = 'E:\\casDataCache;F:\\casDataCache'`

env.CASUSERIGNORECASE='ON'

when in effect (specified using any value), causes the CAS server to ignore the letter casing for user names during authentication, group lookup, and process launch. Always specify `env.CASUSERLOWERCASE` whenever specifying `env.CASUSERIGNORECASE`, unless instructed otherwise by SAS Technical Support.

The typical scenario for declaring `env.CASUSERIGNORECASE` is when users run their CAS sessions under their own host account and the user authentication system is configured to be case-insensitive and contains uppercase or mixed case user names. For more information, see "The CASHostAccountRequired Custom Group" on page 365.

Valid in	casconfig_usermods.lua file
Category	Administration
Default	off
Restrictions	Applies to Linux only.
	`env.CASUSERIGNORECASE` is case sensitive.
Requirement	Use with `env.CASUSERLOWERCASE`.
Note	To turn off `env.CASUSERIGNORECASE`, remove its definition.
Example	In this example, env.CASUSERIGNORECASE is in effect: `env.CASUSERIGNORECASE='on'`

env.CASUSERLOWERCASE='ON'

when in effect (specified using any value), causes the CAS server to convert user names to lower letter casings during group lookup. `env.CASUSERLOWERCASE` is typically used in conjunction with `env.CASUSERIGNORECASE`.

The typical scenario for declaring `env.CASUSERLOWERCASE` is when users run their CAS sessions under their own host account and the user authentication system is configured to be case-insensitive and contains uppercase or mixed case user names. For more information, see "The CASHostAccountRequired Custom Group" on page 365.

Valid in	casconfig_usermods.lua file
Category	Administration
Default	off
Restrictions	Applies to Linux only.
	`env.CASUSERLOWERCASE` is case sensitive.
Requirement	Use with `env.CASUSERIGNORECASE`.
Note	To turn off `env.CASUSERLOWERCASE`, remove its definition.
Example	In this example, env.CASUSERLOWERCASE is in effect: `env.CASUSERLOWERCASE='on'`

env.CONSUL_HTTP_ADDR='https://*SAS-Configuration-Server-machine*:*port*'

specifies the SAS Configuration Server (Consul) machine and port used by the CAS Consul interface.

Valid in	casconfig_usermods.lua file
Category	Server
Restrictions	Applies to Linux only.
	`env.CONSUL_HTTP_ADDR` is case sensitive.
Example	`CONSUL_HTTP_ADDR='my-SAS-Configuration-Server.example.com:8501'`

env.CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT='TRUE' | 'FALSE'

when specified, enables CAS resource management policies by turning on communication with SAS Configuration Server (Consul).

Valid in	casconfig_usermods.lua file
Category	Server
Default	Off (FALSE)
Restrictions	Applies to Linux only.
	`env.CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT` is case sensitive.
Note	When env.CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT is false, the GLIBC environment variable, MALLOC_ARENA_MAX, should exceed the number of concurrent sessions.
See	For more information, see "CAS Resource Management".
Example	`env.CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT='true'`

env.CAS_ENABLE_REMOTE_SAVE=TRUE

specifies whether CAS saves blocks on remote HDFS worker nodes.

Valid in	casconfig_usermods.lua file
Category	Data
Restrictions	Applies to Linux only.
	`env.CAS_ENABLE_REMOTE_SAVE` is case sensitive.
Note	Removing `env.CAS_ENABLE_REMOTE_SAVE=true` causes CAS not to save blocks on remote HDFS worker nodes.
Example	In this example, CAS saves blocks on remote HDFS worker nodes: `env.CAS_ENABLE_REMOTE_SAVE=true`

env.CAS_HEARTBEAT_LOST_TIMEOUT='*interval*'

specifies the interval (in seconds) since the last heartbeat received from a CAS worker node before the controller treats the node as lost.

Smaller intervals detect machines that silently leave the network more quickly. Larger intervals are more tolerant of machines that might be exceptionally overloaded.

Valid in	casconfig_usermods.lua file
Category	Server
Default	120 seconds
Range	60 – (no upper limit) seconds
Restrictions	Applies to Linux only.
	`env.CAS_HEARTBEAT_LOST_TIMEOUT` is case sensitive.
Example	`CAS_HEARTBEAT_LOST_TIMEOUT='300'`

env.CAS_INSTALL='*install-path*'

specifies the installation directory for CAS.

Valid in	casconfig_usermods.lua file
Category	Environment
Restriction	`env.CAS_INSTALL` is case sensitive.
Examples	Here is an example showing the CAS installation directory on Linux: `env.CAS_INSTALL='/opt/sas/viya/home/SASFoundation'`
	Here is an example showing the CAS installation directory on Windows: `env.CAS_INSTALL='C:\\Program Files\\SAS\\Viya\\SASFoundation'`

env.CAS_LICENSE='*path/license-file*'

specifies the path and filename that contains the CAS license.

After CAS deployment, env.CAS_LICENSE is set to **`/opt/sas/viya/config/etc/cas/default/sas_license.txt`**. The deployment process creates a Linux symbolic link between sas_license.txt and the actual SAS license file. You must change the symbolic link whenever the name of the license file changes. For more information, see “Apply New Licenses Manually” on page 183.

Valid in	casconfig_usermods.lua file

Category	Administration
Restriction	`env.CAS_LICENSE` is case sensitive.
Examples	Here is an example showing the location of the CAS license on Linux: `env.CAS_LICENSE='/opt/sas/viya/config/etc/cas/default/SASViyaV0300_09JB84_Linux_x86-64.jwt'`
	Here is an example showing the location of the CAS license on Windows: `env.CAS_LICENSE='C:\\ProgramData\\SAS\\Viya\\etc\\cas\\default\\SASViyaV0300_09JB84_Linux_x86-64.jwt'`

env.CAS_REMOTE_HADOOP_PATH='*SASHDAT-executables-directory*'

specifies the path to the plug-in location when CAS is using an HDFS caslib to a remote HDFS cluster.

Valid in	casconfig_usermods.lua file
Category	Environment
Default	If not specified, defaults to **`$HADOOP_HOME/bin`**
Restrictions	Applies to Linux only.
	`env.CAS_REMOTE_HADOOP_PATH` is case sensitive.
Note	Might be needed to accommodate a nonstandard Hadoop plug-in.
Example	`env.CAS_REMOTE_HADOOP_PATH='$HADOOP_HOME/bin'`

env.CAS_START_MONITOR_UI='TRUE' | 'FALSE'

specifies whether CAS Server Monitor is turned on ('TRUE') or off ('FALSE').

Valid in	casconfig_usermods.lua file
Category	Administration
Default	Off ('FALSE').
Restriction	`env.CAS_START_MONITOR_UI` is case sensitive.
Note	On programming-only deployments, `env.CAS_START_MONITOR_UI` is on ('TRUE'). On full deployments, `env.CAS_START_MONITOR_UI` is off ('FALSE').
Example	In this example, CAS Server Monitor is turned on: `env.CAS_START_MONITOR_UI='true'`

env.CAS_VIRTUAL_HOST= '*host-name*'

The external host or machine name for the controller.

Use `env.CAS_VIRTUAL_HOST` when an external HTTP client needs to use an external address that differs from the actual host name known by the operating system. A common use is when the controller machine is behind a reverse proxy server.

Valid in	casconfig_usermods.lua file
Category	Network
Restriction	`env.CAS_VIRTUAL_HOST` is case sensitive.
Example	`env.CAS_VIRTUAL_HOST='my_machine'`

env.CAS_VIRTUAL_PATH='*URL-path-suffix*'

Use this environment variable when external HTTP clients must reach the CAS controller through a reverse proxy server. This identifies the path portion of the URL for the reverse proxy.

Valid in	casconfig_usermods.lua file
Category	Network
Restriction	`env.CAS_VIRTUAL_PATH` is case sensitive.
Example	`env.CAS_VIRTUAL_PATH='/cas-qstgrd-default-http'`

env.CAS_VIRTUAL_PORT=*port*

The external port number for the controller.

Use `env.CAS_VIRTUAL_PORT` when an external HTTP client needs to use a port that differs from the actual port that is local to the controller machine. A common use is when the controller machine is behind a reverse proxy server.

Valid in	casconfig_usermods.lua file
Category	Network
Restriction	`env.CAS_VIRTUAL_PORT` is case sensitive.
Example	`env.CAS_VIRTUAL_PORT=5580`

env.CAS_VIRTUAL_PROTOCOL='http | https'

Use this environment variable when external HTTP clients must reach the CAS controller through a reverse proxy server. This identifies the protocol portion of the URL for the reverse proxy.

Valid in	casconfig_usermods.lua file
Category	Network
Restriction	`env.CAS_VIRTUAL_PROTOCOL` is case sensitive.
Example	`env.CAS_VIRTUAL_PROTOCOL='https'`

env.HADOOP_HOME='*path*'

specifies the standard HADOOP_HOME variable used by Hadoop.

Valid in	casconfig_usermods.lua file
Category	Data
Restrictions	Applies to Linux only.
	`env.HADOOP_HOME` is case sensitive.
Example	`env.HADOOP_HOME='/opt/hadoop'`

env.HADOOP_NAMENODE='*machine-name* :[*machine-name*]'

identifies which machines in the Hadoop cluster are NameNodes. There can be up to two Hadoop NameNodes. Separate machine names with a colon (:). *Machine-name* can be a name, fully qualified domain name, or an IP address for a machine.

Valid in	casconfig_usermods.lua file

Category	Data
Restrictions	Applies to Linux only. `env.HADOOP_NAMENODE` is case sensitive.
Example	`env.HADOOP_NAMENODE='my_namenode1:my_namenode2'`

env.SAS_RNG_METHOD='*random-number-generator*'

specifies the random number generator (RNG) used when running CAS actions and procedures.

Examples of valid RNGs are: `'PCG'` (permuted congruential generator) or `'TF2'` (Threefry, 2x64-bit counter-based). The complete set of values for *random-number-generator* is listed under the CALL STREAMINIT Routine in the *SAS Functions and CALL Routines: Reference*.

Valid in	cas_usermods.settings file
Category	Data
Default	When env.SAS_RNG_METHOD is not specified, the default RNG is **`'MTHYBRID'`** (Hybrid 1998/2002 32-bit Mersenne twister).
Example	`env.SAS_RNG_METHOD='PCG'`

env.TKTXTANIO_BINDAT_DIR='*install-path*'

specifies the installation directory for SAS linguistic binary files required to perform text analysis.

Note: This environment variable is valid only on native operating systems such as Linux.

Note: TKTGDat.sh contains the SAS linguistic binary files required to perform text analysis in SAS LASR Analytic Server with SAS Visual Analytics and to run PROC HPTMINE and HPTMSCORE with SAS Text Miner.

Valid in	casconfig_usermods.lua or cas_usermods.settings file
Category	Data
Restriction	`env.TKTXTANIO_BINDAT_DIR` is case sensitive.
Examples	Here is an example on Linux: `env.TKTXTANIO_BINDAT_DIR='/opt/sas/viya/home/SASFoundation/utilities/TKTGDat'`
	Here is an example on Windows: `env.TKTXTANIO_BINDAT_DIR='C:\\Program Files\\SAS\\Viya\\SASFoundation\\utilities\\TKTGDat'`

LD_LIBRARY_PATH=*path* :[[*path*] …]

specifies the path to search for additional shared libraries.

Separate multiple paths with a colon (:).

CAUTION! SAS Cloud Analytic Services ignores any instance of LD_LIBRARY_PATH found in the server configuration file. Specify LD_LIBRARY_PATH in the cas_usermods.setting file only. Or, if your site uses Ansible, in vars.yml.

Valid in	cas_usermods.settings file
Category	Data
Restrictions	`LD_LIBRARY_PATH` is case sensitive.

	Applies to Linux only.
Notes	Be careful when specifying `LD_LIBRARY_PATH`. The path order can affect the operation of some applications.
	The `LD_LIBRARY_PATH` export statement must be on a single line without wrapping.
Example	`export LD_LIBRARY_PATH=/var/my_libs:/share/groups_libs:$LD_LIBRARY_PATH`

Grouped by Categories

Administration Variables

- env.CAS_LICENSE on page 539
- env.CAS_START_MONITOR_UI on page 540
- env.CASUSERIGNORECASE on page 537
- env.CASUSERLOWERCASE on page 537

Data Variables

- env.CAS_CONTROLLER_TEMP on page 536
- env.CAS_DISK_CACHE on page 537
- env.HADOOP_HOME on page 541
- env.HADOOP_NAMENODE on page 541
- env.CAS_ENABLE_REMOTE_SAVE on page 538
- env.SAS_RNG_METHOD on page 542
- env.TKTXTANIO_BINDAT_DIR on page 542
- LD_LIBRARY_PATH on page 542

Environment Variables

- env.CAS_INSTALL on page 539
- env.CAS_REMOTE_HADOOP_PATH on page 540

Network Variables

- env.CAS_VIRTUAL_HOST on page 540
- env.CAS_VIRTUAL_PATH on page 541
- env.CAS_VIRTUAL_PORT on page 541
- env.CAS_VIRTUAL_PROTOCOL on page 541

Security Variables

For information about SAS Cloud Analytic Services environment variables for:

- TLS, see "CAS TLS Environment Variables" in *Encryption in SAS Viya: Data in Motion*.
- Authentication, see "CAS_AUTH_METHOD=authinfo | kerberos" on page 335.

Server Variables

- env.CAS_ACTION_THREAD_NICE on page 535
- env.CAS_ADDITIONAL_YARN_OPTIONS on page 536
- env.CAS_CFG_CONSULPATH on page 536
- env.CONSUL_HTTP_ADDR on page 538
- env.CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT on page 538
- env.CAS_HEARTBEAT_LOST_TIMEOUT on page 539

CAS Resource Management Policies

Overview

CAS resource management policies enable you to manage disk cache used by CAS tables through three different table categories:

Table A.4 Table Categories and Policies That Manage Them

Table category	caslib type	caslib example	Policy
Global tables	Global	HPS, PUBLIC	`globalCaslibs`
Global tables	Personal	CASUSER, CASHDFS	`Priority-n`
Session tables*	Personal	CASUSER, CASHDFS	`Priority-n`

* The sum of all tables in each session. Individual caslib limits at the session level are not supported.

> **TIP** With the CAS command line interface, you can create a policy template that can serve as a starting point for creating your own CAS resource management policies. For more information, see “Create Policies from JSON Templates” on page 694.

Global caslib Policy

Because global caslibs apply to all users, global caslibs have a unique policy.

POLICY DEFINITION

```
globalCaslibs
   _ALL_          quota
   [global-caslib  quota]
   [...]
```

global-caslib
specifies a global caslib.

quota
specifies the maximum amount of disk cache space, in bytes, that *global-caslib* can use.

ALL
specifies all global caslibs.

When `_ALL_` is specified, *quota* specifies the total amount of disk cache space that can be used for all global tables, regardless of the caslibs to which the tables belong.

Example

```
globalCaslibs
   _ALL_      400000000
   HPS        200000000
   MyGlobal   100000000

 [CAS-server-priority-n
    cpu               - share
    globalCasuser     - quota
    globalCasuserHdfs - quota
    sessionTables     - quota]
 [...]

priorityAssignments
   user    priority-level
   [...]
```

In this example, the shared global caslibs HPS and MySiteGlobal can use up to 200GB and 100GB of disk cache space, respectively. The maximum amount of disk cache space that all global caslibs can use is 400GB.

Priority-Level Policies

Priority-level policies assign disk cache space quotas and CPU utilization shares based on a user's group membership, or a user's explicit assignment with the `priority-assignment` option. You can define up to five priority-level policies.

POLICY DEFINITION

```
 [CAS-server-priority-n
    cpu               - share
    globalCasuser     - quota
    globalCasuserHdfs - quota
    sessionTables     - quota]

 [...]
```

CAS-server
specifies a CAS server (for example, cas-shared-default or cas-tenant1-default.)

priority-*n*
specifies the priority level for a policy. *n* is a number 1–5.

cpu *share*
specifies the maximum amount of the CPU's capacity that all users in the associated resource management group have access to. *share* is a number that is relative to the sum of all the shares used in a CAS server's policies. The recommended practice is to define shares for all policies to total 100—thus defining percentages.

quota
specifies the maximum amount of disk cache space, in bytes, that the associated table category can use.

globalCasuser
specifies the quota for global tables that are defined in a user's personal caslib (the CASUSER caslib).

globalCasuserHdfs
specifies the quota for global tables that are defined in a user's personal caslib used on distributed server file systems such as DNFS and HDFS. (This is the CASUSERHDFS caslib.)

sessionTables

specifies the quota to be placed on session tables (for example, MyTable).

Example

```
globalCaslibs
   _ALL_         quota
   [global-caslib  quota]
   [...]

priorityLevels
   cas-shared-default-priority-1
      cpu                - 50
      globalCasuser      - 500000000
      globalCasuserHdfs - 500000000
      sessionTables      - 500000000
   cas-shared-default-priority-2
      cpu                - 20
      globalCasuser      - 50000000
      globalCasuserHdfs - 50000000
      sessionTables      - 50000000
   cas-shared-default-priority-3
      cpu                - 30
      globalCasuser      - 10000000
      globalCasuserHdfs - 10000000
      sessionTables      - 10000000

priorityAssignments
   user    priority-level
   [...]
```

In this example, three out of a maximum of five policies are defined for the CAS server, cas-shared-default. The policy, `cas-shared-default-priority-1`, applies to CAS users who are members of the user group with the same name, or to CAS users who are explicitly assigned with the `priority-assignment` option. CAS users for which the `cas-shared-default-priority-1` applies to can use up to 500GB of disk cache space in their personal caslibs for shared global caslibs and shared global caslibs across a distributed server file system (such as DNFS or HDFS). These CAS users can use up to 500GB disk cache space for session tables. CAS sessions running under the `cas-shared-default-priority-1` policy are limited to a 50% share of the CAS server's CPU capacity.

Priority Assignments

Users that are not a member of a resource management user group can be explicitly assigned to a priority-level policy.

POLICY DEFINITION

```
priorityAssignments
   user     priority-level
   [...]
```

user

specifies a CAS user ID that is not a member of a CAS resource management user group (for example, `cas-shared-default-priority-1`).

> **TIP** Instead of a user ID, *user* can be an asterisk (*), that includes any user IDs that do not match any of the specified user IDs or CAS resource management group names.

priority *priority-level*

specifies a priority level that maps to a specific policy. *priority-level* is a number 1–5.

For example, if 4 is specified, *user* is assigned to the `cas-shared-default-priority-4` policy.

Important: If a priority level is specified for which there is no corresponding policy defined, the user has no policy assignment.

Example

```
    globalCaslibs
       _ALL_         quota
       [global-caslib  quota]
       [...]

     [CAS-server-priority-n
       cpu               - share
       globalCasuser     - quota
       globalCasuserHdfs - quota
       sessionTables     - quota]
     [...]

    priorityAssignments
       userA     1
```

In this example, three out of a maximum of five policies are defined for the CAS server, cas-shared-default. UserA is not a member of any of the CAS resource management user groups (for example, `cas-shared-default-priority-1`). Therefore, when UserA starts a CAS session, the `cas-shared-default-priority-1` policy contains the resource definitions.

See Also

"CAS Resource Management"

CAS Configuration and Start-up Logging

How Do I Use Configuration and Start-up Logging?

In addition to its normal server logging, CAS can also log when it processes its configuration and start-up files. This logging feature is similar to the SAS 9 Logging Facility, `log4sas`, and uses the `App.cas.config` logger. When you customize an existing usermods Lua script or write your own Lua script, you can include logging functions to write records to the standard CAS server log files in **`/var/log/sas/viya/cas/default/`**or to custom log files that you define.

Note: For information about managing CAS logging after configuration and start-up file processing, see "Logging: How To" on page 192.

CAS writes configuration and start-up logging records with the following log functions (documented later):

- `log.level='level'`
- `log.file='custom-log-filename'`
- `log.level('custom-message')`

You can add the log functions in the following locations:

- casconfig_usermods.lua
- casstartup_usermods.lua
- Lua scripts in the **`conf.d`** directory
- Lua scripts in the **`start.d`** directory

For more information, see "Standard Configuration Files".

Configuration and Start-up Logging Reference

log.level='[trace | debug | info | warn | error | fatal] | [1 | 2 | 3 | 4 | 5 | 6]'

specifies the lowest logging level for configuration and start-up file logging to write messages in the log file during CAS configuration and start-up file processing. The lowest level is `trace` or `1` (most verbose). The highest level is `fatal` or `6` (least verbose).

Important: The `log.level` function only impacts CAS configuration and start-up logging and does not impact standard CAS server logging.

Default log.level='trace'

Restriction `log.level` is case sensitive.

Examples In this example, log messages that are at the WARN (4) level and higher are written to the log file:

```
log.level='4'
```

In this example, **`log.level`** is used to obtain the current logging level and as an input to the **`log.all('`*`custom-message`*`')`** function. The Lua string concatenation operator, '..' (two dots) is also used:

```
log.all('The current log level is: '..log.level)
```

The following text is written to the log:

```
The current log level is: ERROR
```

log.file='*custom-log-filename*' | '+*custom-log-filename*'

specifies the absolute path and custom log filename in which to write log messages during CAS configuration and start-up file processing. Using a plus sign (+) means that the log messages are appended to an existing log file.

Important: The `log.file` function only impacts CAS configuration and start-up logging and does not impact standard CAS server logging.

> **TIP** You can use various Lua functions in the `log.file` and `log.level('custom-message')` functions. Also, the log functions have built-in tags that, when used, write the CAS server process ID (%P) and the process owner ID (%U). For more information, see the examples that follow and https://www.lua.org/pil/22.1.html.

Restriction `log.file` is case sensitive.

Note When `log.file` is not used, CAS writes configuration and start-up logging to the CAS server log file in `/var/log/sas/viya/cas/default/`.

Examples In this example, CAS configuration and start-up file logging is appended to a pre-existing log file, as indicated with a plus sign (+):

```
log.file='+/var/log/sas/viya/cas/default/config_startup.log'
```

In this example, CAS configuration and start-up file logging is written to a unique file that contains the current date and the CAS server process ID. The date portion of the log filename is obtained using the Lua **`os.date`** function with date tags (**`%Y`**, **`%m`**, **`%d`**). The CAS server process ID is obtained with a built-in **`log.file`** function tag, **`%P`**. The date and process ID are joined using the Lua string concatenation operator, '..' (two dots):

```
log.file=os.date('+/var/log/sas/viya/cas/default/%Y-%m-%d_controller-1_') .. '_%P.log'
```

Here is an example of the resulting log filename: **`2018-06-21_controller-1_11634.log`**

log.[all | trace | debug | info | warn | error | fatal]('*custom-message*')

writes a custom message (a string) to the log file at the specified level during CAS configuration and start-up file processing. When `ALL` is specified, the current logging level is ignored and the custom message is always written to the log file.

> **TIP** You can use various Lua functions in the `log.level('custom-message')` and `log.file` functions. Also, the log functions have built-in tags that, when used, write the CAS server process ID (%P) and the process owner ID (%U). For more information, see the examples that follow and https://www.lua.org/pil/22.1.html.

Restriction `log.level` is case sensitive.

Examples In this example, the following custom message is written to the log file when the current logging level is set to INFO (3) or lower:

```
log.info('CAS writes this message when the current log level is INFO or lower.')
```

Therefore, any higher level custom messages (WARN, ERROR, and FATAL) are also written to the log. But, any lower level custom messages (TRACE or DEBUG) are not written to the log.

In this example, the following custom message is written to the log, using the Lua **`os.time`** function to obtain the time at which the custom message is logged:

```
log.debug(os.time('Debug message at time: %X'))
```

Here is an example of what is written to the log:

```
Debug message at time: 14:22:38
```

In this example, the following custom message is always written, regardless of the log level:

```
log.all('This is a message written at any logging level.')
```

gridmon.sh Commands

Overview

This section describes commands that you can use to operate gridmon.sh. For usage information, see “Use gridmon.sh (Linux)”.

This section is organized into these sub-sections:

- “Global Commands ”
- “Job Mode Commands ”
- “Machine Mode Commands ”
- “Disk Mode Commands ”
- “Show Ranks Menu Commands ”
- “Show Details Menu Commands ”
- “Details Menu Commands ”

Global Commands

Note: Menu options that produce lengthy results redirect the output to your vi editor. Closing vi returns to gridmon.

Table A.5 Global Commands

Command	Description
q	Exits gridmon.sh.
Up and Down arrows Page Up and Page Down	Moves through the list of jobs, machines, or disks.
Backspace Escape	Cancels current menu, prompt, or sub-mode.
?	Shows help information for gridmon.sh.

For usage information, see “Use gridmon.sh (Linux)”.

Job Mode Commands

Table A.6 Job Mode Commands

Command	Description
j	Runs gridmon.sh in job mode.
Left and Right arrows	Changes the column for sorting the list.
h Home	Moves to the top of the list.
Enter	Shows the menu option for the selected job.

For usage information, see “Use gridmon.sh (Linux)”.

Machine Mode Commands

Table A.7 Machine Mode Commands

Command	Description
m	Runs gridmon.sh in machine mode.
Enter	Shows menu options for the selected machine

For usage information, see “Use gridmon.sh (Linux)”.

Disk Mode Commands

Table A.8 *Disk Mode Commands*

Command	Description
d	Runs gridmon.sh in disk mode.
Enter	Shows selected disk use on machines where the disk is present.

For usage information, see “Use gridmon.sh (Linux)”.

Show Ranks Menu Commands

When gridmon.sh is in job mode, you display the **Show Ranks** menu by pressing `Enter` from the main window.

Table A.9 *Show Ranks Menu Commands*

Command	Description
Show Ranks	Displays all the ranks belonging to the job and the machines on which they are running.
Kill job	Kill the selected job.
Kill jobs with user: ***user-ID***	Kills all jobs of the selected user.
Kill jobs with user: ***user-ID*** **ID:** ***process-ID***	Kills all jobs of the selected user and specific ID.
Kill jobs at least this old	Kills all jobs at least as old as the selected job.
Stack Trace all Ranks	Runs the gstack application on all processes in this job and collects results. gstack displays its results in your vi editor.

For usage information, see “Use gridmon.sh (Linux)”.

Show Details Menu Commands

When gridmon.sh is in job mode, you display the **Show Details** menu when you press `Enter` from the Ranks window (**Enter** ⇨ **Show Ranks**).

Table A.10 Show Details Menu Commands

Command	Description
Show Details	Shows process ID, CPU use, virtual memory, and if not zero, the following fields: ■ **DFSSize**: Disk space in CAS_DISK_CACHE owned by the current process. ■ **HDFSSize**: Disk space mapped from HDFS. ■ **DNFSSize**: Disk space mapped from DNFS. ■ **Global FSSize**: Disk space in CAS_DISK_CACHE for global tables, owned by the main server process. ■ **CGroup Limit**: Size of memory cgroup, as specified by cas.MEMORYSIZE. ■ **CGroup Usage**: Amount of the CGroup memory that is in use by all processes belonging to this server on the current machine. ■ **Faults/s**: The number of page faults per second for the process, most commonly caused by paging in table data. (Faults can help you determine whether the process is paging.)
Kill Rank	Kills the selected rank or process.
Stack Trace	Runs the gstack application on all processes in this job and collects results. gstack displays its results in your vi editor.
Process Limits	Displays the contents of **/proc/*pid*/limits**.
FileHandle Count	Counts the files owned by the process.
FileHandle List	Lists the files owned by the process.
Environment	Displays the process's environment handles from **/proc/*pid*/environ**.
List Memory Maps	Shows the process's memory maps from **/proc/*pid*/maps**.
Numa Stats	Shows the output from the Linux `numastat` command for this process.
Show CGroups	Shows the Linux cgroups that this process belongs to.
Xterm*	Starts an Xterm on the selected machine.
Perf Top*	Runs the perf top application on this process in a new Xterm window. **Note:** The perf package must be installed.

Command	Description
Attach Debugger*	Attaches a debugger to the running process. Requires a new X window. Note: Attach Debugger is for use only when directed by SAS Technical Support or by SAS R&D.

* Requires that an X Server be running on the CAS controller machine.

For usage information, see "Use gridmon.sh (Linux)".

Details Menu Commands

When gridmon.sh is in machine mode, you display the **Details** menu by pressing `Enter` from the main window.

Table A.11 Details Menu Commands

Command	Description
Details	Displays information about the machine such as CPU utilization, free memory, and total memory.
Top	Runs the top application on all processes in this job and collects results. Top displays its results in your vi editor.
Xterm*	Starts an Xterm on the selected machine.
Perf Top*	Runs the perf top application on this process in a new Xterm window.

* Requires that an X Server be running on the CAS controller machine.

For usage information, see "Use gridmon.sh (Linux)".

SAS Cloud Analytic Services: Interfaces

There are several interfaces that you can use to administer a CAS server. The following table lists these interfaces and the shading indicates the relative amount of CAS administration that each covers:

Table A.12 Interfaces to CAS Administration

◔	CAS Server Properties action set	A programmatic interface for CASL (the CAS procedure), Python, Lua, and R. Used to display server option values.
◔	Ansible	A software orchestration tool that provides a straightforward approach to deploying and provisioning SAS Viya.
◒	Administrative scripts	Scripts used to operate CAS server, change the process owner account, and to convert from single- to multi-machine CAS.
◕	Command-line interface	A command-line interface that enables you to perform CAS administration.

◕	SAS Environment Manager	A graphical enterprise web application used to modify and view a subset of server properties and to adjust caslib management privileges.
●	CAS Server Monitor	A graphical web application that is embedded in the CAS server. Used to view server information and to manage sessions, nodes, and caslib management privileges.

28

SAS Server Contexts

SAS Viya Server Contexts: Overview

Note: A programming-only on page 17deployment does not use server contexts.

To learn about SAS Viya server contexts, see “Server Contexts: Concepts” on page 557.

To create server contexts, see “Server Contexts: How To” on page 555.

Server Contexts: How To

Introduction

These instructions explain how to view and modify server contexts using SAS Environment Manager.

Navigation

In the applications menu (≡), under **Administration**, select **Manage Environment**. In the navigation bar, click 🔧.

The Contexts page is an advanced interface that is available to SAS Administrators only. If you are a SAS Administrator and the Contexts page is unavailable to you, then the programming run-time servers have not been deployed in your SAS Viya environment.

Create a Context

1. Under **View**, select the type of context that you want to create.

2. On the top left side of the Contexts page, click ✱.

3. (Required) Enter a name for your context.

   Names must not be longer than 40 characters and can consist of any alphanumeric and special characters.

4. If you are creating a compute context, skip to Step 7. Otherwise, enter values for creating a launcher context:

   - **Description**

     Enter a description of the context that you are creating.

   - **Port Range**

     Enter a range of ports. The SAS Launcher Server selects a port in the specified range in order to run the compute server.

   - **Environment Variables**

     Click + and add the environment variable and its value that you want the launcher server to use when running the compute server.

5. Click **Advanced** to override the default server deployment settings that are used by the launcher service at launch time. Next, check the box and provide values for all the fields.

6. When you are finished, click **Save** to create the launcher context

7. Enter values for creating a compute context:

   - **Description**

     Enter a description of the context that you are creating.

   - (Required) **Launcher context**

     Select a launcher context with which to run the SAS Compute Server.

   - (Required) **Identity type**

     Select one of the following:

     - Select **Authenticated users** in order for any authenticated user to use this context.
     - Select **Identities** by clicking 👤, and then select one or more users or groups to use this context.

8. To add any SAS options or additional autoexec file settings that the compute server processes use at start-up, select **Advanced** and enter this information in their respective fields.

   For more information, see Customizing Your SAS Session By Using Configuration and Autoexec Files.

9. When you are finished, click **Save** to create the compute context.

Edit a Context

1. Under **View**, select the type of context that you want to edit.

2. On the top right side of the Contexts page, click ✎ or right-click on the selected context and select **Edit** .

3 Make your modifications, and click **Save** when you are finished.

Delete a Context

1 Under **View**, select the type of context that you want to delete.

2 On the left side of the Contexts page, click 🗑 or right-click on the selected context and select **Delete** .

3 Click **Yes** to confirm the deletion.

View a Context

1 Under **View**, select the type of context that you want to view.

2 The **Basic Properties** are displayed on the right side of the **Contexts** page.

Server Contexts: Concepts

Compute Contexts

A SAS Compute Server is run under a compute context. (Contexts are analogous to SAS 9 SAS Application Servers.) A *compute context* is a specification that contains the information that is needed to run a compute server.

The information that is contained in a compute context is the user identity and any SAS options or autoexec file parameters to be used when starting the server.

Launcher Contexts

The server that starts the compute server, SAS Launcher Server, itself requires a context.

A *launcher context* is a specification that enables SAS administrators to apply environmental and access constraints on processes that are run by a launcher server.

In the following example, the Launcher context (on the right) contains a range of ports for server access and an environment variable that defines the location of the disk cache. The Compute context (on the left) includes the launcher context settings that are applied automatically when the launcher starts the compute server.

Figure A.1 Context Types

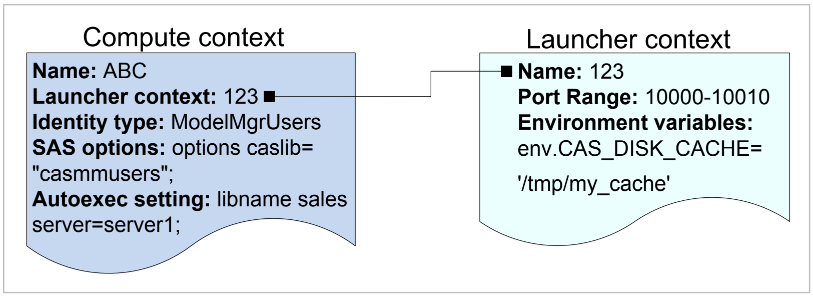

Name: ABC
Launcher context: 123 ■
Identity type: ModelMgrUsers
SAS options: options caslib= "casmmusers";
Autoexec setting: libname sales server=server1;

■ **Name:** 123
Port Range: 10000-10010
Environment variables: env.CAS_DISK_CACHE= '/tmp/my_cache'

SAS Viya Server Contexts: Interfaces

SAS Viya server contexts can be managed using either of two interfaces. The following table lists these interfaces and the shading indicates the relative amount of SAS server contexts administration that each covers:

Table A.1 *Interfaces to SAS Viya Server Contexts*

◕	SAS Environment Manager	Graphical enterprise web application
●	Command-line interface	Command-line interface

29

Programming Run-Time Servers

Programming Run-Time Servers: Overview

A programming run-time environment includes several SAS Viya servers. The following table lists the servers (and services, where applicable) and indicates which are available in a programming-only on page 17deployment:

Server	Full deployment	Programming-only deployment
“SAS Compute Server and Compute Service”	✓	
“SAS Launcher Server and Launcher Service”	✓	
“SAS Workspace Server and SAS Object Spawner”	✓	✓
“Embedded Web Application Server”	✓	✓
“SAS/CONNECT Server and SAS/CONNECT Spawner”	✓	✓

In the following diagram, the highlighted box shows the relationship of the programming run-time servers to other components in the SAS Viya environment in a full deployment:

Figure A.1 *SAS Viya Programming Run-Time Servers (Full Deployment)*

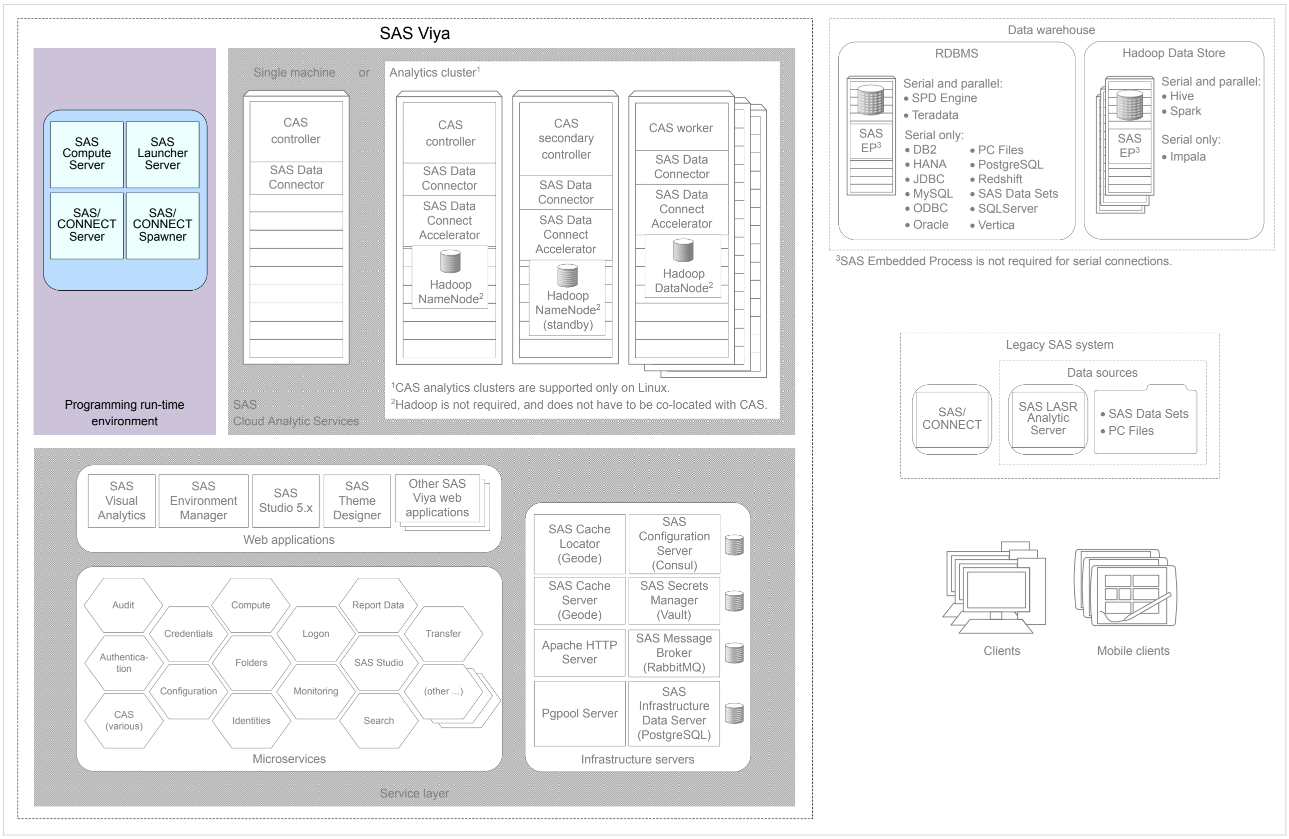

In the following diagram, the highlighted box shows the relationship of the programming run-time servers to other components in the SAS Viya environment in a programming-only deployment:

Figure A.2 *SAS Viya Programming Run-Time Servers (Programming-Only Deployment)*

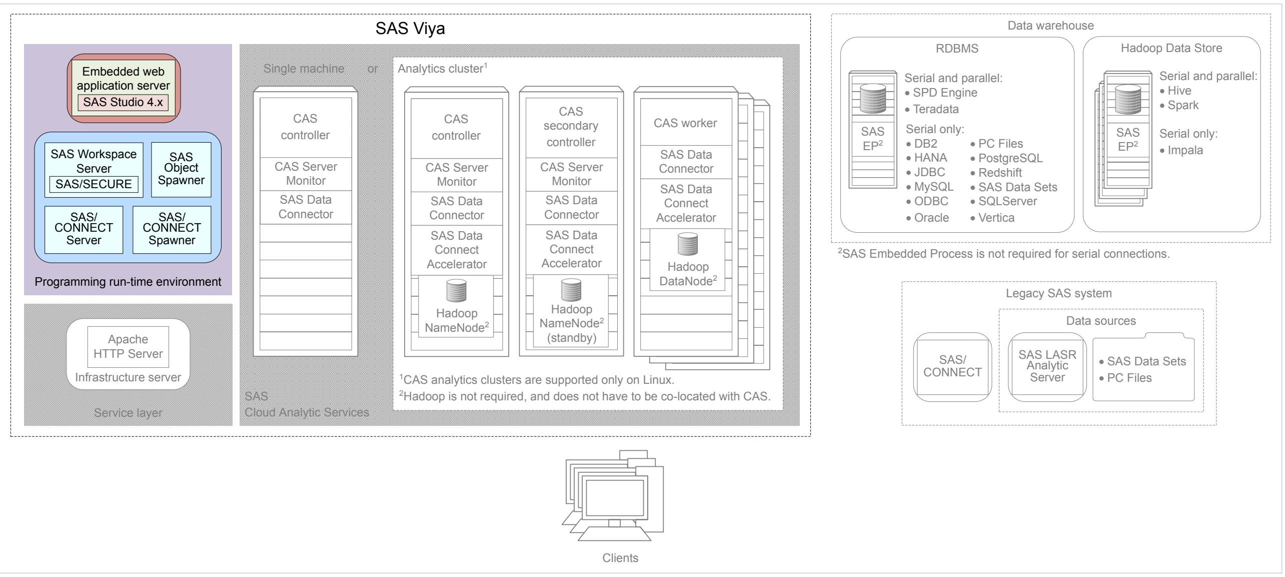

SAS Compute Server and Compute Service

Overview

The Compute service enables clients to submit SAS programs and stored procedures in the form of jobs for processing. The SAS Compute Server implements the Compute service. For more information, see "Concepts" on page 564.

Operate the Compute Service (Linux)

SAS Viya provides a script in **/etc/init.d** that you use to stop, start, restart, and check the status of the compute service. The script is named `sas-viya-compute-default`.

Syntax

How you run `sas-viya-compute-default` depends on your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl status | stop | start | restart sas-viya-compute-default

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-compute-default status | stop | start | restart

Usage Notes and Tips

- You must be logged on to the machine where the compute service resides. Also, you must have sudo privileges to run this script.
- On multi-tenant SAS Viya systems, the script is named `sas-tenant-ID-sas-viya-compute-default`.
- There is another script that you can use to manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" on page 462.

  Note: There is a sequence for starting and stopping SAS Viya servers and services. You must follow this sequence to avoid operational issues. For more information, see "Read This First: Start and Stop Servers and Services" on page 458 .

- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-compute-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call do not use.

CAUTION! On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, do not mix the System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-compute-default` on Red Hat Enterprise Linux 7.*x* with the `service` command, and later attempt to shut down the compute service using the `systemctl` command, the compute service stops responding and does not shut down.

Examples

- To check status of the compute service on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-compute-default
```

- To stop the compute service on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-compute-default stop
```

- To start the compute service on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-compute-default
```

- To restart the compute service on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-compute-default restart
```

Operate the Compute Service (Windows)

Using the Microsoft Management Console (MCC) Services Snap-In, you can start, stop, and restart SAS Compute Service.

Figure A.3 *SAS Compute Service in the Services Snap-In*

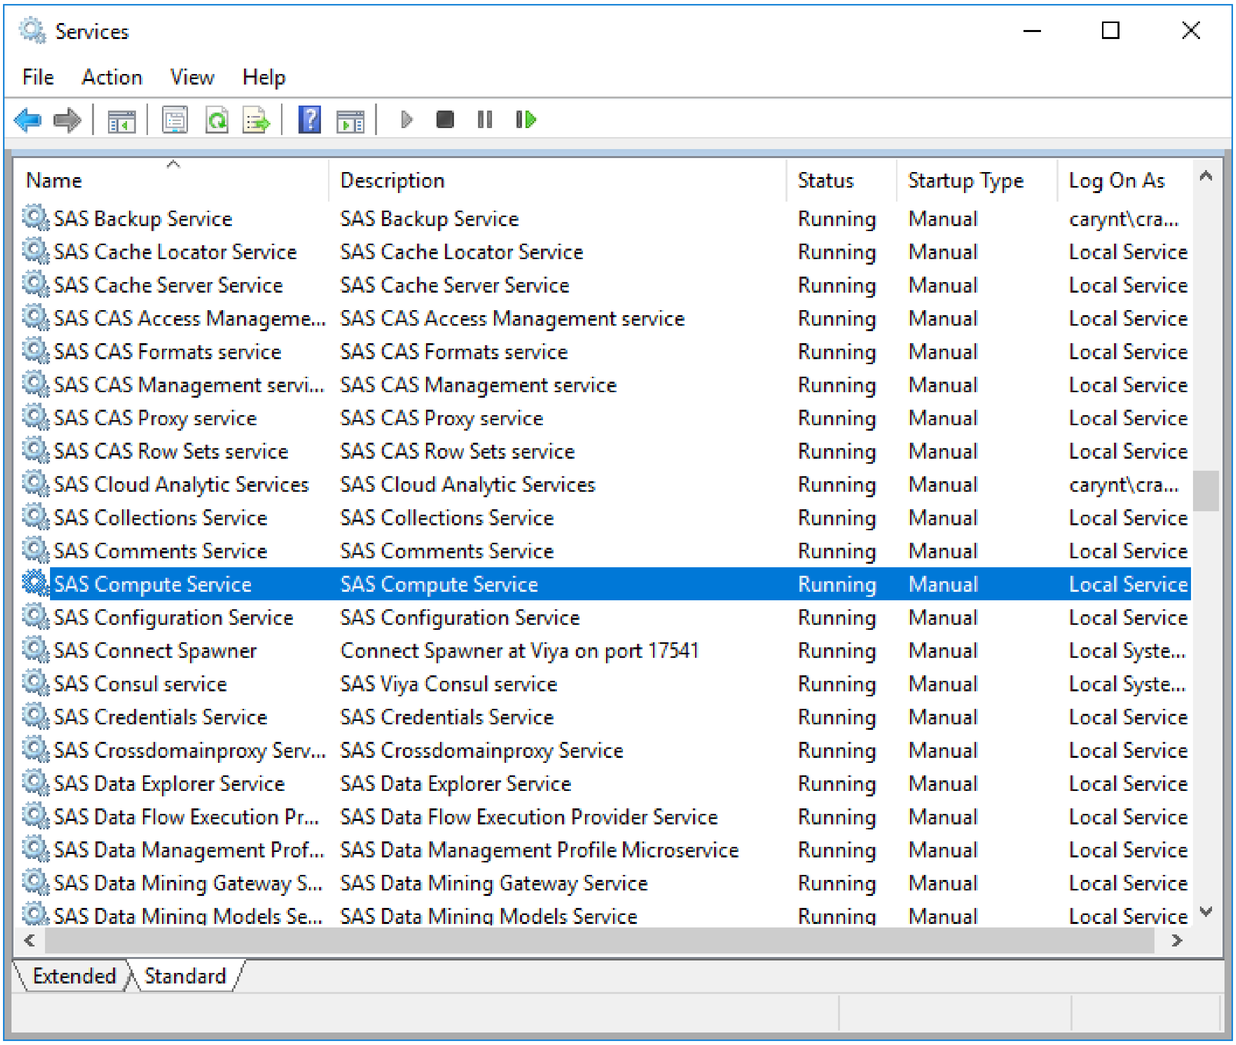

Name	Description	Status	Startup Type	Log On As
SAS Backup Service	SAS Backup Service	Running	Manual	carynt\cra...
SAS Cache Locator Service	SAS Cache Locator Service	Running	Manual	Local Service
SAS Cache Server Service	SAS Cache Server Service	Running	Manual	Local Service
SAS CAS Access Manageme...	SAS CAS Access Management service	Running	Manual	Local Service
SAS CAS Formats service	SAS CAS Formats service	Running	Manual	Local Service
SAS CAS Management servi...	SAS CAS Management service	Running	Manual	Local Service
SAS CAS Proxy service	SAS CAS Proxy service	Running	Manual	Local Service
SAS CAS Row Sets service	SAS CAS Row Sets service	Running	Manual	Local Service
SAS Cloud Analytic Services	SAS Cloud Analytic Services	Running	Manual	carynt\cra...
SAS Collections Service	SAS Collections Service	Running	Manual	Local Service
SAS Comments Service	SAS Comments Service	Running	Manual	Local Service
SAS Compute Service	SAS Compute Service	Running	Manual	Local Service
SAS Configuration Service	SAS Configuration Service	Running	Manual	Local Service
SAS Connect Spawner	Connect Spawner at Viya on port 17541	Running	Manual	Local Syste...
SAS Consul service	SAS Viya Consul service	Running	Manual	Local Syste...
SAS Credentials Service	SAS Credentials Service	Running	Manual	Local Service
SAS Crossdomainproxy Serv...	SAS Crossdomainproxy Service	Running	Manual	Local Service
SAS Data Explorer Service	SAS Data Explorer Service	Running	Manual	Local Service
SAS Data Flow Execution Pr...	SAS Data Flow Execution Provider Service	Running	Manual	Local Service
SAS Data Management Prof...	SAS Data Management Profile Microservice	Running	Manual	Local Service
SAS Data Mining Gateway S...	SAS Data Mining Gateway Service	Running	Manual	Local Service
SAS Data Mining Models Se...	SAS Data Mining Models Service	Running	Manual	Local Service

See “General Servers and Services: Operate (Windows)” on page 463 for details.

Lock Down the SAS Compute Server

Using the “LOCKDOWN System Option” on page 589 and the “LOCKDOWN Statement” on page 590, you can limit access to files and to specific SAS features in a SAS Compute Server session in a multi-tenant environment.

To lock down a compute server:

1 With administrator privileges, log on to the machine that contains the compute server.

2 By default, SAS adds certain predefined paths from the SAS configuration file to the lockdown path list (a whitelist that contains all the paths that are accessible to the compute server). To add more paths to the lockdown whitelist, go to `/opt/sas/viya/config/etc/compsrv/default/autoexec_usermods.sas` and add the lockdown path statements. For more information, see LOCKDOWN Statement Details on page 592.

   Note: For Windows, add the lockdown path list to `\ProgramData\SAS\Viya\etc\compsrv\default\autoexec_usermods.sas`.

   Note: A path that is declared in the whitelist does not mean that an arbitrary user can read any file in that path. Host permissions on physical files and directories always take precedence over the whitelist.

   Changes to the autoexec_usermods.sas file are automatically included when the compute server scripts run. Your changes will take effect the next time SAS starts a compute server session.

   TIP For a suggestion about how to implement the whitelist, see "Example 2: Hiding the Whitelist By Locating the Path outside the Whitelist" on page 593.

3 To enable lockdown, set the environment variable COMPUTESERVER_LOCKDOWN_ENABLE to 1 in the sysconfig file `/opt/sas/viya/config/etc/sysconfig/compsrv/default/sas-compsrv`.

   Setting this variable enables the -lockdown option in the start-up script.

   Note: You can change the default path during installation and install to a different folder.

4 If your site uses SAS Studio, set webdms.showSystemRoot=false.

   For more information, see "Update SAS Studio Configuration Properties" on page 80.

Concepts

SAS Compute Server

The SAS Compute Server enables clients to submit SAS programs and stored procedures in the form of jobs for processing using the SAS language. For every job that is processed, the compute server writes a logging message to a SAS log. If the job produces ODS results, output data sets, files, and so on, the output is associated with the job.

Note: The SAS Compute Server does not support X commands, which enable execution of operating system commands from within SAS.

Compute servers are launched by a SAS Launcher Server.

Compute Service

The Compute service is a SAS Viya microservice that provides API endpoints for requesting a SAS Compute Server session. The compute service also provides API endpoints for creating and managing compute contexts, specifications that contain all the information that is needed to run a compute server.

The launcher service provides a specification to the launcher server called a launcher context, that enables the SAS administrator to apply constraints for how the launcher server starts a compute server.

How It Works

The following figure describes how a SAS client submits code to the SAS Compute Server.

Figure A.4 How a Client Submits Code to the SAS Compute Server

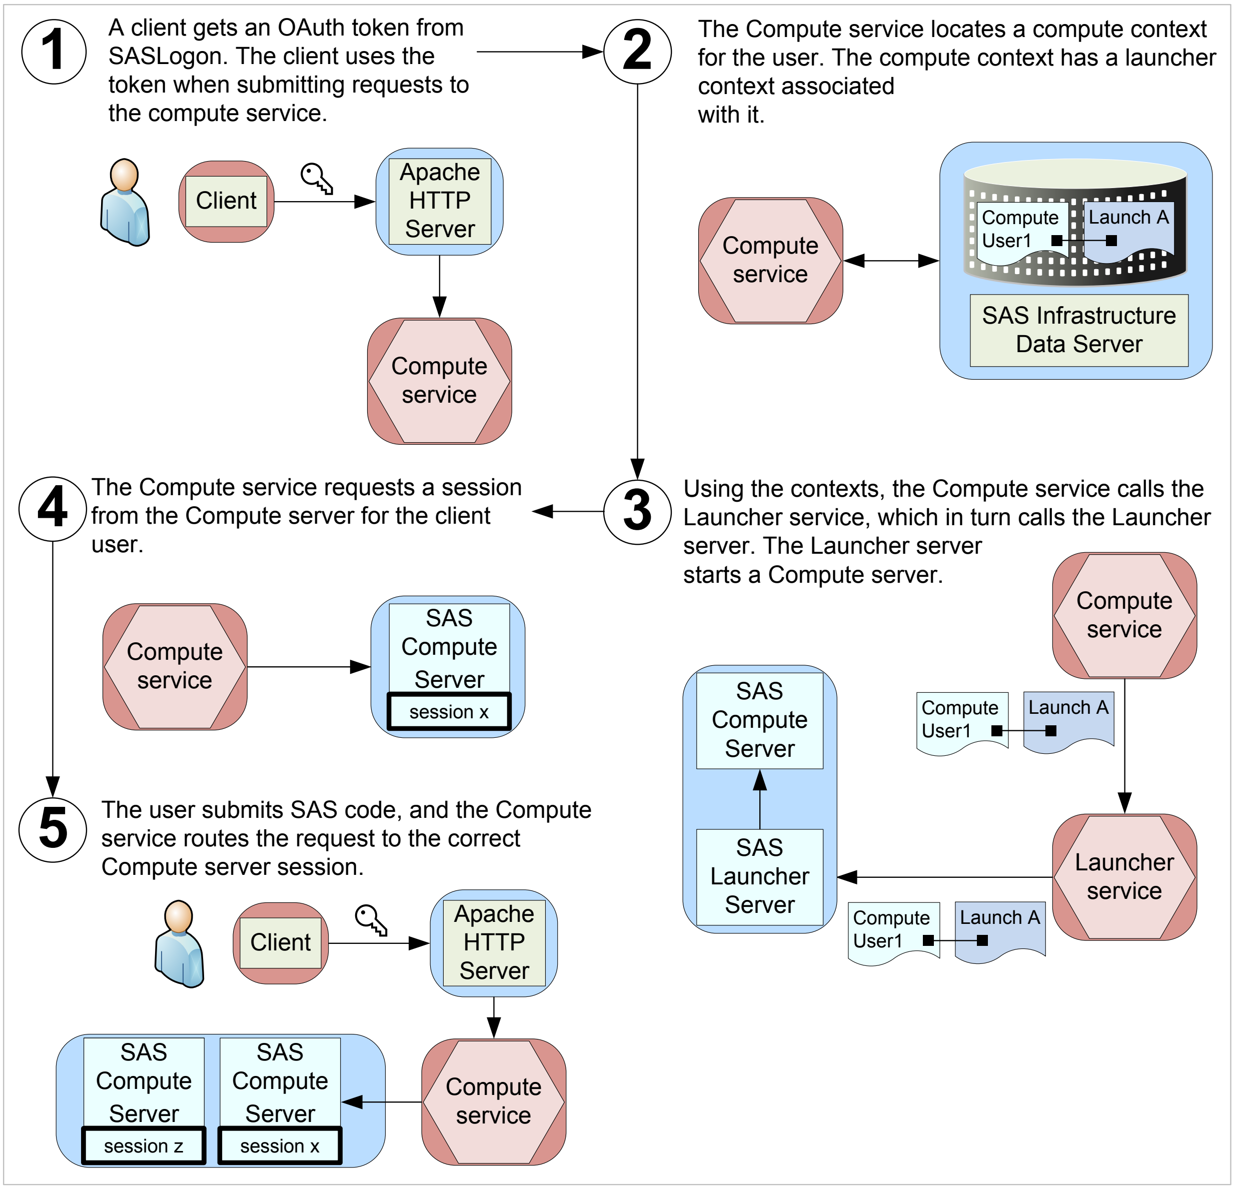
1 A client gets an OAuth token from SASLogon. The client uses the token when submitting requests to the compute service.

2 The Compute service locates a compute context for the user. The compute context has a launcher context associated with it.

3 Using the contexts, the Compute service calls the Launcher service, which in turn calls the Launcher server. The Launcher server starts a Compute server.

4 The Compute service requests a session from the Compute server for the client user.

5 The user submits SAS code, and the Compute service routes the request to the correct Compute server session.

Fault Tolerance

You are able to deploy SAS Compute Servers for fault tolerance. You can deploy multiple SAS Launcher Servers on multiple compute server machines, and the Launcher service randomly routes client requests among the registered Launcher servers.

Only machine-level fault tolerance is supported. If a machine goes down, and you have other machines running Launcher and Compute servers, then fault tolerance is applied. If an individual Launcher or Compute Server process abnormally terminates, then no fault tolerance is applied.

Log Files

Compute service

Log files for the compute service are located in `/opt/sas/viya/config/var/log/compute/default`.

On multi-tenant systems, log files for the compute service are located in **/opt/sas/*tenant-ID*/config/var/log/compute/default**.

In Windows, the log files for the compute service are located in **\ProgramData\SAS\Viya\var\log\compute\default**.

Compute Server

Compute servers and their logs are located where the launcher servers are running. Each compute server generates its own log. Log files are owned by the account under which the server was launched. This is useful in locating the file for a specific user.

Log files for the compute server are located in **/opt/sas/viya/config/var/log/compsrv/default**.

On multi-tenant systems, log files for the compute server are located in **/opt/sas/*tenant-ID*/config/var/log/compsrv/default**.

In Windows, the log files for the compute server are located in **\ProgramData\SAS\Viya\var\log\compsrv\default**.

SAS Launcher Server and Launcher Service

Overview

The SAS Launcher Server runs processes in a SAS Viya environment. The Launcher service is a SAS Viya microservice that provides API endpoints for how the launcher server runs a process.

Operate the Launcher Service (Linux)

SAS Viya provides a script in **/etc/init.d** that you use to stop, start, restart, and check the status of the launcher service. The script is named, `sas-viya-launcher-default`.

Syntax

How you run `sas-viya-launcher-default` depends on your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl status | stop | start | restart sas-viya-launcher-default
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-launcher-default status | stop | start | restart
  ```

Usage Notes and Tips

- You must be logged on to the machine where the launcher service resides. Also, you must have sudo privileges to run this script.
- On multi-tenant SAS Viya systems, the script is named `sas-tenant-ID-sas-viya-launcher-default`.
- There is another script that you can use to manage and view the running state of all SAS Viya services. For more information, see “Start and Stop All Servers and Services” on page 462.

  Note: There is a sequence for starting and stopping SAS Viya servers and services. You must follow this sequence to avoid operational issues. For more information, see “Read This First: Start and Stop Servers and Services” on page 458 .

- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-launcher-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-launcher-default` on Red Hat Enterprise Linux 7.**x** with the `service` command, and later attempt to shut down the launcher server using the `systemctl` command, the launcher server stops responding and does not shut down.

- The launcher server and launcher service support Kerberos on Linux. For more information, see "Configure Kerberos for SAS Launcher Server" on page 294.

Examples

- To check status of the launcher service on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-launcher-default
```

- To stop the launcher service on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-launcher-default stop
```

- To start the launcher service on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-launcher-default
```

- To restart the launcher service on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-launcher-default restart
```

Operate the Launcher Service (Windows)

Using the Microsoft Management Console (MCC) Services Snap-In, you can start, stop, and restart SAS Launcher Service.

Figure A.5 SAS Launcher Service in the Services Snap-In

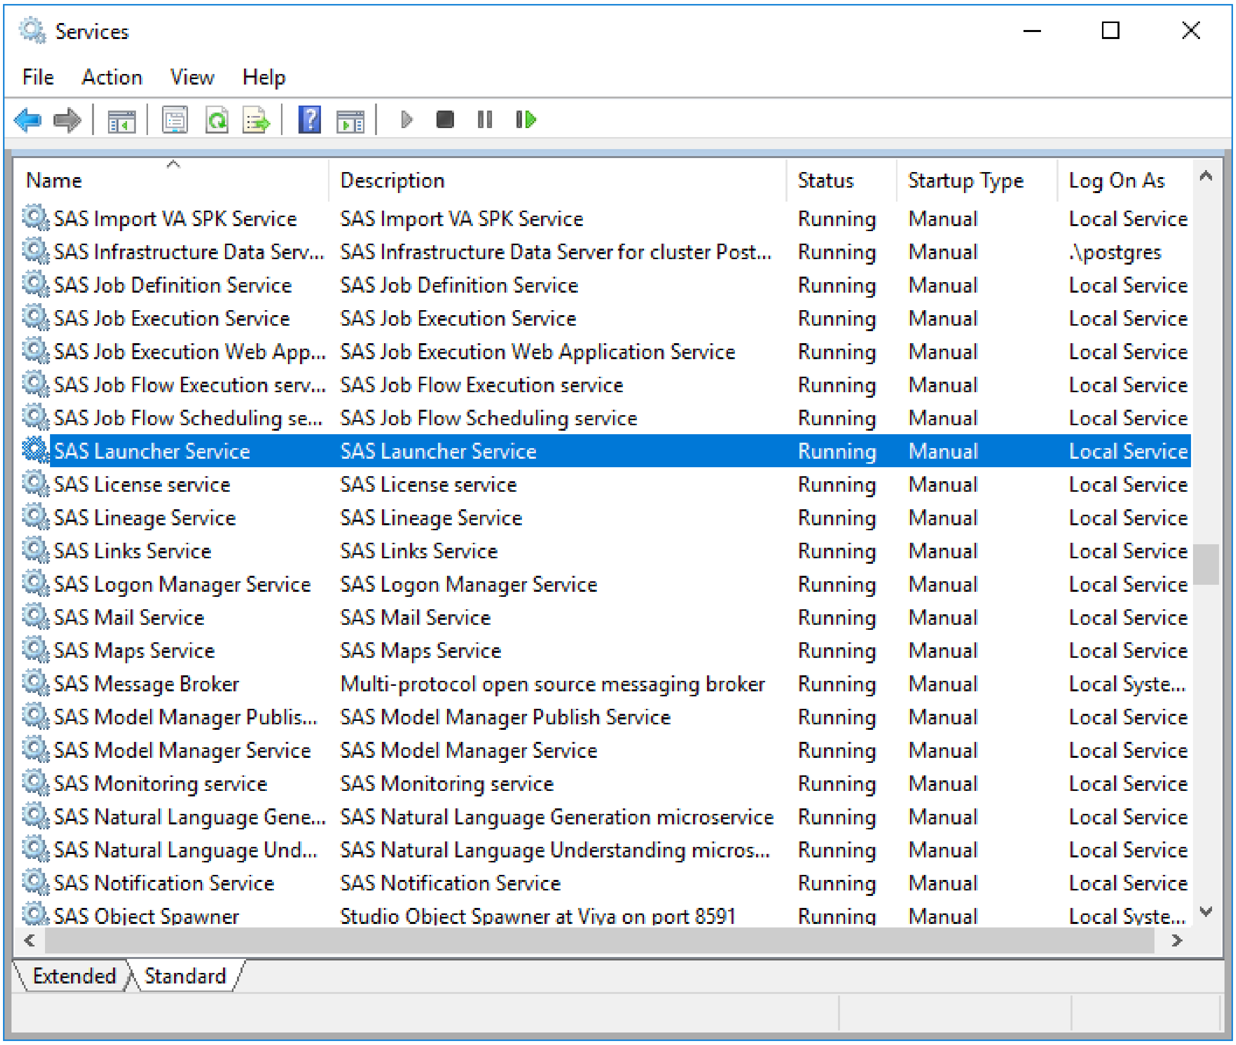

Name	Description	Status	Startup Type	Log On As
SAS Import VA SPK Service	SAS Import VA SPK Service	Running	Manual	Local Service
SAS Infrastructure Data Serv...	SAS Infrastructure Data Server for cluster Post...	Running	Manual	.\postgres
SAS Job Definition Service	SAS Job Definition Service	Running	Manual	Local Service
SAS Job Execution Service	SAS Job Execution Service	Running	Manual	Local Service
SAS Job Execution Web App...	SAS Job Execution Web Application Service	Running	Manual	Local Service
SAS Job Flow Execution serv...	SAS Job Flow Execution service	Running	Manual	Local Service
SAS Job Flow Scheduling se...	SAS Job Flow Scheduling service	Running	Manual	Local Service
SAS Launcher Service	SAS Launcher Service	Running	Manual	Local Service
SAS License service	SAS License service	Running	Manual	Local Service
SAS Lineage Service	SAS Lineage Service	Running	Manual	Local Service
SAS Links Service	SAS Links Service	Running	Manual	Local Service
SAS Logon Manager Service	SAS Logon Manager Service	Running	Manual	Local Service
SAS Mail Service	SAS Mail Service	Running	Manual	Local Service
SAS Maps Service	SAS Maps Service	Running	Manual	Local Service
SAS Message Broker	Multi-protocol open source messaging broker	Running	Manual	Local Syste...
SAS Model Manager Publis...	SAS Model Manager Publish Service	Running	Manual	Local Service
SAS Model Manager Service	SAS Model Manager Service	Running	Manual	Local Service
SAS Monitoring service	SAS Monitoring service	Running	Manual	Local Service
SAS Natural Language Gene...	SAS Natural Language Generation microservice	Running	Manual	Local Service
SAS Natural Language Und...	SAS Natural Language Understanding micros...	Running	Manual	Local Service
SAS Notification Service	SAS Notification Service	Running	Manual	Local Service
SAS Object Spawner	Studio Object Spawner at Viya on port 8591	Running	Manual	Local Syste...

See “General Servers and Services: Operate (Windows)” on page 463 for details.

View Launcher Server Properties

To access the **Servers** window from SAS Environment Manager:

1 In the applications menu (≡), under **Administration**, select **Manage Environment**.

2 In the vertical navigation bar, click 🖥.

3 Select the server whose properties you want to view, and then click ▤.

Concepts

SAS Launcher Server

The SAS Launcher Server starts processes, stops processes, and checks the status of processes in a SAS Viya environment.

For information about clustering, see “Fault Tolerance”.

Launcher Service

The launcher service is a SAS Viya microservice that provides API endpoints for how the launcher server runs a process. These API endpoints are used to create and manage launcher contexts.

Troubleshooting

Failure to launch Compute server sessions

Explanation:

Here are some reasons why a Compute server fails to launch:

- The user account under which the client is running does not have a home directory on the machine where the Compute server resides.
- Client users in a multi-tenant environment have to be a member of the sas group on the machine where the Compute server resides.
- Kerberos is present without valid credentials.

Resolution:

Check for the preceding issues in logs for the client application, Compute service , and Launcher service.

Log Files

Log files for the launcher service are located in **`/opt/sas/viya/config/var/log/launcher/default`**.

On multi-tenant systems, log files for the launcher service are located in **`/opt/sas/`*`tenant-ID`*`/config/var/log/launcher/default`**.

For Windows, the log files for the launcher service are located in **`\ProgramData\SAS\Viya\var\log\launcher\default`**.

SAS Workspace Server and SAS Object Spawner

Overview

The SAS Workspace Server enables client programs to access SAS libraries, to perform tasks by using the SAS language, and to retrieve the results. One or more SAS Workspace Servers are initialized by the SAS Object Spawner.

How To

Operate (Linux)

SAS Viya provides a script in **`/etc/init.d`** that you use to stop, start, restart, and check the status of the SAS Object Spawner. The script is named, `sas-viya-spawner-default`.

Syntax

How you run `sas-viya-spawner-default` depends on your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

`sudo systemctl status | stop | start | restart sas-viya-spawner-default`

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  `sudo service sas-viya-spawner-default status | stop | start | restart`

Usage Notes and Tips

- You must be logged on to the machine where the object spawner resides. Also, you must have sudo privileges to run this script.
- On multi-tenant SAS Viya systems, the script is named `sas-tenant-ID-sas-viya-spawner-default`.
- There is another script that you can use to manage and view the running state of all SAS Viya services. For more information, see “Start and Stop All Servers and Services” on page 462.

  Note: There is a sequence for starting and stopping SAS Viya servers and services. You must follow this sequence to avoid operational issues. For more information, see “Read This First: Start and Stop Servers and Services” on page 458 .
- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-spawner-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Server 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-spawner-default` on Red Hat Enterprise Server 7.*x* with the `service` command, and later attempt to shut down the object spawner using the `systemctl` command, the object spawner stops responding and does not shut down.

Examples

- To check status of the object spawner on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-spawner-default
```

- To stop the object spawner on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-spawner-default stop
```

- To start the object spawner on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-spawner-default
```

- To restart the object spawner on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-spawner-default restart
```

Operate (Windows)

Using the Microsoft Management Console (MCC) Services Snap-In, you can start, stop, and restart SAS Object Spawner.

Figure A.6 SAS Object Spawner in the Services Snap-In

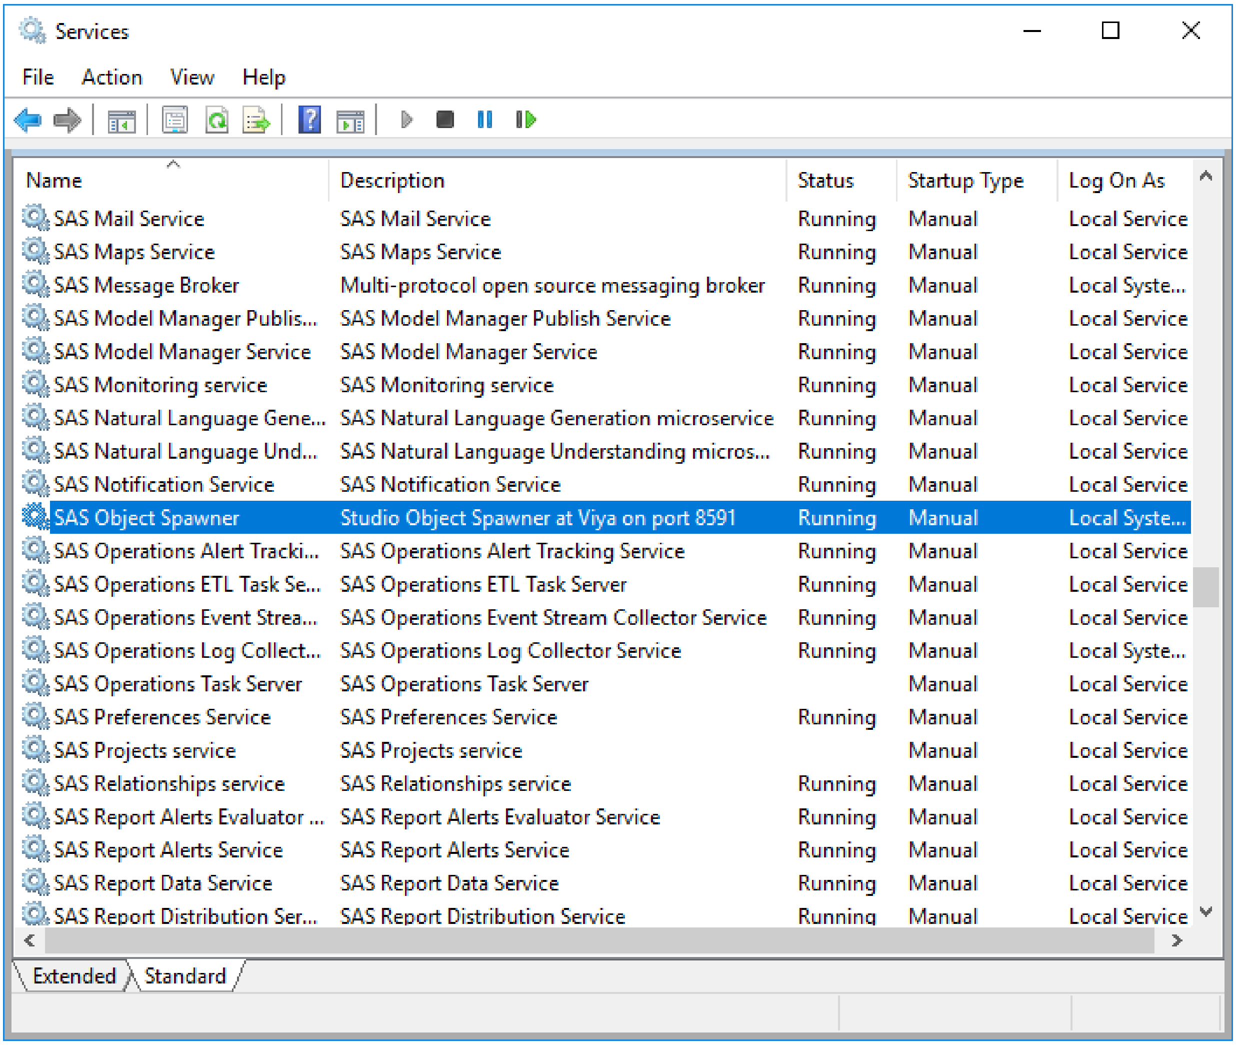

Name	Description	Status	Startup Type	Log On As
SAS Mail Service	SAS Mail Service	Running	Manual	Local Service
SAS Maps Service	SAS Maps Service	Running	Manual	Local Service
SAS Message Broker	Multi-protocol open source messaging broker	Running	Manual	Local Syste...
SAS Model Manager Publis...	SAS Model Manager Publish Service	Running	Manual	Local Service
SAS Model Manager Service	SAS Model Manager Service	Running	Manual	Local Service
SAS Monitoring service	SAS Monitoring service	Running	Manual	Local Service
SAS Natural Language Gene...	SAS Natural Language Generation microservice	Running	Manual	Local Service
SAS Natural Language Und...	SAS Natural Language Understanding micros...	Running	Manual	Local Service
SAS Notification Service	SAS Notification Service	Running	Manual	Local Service
SAS Object Spawner	Studio Object Spawner at Viya on port 8591	Running	Manual	Local Syste...
SAS Operations Alert Tracki...	SAS Operations Alert Tracking Service	Running	Manual	Local Service
SAS Operations ETL Task Se...	SAS Operations ETL Task Server	Running	Manual	Local Service
SAS Operations Event Strea...	SAS Operations Event Stream Collector Service	Running	Manual	Local Service
SAS Operations Log Collect...	SAS Operations Log Collector Service	Running	Manual	Local Syste...
SAS Operations Task Server	SAS Operations Task Server		Manual	Local Service
SAS Preferences Service	SAS Preferences Service	Running	Manual	Local Service
SAS Projects service	SAS Projects service		Manual	Local Service
SAS Relationships service	SAS Relationships service	Running	Manual	Local Service
SAS Report Alerts Evaluator ...	SAS Report Alerts Evaluator Service	Running	Manual	Local Service
SAS Report Alerts Service	SAS Report Alerts Service	Running	Manual	Local Service
SAS Report Data Service	SAS Report Data Service	Running	Manual	Local Service
SAS Report Distribution Ser...	SAS Report Distribution Service	Running	Manual	Local Service

See "General Servers and Services: Operate (Windows)" on page 463 for details.

Enable X Commands (Linux)

Because clients can use host commands to perform potentially harmful operations such as file deletion, by default, X commands are disabled for the SAS Object Spawner. However, to enable X commands, follow these steps:

1 Log on to the machine on which the object spawner resides.

2 Using a text editor, open `/opt/sas/viya/config/etc/spawner/default/spawner_usermods.sh`.

3 Add the following line, save, and close the spawner_usermods.sh file:

```
USERMODS="$JREOPTIONS -allowxcmd"
```

4 Restart the object spawner:

```
sudo service sas-viya-spawner-default restart
```

Set umask or ulimit Values (Linux)

In many circumstances, it might be desirable to control the permissions of files created from SAS sessions, or to set process limits for SAS sessions. To set permissions on files created from SAS sessions on Linux, the umask command can be used. The ulimit command is used to set process limits. The location from which these commands are executed will affect the scope of the settings.

1 Log on to the machine on which the SAS Workspace Server or the SAS/CONNECT Server resides. Log on as the SAS install user or log on with sudo privileges.

2 Using a text editor, open one of the following files, as appropriate:

   - For the SAS Workspace Server:

     /opt/sas/viya/config/etc/workspaceserver/default/workspaceserver_usermods.sh

   - For the SAS/CONNECT Server:

     /opt/sas/viya/config/etc/connectserver/default/connectserver_usermods.sh

   - For the SAS Workspace Server, the SAS/CONNECT Server, and all SAS instances:

     /opt/sas/spre/home/SASFoundation/bin/sasenv_local

3 Add your umask and ulimit values, and save the file.

   Your changes take effect the next time the server or servers are launched.

> **TIP** The umask and ulimit settings can be set for all users (or values can be set conditionally for each user), for collections of users, or for all members of a given Linux group. For more information, see "Examples of umask and ulimit Settings".

Examples of umask and ulimit Settings

In the following example, the umask command creates all files for all users with effective permissions of rw-r--r-- (owner:read and write; group:read; other:read):

```
umask 022
```

In the following example, umask is set for user joe00001 only:

```
if [ "$LOGNAME" = joe00001 ]
then
umask 022
fi
```

In the following example, ulimits are set according to user ID or group membership.

```
# determine primary group membership of user
# GP=`groups $LOGNAME | awk '{ print $1 }'`
# assign new ulimit based on userid or group membership as desired
if [ "$LOGNAME" = joe00001 ]
   then
     MAXSIZE=4096
     umask 022
elif [ "$LOGNAME" = fred0002 -o "$GP" = saspower ]
   then
     MAXSIZE=8192
     umask 077
elif [ "$GP" = sasuser ]
   then
```

```
     MAXSIZE=6144
else
     MAXSIZE=8192
fi

export MAXSIZE

ulimit -f $MAXSIZE
```

Lock Down SAS Workspace Servers

Using the "LOCKDOWN System Option" on page 589 and the "LOCKDOWN Statement" on page 590, you can limit access to files and to specific SAS features in a SAS Workspace Server session that executes in a batch mode or a server processing mode in a multi-tenant environment.

To lock down one or more workspace servers:

1. With administrator privileges, log on to the machine that contains the workspace server.

2. Create a lockdown path list (a whitelist) that contains all the paths that are accessible to the server, and add it to **/opt/sas/viya/config/etc/workspaceserver/default/autoexec_usermods.sas**.

   Note: For Windows, add the lockdown path list to **\ProgramData\SAS\Viya\etc\workspaceserver\default\autoexec_usermods.sas**

   Note: A path that is declared in the whitelist does not mean that an arbitrary user can read any file in that path. Host permissions on physical files and directories always take precedence over the whitelist. SAS adds certain predefined paths from the SAS configuration file by default. For more information, see LOCKDOWN Statement Details on page 592.

   Changes to the autoexec_usermods.sas file are automatically included when the workspace server scripts run. Your changes will take effect the next time SAS starts a workspace server session.

   TIP For a suggestion about how to implement the whitelist, see "Example 2: Hiding the Whitelist By Locating the Path outside the Whitelist" on page 593.

3. To enable lockdown, set the environment variable WORKSPACESERVER_LOCKDOWN_ENABLE to 1 in the sysconfig file **/opt/sas/viya/config/etc/sysconfig/workspaceserver/default/sas-workspaceserver.**

   Setting this variable enables the -lockdown option in the start-up script.

   Note: For Windows, add the -lockdown system option to the configuration file that is located in the server's configuration directory **\ProgramData\SAS\Viya\etc\workspaceserver\default\sasv9_usermods.cfg**.

4. If your site uses SAS Studio, set webdms.showSystemRoot=false.

   For more information, see "Update SAS Studio Configuration Properties" on page 80.

5. If your site uses SAS/CONNECT, see "Lock Down the SAS/CONNECT Server" on page 581.

TIP To limit the paths that are available to non-administrators when they create or edit a caslib, see "Paths List" on page 500.

Restricting SAS System Options

You can restrict SAS system options so that they cannot be changed by a user. An option can be restricted globally, by group, or by user.

Global Restrictions

Create the **`/opt/sas/spre/home/SASFoundation/misc/rstropts/rsasv9.cfg`** file and add options to this file.

Group Restrictions

Create the **`/opt/sas/spre/home/SASFoundation/misc/rstropts/groups/`*`groupname`*`_rsasv9.cfg`** file and add options to this file.

For example, for user smith in the group staff, the filename would be **`staff_rsasv9.cfg`**.

User Restrictions

Create the **`/opt/sas/spre/home/SASFoundation/misc/rstropts/users/`*`username`*`_rsasv9.cfg`** file and add options to this file.

For example, for user smith, the filename would be **`smith_rsasv9.cfg`**.

Concepts

SAS Workspace Server

The SAS Workspace Server enables client programs to access SAS libraries, to perform tasks by using the SAS language, and to retrieve results. Each workspace server process is owned by the client user that made the server request.

SAS Object Spawner

SAS Object Spawners interact with SAS by creating a server process for each client connection. SAS Workspace Servers are initialized by the SAS Object Spawner. An object spawner runs on the same machine as the workspace server, listens for requests, and launches the servers as necessary.

SAS Workspace Servers and SAS Cloud Analytic Services

In a SAS Viya environment, you can set up your autoexec.sas file to start a CAS session automatically. If you opt for automatic CAS session start-up, SAS uses that CAS session whenever it needs to communicate with SAS Cloud Analytic Services.

Many SAS procedures that are used in a SAS Viya deployment (such as PROC CARDINALITYand PROC NNET) use the CAS engine to communicate with CAS. The CAS engine uses the CAS session. In this context, the workspace server is used to interpret your SAS program and to determine how to run the lower-level actions in CAS.

Use of the SESSREF= DATA statement option in a SAS program is another method to inform the workspace server that CAS is being used. To run a DATA step in CAS, you must use a libref from the CAS engine, and you must specify the CAS session name in the SESSREF= option. When the workspace server interprets these language elements, it knows to run your DATA step in CAS.

In a SAS Viya environment, the workspace server is also used to do some work outside of CAS. Here are two examples:

- When creating graphics with procedures like PROC SGPLOT, although the data might be read from CAS with a CAS engine libref, the graphics are created with the workspace server.

- When processing data with the INFILE statement, the INPUT statement, and related DATA step statements and functions, the workspace server reads the contents of external files before the data can be transferred to CAS for analysis.

SAS Object Spawner Invocation

The SAS Object Spawner uses an suid root program, called elssrv, to launch processes under the identity of the requesting client. The user ID must be root in order to switch the identity to another user.

When launching a SAS Workspace Server, the client provides host credentials for the user who is requesting the SAS process (for example, a query or an ETL process) via the spawner. The spawner host authenticates the client and receives confirmation of valid credentials from sasauth. In addition, sasauth returns the UNIX uid and the list of groups. The suid root program launches the workspace server under this identity so that the process runs with the host authority of the requesting client.

Embedded Web Application Server

Overview

The embedded Apache Tomcat server that is used in all of the SAS Viya web applications provides the execution environment for SAS Studio.

How To

Operate (Linux)

SAS Viya provides a script in **`/etc/init.d`** that you use to stop, start, restart, and check the status of SAS Studio. The script is named, `sas-viya-sasstudio-default`.

Syntax

How you run `sas-viya-sasstudio-default` depends on your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl status | stop | start | restart sas-viya-sasstudio-default
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasstudio-default status | stop | start | restart
  ```

Usage Notes and Tips

- You must be logged on to the machine where the embedded web application server resides. Also, you must have sudo privileges to run this script.
- On multi-tenant SAS Viya systems, the script is named `sas-tenant-ID-sas-viya-sasstudio-default`.
- There is another script that you can use to manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" on page 462.

  Note: There is a sequence for starting and stopping SAS Viya servers and services. You must follow this sequence to avoid operational issues. For more information, see "Read This First: Start and Stop Servers and Services" on page 458 .
- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-sasstudio-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-sasstudio-default` on Red Hat Enterprise Linux 7.*x* with the `service` command, and later attempt to shut down SAS Studio using the `systemctl` command, SAS Studio stops responding and does not shut down.

Examples

- To check status of SAS Studio on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-sasstudio-default
```

- To stop SAS Studio on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasstudio-default stop
```

- To start SAS Studio on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-sasstudio-default
```

- To restart SAS Studio on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasstudio-default restart
```

Operate (Windows)

Using the Microsoft Management Console (MCC) Services Snap-In, you can start, stop, and restart SAS Studio.

Figure A.7 SAS Studio in the Services Snap-In

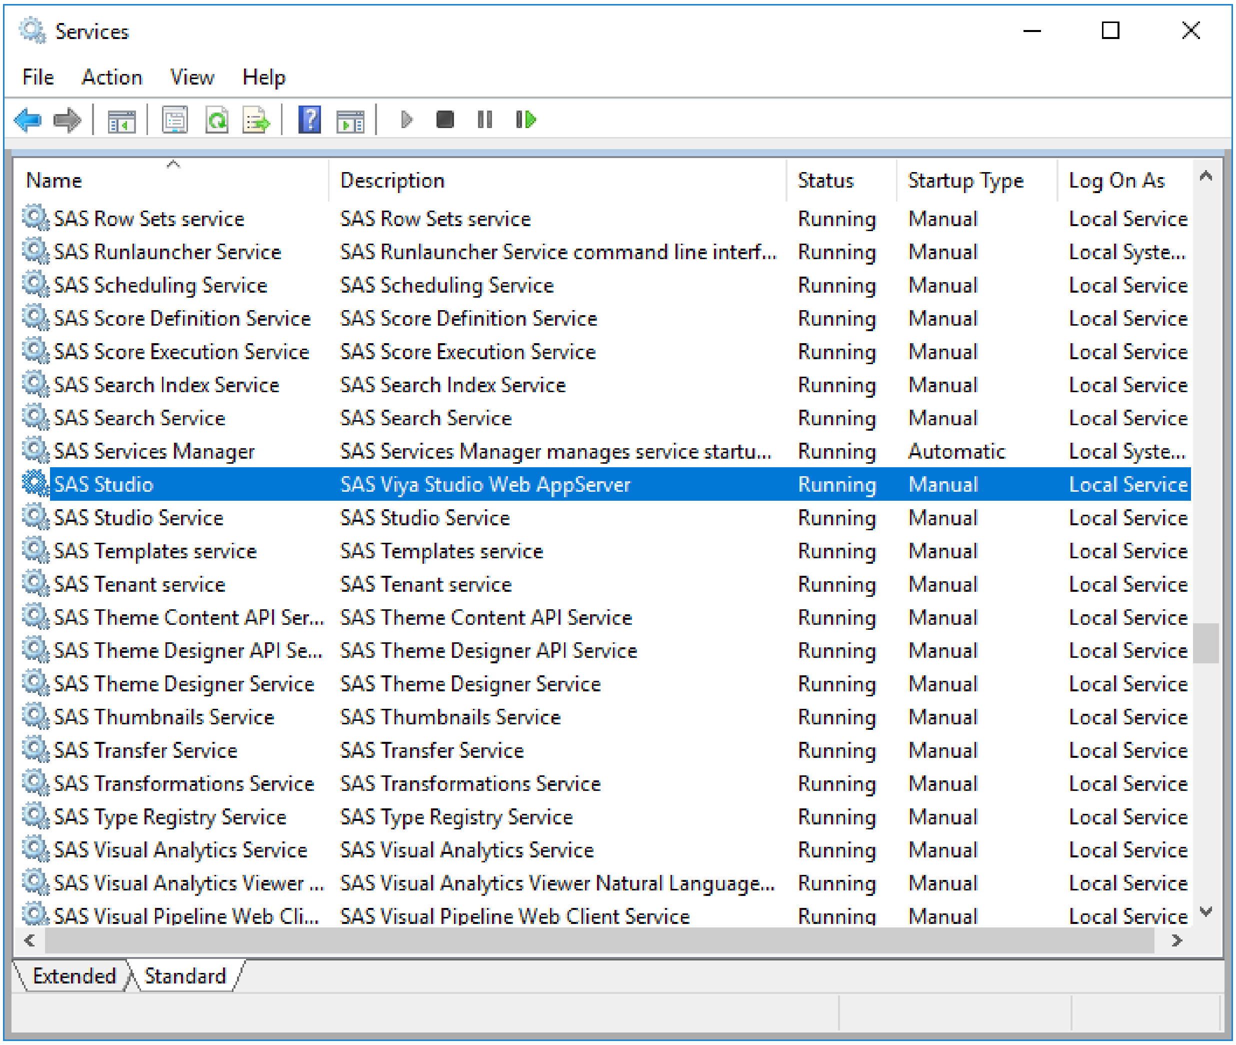

Name	Description	Status	Startup Type	Log On As
SAS Row Sets service	SAS Row Sets service	Running	Manual	Local Service
SAS Runlauncher Service	SAS Runlauncher Service command line interf...	Running	Manual	Local Syste...
SAS Scheduling Service	SAS Scheduling Service	Running	Manual	Local Service
SAS Score Definition Service	SAS Score Definition Service	Running	Manual	Local Service
SAS Score Execution Service	SAS Score Execution Service	Running	Manual	Local Service
SAS Search Index Service	SAS Search Index Service	Running	Manual	Local Service
SAS Search Service	SAS Search Service	Running	Manual	Local Service
SAS Services Manager	SAS Services Manager manages service startu...	Running	Automatic	Local Syste...
SAS Studio	SAS Viya Studio Web AppServer	Running	Manual	Local Service
SAS Studio Service	SAS Studio Service	Running	Manual	Local Service
SAS Templates service	SAS Templates service	Running	Manual	Local Service
SAS Tenant service	SAS Tenant service	Running	Manual	Local Service
SAS Theme Content API Ser...	SAS Theme Content API Service	Running	Manual	Local Service
SAS Theme Designer API Se...	SAS Theme Designer API Service	Running	Manual	Local Service
SAS Theme Designer Service	SAS Theme Designer Service	Running	Manual	Local Service
SAS Thumbnails Service	SAS Thumbnails Service	Running	Manual	Local Service
SAS Transfer Service	SAS Transfer Service	Running	Manual	Local Service
SAS Transformations Service	SAS Transformations Service	Running	Manual	Local Service
SAS Type Registry Service	SAS Type Registry Service	Running	Manual	Local Service
SAS Visual Analytics Service	SAS Visual Analytics Service	Running	Manual	Local Service
SAS Visual Analytics Viewer ...	SAS Visual Analytics Viewer Natural Language...	Running	Manual	Local Service
SAS Visual Pipeline Web Cli...	SAS Visual Pipeline Web Client Service	Running	Manual	Local Service

See “General Servers and Services: Operate (Windows)” on page 463 for details.

Configure Mail

To use the email functionality in SAS Studio, an SMTP server and the following information is required:

- *fully-qualified-SMTP-server-name*

  The fully qualified host name of the SMTP server for the outbound mail (for example, `my_mail_server.example.com`).

- *SMTP-server-port*

  The port for the SMTP server (for example, 25).

- *site-administrator-email-address*

  The user name that accesses the SMTP server.

  This user name is not necessarily the person who is sending the mail.

- *site-administrator-password*

  The password for the user name that accesses the SMTP server.

- *company-domain*

The domain name for your site (for example, `my_company.example.com`).

To configure SAS Studio for SMTP email, follow these steps:

1 Log on to the machine on which the embedded web application server resides.

2 Using a text editor, open **`/opt/sas/viya/config/etc/sasstudio/default/init_usermods.properties`**.

3 Add the following lines, save, and close the init_usermods.properties file:

```
webdms.SMTP.hostName=fully-qualified-SMTP-server-name
```

```
webdms.SMTP.port=SMTP-server-port
```

```
webdms.SMTP.user=site-administrator-email-address
```

```
webdms.SMTP.password=site-administrator-password
```

```
webdms.domain=company-domain
```

4 Restart the embedded web application server:

```
sudo service sas-viya-sasstudio-default restart
```

When sending email, the sender address is derived from the user name that logged on to SAS Studio and the value of the `webdms.domain` property in the appserver_usermods.sh file. For example, if the user name is test, the sender address would be test@your-company.com.

SAS/CONNECT Server and SAS/CONNECT Spawner

Overview

SAS/CONNECT software provides the essential tools for sharing data and processing power across multiple computing environments:

- For users of SAS 9.4 and earlier versions, SAS/CONNECT enables you to use SAS Viya functionality and features.

  For more information, see “SAS 9 and SAS Viya” on page 24.

- For SAS Viya users who might also have SAS 9, SAS/CONNECT provides parallel processing for CAS procedures.

  For more information, see SAS/CONNECT for SAS Viya *User’s Guide*

In a full deployment of SAS Viya on Linux, SAS/CONNECT is secure by default. In a programming-only deployment on Linux and Windows, you must configure security using Transport Layer Security (TLS). See “Use SAS/CONNECT with TLS Enabled to Import Data” in *Encryption in SAS Viya: Data in Motion*.

How To

Operate (Linux)

SAS Viya provides a script in **`/etc/init.d`** that you use to stop, start, restart, and check the status of SAS/CONNECT Spawner. The script is named, `sas-viya-consul-default`.

Syntax

How you run `sas-viya-connect-default` depends on your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  `sudo systemctl status | stop | start | restart sas-viya-connect-default`

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  `sudo service sas-viya-connect-default status | stop | start | restart`

Usage Notes and Tips

- You must be logged on to the machine where the spawner resides. Also, you must have sudo privileges to run this script.
- On multi-tenant SAS Viya systems, the script is named `sas-tenant-ID-sas-viya-connect-default`.

  An example is sas-tenant1-connect-default.

- There is another script that you can use to manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" on page 462.

  Note: There is a sequence for starting and stopping SAS Viya servers and services. You must follow this sequence to avoid operational issues. For more information, see "Read This First: Start and Stop Servers and Services" on page 458 .

- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-connect-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-connect-default` on Red Hat Enterprise Linux 7.*x* with the `service` command, and later attempt to shut down the spawner using the `systemctl` command, the configuration server stops responding and does not shut down.

Examples

- To check status of the spawner on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-connect-default
```

- To stop the spawner on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-connect-default stop
```

- To start the spawner on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-connect-default
```

- To restart the spawner on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-connect-default restart
```

Operate (Windows)

Using the Microsoft Management Console (MCC) Services Snap-In, you can start, stop, and restart SAS Connect Spawner.

Figure A.8 SAS Connect Spawner in the Services Snap-In

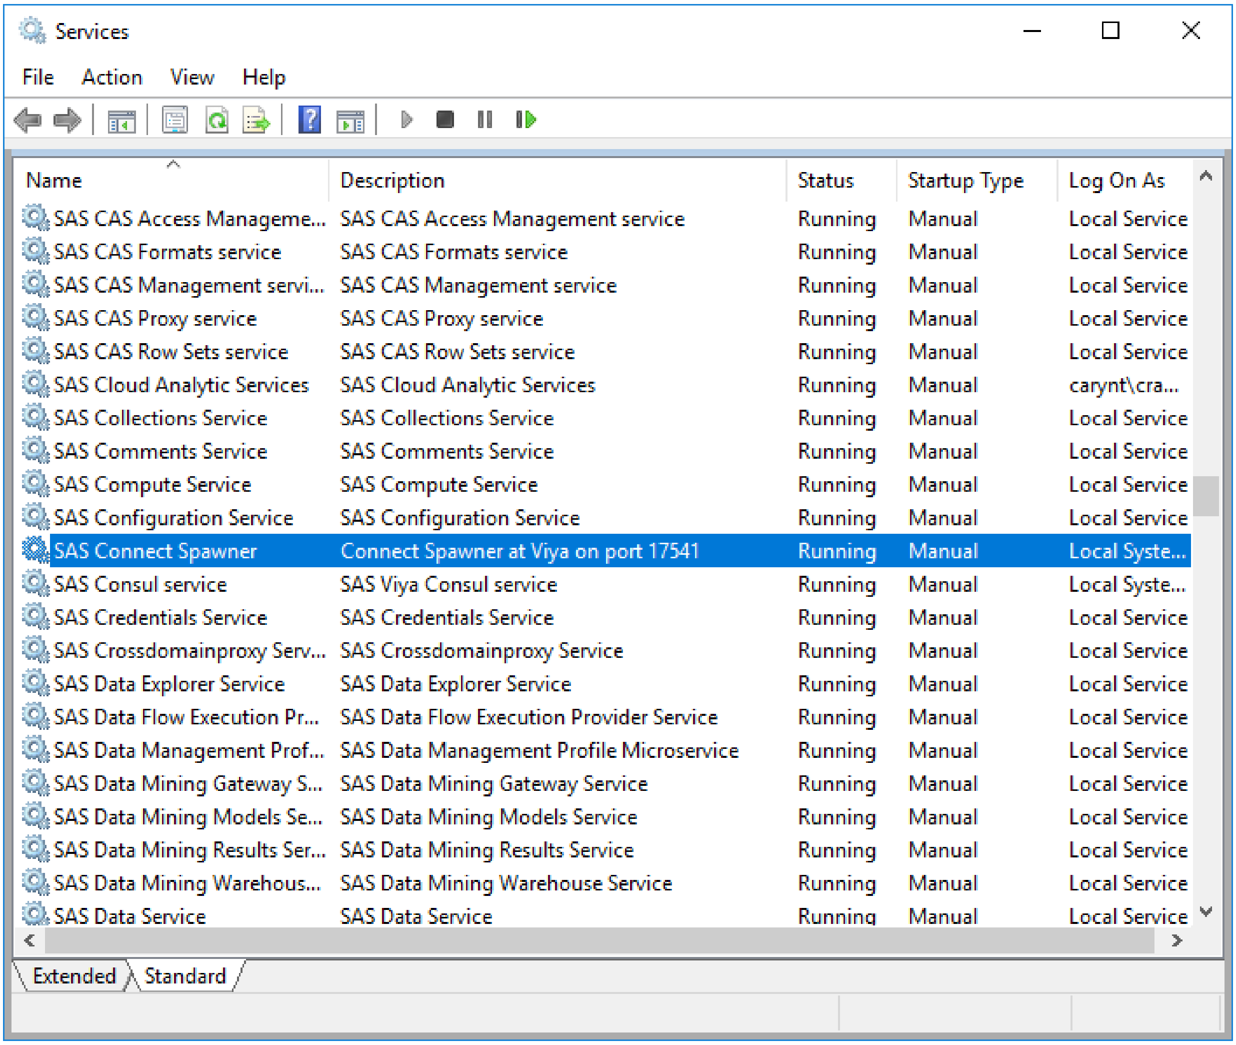

Name	Description	Status	Startup Type	Log On As
SAS CAS Access Manageme...	SAS CAS Access Management service	Running	Manual	Local Service
SAS CAS Formats service	SAS CAS Formats service	Running	Manual	Local Service
SAS CAS Management servi...	SAS CAS Management service	Running	Manual	Local Service
SAS CAS Proxy service	SAS CAS Proxy service	Running	Manual	Local Service
SAS CAS Row Sets service	SAS CAS Row Sets service	Running	Manual	Local Service
SAS Cloud Analytic Services	SAS Cloud Analytic Services	Running	Manual	carynt\cra...
SAS Collections Service	SAS Collections Service	Running	Manual	Local Service
SAS Comments Service	SAS Comments Service	Running	Manual	Local Service
SAS Compute Service	SAS Compute Service	Running	Manual	Local Service
SAS Configuration Service	SAS Configuration Service	Running	Manual	Local Service
SAS Connect Spawner	Connect Spawner at Viya on port 17541	Running	Manual	Local Syste...
SAS Consul service	SAS Viya Consul service	Running	Manual	Local Syste...
SAS Credentials Service	SAS Credentials Service	Running	Manual	Local Service
SAS Crossdomainproxy Serv...	SAS Crossdomainproxy Service	Running	Manual	Local Service
SAS Data Explorer Service	SAS Data Explorer Service	Running	Manual	Local Service
SAS Data Flow Execution Pr...	SAS Data Flow Execution Provider Service	Running	Manual	Local Service
SAS Data Management Prof...	SAS Data Management Profile Microservice	Running	Manual	Local Service
SAS Data Mining Gateway S...	SAS Data Mining Gateway Service	Running	Manual	Local Service
SAS Data Mining Models Se...	SAS Data Mining Models Service	Running	Manual	Local Service
SAS Data Mining Results Ser...	SAS Data Mining Results Service	Running	Manual	Local Service
SAS Data Mining Warehous...	SAS Data Mining Warehouse Service	Running	Manual	Local Service
SAS Data Service	SAS Data Service	Running	Manual	Local Service

See “General Servers and Services: Operate (Windows)” on page 463 for details.

Set Configuration Options

For Linux, you can use any of these methods to set options such as encryption options to invoke SAS/CONNECT spawner:

- `/opt/sas/viya/config/etc/connect/default/connect_usermods.sh`
- the SASCMD option
- `/opt/sas/viya/config/etc/sysconfig/connect/default/sas-connect`

For Windows, use `\ProgramData\SAS\Viya\etc\connect\default\connect_usermods.bat`

For Linux, you can use any of these methods to set options to invoke SAS/CONNECT server :

- `/opt/sas/viya/config/etc/connectserver/default/connectserver_usermods.sh`
- `/opt/sas/viya/config/etc/sysconfig/connectserver/default/sas-connectserver`

For Windows, you can use `\ProgramData\SAS\Viya\etc\connectserver\default\connectserver_usermods.bat`.

Your changes take effect the next time the SAS/CONNECT server or spawner is restarted.

Lock Down the SAS/CONNECT Server

Using the "LOCKDOWN System Option" on page 589 and the "LOCKDOWN Statement" on page 590, you can limit access to files and to specific SAS features in a SAS/CONNECT server session in a multi-tenant environment.

To lock down your SAS/CONNECT server:

1 With administrator privileges, log on to the machine that contains the SAS/CONNECT server.

2 If you have not done so already, create a lockdown path list (a whitelist) that contains all the paths that are accessible to the server, and add it to `/opt/sas/viya/config/etc/connectserver/default/autoexec_usermods.sas`.

   Note: For Windows, add the lockdown path list to `\ProgramData\SAS\Viya\etc\connectserver\default\autoexec_usermods.sas`.

   Changes to the autoexec_usermods.sas file are automatically included when the SAS/CONNECT server scripts run. Your changes will take effect the next time SAS starts a SAS/CONNECT server session.

   > **TIP** If you have already locked down your SAS/CONNECT server, then this step is unnecessary.

   Note: A path declared in the whitelist does not mean that an arbitrary user can read any file in that path. Host permissions on physical files and directories always take precedence over the whitelist. SAS adds certain predefined paths from the SAS configuration file by default. For more information, see LOCKDOWN Statement Details on page 592.

   > **TIP** For a suggestion about how to implement the whitelist, see "Example 2: Hiding the Whitelist By Locating the Path outside the Whitelist" on page 593.

3 To enable lockdown, set the environment variable `CONNNECTSERVER_LOCKDOWN_ENABLE` to 1 in the sysconfig file `/opt/sas/viya/config/etc/sysconfig/connectserver/default/sas-connectserver`.

   Setting this variable enables the –lockdown option in the start-up script.

   Note: For Windows, add the -lockdown system option to the configuration file that is located in the server's configuration directory `\ProgramData\SAS\Viya\etc\connectserver\default\sasv9_usermods.cfg`.

   Note: Do not start the SAS/CONNECT spawner using the `-SHELL` option. As long as the -SHELL option is *not* specified, the `-NOXCMD` option is added by default to the server's invocation parameters. `-NOXCMD` prevents clients from executing X commands from their SAS sessions to access system files.

4 If your site uses SAS Studio, set webdms.showSystemRoot=false.

   For more information, see "Update SAS Studio Configuration Properties" on page 80.

Concepts

SAS/CONNECT software provides the essential tools for sharing data and processing power across multiple computing environments.

Note: SAS/CONNECT is ordered and licensed separately from other SAS Viya products.

SAS code uses these tools to perform tasks such as the following:

- dividing time-consuming tasks into multiple units of work and executing these units in parallel

- moving data from a client machine to a server machine (including legacy data from SAS 9), or vice versa, so that the data is on the same machine as the code processing it.

Reference

TCPPORTFIRST System Option

TCPPORTFIRST=<*port-number*>
TCPPORTLAST=<*port-number*>

Restricts the range of TCP/IP ports that clients can use to remotely access servers. Within the range of 0 through 32767, assign a beginning value to TCPPORTFIRST and an ending value to TCPPORTLAST. To restrict the range of ports to only one port, set the values for TCPPORTFIRST and TCPPORTLAST to the same number. Consult with your network administrator for advice about these settings.

When -NOINHERITANCE is on, you can set TCPPORTFIRST and TCPPORTLAST in a SAS start-up command or in the configuration file.

This applies in the noinheritance case only. When socket inheritance is enabled, the child SAS/CONNECT server does not start up as a listening port.

Server	Optional
Range	0–32767

Example	In the example below, the server is restricted to the TCP/IP ports 4020 through 4050:
	`options tcpportfirst=4020;`
	`options tcpportlast=4050;`

Server Environment Variables

The following SAS/CONNECT Server environment variables are available for configuring your TCP/IP connections. Place them in the **/opt/sas/viya/config/etc/connectserver/default/connectserver_usermods.sh** script file. For information about configuring environment variables in a Linux environment, see Defining Environment Variables in UNIX Environments.

CONNECTWDWAIT=<*seconds*>

Specify to limit the possibility that a client session disconnect might orphan a runaway DMR mode session. To ensure the responsiveness of the spawner, SAS starts a “watchdog” thread to monitor the connection. The default interval is five seconds. If a disconnect occurs, CONNECTWDWAIT checks 18 times and then terminates the DMR thread (for a default elapsed time of 90 seconds). Setting the CONNECTWDWAIT value to zero means that the process does not monitor the connection.

Defaults	interval: 5 seconds
	total elapsed time: 90 seconds
Examples	In the following example, the option is set to 10, so the process waits 180 seconds, and then terminates the thread: `set CONNECTWDWAIT=10`
	In the following example, the option is set to 0, so the process does not monitor the connection: `set CONNECTWDWAIT=0`

TCPLISTENTIME=<*seconds*> | <MIN> | <MAX>

Specifies the amount of time that a SAS/CONNECT server listens for a SAS/CONNECT client to connect before terminating the server session. It enables you to control idle and unresponsive sign-on connections by specifying how long (in seconds) a server “listens” for a response from the client during sign-on before it exits automatically. The default value for a session time-out is 0 (meaning, no time limit). The maximum value is 600 seconds.

Client	Optional
Defaults	0 (no time limit)
	MIN (minimum value is 0)
Examples	TCPLISTENTIME=MIN
	TCPLISTENTIME=1
	TCPLISTENTIME=90
	TCPLISTENTIME=MAX (maximum value is 600)

Note: If `TCPLTO` is already set, it will have precedence over `TCPLISTENTIME`.

TCP_POLL_INTERVAL=<*seconds*>

Specify to ensure responsiveness of SAS spawners and servers to various conditions outside of normal request processing. When idle, servers and spawners periodically awaken to check for requests. The interval in seconds for this check is governed by the TCP_POLL_INTERVAL environment variable. Generally, the default setting of 60 seconds should be acceptable.

A value of zero means the server remains idle and awakens for request processing only.

Example In the following example, the option is set to 50, so the process checks every 50 seconds for a connection:

```
TCP_POLL_INTERVAL=50
```

TCPMSGLEN=<*size*>

Specifies the size of the buffer (in bytes) that the TCP/IP access method uses for breaking up a message that it sends to or receives from the SAS/CONNECT application layer during a SAS/CONNECT session. The application layer uses a message size that is stored in the TBUFSIZE option that you can specify in the SIGNON statement or as a SAS option.

If TBUFSIZE is larger than TCPMSGLEN, the TCP/IP access method breaks the message into a buffer whose size is defined by TCPMSGLEN, and issues the number of send and receive messages that are necessary to complete the message transaction.

The value for TCPMSGLEN must be set at both the client and the server. If the values that are set for TCPMSGLEN at the client and at the server are different, the smaller value of the two is used during the SAS/CONNECT session. If the TCPMSGLEN environment variable is not set, SAS uses the TCP stack’s default size and allows autotuning if implemented by the stack.

Client	Optional
Server	Optional
Example	`set TCPMSGLEN=65536`

CONNECTKEEPALIVE=<*seconds*>

Prevents a SAS/CONNECT client connection to the SAS/CONNECT server from being terminated.

Setting this environment variable in the server session prevents firewalls from terminating a connection between a client and server when there are long periods of inactivity on the connection. A keepalive packet is sent from a thread that is started by the server session for the specified number of seconds.

Example The value of 5 causes the keepalive packet to be sent every 5 seconds to prevent connection termination.

```
set CONNECTKEEPALIVE=5
```

Spawner General Options

-CLEARTEXT

Allows sign-ons from clients that do not support user ID and password encryption. This option allows clients that are running older releases (prior to SAS 6.09E and SAS 6.11 TS040, which do not support user ID and password encryption) to sign on to the spawner program. Use this option only when absolutely necessary because credentials are transmitted unencrypted. The default encodes all communications.

Default -NOCLEARTEXT

-DEBUG

Turns on debug level output.

-HELP

Specifies to print the Help message.

-LOG | -LOGFILE <*filename*>

Specifies the filename to use for spawner log output if you are not using the -LOGCONFIGLOC option. The -LOG option should not be used with the -LOGCONFIGLOC option. If both options are specified, then the -LOGCONFIGLOC option takes precedence.

You can specify the -DEBUG or -TRACE options with the -LOG <filename> option to cause the spawner to send detailed log messages to a log file.

Example In this example, the following option is enclosed in double quotation marks and added after USERMODS= in **`/opt/sas/viya/config/etc/connect/default/connect_usermods.sh`**. When the spawner starts, it sends debug-level log messages to a file named sas-connect.log:

```
USERMODS="-log /var/log/sas/viya/connect/default/sas-connect.log"
```

-LOGCONFIGLOC <*filename*>

Enables the SAS logging facility for SAS servers and names the location of the configuration file that is used by the SAS logging facility to create spawner log output. The configuration file is an XML file that specifies and configures loggers and appenders for the SAS/CONNECT spawner.

The file specification that defines the location of the XML configuration file must be a valid filename or a path and filename for your operating environment. If the path contains spaces, enclose the file specification in quotation marks.

Note If LOGCONFIGLOC is specified, spawner messages are routed by default to the App.Connect.Spawner logger.

-NOINHERITANCE

Disables socket inheritance.

Socket inheritance enables SAS/CONNECT servers to use the socket connection that is established between the SAS/CONNECT client and the spawner. Socket inheritance saves resources and is easier to configure when clients connect to a server that is within a firewall.

Default Socket inheritance is on.

-NOSCRIPT

Prevents sign-on from clients that use scripts, and allows sign-on only from clients that do not use scripts.

-NOSCRIPT can be useful if you want to limit SAS start-up commands to the use of the -SASCMD option. Specifying -NOSCRIPT restricts clients from specifying additional options in SAS start-up commands or script files.

Requirement Must be used with -SASCMD

-SASCMD | -CMD <*command*>

Specifies the SAS command or a command file that starts a SAS session when you sign on without a script. If the client does not specify a script file at sign-on, the -SASCMD option must be specified when starting the spawner.

Example In this example, the following option is enclosed in double quotation marks and added after USERMODS= in **`/opt/sas/viya/config/etc/connect/default/connect_usermods.sh`**. When the spawner starts, it uses a command file named mystartup:

```
 USERMODS="-sascmd '/u/username/mystartup'"
```

Here is a sample command file named mystartup:

```
#!/bin/ksh
#---------------------------------
# mystartup
#---------------------------------
. ~/.profile
sas -noterminal -nosyntaxcheck $*
#-----------------------------
```

The $* positional parameter enables you to specify additional SAS options when you invoke SAS. In addition, $* also allows the options that the spawner adds automatically, like -DMR, to be included in the server session.

-SASDAEMONSERVICE <*service-name* | *port*>

Specifies the service name or port number that the SAS/CONNECT spawner uses to listen for child SAS/CONNECT server process connections.

If you use a service, its name must be configured in the SERVICES file on the computer that the SAS/CONNECT server session runs on.

-SERVICE <*service-name* | *port*>

Specifies the service name or port number to use to listen for client connections.

The -SERVICE option values that are used to start the spawner determine what is used by the client to sign on.

Note If the -SERVICE option is not specified, the spawner listens on Telnet port (23).

Example In this example, the following option is enclosed in double quotation marks and added after USERMODS= in **`/opt/sas/viya/config/etc/connect/default/connect_usermods.sh`**. When the spawner starts, it uses port 5020 for the -SERVICE option during spawner start-up:

```
USERMODS="-service 5020"
```

The client can then sign on by specifying the explicit port-number in the SIGNON statement:

```
%let myHost=<spawner-host> 5020;
signon myHost user='myuserid' password='mypassword';
```

-SHELL

Specifies that the started SAS/CONNECT servers allow X commands.

Without specifying the -SHELL option to the spawner, X command processing is disabled by default.

-SSPI

Identifies support for the Security Support Provider Interface for single sign-on connections to the spawner. To enable SSPI authentication, you must specify -SSPI in the spawner start-up command.

Default	-NOSSPI

-TRACE | -VERBOSE
Turns on trace level output.

Spawner Security Options

SAS/CONNECT Spawner uses the “SAS System Options for Encryption” in *Encryption in SAS Viya: Data in Motion* .

Server Configuration Files

Configuration Home Directory

The SAS Viya deployment process creates a configuration home directory for each server instance.

For Linux, the location for all services:

```
/opt/sas/viya/config/etc/compsrv/default
/opt/sas/viya/config/etc/workspaceserver/default
/opt/sas/viya/config/etc/spawner/default
/opt/sas/viya/config/etc/connectserver/default
/opt/sas/viya/config/etc/connect/default
```

For Windows, the location for all services:

```
\ProgramData\SAS\Viya\etc\compsrv\default
\ProgramData\SAS\Viya\etc\workspaceserver\default
\ProgramData\SAS\Viya\etc\spawner\default
\ProgramData\SAS\Viya\etc\connectserver\default
\ProgramData\SAS\Viya\etc\connect\default
```

Note: Linux and Windows support the same services. The last directory in the path, *default*, is the deployment instance for the server.

Server Configuration Files

Each of the following SAS Viya programming run-time servers uses one or more server configuration files, as appropriate.

- SAS Compute Server
- SAS Workspace Server
- SAS Object Spawner
- SAS/CONNECT Server
- SAS/CONNECT Spawner

Table A.1 Server Configuration Files

Standard File Name	Description
autoexec.sas	ContainsSAS statements that are executed immediately after SAS initializes all components of the SAS Application Server. Do not modify this file. If you need to make changes, modify the appserver_autoexec_usermods.sas file that is in the same directory.
autoexec_deployment.sas	Contains server configuration settings that are created during deployment by Ansible from vars.yml. During updates, user configuration settings are overwritten. Do not modify this file. If you need to make changes, modify the sasv9_usermods.cfg file that is in the same directory.
autoexec_usermods.sas	Contains modifications made by the SAS administrator. Using autoexec_usermods.sas ensures that your modifications are not overwritten when you update SAS Viya.
sasv9.cfg	Specifies start-up options for the server and contains calls to other files that are listed in this table. Do not modify this file. If you need to make changes, modify the sasv9_usermods.cfg file that is in the same directory.
sasv9_deployment.cfg	Specifies start-up options for the server and contains calls to other files that are listed in this table that are created during deployment by Ansible from vars.yml. Do not modify this file. If you need to make changes, modify the sasv9_usermods.cfg file that is in the same directory.
sasv9_usermods.cfg	Contains modifications made by the SAS administrator. Using sasv9_usermods.cfg ensures that your modifications are not overwritten when you update SAS Viya.
logconfig.xml	Specifies the logging configuration for the server or the spawner.
logconfig.trace.xml	Contains alternative logging configuration settings for high-level logging messages (for example, DEBUG and TRACE messages) that can be used by SAS Technical Support to help resolve server issues. The messages are written to the server or spawner rolling log file.
logconfig.arm.xml	Specifies logging configuration files for servers that include specifications for collecting Application Response Measurement (ARM) log information and sending it to a log file. ARM logging is used to collect performance-related events.
logconfig.trace.arm.xml	Contains alternative ARM log configuration settings for high-level logging information. The messages are written to the server rolling log file.
sasenv_deployment	Contains server environmental variable settings that are created during deployment by Ansible from vars.yml. During updates, user configuration settings are overwritten. Do not modify this file. Add local environmental variable settings in the sasenv_local file in the `/opt/sas/viya/config/etc/`*`server`*`/default` directory.
■ connect.sh ■ connectserver.sh ■ spawner.sh ■ workspaceserver.sh	The connect.sh and spawner.sh scripts start the SAS/CONNECT server and the SAS Workspace Server, respectively. The spawners accept connections from the clients to start SAS servers. Do not modify these files. If you need to make changes, modify the *server-spawner*_usermods.sh file that is in the same directory.

Standard File Name	Description
■ connect_usermods.sh ■ connectserver_usermods.sh ■ spawner_usermods.sh ■ workspaceserver_usermods.sh ■ sas-compsrv	Contain modifications made by the SAS administrator to the configurations for these application servers. Using *server-spawner*_usermods.sh ensures that your modifications are not overwritten when you update SAS Viya. Note: For the Compute Server, the file that is used for modifications is **`sas-compsrv`**, which resides in the sysconfig directory.

Configuring SAS to Run External Languages

Configuring SAS to Run Python

If your installation includes SAS Micro Analytic Service and you have the May 2019 release of SAS Viya 3.4, you can enable Python code to run in SAS in lockdown mode. You can also configure SAS to run Python code using PROC FCMP.

For more information, see the following resources:

- Enabling Python Code while in Lockdown Mode on page 593
- Working with Python and SAS Micro Analytic Service in *SAS Micro Analytic Service Programming and Administration Guide*.

Python Requirements

These requirements must be met before PROC FCMP can be used to run Python code.

1. Install Python. Python version v2.7 or later is recommended for use with PROC FCMP.
2. Set the MAS_M2PATH environment variable to specify the absolute path to the mas2py.py file. The mas2py.py file is used to execute Python code within a Python process that is launched by SAS Micro Analytic Service. Here is an example:
   - UNIX:

     ```
     export MAS_M2PATH="/opt/sas/spre/home/SASFoundation/misc/embscoreeng/
     mas2py.py"
     ```

   - Windows:

     ```
     set MAS_M2PATH="C:\Program Files\SAS\SPRE\SASFoundation\misc\embscoreeng
     \mas2py.py"
     ```

3. Set the MAS_PYPATH environment variable to specify the absolute path to the Python executable. Here is an example:
   - UNIX:

     ```
     export MAS_PYPATH="/bin/python"
     ```

   - Windows:

```
set MAS_PYPATH="c:\python\python.exe"
```

References

LOCKDOWN System Option

Enables the ability to limit access to files and to specific SAS Viya features for a SAS Viya session that is executing in batch or server processing mode.

LOCKDOWN

enables the ability to limit access to files and to specific SAS features for a SAS session that is executing in batch mode or server processing mode.

Valid in	SAS 9.4: Configuration file, SAS invocation
	SAS Viya: Configuration file, SAS invocation, SASV9_OPTIONS environment variable
Category	Environment Control: Initialization and Operation
PROC OPTIONS GROUP=	EXECMODES
Default	NOLOCKDOWN
Restriction	This version of the LOCKDOWN system option is for SAS Viya only. .
Requirement	XCMD must be disabled to use the LOCKDOWN option.
Note	This option can be restricted by a site administrator. For more information, see Restricted Options in *SAS Intelligence Platform: Administration / Application Server Administration Guide*.
See	For XCMD, see XCMD System Option: UNIX in SAS 9.4 and SAS Viya Programming Documentation / SAS Companion for UNIX Environments.
	For SAS 9.4, see Locked-Down Servers in *SAS Intelligence Platform: Administration / Application Server Administration Guide*.
	For SAS Viya, see SAS Workspace Server and SAS Object Spawner: How to on page 569 and SAS Compute Server on page 562 .

In addition to LOCKDOWN, security administrators also rely on the NOXCMD system option. For more information, see XCMD System Option: UNIX in *SAS Companion for UNIX Environments*.

When the LOCKDOWN option is specified for a SAS session, SAS enters a locked-down state at a lockdown point. A SAS session in the locked down state has following restrictions:

- limited file system access

  All access to local files and directories is validated through the lockdown path list. The lockdown path list specifies which host file resources are available when a SAS session is in the locked-down state. This list includes the default system directories and files.

- limited SAS language features

  The following SAS language features are disabled:

  - □ DATA step Java Object **javaobj**
  - □ PROC JAVAINFO
  - □ SAS functions: ADDR, ADDRLONG, PEEK, PEEKLONG, PEEKC, PEEKCLONG, POKE, POKELONG, and MODULE

LOCKDOWN does not take effect in a SAS session until after the lockdown point has been reached.

The lockdown point has been reached during SAS execution when the following tasks have been completed in order to establish a user's SAS environment:

- SAS session initialization
- AUTOEXEC execution
- INITSTMT execution

During initialization of the user's SAS environment, all paths and files are available and work as designed (for example, SASHELP, WORK, LOG, and so on). AUTOEXEC predefined libraries also work as designed. When initialization is complete, SAS is put in the locked-down state with limited file system access.

LOCKDOWN Statement

Secures the SAS Viya workspace server or the SAS/CONNECT server or SAS Compute Server on SAS Viya by restricting access from within a server process to the host operating environment.

Summary

Valid in	AUTOEXEC file or INITSTMT= system option
Category	Control
Type	Executable
Restriction	This version of the LOCKDOWN statement is for SAS Viya only.
See	For information about LOCKDOWN on SAS 9.4, see LOCKDOWN Statement in *SAS Intelligence Platform: Administration / Application Server Administration Guide*.

Syntax

LOCKDOWN ENABLE_AMS= *access-method-1* < … *access-method-n* >;

LOCKDOWN FILE= *file-ref*;

LOCKDOWN PATH= '*pathname-1*' < '… *pathname-n* '>;

LOCKDOWN LIST;

Arguments

ENABLE_AMS=*access-method-1* < … *access-method-n* >;

By default, when SAS is in a locked-down state, these access methods are not available:

- EMAIL
- FTP
- HADOOP
- HTTP

- SOCKET
- TCPIP
- URL
- PYTHON
- PYTHON_EMBED

Likewise, the HTTP and HADOOP procedures are also not available. These procedures are initially disabled, preventing users from potentially generating spam or accessing files on a server that is not configured for authentication.

ENABLE_AMS allows administrators to re-enable those access methods and procedures that are disabled by default in the locked-down state.

Access method names are case-insensitive. Duplicate names are ignored. Separate each name with a space, and do not use quotation marks.

URL/HTTP and SOCKET/TCPIP are aliased name pairs. If one is re-enabled, the other is also automatically re-enabled. For example, if URL is re-enabled, HTTP is also re-enabled.

If either URL or HTTP is re-enabled, PROC HTTP is also automatically re-enabled.

If HADOOP is re-enabled, PROC HADOOP is re-enabled.

ENABLE_AMS= can be included only in an AUTOEXEC file or an INITSTMT= option.

FILE=*file-ref*

Names a file that contains a list of valid directories and files to be added into the lockdown path list. SAS automatically adds subdirectories of any valid directories in the list to the lockdown path list.

The lockdown path list specifies which files and directories a SAS session can access when in a locked-down state. (This list is often referred to as a whitelist.)

A lockdown file can contain multiple lines. Each line specifies a path string. A path string on each line can be one pathname or a concatenation of pathname specifications. If a line contains only one pathname, the pathname can contain spaces or the host-supported pathname characters. This pathname does not need to be enclosed in quotation marks.

If a line contains a concatenation specification, enclose the entire group of concatenated path specifications in parentheses. Enclose each pathname in quotation marks. Use a space to separate each pathname specification. Do not use newline characters in a path string. For more information, see PATH = 'pathname' on page 591.

If DBCS characters are included in the path list in the lockdown text file, add the ENCODING option to the FILENAME statement. This addition ensures that the lockdown path list is transcoded properly to the SAS session encoding. For example, suppose you have a text file, lockdown-text-file.txt, that includes a path list with DBCS characters, add the following code to the server's AUTOEXEC file appserver_autoexec.sas:

```
filename myfile "c:\lockdown-text-file.txt" encoding='euc-cn';
lockdown file=myfile;
lockdown list;
filename myfile clear;
```

PATH= '*pathname*'

Specifies one or multiple paths to be added to the lockdown path list. Enclose pathnames in quotation marks. (Single quotation marks are preferable to avoid conflicts with SAS macro variables.) Use a space to separate multiple pathnames.

The lockdown path list specifies which files and directories a SAS session can access when in a locked-down state. (This list is often referred to as a whitelist.) SAS automatically adds subdirectories of any valid directories in the list to the lockdown path list.

A pathname can be relative or absolute. A pathname can also include operating system environment variables, SAS environment variables such as !SASROOT, the current working directory (.), and other host-specific character substitutions that are typically allowed in pathnames. For example, a UNIX path accepts ~/.

Note: SAS environment variables can be specified only at the beginning of a pathname.

LIST

Prints the current valid paths in the lockdown path list to the SAS log.

This list is an optimized path list. Subdirectories and duplicate pathnames are not shown.

Details

The LOCKDOWN statement enables you to limit access to local files and to specific SAS features for a SAS session that executes in a server session or in batch processing mode.

To use the LOCKDOWN statement, you must specify the LOCKDOWN system option in the sasv9_usermods.cfg for the specific server. For more information, see “LOCKDOWN System Option” on page 589.

A SAS server in the locked-down state validates all access to the host file system through the lockdown path list.

The lockdown path list contains all the paths that are accessible to a particular server. (This list is often referred to as a whitelist.) SAS does not verify the existence of any path in the lockdown path list. The operating system permissions on directories that are included in the lockdown whitelist are still in effect.

A path that is declared in the whitelist does not mean that an arbitrary user can read any file in that path. When the LOCKDOWN statement is not used, host permissions on physical files and directories always take precedence over the whitelist.

For details about implementing the whitelist, see Hiding the Whitelist By Locating the Path outside the Whitelist on page 593.

The lockdown path list is established during SAS initialization and finalized at the lockdown point.

There are two types of paths in the lockdown path list: default lockdown directories (paths from SAS configuration files) and user directories and files, which are added using LOCKDOWN statements in a server autoexec file.

Any modifications made to the whitelist (including defining a new stored process repository) do not affect servers that are currently running. After making changes, you can stop or quiesce any currently running processes on the affected pooled workspace server or stored process server. Your changes take effect the next time a server session is started.

For more information, see “Lock Down SAS Workspace Servers” on page 573 , “Lock Down the SAS/CONNECT Server” on page 581, and “Lock Down the SAS Compute Server” on page 563 .

By default, SAS adds the following predefined paths from the SAS configuration file and the SAS server metadata definition to the whitelist:

- stored process repository paths
- SASROOT path
- current working directory (not the User home directory)
- SASAUTOS option path or environment variable path
- UTILLOC option path
- LOGPARM defined LOG file path
- MAPS option path
- TEXTURELOC option path

- FONTSLOC option path
- JAVA_HOME/lib/fonts
- SASINITIALFOLDER (Windows only)
- “%SYSTEMROOT%\fonts” (Windows only) path

Note: Pre-assigned library paths are not automatically added to the lockdown list. Users can access pre-assigned libraries through librefs that are established through metadata or the AUTOEXEC files.

Other valid paths can be added in LOCKDOWN statements in the server’s AUTOEXEC files or the INITSTMT statement. Statements are typically added in the server autoexec_usermods.sas file.

Different types of SAS servers might require different lockdown path lists. For example, you might want a workspace server to have access only to an individual user’s home directory. Alternatively, a stored process server might require access to locations that require greater privileges. Librefs that are created by pre-assigned libraries in the metadata or autoexec file are also available to users.

Examples

Example 1: Enabling Access to a User’s Files via Servers

The servers can be any one of Workspace Server, SAS/CONNECT Server, or Compute Server.

For a server that is launched under the client user’s personal account, an administrator can enable access to the user’s personal files but prohibit access to other users. It is not necessary to establish separate server definitions or implement special logic to submit a customized LOCKDOWN statement in the server’s autoexec file. Instead, an administrator can specify a special form of the lockdown path that refers to the home directory of the user under whose account the server was launched. Here are the declarations for host platforms that the server can run on.

- The declaration for Linux : `lockdown path= '~';`
- The declaration for Windows: `lockdown path= &apos;?FOLDERID_Profile&apos;`

Example 2: Hiding the Whitelist By Locating the Path outside the Whitelist

The SAS administrator can hide the contents of the whitelist by locating the whitelist file in a path that is not on the whitelist.

In your whitelist file, define all valid paths. For example, in whitelist.txt, add this code:

```
/valid/path1
/valid/path2
```

Modify the server autoexec file by adding a `LOCKDOWN file=` statement to point to the whitelist file. In this example, **`/opt/sas/viya/config/etc/lockdown`** is *not* defined in the lockdown whitelist:

```
filename lkdn "/opt/sas/viya/config/etc/lockdown/whitelist.txt";
lockdown file = lkdn;
```

Example 3: Enabling a URL while in Lockdown Mode

In this example, an administrator allows users access to certain URL sites while SAS is in LOCKDOWN mode. To do this, the administrator adds the following global statement to the server autoexec file:

```
lockdown enable_ams = URL;
```

Example 4: Enabling Python Code while in Lockdown Mode

To enable Python code to run in SAS in lockdown mode, add the following line to the server autoexec file:

```
lockdown enable_ams=PYTHON;
```

To enable Python code to run in SAS by directly submitting through a SUBMIT block, add the following line to the server autoexec file:

```
lockdown enable_ams=PYTHON_EMBED;
```

The access methods list is used for re-enabling those access methods that are usually restricted when SAS is in lockdown mode. PYTHON and PYTHON_EMBED are added to the existing list of access methods.

See Also

Types of Sign-ons in *SAS/CONNECT User's Guide*

"LOCKDOWN System Option" on page 589

30

Infrastructure Servers

Infrastructure Servers: Overview

Note: A programming-only on page 17deployment uses only one of the infrastructure servers—Apache HTTP Server.

SAS Viya contains these infrastructure servers:

- "SAS Configuration Server"
- "SAS Secrets Manager (Linux)"
- "SAS Infrastructure Data Server"
- "SAS Message Broker"
- "SAS Cache Locator and Cache Server"
- "Apache HTTP Server"

The following diagram identifies the infrastructure service components of a SAS Viya full deployment:

Figure A.1 SAS Viya Infrastructure Servers (Full Deployment)

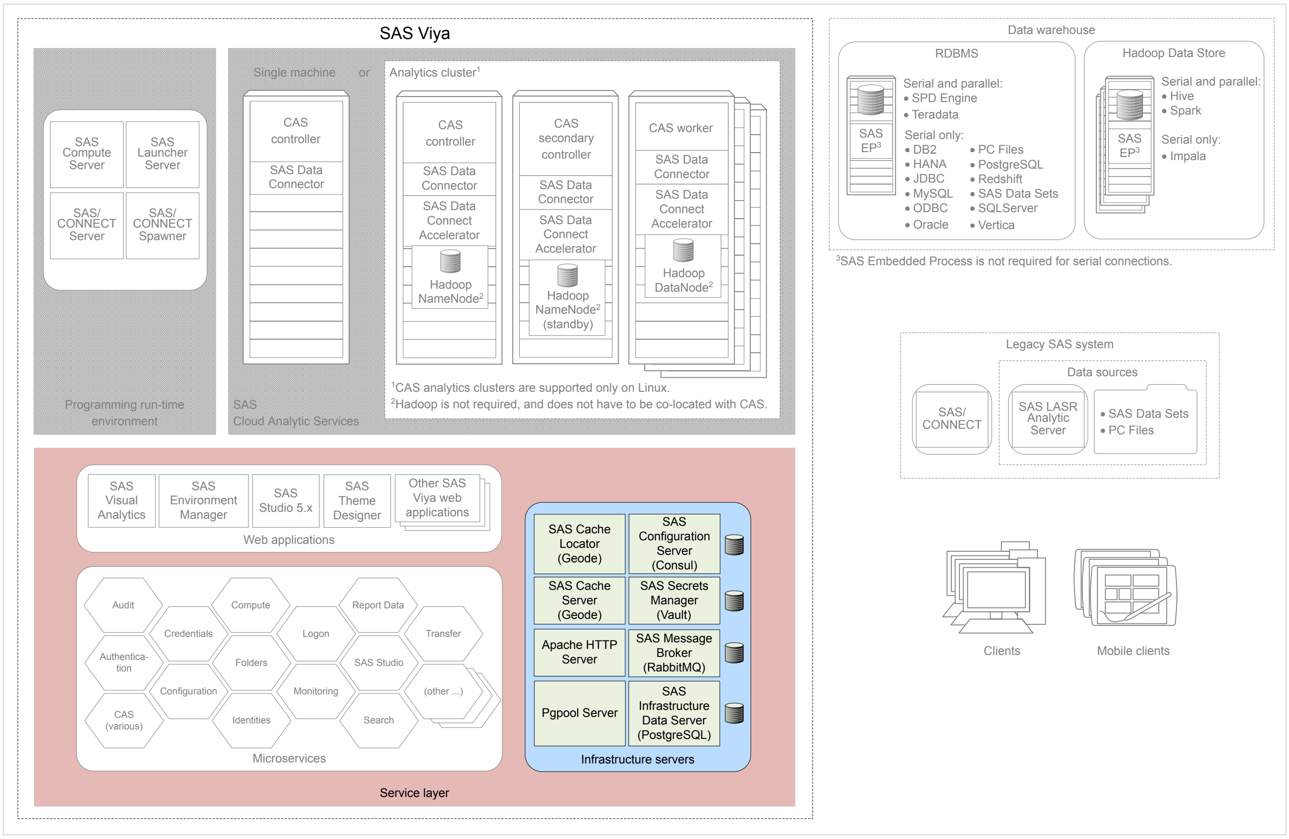

Note: SAS Secrets Manager does not run on Windows.

The following diagram shows that a programming-only deployment uses only the Apache HTTP Server from the SAS Viya server layer:

Figure A.2 *SAS Viya Infrastructure Servers (Programming-only Deployment)*

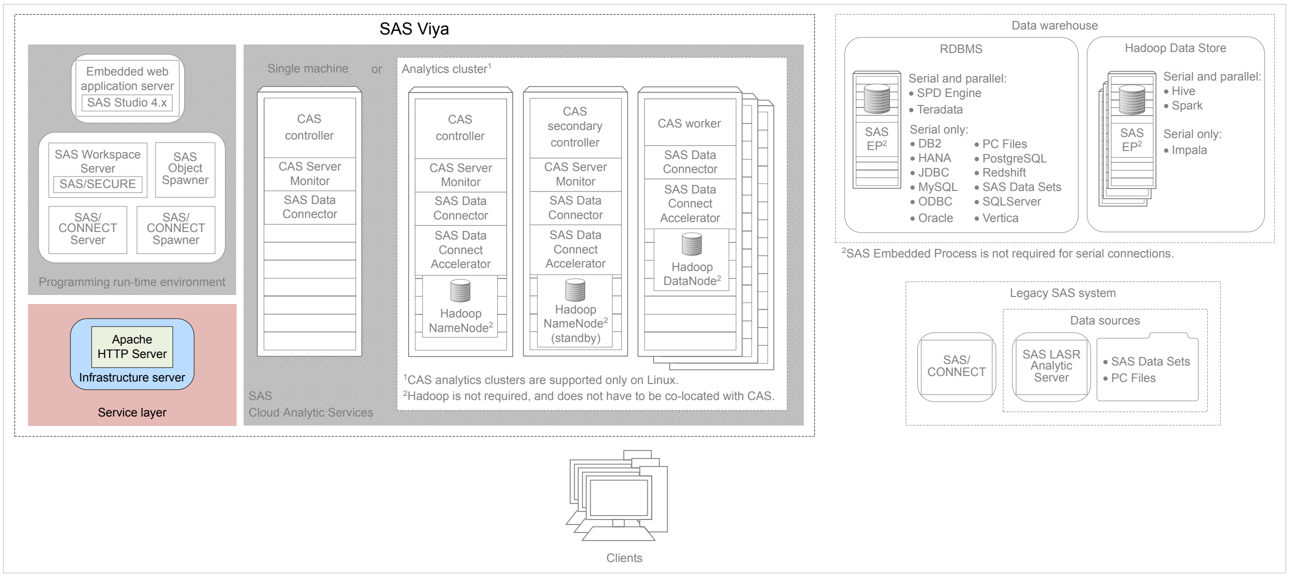

SAS Configuration Server

Overview

SAS Configuration Server is based on HashiCorp Consul. SAS Configuration Server acts as a service configuration registry that serves as a central repository for configuration data, service discovery, and health status.

Note: A programming-only deployment on page 1does not use SAS Configuration Server.

Operate (Linux)

SAS Viya provides a script in **`/etc/init.d`** that you use to stop, start, restart, and check the status of SAS Configuration Server. The script is named, `sas-viya-consul-default`.

Syntax

How you run `sas-viya-consul-default` depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status | stop | start | restart sas-viya-consul-default
```

- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-consul-default status | stop | start | restart
```

Usage Notes and Tips

- You must be logged on to the machine where SAS Configuration Server resides. Also, you must have root-level privileges to run this script.
- For multi-machine deployments, run `sas-viya-consul-default` on every SAS Viya machine. Start SAS Configuration Server *first*. Stop SAS Configuration Server *last*.
- If there are multiple instances of SAS Configuration Server, start them in this sequence:

  - First, start `sas-viya-consul-default` on all machines in the `[consul]` host group.

    The `[consul]` host group is located in the Ansible playbook inventory.ini file and it defines which machines host the SAS Configuration Server instances.

  - Next, start `sas-viya-consul-default` on all other machines in the deployment, which launches the agent processes for SAS Configuration Server.

- If there are multiple instances of SAS Configuration Server, stop them in this sequence:

  - First, stop `sas-viya-consul-default` on all machines not in the `[consul]` host group.

    The `[consul]` host group is located in the Ansible playbook inventory.ini file.

  - Next, stop `sas-viya-consul-default` on machines in the `[consul]` host group.

- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" on page 462.

- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-consul-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-consul-default` on RHEL 7.x with the `service` command, and later attempt to shut down SAS Configuration Server using the `systemctl` command, the configuration server stops responding and does not shut down.

Examples

- To check status of SAS Configuration Server on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-consul-default
```

- To stop SAS Configuration Server on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-consul-default stop
```

- To start SAS Configuration Server on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-consul-default
```

- To restart SAS Configuration Server on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-consul-default restart
```

Operate (Windows)

Using the Services snap-in in the Microsoft Management Console, you can start, stop, and restart SAS Configuration Server (Consul).

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services".

Figure A.3 SAS Configuration Server in the Services snap-in

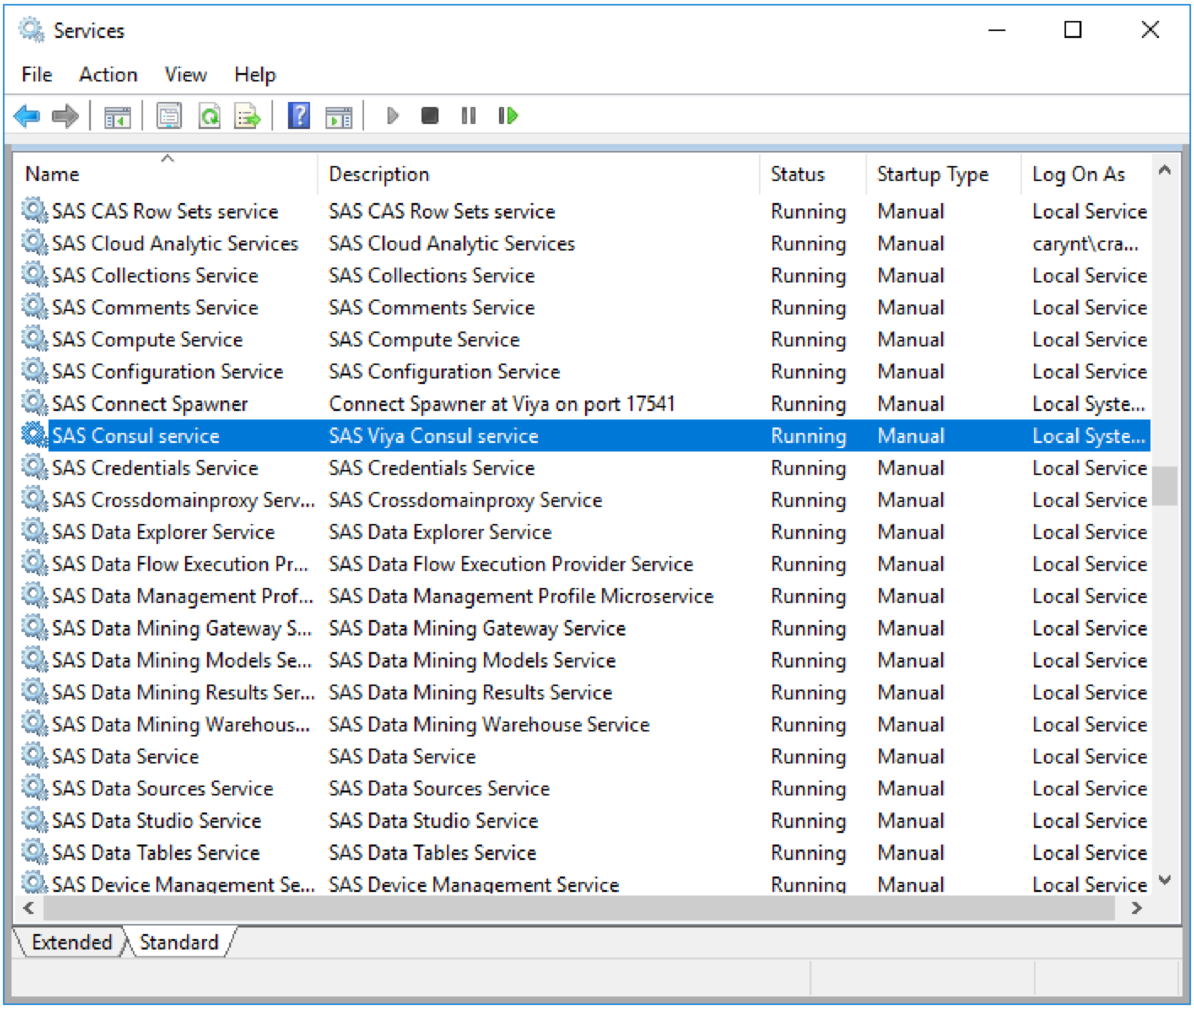

Name	Description	Status	Startup Type	Log On As
SAS CAS Row Sets service	SAS CAS Row Sets service	Running	Manual	Local Service
SAS Cloud Analytic Services	SAS Cloud Analytic Services	Running	Manual	carynt\cra...
SAS Collections Service	SAS Collections Service	Running	Manual	Local Service
SAS Comments Service	SAS Comments Service	Running	Manual	Local Service
SAS Compute Service	SAS Compute Service	Running	Manual	Local Service
SAS Configuration Service	SAS Configuration Service	Running	Manual	Local Service
SAS Connect Spawner	Connect Spawner at Viya on port 17541	Running	Manual	Local Syste...
SAS Consul service	SAS Viya Consul service	Running	Manual	Local Syste...
SAS Credentials Service	SAS Credentials Service	Running	Manual	Local Service
SAS Crossdomainproxy Serv...	SAS Crossdomainproxy Service	Running	Manual	Local Service
SAS Data Explorer Service	SAS Data Explorer Service	Running	Manual	Local Service
SAS Data Flow Execution Pr...	SAS Data Flow Execution Provider Service	Running	Manual	Local Service
SAS Data Management Prof...	SAS Data Management Profile Microservice	Running	Manual	Local Service
SAS Data Mining Gateway S...	SAS Data Mining Gateway Service	Running	Manual	Local Service
SAS Data Mining Models Se...	SAS Data Mining Models Service	Running	Manual	Local Service
SAS Data Mining Results Ser...	SAS Data Mining Results Service	Running	Manual	Local Service
SAS Data Mining Warehous...	SAS Data Mining Warehouse Service	Running	Manual	Local Service
SAS Data Service	SAS Data Service	Running	Manual	Local Service
SAS Data Sources Service	SAS Data Sources Service	Running	Manual	Local Service
SAS Data Studio Service	SAS Data Studio Service	Running	Manual	Local Service
SAS Data Tables Service	SAS Data Tables Service	Running	Manual	Local Service
SAS Device Management Se...	SAS Device Management Service	Running	Manual	Local Service

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services".

Concepts

What Is SAS Configuration Server?

SAS Configuration Server is based on HashiCorp's Consul. Consul is a distributed, highly available registry that contains service configuration data and availability and overall performance (health) information.

Configuration data resides in SAS Configuration Server as key-value pairs. This data is used by SAS Viya microservices at start-up to load default values and to discover any service dependencies.

During run time, whenever a service's properties change, the service is notified, and it rereads its properties from SAS Configuration Server. (The exceptions are noted in "What Services Must Be Restarted?".)

Each service registers its health checks when it starts. The Monitoring system periodically queries the status of the health checks.

How Does the SAS Configuration Service Work with SAS Configuration Server?

For information about how the SAS Configuration Service works with SAS Configuration Server, see "How SAS Viya Configuration Works" on page 86.

Log Files

SAS Configuration Server log files are located in **`/opt/sas/viya/config/var/log/consul/default`**.

SAS Secrets Manager (Linux)

Overview

SAS Secrets Manager is based on HashiCorp Vault. SAS Secrets Manager uses Vault to store and generate secrets such as Transport Layer Security (TLS) certificates.

Important: SAS Secrets Manager is not deployed on Windows.

Note: A programming-only deployment does not use SAS Secrets Manager. For more information, see "Deployment Types" on page 1.

Operate

SAS Viya provides a script in **`/etc/init.d`** that you use to stop, start, restart, and check the status of SAS Secrets Manager. The script is named, `sas-viya-vault-default`.

Syntax

How you run `sas-viya-vault-default` depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

  ```
  sudo systemctl status | stop | start | restart sas-viya-vault-default
  ```

- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

  ```
  sudo service sas-viya-vault-default status | stop | start | restart
  ```

Usage Notes and Tips

- You must be logged on to the machine where SAS Secrets Manager resides. Also, you must have root-level privileges to run this script.
- For multi-machine deployments, run `sas-viya-vault-default` on every SAS Viya machine that also contains SAS Configuration Server (Consul). SAS Secrets Manager (Vault) is always deployed on the same machine as the Configuration server. (Machines that contain Configuration agents do not have SAS Secrets Manager.)

  Start or restart SAS Secrets Manager immediately *after* you run SAS Configuration Server.

  Stop SAS Secrets Manager immediately *before* you stop SAS Configuration Server.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" on page 462.

- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-vault-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-vault-default` on RHEL 7.x with the `service` command, and later attempt to shut down SAS Secrets Manager using the `systemctl` command, secrets manager stops responding and does not shut down.

Examples

- To check status of SAS Secrets Manager on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-vault-default
```

- To stop SAS Secrets Manager on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-vault-default stop
```

- To start SAS Secrets Manager on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-vault-default
```

- To restart SAS Secrets Manager on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-vault-default restart
```

Concepts

What Is SAS Secrets Manager?

SAS Secrets Manager is based on HashiCorp Vault. Vault is a distributed, highly available server used to manage secrets. A *secret* is information that you want to secure, such as keys, passwords, certificates, and so on. Vault provides a secure interface to secrets, in addition to access control, and audit logging.

Here are some features of secrets manager and examples of how SAS Viya uses them:

- On-demand generation of secrets

  Secrets manager generates TLS certificates for SAS Viya servers at start-up.

- Secure storage for secrets

  Microservices use secure storage so that multiple microservice instances running on the same machine do not request multiple TLS certificates.

- Encrypt and decrypt data without storing it

  SAS Compute Server uses this feature when it sends a password to child processes.

- Revocation of secrets

  SAS Viya services use this feature when rotating security artifacts. (For example, services use vault tokens to request TLS certificates).

For more information, see “Concepts” in *Encryption in SAS Viya: Data in Motion*.

Dependency on SAS Configuration Server (Consul)

SAS Secrets Manager is installed on the same machines where SAS Configuration Server (Consul) resides. SAS Configuration Server contains a namespace where secrets manager stores secrets in encrypted form, which enables all instances of secrets manager access to consistent data. Also, secrets manager relies on the configuration server for locking and leader election. Therefore, in order for SAS Secrets Manager to be operational, the configuration server must be running.

In multiple-machine, fault tolerant deployments, SAS Secrets Manager has a primary (leader) and one or more standbys (hot standbys). For information about SAS Secrets Manager topology, see “Fault Tolerance in SAS Viya (Linux)” on page 465.

TTL Precedence Rules

The SAS Secrets Manager (Vault) time to live (TTL) properties have certain rules that you must follow. Failure to follow these rules can cause secrets manager not to start.

Figure A.4 Time to Live Properties Precedence Rules

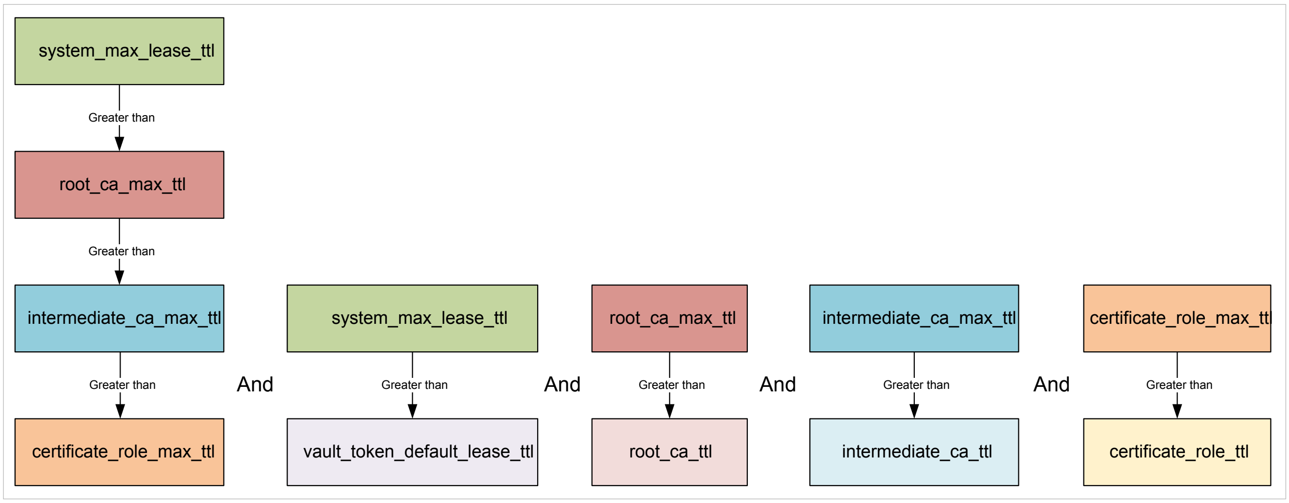

Log Files

SAS Secrets Manager log files are located in **/opt/sas/viya/config/var/log/vault/default**.

SAS Infrastructure Data Server

Overview

SAS Infrastructure Data Server is based on PostgreSQL version 9 and is configured specifically to support SAS software. SAS Infrastructure Data Server stores user content, such as reports, custom groups, comments, authorization rules, selected source definitions, attachments, audit records, and user preferences.

Note: A programming-only deployment on page 1does not use SAS Infrastructure Data Server.

How To (Cluster)

Operate a Cluster (Linux)

SAS Viya provides a script in **/etc/init.d** that you use to stop, start, restart, and check the status of a SAS Infrastructure Data Server cluster. The script is named, `sas-viya-sasdatasvrc-postgres`.

Syntax

How you run `sas-viya-sasdatasvrc-postgres` depends on your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl status | stop | start | restart sas-viya-sasdatasvrc-postgres

  Important: When using the status argument, `sas-viya-sasdatasvrc-postgres` does not display the data server node list. For this reason, we recommend that, on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, you make a direct call to `sas-viya-sasdatasvrc-postgres` when querying the data server status (for example, **cd /etc/init.d**, then **sudo ./sas-viya-sasdatasvrc-postgres status**.

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-postgres status | stop | start | restart

Usage Notes and Tips

- You must be logged on to the machine where the Pgpool server resides. Also, you must have root-level privileges to run this script.
- On multi-tenant SAS Viya systems, the script is named `sas-tenant-ID-sas-viya-sasdatasvrc-postgres`.
- In the list of running processes you see data server-related processes with **ct** in their names (for example, sas-viya-sasdatasvrc-postgres-node0-**ct**-pg_hba). These are processes that keep the data server configuration files in-sync with SAS Configuration Server (Consul). For more information, see "Synchronizing Configuration Files with SAS Configuration Server" on page 633.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" on page 462.
- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-sasdatasvrc-postgres`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-sasdatasvrc-postgres` on RHEL 7.*x* with the `service` command, and later attempt to shut down the data server cluster using the `systemctl` command, the data server stops responding and does not shut down.

Examples

- To check status of the data server cluster (and see the nodes list) on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
cd /etc/init.d

sudo ./sas-viya-sasdatasvrc-postgres status
```

- To check status of the data server cluster (and not see the nodes list) on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-sasdatasvrc-postgres
```

- To stop the data server cluster on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres stop
```

- To start the data server cluster on Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl start sas-viya-sasdatasvrc-postgres
```

- To restart the data server cluster on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres restart
```

Operate a Cluster (Windows)

Using the Services snap-in in the Microsoft Management Console, you can start, stop, and restart SAS Infrastructure Data Server.

Figure A.5 *SAS Infrastructure Data Server in the Services snap-in*

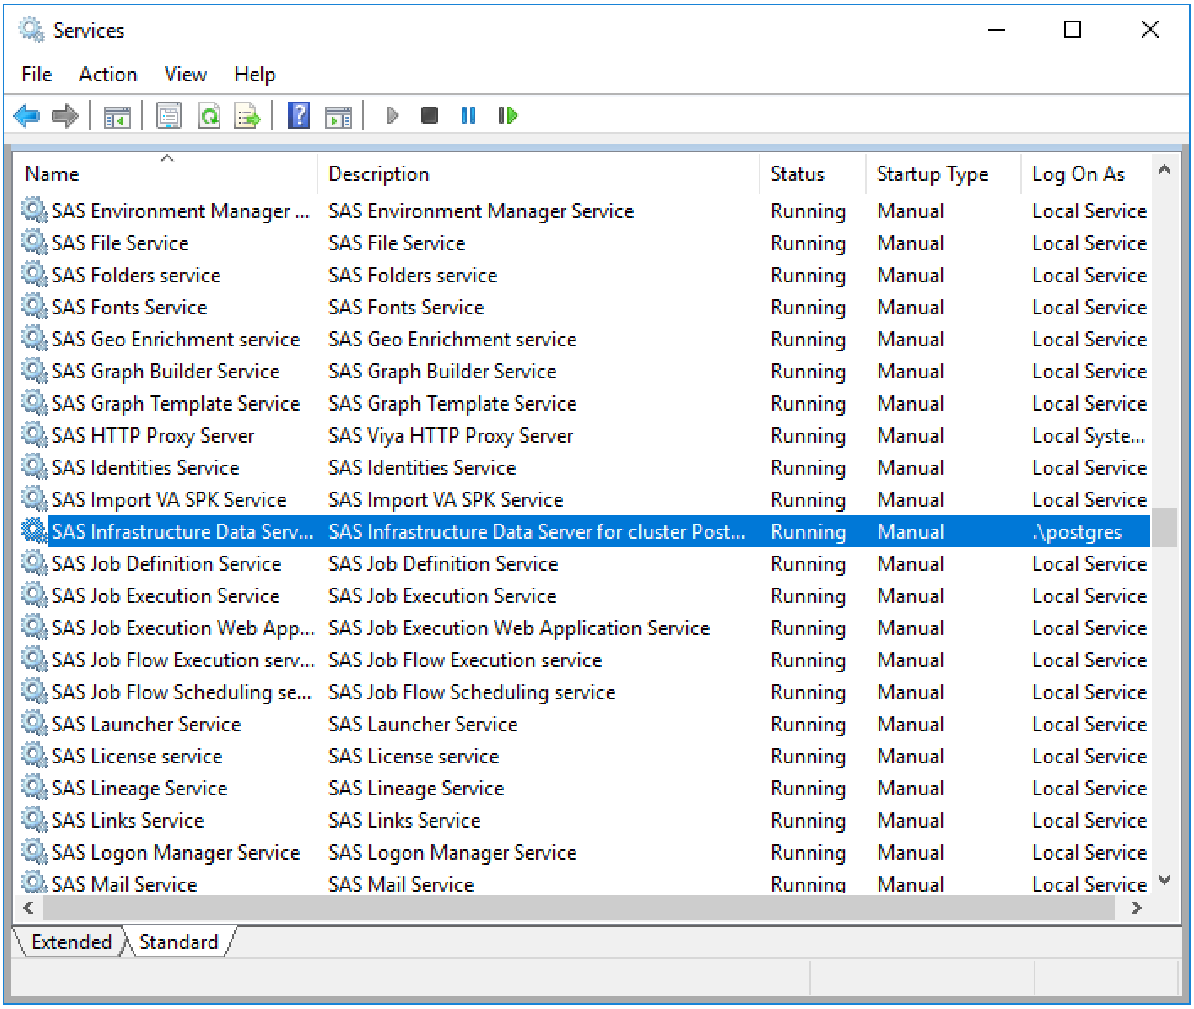

Name	Description	Status	Startup Type	Log On As
SAS Environment Manager ...	SAS Environment Manager Service	Running	Manual	Local Service
SAS File Service	SAS File Service	Running	Manual	Local Service
SAS Folders service	SAS Folders service	Running	Manual	Local Service
SAS Fonts Service	SAS Fonts Service	Running	Manual	Local Service
SAS Geo Enrichment service	SAS Geo Enrichment service	Running	Manual	Local Service
SAS Graph Builder Service	SAS Graph Builder Service	Running	Manual	Local Service
SAS Graph Template Service	SAS Graph Template Service	Running	Manual	Local Service
SAS HTTP Proxy Server	SAS Viya HTTP Proxy Server	Running	Manual	Local Syste...
SAS Identities Service	SAS Identities Service	Running	Manual	Local Service
SAS Import VA SPK Service	SAS Import VA SPK Service	Running	Manual	Local Service
SAS Infrastructure Data Serv...	SAS Infrastructure Data Server for cluster Post...	Running	Manual	.\postgres
SAS Job Definition Service	SAS Job Definition Service	Running	Manual	Local Service
SAS Job Execution Service	SAS Job Execution Service	Running	Manual	Local Service
SAS Job Execution Web App...	SAS Job Execution Web Application Service	Running	Manual	Local Service
SAS Job Flow Execution serv...	SAS Job Flow Execution service	Running	Manual	Local Service
SAS Job Flow Scheduling se...	SAS Job Flow Scheduling service	Running	Manual	Local Service
SAS Launcher Service	SAS Launcher Service	Running	Manual	Local Service
SAS License service	SAS License service	Running	Manual	Local Service
SAS Lineage Service	SAS Lineage Service	Running	Manual	Local Service
SAS Links Service	SAS Links Service	Running	Manual	Local Service
SAS Logon Manager Service	SAS Logon Manager Service	Running	Manual	Local Service
SAS Mail Service	SAS Mail Service	Running	Manual	Local Service

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services".

Recover a Failed Single Node PostgreSQL Cluster

The SAS Infrastructure Data Server cluster is considered to be failed under these conditions: when it fails to start, and when all its data nodes are marked as `unhealthy` in the cluster definition file **`/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/pool.cdf`**.

Common causes for cluster failure include a power failure, network connectivity issues, or a machine reboot. Another cause of cluster failure can be from lack of system resources. Examples are disk space, memory, number of processes, number of open files, ports, semaphores, and shared memory. In such cases, the data server logs contain failure information.

The cluster will not start until the problem that caused the cluster failure has been fixed, and the server nodes are marked as `healthy` in pool.cdf.

On the pgpool server machine, you can manually update pool.cdf, or you can run the repair_postgres_nodes.sh script. After the script updates pool.cdf, it attempts to start the cluster.

1. Before attempting any cluster repair procedures, do the following:
   - Examine the integrity of the data on the data server.

     One or more failovers might have occurred. Therefore, examine the server with the most current data.
   - Back up the data server data directories.
2. Fix the problem that caused the cluster to fail.
3. Verify the log files on both the Postgres and pgpool nodes for any errors.
   - Postgres log file:

     `/opt/sas/viya/config/var/log/sasdatasvrc/postgres/node`*`n`* , where *n* is the node number.
   - pgpool log file:

     `/opt/sas/viya/config/var/log/sasdatasvrc/postgres/pgpool0`
4. Stop the SAS Infrastructure Data Server cluster using the command that is appropriate for your operating system:
   - Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

     ```
     sudo systemctl stop sas-viya-sasdatasvrc-postgres
     ```
   - Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

     ```
     sudo service sas-viya-sasdatasvrc-postgres stop
     ```
5. Determine whether there are any leftover processes on one or more Postgres and pgpool nodes. The command lists only those leftover processes that belong to Postgres and pgpool. For command syntax details, see the Linux ps and egrep manual pages.

   ```
   ps -ef | egrep -i 'pgpool:|postgres:|all|bin/pgpool|bin/postgres'
   ```
6. Remove one or more leftover processes on the Postgres and pgpool nodes using the PID, as appropriate.

   ```
   kill PID
   ```

   Note: Make sure that you remove only the pgpool parent process. When the parent process is removed, the child processes are removed by default.

Do not use the kill -9 command. This command removes only the parent process and leaves behind the child processes.

7 As the SAS install user (sas) or the user who has root-level privileges, sign in to the pgpool server machine.

8 Choose one of the following methods to update pool.cdf:

Note: Each method attempts to restart the cluster after pool.cdf has been modified.

- Run the repair_postgres_nodes.sh script to mark all cluster nodes as healthy in pool.cdf:

  ```
  sudo /opt/sas/viya/home/libexec/sasdatasvrc/script/maintenance/
  repair_postgres_nodes.sh
  ```

- Using a text editor, open the cluster definition file and mark all of the data server nodes as healthy:

  ```
  vi /opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/pool.cdf
  ```

  Change `node0=unhealthy` to `node0=healthy`.

9 Check the value of `network.databaseTraffic.enabled` is `true` by running the following commands:

Note: Specify each of these commands on a single line. Multiple lines are used here to improve readability

```
source /opt/sas/viya/config/consul.conf

export CONSUL_HTTP_TOKEN=$(sudo cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
tokens/consul/default/client.token)

/opt/sas/viya/home/bin/sas-bootstrap-config kv read --recurse config| grep network.databaseTraffic.enabled
```

10 If the value is true, verify that the SAS Configuration Server and SAS Secret Manager (vault) are alive and accessible:

```
/opt/sas/viya/home/bin/consul members
/opt/sas/viya/home/bin/vault status
```

Here is an example:

```
/opt/sas/viya/home/bin/consul members
```

```
node_id     | Address     | status | Type    |  Build | Protocol | DC    | Segment
------------+------------+---------+--------+--------+----------+------+----------
machine1 | ip-address  | alive   | server | 1.0.6  | 2         | viya | <all>
machine2 | ip-address  | alive   | client | 1.0.6  | 2         | viya | <default>
machine3 | ip-address  | alive   | client | 1.0.6  | 2         | viya | <default>
machine4 | ip-address  | alive   | client | 1.0.6  | 2         | viya | <default>
```

```
/opt/sas/viya/home/bin/vault status
```

```
Seal Type: <seal-type>
    Sealed: false
    Key Shares: 1
    Key Threshold: 1
    Unseal Progress: 0
    Unseal Nonce:
    Version: <version-number>
    Cluster Name: <vault-cluster-name>
    Cluster ID: <cluster-id>

    High-Availability Enabled: true
            Mode: active
            Leader Cluster Address: <cluster-address>
```

11 If the value is false, verify that at least the SAS Configuration Server is alive and accessible so that the cluster can be started.

```
/opt/sas/viya/home/bin/consul members
```

12 Start the cluster using the command that is appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl start sas-viya-sasdatasvrc-postgres
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-postgres start
  ```

13 Check the status of the data server cluster.

A **status** of `up` in the cluster status list indicates that the node is connected and is an active part of the cluster. There should be only one primary server, with no standby servers or one or more standby servers, as appropriate.

Failback a Cluster

Failback refers to the restoration of the high availability (HA) SAS Infrastructure Data Server cluster to its original configuration before the failover.

A PostgreSQL *original configuration* is indicated when the cluster status list displays the following:

- node0 has the **role** of `primary`
- all other nodes have a **role** of `standby`
- all nodes have a **status** of `up`

To failback a SAS Infrastructure Data Server cluster to its original configuration before the failover, follow these steps:

1 Make sure that the following servers are running and are accessible:

- SAS Configuration Server (Consul)
- Pgpool server
- SAS Infrastructure Data Server (PostgreSQL) cluster

2 As the SAS install user (sas) or the user who has root-level privileges, sign in to the pgpool server machine.

3 To ensure that the cluster is in a failover condition, check its status by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
cd /etc/init.d
```

```
sudo ./sas-viya-sasdatasvrc-service-name status
```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name status
  ```

Verify that node0 has a **status** of `down` and a **role** of `standby`.

Here is an example of running the command on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution) where the *service-name* is named **`postgres2`**:

```
sudo service sas-viya-sasdatasvrc-postgres2 status
```

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | down   | 0.250000  | standby | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | primary | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

4 In the cluster status list, if node0 has a **status** of `up` and a **role** of `standby`, you can go directly to Step 7.

5 To recover (start) node0, start the node using the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl start sas-viya-sasdatasvrc-service-name-node0
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name-node0 start
  ```

Here is an example Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres2-node0 start
```

Here is typical output:

```
Starting sas-viya-sasdatasvrc-postgres2-node0 service...

                                                   [  OK  ]
```

6 Check the cluster status to verify that the node has started using the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  cd /etc/init.d
  ```

  ```
  sudo ./sas-viya-sasdatasvrc-service-name status
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name status
  ```

Verify that node0 has a **status** of `up`.

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

`sudo service sas-viya-sasdatasvrc-postgres2 status`

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | standby | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | primary | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

7 Now, stop the current primary node to failback to the original primary node by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  `sudo systemctl stop sas-viya-sasdatasvrc-`*`service-name-node-name`*

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  `sudo service sas-viya-sasdatasvrc-`*`service-name-node-name`*` stop`

Here is an example where *node-name* is named **`node1`** on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

`sudo service sas-viya-sasdatasvrc-postgres2-node1 stop`

Here is typical output:

```
Stopping sas-viya-sasdatasvrc-postgres2-node1 service...
                                                      [  OK  ]
```

8 Check the status of the cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  `cd /etc/init.d`

  `sudo ./sas-viya-sasdatasvrc-`*`service-name`*` status`

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  `sudo service sas-viya-sasdatasvrc-`*`service-name`*` status`

Verify that **`node0`** has a **status** of `up` and a **role** of `primary`.

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution) where *service-name* is named **`postgres2`**:

`sudo service sas-viya-sasdatasvrc-postgres2 status`

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | primary | 1          | false             | 0
 1       | machine2 | 5452 | down   | 0.250000  | standby | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

Note: The remaining running nodes of the cluster initially show a **status** of `down` while replication is established for the new primary node. Continue to monitor the cluster status until all running nodes have a **status** of `up`.

9 Now, recover (start) the stopped node by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  sudo systemctl start sas-viya-sasdatasvrc-*service-name-node-name*

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name-node-name* start

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution) where the previous primary node is named **node1**:

sudo service sas-viya-sasdatasvrc-postgres2-node1 start

Here is typical output:

```
Starting sas-viya-sasdatasvrc-postgres2-node1 service...

                                                     [  OK  ]
```

10 Re-check the status of the cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  cd /etc/init.d

  sudo ./sas-viya-sasdatasvrc-*service-name* status

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  sudo service sas-viya-sasdatasvrc-*service-name* status

Verify that the cluster has returned to its initial configuration:

- **node0** has the **role** of `primary`
- all other nodes have a **role** of `standby`
- all nodes have a **status** of `up`

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution) where the data server service is named **postgres2**:

sudo service sas-viya-sasdatasvrc-postgres2 status

Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | primary | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | standby | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

Add a Cluster (Ansible)

1 Sign on your Ansible controller with administrator privileges, and locate the file, **`/playbook/vars.yml`**.

2 Using a text editor, open vars.yml and locate the `INVOCATION_VARIABLES` section.

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

3 Copy and paste an existing cluster definition.

In this example, the new cluster is being added to Machine2:

```
INVOCATION_VARIABLES:
   Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres

   Machine2:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

4 Configure the new cluster definition for the pgpool server (pgpoolc) and the data server nodes (sasdatasvrc):

- Pgpool server definition parameters:
  - □ PCP_PORT

    Specifies the pcp port for the pgpool instance.
  - □ PGPOOL_PORT

    Specifies the pgpool port. This is the primary port through which all databases connect.
  - □ SANMOUNT

    Specifies the location of the data files. This path is typically the same value as the other data nodes.
  - □ SERVICE_NAME

    Specifies the unique service name for the data server cluster. SERVICE_NAME should be the same for the pgpool server and all nodes in the cluster.
- Data server node definition parameters:
  - □ NODE_NUMBER

    Specifies the sequential node identifier. The primary node is 0. Standby nodes start at 1 and are incremented sequentially.
  - □ NODE_TYPE

    Specifies the type of node that you are adding. The primary node should have a value of **P**. Standby nodes should have a value of **S**.
  - □ PG_PORT

    Specifies the Postgres database port. The pgpool server communicates with the database on this port. Clients use the PGPOOL_PORT. The port must be available for use on the deploy target.
  - □ SANMOUNT

    Specifies the location of the data files. This path is typically the same value as the other data nodes.
  - □ SERVICE_NAME

    Specifies the unique service name for the data server cluster. SERVICE_NAME should be the same for the pgpool server and all the nodes in the cluster.

Here is an example:

```
INVOCATION_VARIABLES:
  Machine2:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres2
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres2
    - NODE_NUMBER: '1'
```

```
      NODE_TYPE: S
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres2
```

5 If the cluster machines that you are adding are already a part of your SAS Viya system, and are already in `[pgpoolc]` and `[sasdatasvrc]` host groups, then go to Step 6.

Otherwise, add the machines to your Ansible inventory.ini file at the top of the file in the target list, and in the `[pgpoolc]` and `[sasdatasvrc]` host groups, respectively.

6 Run your Ansible playbook using the sitedefault.yml file.

Here is an example:

```
ansible-playbook site.yml
```

For a complete list of playbook commands, see "Deploy the Software" in *SAS Viya for Linux: Deployment Guide*.

Delete a Node or a Cluster

CAUTION! Do not delete a primary node unless you plan to delete the entire cluster. Deleting a primary node would increase chances of introducing data corruption. To delete the primary node, failover the node to a standby node, and wait for all remaining nodes to indicate that they are available. When the nodes are available, it is safe to delete the former primary node. Do not delete a pgpool node without first moving the pgpool node of the cluster. Failure to do so will make the cluster unusable. If you choose to delete the node data using the `-d` option, its data files are deleted. Use caution when deciding to use the `-d` option.

1 As root or with an account that has root-level privileges, sign in to the machine where the node that you want to remove resides.

2 Change the directory to **`/opt/sas/viya/home/libexec/sasdatasvrc/script`**.

3 Run the sds_delete_node.sh script with the following options:

Note: When the sds_delete_node.sh script runs, it stops the cluster.

- -s *service-name*
- -n *node-name*
- -d y | n

CAUTION! A yes (y) value specifies that the script deletes the node or the cluster data files.

- -c *absolute-path*/sds_env_var.sh

Here is an example:

```
sudo ./sds_delete_node.sh -s postgres -n node1 -d y
-c /opt/sas/viya/config/etc/sasdatasvrc/postgres/node1/sds_env_var.sh
```

Every time the script runs, it generates a new log file in **`/tmp/sds_uninstall_log`**.

4 After the script runs, be sure to delete the node or the cluster definition in the `INVOCATION_VARIABLES` section of vars.yml. For more information, see "Add a Node (Ansible)" on page 620.

Recover a Cluster When All Nodes Are Stopped in a Multi-Node High Availability Cluster

Consider this example: in a three-node cluster where the primary is on node0, someone does the following:

1 Stops node0, and waits for fail-over to node1 to complete.

2 Stops node1, and waits for fail-over to node2 to complete.

3 Stop node2, and there are no more nodes to fail over to.

As a result of these actions, there is no primary node. The cluster is unable to be stopped or restarted.

To recover such a cluster, follow these steps:

1 As the root user or a user privileges, sign in to the SAS Configuration Server (Consul) machine.

2 Verify that the SAS Configuration Server and SAS Secret Manager (Vault) are running, appropriate for your operating system:

   - Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

     ```
     sudo systemctl status sas-viya-consul-default
     ```

     ```
     sudo systemctl status sas-viya-vault-default
     ```

   - Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

     ```
     sudo service sas-viya-consul-default status
     ```

     ```
     sudo service sas-viya-vault-default status
     ```

3 If you have made the SAS Viya service layer highly available, do the following:

   - Verify that all the configuration servers are alive and accessible:

     Note: Specify each of these commands on a single line. Multiple lines are used here to improve readability

     ```
     source /opt/sas/viya/config/consul.conf

     /opt/sas/viya/home/bin/consul members

     export CONSUL_HTTP_TOKEN=$(sudo cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
     tokens/consul/default/client.token)

     /opt/sas/viya/home/bin/sas-bootstrap-config kv read config/postgres/sas.dataserver.conf/
     common/max_connections
     ```

   - Verify that the secret manager leader and its hot standbys are alive and accessible:

     ```
     /opt/sas/viya/home/bin/vault status
     ```

   Important: Bring up any nodes that are down, before you continue.

4 Log on to the primary SAS Infrastructure Data Server node machine as the SAS Installer user (sas), and set up the necessary environment variables and the Consul token:

   ```
   source /opt/sas/viya/config/consul.conf
   ```

   ```
   export CONSUL_HTTP_TOKEN=$(sudo cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
   tokens/consul/default/client.token)
   ```

5 Verify that there is one (and only one) primary node:

   ```
   /opt/sas/viya/home/bin/sas-bootstrap-config catalog service postgres
   ```

   CAUTION! Failure to locate the correct primary node, might cause data loss.

6 If two nodes are listed as primary, then locate the correct primary node. The node directory that contains the file, recovery.done, is the correct primary node. Check all the node directories (**/opt/sas/viya/config/data/sasdatasvrc/postgres/node0...*n***).

7. After you have determined which node is the correct primary node, de-register the Postgres service for the correct primary node.

   In this example, the correct primary node is node2:

   ```
   /opt/sas/viya/home/bin/sas-bootstrap-config agent service deregister postgres-datanode2
   ```

8. Create a temporary JSON file to re-register node2 as a primary node using the tag **`primary`**:

   a. Start by making a copy of the file, service_node_registration.json.

      In this example, the correct primary node is node2:

      ```
      cd /opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/
      ```

      ```
      cp /opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/service_node_registration.json
      /opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/service_node_registration.json.primary.tmp
      ```

   b. Using a text editor, open the copied JSON file (service_node_registration.json.primary.tmp) and add **`primary`** to the tag list:

      ```
        {
            "id": "postgres-datanode2",
            "name": "postgres",
            "tags": [
               "postgres", "primary"
            ],
            "address": "machine3.example.com",
            "port": 5432,
            "checks": [
              {
                "script": "/opt/sas/viya/home/libexec/sasdatasvrc/script/sds_consul_health_check.sh -config_path
                      /opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/sds_env_var.sh",
                "interval": "30s",
                "timeout": "5s"
              }
            ]
          }
      ```

9. Re-register the correct node as a primary node.

   In this example, the correct node is node2:

   ```
   /opt/sas/viya/home/bin/sas-bootstrap-config agent service register --json
   /opt/sas/viya/config/etc/sasdatasvrc/postgres/node2/service_node_registration.json.primary.tmp
   ```

10. As you did in Step 5, check the PostgreSQL service for nodes and their tags:

    ```
    /opt/sas/viya/home/bin/sas-bootstrap-config catalog service postgres
    ```

11. Repeat Step 8 – Step 10 to correct the wrong primary node, if any.

12. Stop the SAS Infrastructure Data Server cluster using the command appropriate for your operating system:

    - Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

      `sudo systemctl stop sas-viya-sasdatasvrc-postgres`

    - Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

      `sudo service sas-viya-sasdatasvrc-postgres stop`

13. Using a text editor, open **`/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/pool.cdf`**, and specify the correct primary node as **`healthy`** and all standby nodes as **`unhealthy`**.

In this example, node2 is the correct primary node:

```
cluster_definition_file_version=10.0.1
  pgpool_server_version=3.6.6
  postgresql_server_version=9.4.17
  pgpool0=healthy
  node0=unhealthy
  node1=unhealthy
  node2=healthy
  pgpoolConfigured
```

14 Using a text editor, open the PostgreSQL and Pgpool configuration files and verify that all parameters have proper values and that there are no missing values.

- Pgpool configuration file:

  ```
  /opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pgpool.conf
  ```

- PostgreSQL configuration files:

  ```
  /opt/sas/viya/config/data/sasdatasvrc/postgres/noden/postgresql.conf
  ```

> **TIP** Pgpool does not recognize a primary node if back-end parameters are missing in the pgpool.conf file: `backend_data_directoryn`, `backend_hostnamen`, `backend_portn`, `backend_weightn`, where *n* is the node number.

15 If you find any missing values in the configuration files that you verified in Step 14, then refresh those configuration files by stopping and restarting the corresponding consul-template services.

Here is an example for node0 on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres-node0-ct-postgresql stop

sudo service sas-viya-sasdatasvrc-postgres-node0-ct-postgresql start

sudo service sas-viya-sasdatasvrc-postgres-node0-ct-pg_hba stop

sudo service sas-viya-sasdatasvrc-postgres-node0-ct-pg_hba start
```

Here is an example for Pgpool on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pgpool stop

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pgpool start

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pool_hba stop

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pool_hba start

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pcp stop

sudo service sas-viya-sasdatasvrc-postgres-pgpool0-ct-pcp start
```

16 After you have verified that the PostgreSQL and Pgpool configuration files are correct, start the cluster using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl start sas-viya-sasdatasvrc-postgres
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-postgres start
  ```

17 Check the status of the cluster using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
cd /etc/init.d
```

```
sudo ./sas-viya-sasdatasvrc-postgres status
```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-postgres status
  ```

In this example, the two nodes are down:

```
PGPool is running with PID=21226
Checking Postgresql nodes status...
 node_id |       hostname       | port | status | lb_weight |  role   | select_cnt | load_balance_node |
replication_delay
---------+---------------------+------+--------+-----------+---------+------------+-------------------
+-------------------
 0       | machine1.example.com | 5432 | down   | 0.333333  | standby | 0          | false             | 0
 1       | machine2.example.com | 5432 | down   | 0.333333  | standby | 0          | false             | 0
 2       | machine3.example.com | 5432 | up     | 0.333333  | primary | 159        | true              | 0
(3 rows)
```

18 Recover the stopped nodes. Starting the node automatically recovers it and attaches it back to the cluster.

In this example, node0 and node1 are the stopped nodes that must be recovered. The commands used here are for Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres-node0 start
```

```
sudo service sas-viya-sasdatasvrc-postgres-node1 start
```

19 Re-check the status of the cluster.

(For the necessary command, see Step 17.)

In this example, the two nodes that were down are now up:

```
PGPool is running with PID=21226
Checking Postgresql nodes status...
 node_id |       hostname       | port | status | lb_weight |  role   | select_cnt | load_balance_node |
replication_delay
---------+---------------------+------+--------+-----------+---------+------------+-------------------
+-------------------
 0       | machine1.example.com | 5432 | up     | 0.333333  | standby | 0          | false             | 0
 1       | machine2.example.com | 5432 | up     | 0.333333  | standby | 0          | false             | 0
 2       | machine3.example.com | 5432 | up     | 0.333333  | primary | 159        | true              | 0
(3 rows)
```

How To (Nodes)

Check the Status of a Node (Linux)

SAS Viya uses the operating system's default init system or systemd command to launch a script that can check the status of a SAS Infrastructure Data Server node. This script, `sas-viya-sasdatasvrc-postgres-noden`, resides in **`/etc/init.d`**.

Note: As the SAS install user (sas) or the user who has root-level privileges, you must be signed in to the machine where the node resides.

Note: Each node script is numbered, starting at zero (0).

Run the script, using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

```
sudo systemctl status sas-viya-sasdatasvrc-postgres-noden
```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-postgres-noden status
  ```

Note: On Red Hat Enterprise Linux 7.*x*, the first time you run the `sas-viya-sasdatasvrc-postgres-noden` command, you must run the command twice. Red Hat Enterprise Linux 7.*x* has backward compatibility for the system V service, and it does not have initial systemd unit files. The first time the `sas-viya-sasdatasvrc-postgres-noden` command runs, it builds the service unit file from the system V init file. Therefore, until the unit file is built, the `sas-viya-sasdatasvrc-postgres-noden` commands do not function properly.

Stop a Node (Linux)

SAS Viya uses the operating system's default init system or systemd command to launch a script that can stop a SAS Infrastructure Data Server node. This script, `sas-viya-sasdatasvrc-postgres-noden`, resides in **`/etc/init.d`**.

CAUTION! The act of stopping individual nodes changes the cluster state. Stopping the primary node causes a failover to occur in a cluster of two or more nodes. In addition, stopping a standby node removes the node from the active cluster. Stopped nodes must be recovered in order for them to added back to the cluster. The recovery of stopped nodes occurs automatically during node start-up. During failover of the primary node (0), other healthy standby nodes (2 and 3) go through a process of "following" the new primary node (1). During failover, nodes 2 and 3 briefly detach from the cluster and display a status of 3 (unhealthy). Wait several minutes and then recheck the cluster status. Eventually, nodes 2 and 3 should re-attach to the cluster and display a **status** of `up` (healthy).

Note: As the SAS install user (sas) or the user who has root-level privileges, you must be signed in to the machine where the node resides.

Note: Each node script is numbered, starting at zero (0).

Note: On Red Hat Enterprise Linux 7.*x*, the first time you run the `sas-viya-sasdatasvrc-postgres-noden` command, you must run the command twice. Red Hat Enterprise Linux 7.*x* has backward compatibility for the system V service, and it does not have initial systemd unit files. The first time the `sas-viya-sasdatasvrc-postgres-noden` command runs, it builds the service unit file from the system V init file. Therefore, until the unit file is built, the `sas-viya-sasdatasvrc-postgres-noden` commands do not function properly.

Run the script, using the command appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl stop sas-viya-sasdatasvrc-postgres-noden
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-postgres-noden stop
  ```

Start a Node (Recover a Node) (Linux)

A SAS Infrastructure Data Server data node is considered to be "unhealthy" when it has a **status** of `down` in the cluster status list. If a PostgreSQL data node has been stopped or has been taken offline, the pgpool server removes this node from the cluster.

When you restart an unhealthy node, pgpool server automatically initiates the node recovery process. To recover an unhealthy data node, follow these steps:

1 Make sure that the following servers are running and accessible:

  - SAS Configuration Server (Consul)

- pgpool server
- SAS Infrastructure Data Server (PostgreSQL) cluster

2 As the SAS install user (sas) or the user who has root-level privileges, sign in to the pgpool server machine, and run the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  `cd /etc/init.d`

  `sudo ./sas-viya-sasdatasvrc-`*`service-name`*` status`

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  `sudo service sas-viya-sasdatasvrc-`*`service-name`*` status`

Verify that the unhealthy data node has a **status** of `down` and a **role** of `standby`.

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution), where the data server service is named **`postgres2`**:

`sudo service sas-viya-sasdatasvrc-postgres2 status`

Here is typical output:

Note: In this example, the unhealthy node is **`node0`**.

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | down   | 0.250000  | standby | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | primary | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

3 As the SAS install user (sas) or the user who has root-level privileges, sign in to the machine that contains the unhealthy data node.

4 Make sure that the unhealthy node is stopped by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  `sudo systemctl stop sas-viya-sasdatasvrc-`*`service-name-node-name`*

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  `sudo service sas-viya-sasdatasvrc-`*`service-name-node-name`*` stop`

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution), where the unhealthy node is named **`node0`**.

`sudo service sas-viya-sasdatasvrc-postgres2-node0 stop`

Here is typical output:

```
Service sas-viya-sasdatasvrc-postgres2-node0 is not running.
                                                         [  OK  ]
```

5 Recover the node as a standby server by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl start sas-viya-sasdatasvrc-service-name-node0
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name-node0 start
  ```

The pgpool server automatically starts the unhealthy node.

A node **status** of `up` indicates that the node is connected and is an active part of the cluster. There should be only one server with a **role** of `primary`, with zero or more servers with a **role** of `standby`.

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres2-node0 start
```

Here is typical output:

```
Starting sas-viya-sasdatasvrc-postgres2-node0 service...

                                                   [  OK  ]
```

6 Make sure that the node has been successfully added to the cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  cd /etc/init.d
  sudo ./sas-viya-sasdatasvrc-service-name status
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name status
  ```

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres2 status
```

Here is typical output:

Note: In this example, the previously unhealthy node (`node0`) has a **status** of `up`.

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | standby | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | primary | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

Note: If starting (recovering) a node fails, refer to the Troubleshooting section on page 634.

Add a Node (Ansible)

Adding a data node to your SAS Infrastructure Data Server cluster consists of modifying the vars.yml file and running your Ansible playbook.

1 With administrator privileges, sign in to your Ansible controller , and locate the file, `/playbook/vars.yml`.

2 Using a text editor, open vars.yml and locate the INVOCATION_VARIABLES section.

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

3 Copy an existing node definition and place it under the deploy target on which the node will be configured.

Here is an example:

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

4 Configure the node definition in order to meet the requirements of the cluster:

- NODE_NUMBER

  Specifies the sequential node identifier. Standby nodes start at 1 and are incremented sequentially.

  For example, if you have only a primary node, the node that you are adding should have a NODE_NUMBER of 1. If the last standby node in your cluster has the value of 1, the node that you are adding should have a NODE_NUMBER of 2.

- NODE_TYPE

  Specifies the type of node that you are adding. The only acceptable value is S (standby). After initial deployment, you cannot add a primary node.

- PG_PORT

  Specifies the Postgres database port. The pgpool server communicates with the database on this port. Clients use the PGPOOL_PORT. The port must be available for use on the deploy target.

- SANMOUNT

  Specifies the location of the data files. This path is typically the same value as the other data nodes.

- SERVICE_NAME

  Specifies the service name for the data server cluster. It must be an exact match of the name of the cluster to which you are adding a data node.

Here is an example:

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    - NODE_NUMBER: '1'
      NODE_TYPE: S
      PG_PORT: '5433'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

5 Run your Ansible playbook using the sitedefault.yml file.

Here is an example:

`ansible-playbook site.yml`

For a complete list of playbook commands, see “Deploy the Software” in *SAS Viya for Linux: Deployment Guide*.

Move a Node (Ansible)

Moving a data node to your SAS Infrastructure Data Server cluster consists of modifying the vars.yml file and running your Ansible playbook.

1 With administrator privileges, sign in to your Ansible controller , and locate the file **`/playbook/vars.yml`**.

2 Using a text editor, open vars.yml and locate the `INVOCATION_VARIABLES` section.

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    - NODE_NUMBER: '1'
      NODE_TYPE: S
      PG_PORT: '5433'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

3. Copy the existing node definition and place it under the deploy target to which you want to move the node.

   In this example, the deploy target is Machine2:

   ```
   INVOCATION_VARIABLES:
     Machine2:
       pgpoolc:
       - PCP_PORT: '5430'
         PGPOOL_PORT: '5431'
         SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
         SERVICE_NAME: postgres2
       sasdatasvrc:
       - NODE_NUMBER: '0'
         NODE_TYPE: P
         PG_PORT: '5432'
         SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
         SERVICE_NAME: postgres2
       - NODE_NUMBER: '1'
         NODE_TYPE: S
         PG_PORT: '5433'
         SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
         SERVICE_NAME: postgres2
   ```

4. Configure the node definition to meet the requirements of the cluster:

   - NODE_NUMBER

     Specifies the sequential node identifier. This number should change to fit with the target cluster. For example, if the last standby node in the cluster is 1, the node that you are moving should have a `NODE_NUMBER` of 2. If there is only a primary node in the target cluster, the node that you are moving should have a `NODE_NUMBER` of 1.

   - NODE_TYPE

     Specifies the type of node that you are moving. The only acceptable value is `S` (standby). After initial deployment, you cannot move a primary node.

   - PG_PORT

     Specifies the Postgres database port. The pgpool server communicates with the database on this port. Clients use the `PGPOOL_PORT`. The port must be available for use on the deploy target.

   - SANMOUNT

     Specifies the location of the data files. This path is typically the same value as other data nodes.

   - SERVICE_NAME

     Specifies the service name for the data server cluster.

     Note: Do not change this value.

5. Run your Ansible playbook using the sitedefault.yml file.

   Here is an example:

   `ansible-playbook site.yml`

   For a complete list of playbook commands, see “Deploy the Software” in *SAS Viya for Linux: Deployment Guide*.

Change the Port Number or the Data Directory for a Node (Ansible)

CAUTION! To avoid data corruption, do not change the port number or the data directory on a primary node.

1 As the SAS install user (sas) or the user who has root-level privileges, sign in to the pgpool machine.

2 To stop the node whose port or directory you want to change, run the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl stop sas-viya-sasdatasvrc-service-name-node-name
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name-node-name stop
  ```

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres-node3 stop
```

3 To check the status of the node, run the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  cd /etc/init.d
  ```

  ```
  sudo ./sas-viya-sasdatasvrc-service-name status
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name status
  ```

Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

```
sudo service sas-viya-sasdatasvrc-postgres status
```

Verify that failover has successfully occurred. In the cluster status list, the **status** of the node should be `down`.

4 With administrator privileges, sign in to your Ansible controller, and locate the file`/playbook/vars.yml`.

5 Using a text editor, open vars.yml and locate the `INVOCATION_VARIABLES` section.

```
INVOCATION_VARIABLES:
  Machine1:
    pgpoolc:
    - PCP_PORT: '5430'
      PGPOOL_PORT: '5431'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    sasdatasvrc:
    - NODE_NUMBER: '0'
      NODE_TYPE: P
      PG_PORT: '5432'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
    - NODE_NUMBER: '1'
      NODE_TYPE: S
      PG_PORT: '5433'
      SANMOUNT: '{{ SAS_CONFIG_ROOT }}/data/sasdatasvrc'
      SERVICE_NAME: postgres
```

6 Make the necessary changes to the port number or the data directory in the definition for the node:

- NODE_NUMBER

  Specifies the sequential node identifier.

  Note: Do not change this value.

- NODE_TYPE

  Specifies the type of node: `P` (primary) or `S` (standby).

  Note: Do not change this value.

- PG_PORT

  Specifies the Postgres database port. The pgpool server communicates with the database on this port. Clients use the `PGPOOL_PORT`. The port must be available for use on the deploy target.

- SANMOUNT

  Specifies the location of the data files. This path is typically the same value as other data nodes.

- SERVICE_NAME

  Specifies the service name for the data server cluster.

  Note: Do not change this value.

7 Run your Ansible playbook using the sitedefault.yml file.

Here is an example:

```
ansible-playbook site.yml
```

For a complete list of playbook commands, see "Deploy the Software" in *SAS Viya for Linux: Deployment Guide*.

8 Check the status of the node. In the cluster status list, the **status** of the node should be `up`.

How To (General)

Get Current Passwords (Linux)

1 As the SAS install user (sas) or the user who has root-level privileges, sign in to any SAS Infrastructure Data Server machine.

2 Obtain the security token from the configuration server, and set it as an environment variable, using the appropriate command:

- Install user or root accounts:

```
source /opt/sas/viya/config/consul.conf
```

```
export CONSUL_HTTP_TOKEN=$(cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
tokens/consul/default/client.token)
```

- With root-level privileges (but not as the install user), install accounts:

```
source /opt/sas/viya/config/consul.conf
```

```
export CONSUL_HTTP_TOKEN=$(sudo cat /opt/sas/viya/config/etc/SASSecurityCertificateFramework/
tokens/consul/default/client.token)
```

3 Run the sas-bootstrap-config script for the data server user ID whose password you want to obtain:

- sas

```
/opt/sas/viya/home/bin/sas-bootstrap-config kv read config/postgres/sas.dataserver.pooluser/
```

```
common/sr_check_password
```

- dbmsowner

  ```
  /opt/sas/viya/home/bin/sas-bootstrap-config kv read config/application/sas/database/postgres/password
  ```

Change User Passwords (Linux)

The script, sds_change_user_pw.sh, changes SAS Infrastructure Data Server passwords and synchronizes them with SAS Configuration Server (Consul) and configuration files.

CAUTION! To avoid data loss, change the sas user account password only during a scheduled maintenance when users are not accessing SAS Viya. The data server must be running when you change the sas user's password Changing the password for the database user, sas, causes all nodes on the database cluster to restart.

Note: To change the password, you must know the current password. For more information, see "Get Current Passwords (Linux)".

1 As the SAS install user (sas), sign in to the SAS Infrastructure Data Server Pgpool machine.

   Note: The change user password script requires sudo execution privileges.

2 You can determine the status of your cluster by running the following command, appropriate for your operating system:

   - Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

     ```
     cd /etc/init.d
     sudo ./sas-viya-sasdatasvrc-service-name status
     ```

   - Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

     ```
     sudo service sas-viya-sasdatasvrc-service-name status
     ```

   Before you run the change password script, verify that the cluster is in its initial configuration state (running and healthy):

   - node0 has the **role** of `primary`
   - all other nodes have a **role** of `standby`
   - all nodes have a **status** of `up`

   Here is an example on Red Hat Enterprise Linux 6.*x* (or an equivalent distribution), where the data server service is named **`postgres2`**:

   ```
   sudo service sas-viya-sasdatasvrc-postgres2 status
   ```

   Here is typical output:

```
Checking status of sas-viya-sasdatasvrc-postgres2...

PGPool is running with PID=11445
Checking Postgresql nodes status...
 node_id | hostname | port | status | lb_weight |  role   | select_cnt | load_balance_node | replication_delay
---------+----------+------+--------+-----------+---------+------------+-------------------+-------------------
 0       | machine1 | 5452 | up     | 0.250000  | primary | 1          | false             | 0
 1       | machine2 | 5452 | up     | 0.250000  | standby | 0          | true              | 0
 2       | machine3 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
 3       | machine4 | 5452 | up     | 0.250000  | standby | 0          | false             | 0
(4 rows)
```

3 Locate the data server environment variables file, sds_env_var.sh, and record its location.

By default, sds_env_var.sh resides in **`/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0`**.

4. The script prompts for the following information. Have this information ready when you run the script in a later step:

   - database user name
   - current database password
   - new database password

     Note: Your password must conform to the data server password policy on page 635.

5. Using the location of sds_env_var.sh noted in Step 3, run the script using the following command:

   ```
   sudo -Hu sas /opt/sas/viya/home/libexec/sasdatasvrc/script/sds_change_user_pw.sh
   -config_path
   /opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/sds_env_var.sh
   ```

   TIP If you run the script from the directory where it resides, you might see several `cannot open [No such file or directory]` messages. This is a known issue, and you can safely ignore these messages.

6. Enter the information that you collected in Step 4 as the script prompts you for it.

   After you provide the values in response to the prompts, the script connects to SAS Configuration Server and updates all instances of the database user password that it finds. Changes made in the configuration server are synchronized with the proper SAS Infrastructure Data Server configuration files. Finally, the script issues the necessary SQL commands in the data server to update the permissions for the database user.

7. To validate that your password has successfully changed, connect to the data server first database, **postgres**, using the PostgreSQL interactive terminal, psql:

   ```
   /opt/sas/viya/home/bin/psql -h data-server-machine-name -U user-IDservice-name
   ```

8. When prompted, enter the new password for dbmsowner.

9. Restart all SAS Viya services.

   For more information, see “Start and Stop All Servers and Services” on page 462.

Clean Up after a Hardware Failure (Linux)

If the machine on which the high availability (HA) SAS Infrastructure Data Server cluster runs was stopped unexpectedly, you might need to perform some cleanup steps after you restart the machine.

These steps involve removing any socket-lock files and any PID files that might have become orphaned after the PostgreSQL and pgpool servers were improperly shut down.

1. As the SAS install user (sas) or with root-level privileges, sign in to the pgpool machine.

2. Stop the HA data server cluster by running the following command, appropriate for your operating system:

   - Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

     ```
     sudo systemctl stop sas-viya-sasdatasvrc-service-name
     ```

   - Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

     ```
     sudo service sas-viya-sasdatasvrc-service-name stop
     ```

3 Delete any socket-lock file (in the form .s.PGSQL.**xxxx**) or any PID file (in the form *server*.pid) that corresponds to your HA data server cluster ports.

For the default HA data server instance with one data node, remove the following files:

- **/tmp/.s.PGSQL.5430**
- **/tmp/.s.PGSQL.5431**
- **/tmp/.s.PGSQL.5432**
- **/tmp/.s.PGSQL.5432.lck**
- **/opt/sas/viya/config/data/sasdatasvrc/postgres/node0/postmaster.pid**
- **/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/run/pgpool.pid**

4 Restart the HA data server cluster by running the following command, appropriate for your operating system:

- Red Hat Enterprise Linux 7.*x* (or an equivalent distribution) and SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl start sas-viya-sasdatasvrc-service-name
  ```

- Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service sas-viya-sasdatasvrc-service-name start
  ```

Remove a Persistent Lock on a Database Table (Linux)

Persistent locks on a SAS Infrastructure Data Server database table are caused by an uncommitted transaction or a long running query. To fix this problem, you must identify the process IDs of the client connections that are locking the table and terminate these connections.

1 As the SAS install user (sas) or with root-level privileges, sign in to the pgpool server machine.

2 If you know the PostgreSQL dbmsowner (superuser) password, go to Step 3. Otherwise, follow the steps in “Get Current Passwords (Linux)”.

3 If you have not already done so, install pgAdmin on any machine (including Microsoft Windows) that has access to the machine that is running the pgpool server.

4 In pgAdmin, perform the following steps:

a Create a **New Server Registration** object and specify the following information:

- **Host**: ***machine-name***

  The name of the machine on which the pgpool server resides.

- **Port**: ***pgpool-client-connection-port***

  The default port is `5431`.

- **Maintenance DB**: **SharedServices**

  Do not use the default, `postgres`.

- **Username**: *superuser*

  The database superuser. The default is `dbmsowner`.

- **Password**: string

  The superuser password.

b Connect to the pgpool server.

c Highlight the server name, and choose **Tools** ⇨ **Server Status**.

The status panel shows all the client connections in the top panel. The second panel shows the persistent locks.

d Choose **Actions** ⇨ **Refresh** multiple times in order to determine whether the listed locks are transient or persistent. A transient lock disappears, and a persistent lock remains throughout refreshes.

e (Optional) You can cross-reference the process identifiers (PIDs) for the locked tables with the connection listing in order to identify the client (the application name) that has locked the table.

Note: If you choose this option, open a SAS Technical Support track about this issue. Include the **PID**, **Application Name**, **Connection State**, and **Query** (if applicable) from the top connections section. Also include the PID and the persistent locked table names from the **Lock** section.

f To clear the locks, run the **pg_terminate_backend()** command on each PID that has a persistent lock.

To do this, go back to the main pgAdmin panel. Highlight the **SharedServices** database name in the server **Object Browser** and choose **Tools** ⇨ **Query Tool** to open an SQL query execution window.

g Execute the **pg_terminate_backend(__PID__)** command to close each connection that is associated with a table that has a persistent lock.

Here is an example:

```
SELECT pg_terminate_backend(14826);
SELECT pg_terminate_backend(16697);
SELECT pg_terminate_backend(22246);
```

h Select **Tools** ⇨ **Server Status** and refresh the panel (**Actions** ⇨ **Refresh** from the menu).

If all the persistent locks have been removed from the second **Locks** section, the persistent locks are successfully removed.

i Exit pgAdmin.

Routine Maintenance Tasks

Overview

Routine maintenance for a SAS Infrastructure Data Server consists of the following tasks:

- Adhering to a rigid schedule of performing database backups.
- Performing a re-index, vacuuming, and analyzing each table in the database during a maintenance cycle.
- Inspecting the data server logs periodically to make sure that there is no data corruption.
- Removing large orphaned data objects from the database periodically to free disk space.
- Track PostgreSQL software patches, and apply the patches that contain critical fixes, such as the CVE-2018-10915 security patch.

Re-Index, Vacuum, and Analyze Database Tables (Linux)

Follow these steps to re-index, vacuum, and analyze each table in the SAS Infrastructure Data Server databases. SAS recommends that you perform these steps during a maintenance cycle in order to reduce the chance of a PostgreSQL database command hanging because of a long-term lock on a table. If you encounter a hang condition, to remove the lock, you might need to restart the SAS Infrastructure Data Server.

1 As the SAS install user (sas) or with root-level privileges, sign in to the pgpool server machine.

2 If you know the PostgreSQL dbmsowner (superuser) password, go to Step 3. Otherwise, follow the steps in "Get Current Passwords (Linux)".

3 Run the following commands to set up the PostgreSQL command-line environment:

```
export PATH=/opt/sas/viya/home/bin:$PATH
```

```
export LD_LIBRARY_PATH=/opt/sas/viya/home/lib:/opt/sas/viya/home/
lib64:$LD_LIBRARY_PATH
```

4 (Optional) Stop all SAS Viya services, and run only SAS Infrastructure Data Server.

5 Run the following commands:

Note: For illustration, `5431` is used as the client connection port, **`/opt/sas/viya/home`** is used as the installation directory, and `dbmsowner` is used as the database superuser. Substitute the values that are appropriate for your site.

> **TIP** To prevent having to enter the superuser password multiple times, you can create a **`~/.pgpass`** file.

- Re-index all databases:

```
./reindexdb -a -p 5431 -h localhost -U dbmsowner
```

- Perform a full vacuum and analyze all databases:

```
./vacuumdb -p 5431 -h localhost -U dbmsowner -f -v -z -a
```

Note: If you encounter a hang condition, you might need to restart the SAS Infrastructure Data Server to remove the lock.

If there were no errors in the previous step, then you are done.

If you stopped the all of the SAS Viya services, then you can now stop the SAS Infrastructure Data Server and restart all the SAS Viya services.

Remove Large Orphaned Data Objects (Linux)

Large objects in the SAS Infrastructure Data Server database are stored separately from the tables that reference them. When a particular row is updated or deleted, these objects can become orphaned (unattached) from a table. Periodically, these orphaned large objects must be manually removed to free disk space.

1 Create an SQL command file named lo-cleanup.sql with the following content:

```
DROP FUNCTION IF EXISTS sas_lob_cleanup();
CREATE FUNCTION sas_lob_cleanup() RETURNS VOID AS $function$
DECLARE
   possible_lob_col record;
   possible_oid_row record;
   possible_oid_val bigint;
BEGIN
   DROP TABLE IF EXISTS sas_possible_lobs;

   CREATE TABLE sas_possible_lobs (table_schema TEXT NOT NULL, table_name
   TEXT NOT NULL, column_name TEXT NOT NULL, column_value BIGINT);

   FOR possible_lob_col IN SELECT * FROM information_schema.columns WHERE
   udt_name IN ('int4', 'int8', 'numeric', 'oid', 'text', 'varchar',
   'char', 'lo') AND NOT (table_schema = 'pg_catalog' AND table_name =
   'pg_shdepend') AND NOT (table_schema = 'pg_catalog' AND table_name =
   'pg_largeobject') AND table_name != 'sas_possible_lobs'

   LOOP

      BEGIN

          FOR possible_oid_val IN EXECUTE 'SELECT CAST(' ||
      possible_lob_col.column_name || ' AS BIGINT) FROM ' ||
      possible_lob_col.table_schema || '.' || possible_lob_col.table_name
      || ' WHERE ' || possible_lob_col.column_name || ' IS NOT NULL AND
      CAST(' || possible_lob_col.column_name || ' AS BIGINT) < ((2 ^ 32) -
      1) AND CAST(' || possible_lob_col.column_name || ' AS BIGINT) > 0'

          LOOP

          --raise notice 'successfully cast % % %',
          possible_lob_col.table_schema,
          possible_lob_col.table_name,
          possible_lob_col.column_name;

              INSERT INTO sas_possible_lobs (table_schema, table_name,
          column_name, column_value) VALUES (possible_lob_col.table_schema,
          possible_lob_col.table_name, possible_lob_col.column_name,
          possible_oid_val);

           END LOOP;

           EXCEPTION

               WHEN cannot_coerce THEN

                  --RAISE NOTICE 'error coercing % % %',
          possible_lob_col.table_schema, possible_lob_col.table_name,
          possible_lob_col.column_name;

               WHEN invalid_text_representation THEN

                  --RAISE NOTICE 'error casting % % %',
          possible_lob_col.table_schema, possible_lob_col.table_name,
          possible_lob_col.column_name;

               WHEN others THEN

                  RAISE NOTICE 'unexpected failure';

           END;

   END LOOP;

   SELECT LO_UNLINK(lo.loid) FROM pg_catalog.pg_largeobject lo GROUP BY loid
   HAVING (NOT EXISTS (SELECT 1 FROM public.sas_possible_lobs pl WHERE lo.loid
   = pl.column_value));

   DROP TABLE IF EXISTS sas_possible_lobs;

END;
```

```
$function$ LANGUAGE PLPGSQL;

SELECT sas_lob_cleanup();
```

2 As the SAS install user (sas) or with root-level privileges, sign in to the pgpool server machine

3 Run the following command for each database:

   Note: For illustration, `5431` is used as the client connection port, **`/opt/sas/viya/home`** is used as the installation directory, and `dbmsowner` is used as the database superuser. Substitute the values that are appropriate for your site.

   ```
   psql -p 5431 -h localhost -U dbmsowner -d postgres -a -f lo-cleanup.sql
   ```

   ```
   psql -p 5431 -h localhost -U dbmsowner -d SharedServices -a -f lo-cleanup.sql
   ```

Apply the CVE-2018-10915 Security Patch

A new security patch, CVE-2018-10915, fixes a vulnerability was found in libpq, the default PostgreSQL client library.

Sites that deploy SAS Viya 3.4 on Windows have a newer version of PostgreSQL (version 9.4.19) that contains the fix for this security issue. However, sites running SAS Viya 3.4 or earlier on Red Hat Enterprise Linux and SUSE Linux Enterprise Server, must manually apply patch SAS Infrastructure Data Server with patch CVE-2018-10915.

1 As the SAS install user (sas) or with root-level privileges, sign in to the pgpool server machine.

2 Stop all SAS Viya services.

3 Perform an update using `yum`, or run Ansible with the appropriate playbook.

4 Restart all the SAS Viya services.

Concepts

What is the SAS Infrastructure Data Server?

SAS Infrastructure Data Server is used for transactional storage by SAS middle-tier software. It is also used by some SAS solutions software for user content such as reports, custom groups, comments, authorization rules, selected source definitions, attachments, and user preferences. The server is configured specifically to support SAS software, and is based on PostgreSQL version 9.

By default, the SAS installer account is used to start the server.

The databases that are managed by the server are backed up and restored with the Backup and Recovery Deployment Tool. For more information, see *SAS Viya Administration: Backup and Restore*.

Figure A.6 SAS Infrastructure Data Server Architecture

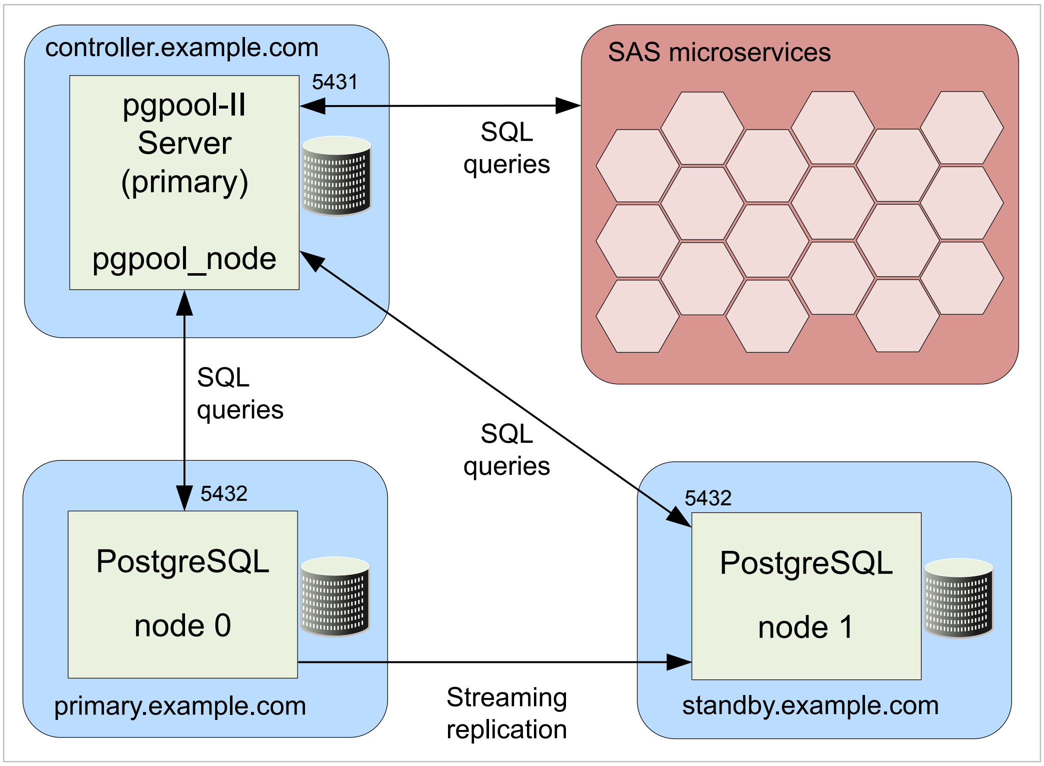

Pgpool-II

SAS provides Pgpool-II (version 3) open-source software to enable you to manage PostgreSQL clusters for high availability (failover management). The Pgpool-II software resides and operates between SAS Infrastructure Data servers and clients. All data connections and database requests are routed through the pgpool service.

Synchronizing Configuration Files with SAS Configuration Server

The SAS Infrastructure Data Server uses a daemon called consul-template to keep configuration files synchronized with configuration properties that are modified in SAS Environment Manager and stored as key-value pairs in SAS Configuration Server (Consul). For more information, see “Non-Spring-Based Servers” on page 86.

Each configuration file for each Pgpool node and PostgreSQL data node of the cluster has an init script to manage the synchronization. The consul-template init scripts for the SAS Infrastructure Data Server cluster are the following:

- For each Pgpool
  - □ sas-viya-sasdatasvrc-postgres-pool0-ct-pgpool
  - □ sas-viya-sasdatasvrc-postgres-pool0-ct-pool_hba
  - □ sas-viya-sasdatasvrc-postgres-pool0-ct-pcp
- For each data server node, where *n* is the node number, such as node0:
  - □ sas-viya-sasdatasvrc-postgres-node*n*-pg_hba
  - □ sas-viya-sasdatasvrc-postgres-node*n*-postgresql

Troubleshooting

psql: server closed the connection unexpectedly. This probably means the server terminated abnormally before or while processing the request.

Explanation:

The SAS Viya environment was shut down abnormally.

Resolution:

Restart the SAS Viya environment using the sas-viya-all-services start command. For more information, see "Start and Stop All Servers and Services" on page 462.

/opt/sas/viya/config/etc/sasdatasvrc/../node.cdf was already marked with 'recoveryInProgress=y'. Exiting from auto-recovery.

Explanation:

PostgreSQL was in the process of recovering the SAS Infrastructure Data Server node when it encountered an error, and stopped the recovery process. Whenever it restarts a data server node, PostgreSQL always inserts the line, `recoveryInProgress=y`, in the node.cdf file to avoid a simultaneous recovery.

Resolution:

1. Review the recovery log to determine what the problem is.

   (The recovery log is located here: **`/opt/sas/viya/config/var/log/sasdatasvrc/cluster/nodex/sds_auto_recovery_node.log`**.)

2. Fix the problem.

3. Remove the following line from the node's node.cdf file:

   ```
   recoveryInProgress=y
   ```

4. Restart (recover) the node.

EDTERROR: missing chunk number 0 for toast value 9558737 in pg_toast_2619
EDTCONTEXT: automatic analyze of table "SharedServices.public.sas_audit"
EDTERROR: could not read block 3062 in file "base/18797/19703": read only 0 of 8192 bytes
***yyyy-mm-dd* EDTERROR: unexpected data beyond EOF in block 0 of relation base/16715/107679**

Explanation:

There is a high probability that your SAS Infrastructure Data Server database is corrupted.

Resolution:

After you correct the cause of the data corruption, and recover the database using a restored backup.

ERROR: Cluster stop failed. Please review the log file. /opt/sas/viya/config/var/log/sasdatasvrc/postgres/pgpool0/sas-viya-sasdatasvrc-postgres-service_YYYYMMDD_####.log [FAILED]
Unexpected response code: 500 ERROR: Unable to read a key Unexpected response code: 500 (rpc error: failed to get conn: dial tcp someotherhost.com:8300: getsockopt: connection refused) ERROR: Unable to list the nodes that provide the service 'postgres'

Explanation:

SAS Infrastructure Data Server fails to stop because it cannot connect to SAS Configuration Server (Consul) even though Consul is running.

This problem occurs in a multi-machine deployment where the primary data server node is not on the same system as the primary Consul server (the server designated in the inventory.ini file). If the primary Consul server has already been shut down, the data server fails to stop, even if a local Consul service is running.

Resolution:

Always stop SAS Infrastructure Data Server before the primary Consul server, regardless of which machine it is located on. For more information, see "Read This First: Start and Stop Servers and Services" on page 458.

Reference

Database

TIP All PostgreSQL data servers have a *first database* named **postgres**. For more information, see Creating a Database in PostgreSQL documentation.

In a SAS Viya deployment, SAS Infrastructure Data Server is configured to manage the SharedServices database. SAS Viya microservices create database schemas within SharedServices.

If your deployment includes SAS solutions software that supports SAS Infrastructure Data Server, more databases might be configured on the server.

Default Users

dbmsowner
: The PostgreSQL database owner and the SAS database administrator user.

sas
: The SAS Viya install user and the account used for SAS Infrastructure Data Server cluster management.

Network Access

SAS Infrastructure Data Server is configured to accept connections on all network interfaces, and it requires password authentication. By default, SAS configures the server to use network port number 5431.

PostgreSQL instances are configured with JDBC data sources that reference the SharedServices database.

Password Policy

The user name and password for the SAS Infrastructure Data Server administrator are specified during deployment. The password can be updated. Passwords for SAS Infrastructure Data Server are subject to the following guidelines:

- The password must not contain any non-alphanumeric characters.

  Examples are underscores (_), hyphens (-), and periods (.).
- The password must be at least six characters long.
- The password can contain alphanumeric characters.
- There are no restrictions for including numbers or mixed-case characters.

Environment Parameters (Linux)

Export the following path in order to execute PostgreSQL and pgpool commands:

```
export LD_LIBRARY_PATH=/opt/sas/viya/home/lib:/opt/sas/viya/home/lib64
```

Configuration Files (Linux)

- `/opt/sas/viya/config/etc/sasdatasvrc/postgres/node0/node.cdf`

- **/opt/sas/viya/config/etc/sasdatasvrc/postgres/pgpool0/pool.cdf**
- **/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pgpool.conf**
- **/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pcp.conf**
- **/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pool_hba.conf**
- **/opt/sas/viya/config/data/sasdatasvrc/postgres/pgpool0/pool_passwd**
- **/opt/sas/viya/config/data/sasdatasvrc/postgres/node0/postgresql.conf**
- **/opt/sas/viya/config/data/sasdatasvrc/postgres/node0/pg_hba.conf**

Configuration File (Windows)

The SAS Infrastructure Data Server configuration file on Windows is **\ProgramData\SAS\Viya\etc\sasdatasvrc\postgres\node0\node.cdf**.

Log Files

SAS Infrastructure Data Server log files are located in the following path, depending on your operating system:

- Linux:

  /opt/sas/viya/config/var/log/sasdatasvrc
- Windows:

  \ProgramData\SAS\Viya\var\log\sasdatasvrc\postgres\node0

SAS Message Broker

Overview

SAS uses a set of event APIs that are dependent on Spring Integration and Spring AMQP to interact with the message broker. The AMQP-compliant message broker that SAS uses is Pivotal's RabbitMQ, version 3. RabbitMQ includes the Erlang platform, version 20.

Note: A programming-only deployment on page 1does not use SAS Message Broker.

Operate (Linux)

SAS Viya provides a script in **/etc/init.d** that you use to stop, start, restart, and check the status of SAS Message Broker. The script is named, sas-viya-rabbitmq-server-default.

Syntax

How you run sas-viya-rabbitmq-server-default depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

  ```
  sudo systemctl status | stop | start | restart sas-viya-rabbitmq-server-default
  ```
- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

  ```
  sudo service sas-viya-rabbitmq-server-default status | stop | start | restart
  ```

Usage Notes and Tips

- You must be logged on to the machine where SAS Message Broker resides. Also, you must have root-level privileges to run this script.
- If there are multiple instances of SAS Message Broker:
  - start them in the reverse sequence in which you stopped them.
  - stop them in the reverse sequence in which you started them.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" on page 462.
- On Linux systems that support systemd, use the `systemctl` command when running `sas-viya-rabbitmq-server-default`. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-rabbitmq-server-default` on RHEL 7.x with the `service` command, and later attempt to shut down SAS Message Broker using the `systemctl` command, the configuration server stops responding and does not shut down.

Examples

- To check status of SAS Message Broker on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-rabbitmq-server-default
```

- To stop SAS Message Broker on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-rabbitmq-server-default stop
```

- To start SAS Message Broker on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-rabbitmq-server-default
```

- To restart SAS Message Broker on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-rabbitmq-server-default restart
```

Operate (Windows)

Using the Microsoft Management Console (MMC) Services snap-in, you can start, stop, and restart SAS Message Broker.

Figure A.7 SAS Message Broker in the Services Snap-In

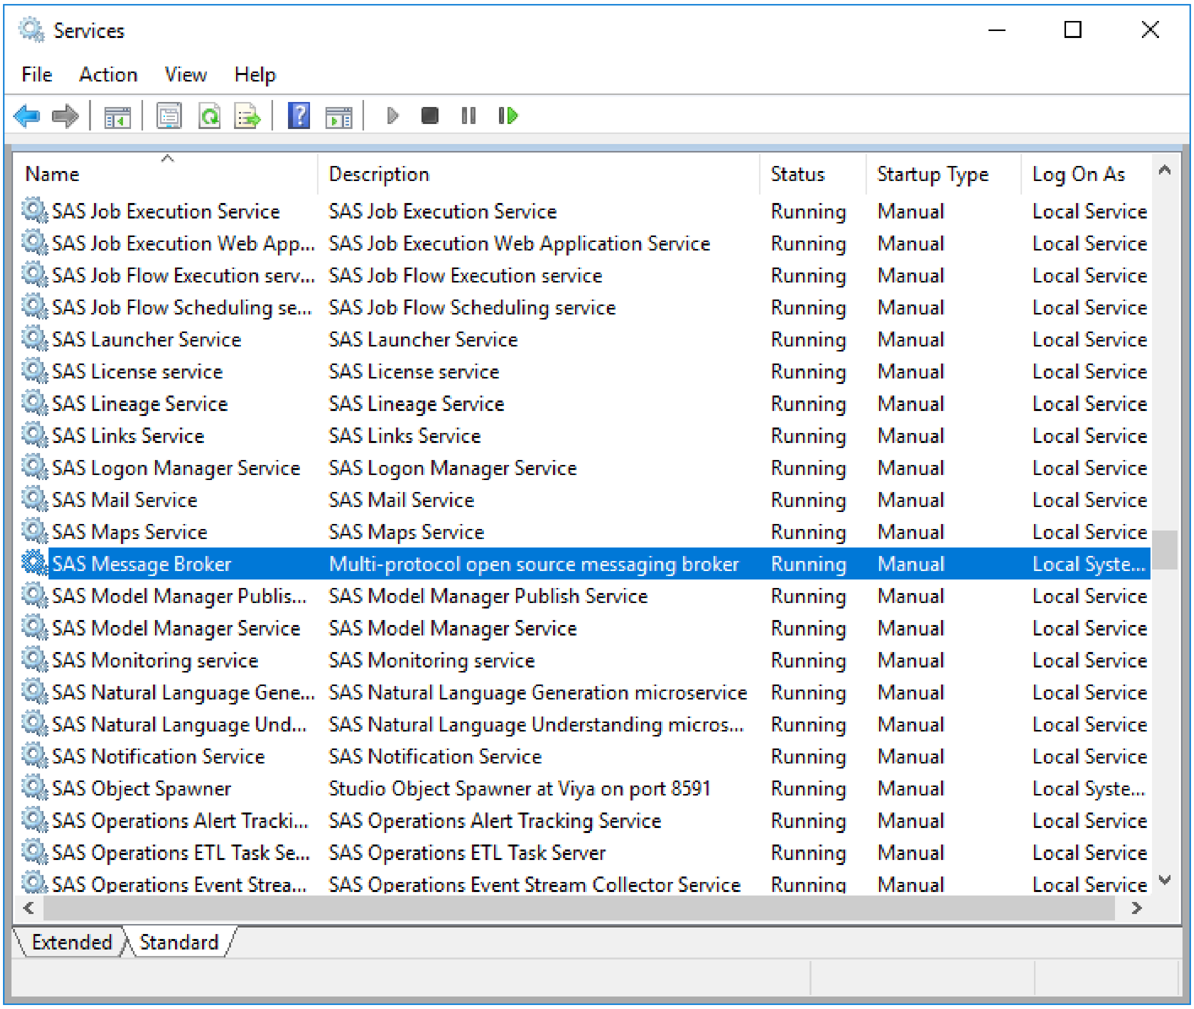

Name	Description	Status	Startup Type	Log On As
SAS Job Execution Service	SAS Job Execution Service	Running	Manual	Local Service
SAS Job Execution Web App...	SAS Job Execution Web Application Service	Running	Manual	Local Service
SAS Job Flow Execution serv...	SAS Job Flow Execution service	Running	Manual	Local Service
SAS Job Flow Scheduling se...	SAS Job Flow Scheduling service	Running	Manual	Local Service
SAS Launcher Service	SAS Launcher Service	Running	Manual	Local Service
SAS License service	SAS License service	Running	Manual	Local Service
SAS Lineage Service	SAS Lineage Service	Running	Manual	Local Service
SAS Links Service	SAS Links Service	Running	Manual	Local Service
SAS Logon Manager Service	SAS Logon Manager Service	Running	Manual	Local Service
SAS Mail Service	SAS Mail Service	Running	Manual	Local Service
SAS Maps Service	SAS Maps Service	Running	Manual	Local Service
SAS Message Broker	Multi-protocol open source messaging broker	Running	Manual	Local Syste...
SAS Model Manager Publis...	SAS Model Manager Publish Service	Running	Manual	Local Service
SAS Model Manager Service	SAS Model Manager Service	Running	Manual	Local Service
SAS Monitoring service	SAS Monitoring service	Running	Manual	Local Service
SAS Natural Language Gene...	SAS Natural Language Generation microservice	Running	Manual	Local Service
SAS Natural Language Und...	SAS Natural Language Understanding micros...	Running	Manual	Local Service
SAS Notification Service	SAS Notification Service	Running	Manual	Local Service
SAS Object Spawner	Studio Object Spawner at Viya on port 8591	Running	Manual	Local Syste...
SAS Operations Alert Tracki...	SAS Operations Alert Tracking Service	Running	Manual	Local Service
SAS Operations ETL Task Se...	SAS Operations ETL Task Server	Running	Manual	Local Service
SAS Operations Event Strea...	SAS Operations Event Stream Collector Service	Running	Manual	Local Service

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see "Start and Stop All Servers and Services".

Concepts

What is SAS Message Broker?

SAS Message Broker is an integral part of the event-driven architecture in which SAS Viya services participate. SAS uses a set of event APIs that are dependent on Spring Integration and Spring AMQP for interacting with the message broker. The AMQP-compliant message broker that SAS uses is Pivotal's RabbitMQ. The SAS event APIs provide a layer of abstraction between the message broker and its clients. The SAS event APIs also prevent code from breaking, which could result if SAS changed its third-party message broker from RabbitMQ to another third-party message broker in the future.

How Does Message Broker Work?

SAS Message Broker accepts messages in a standard format and routes them through exchanges and queues, which provide transaction acknowledgment, message persistence, and redundancy. Message broker exchanges accept messages from publishers and route them to queues, as appropriate. The exchange type controls whether messages are sent to a specific queue, to all associated queues, or only to queues that accept a particular message routing key or that match a key pattern.

SAS Message Broker Reference

Exchanges

SAS Message Broker uses the following exchanges:

- sas.application
- sas.application.backup
- sas.backup.topic
- sas.ledger
- sas.log
- sas.metric
- sas.notification
- sas.search.schema.topic

Configuration Files

SAS Message Broker configuration files are located in **`/opt/sas/viya/config/etc/rabbitmq-server/`**.

Note: Change these configuration files only when instructed to do so by SAS Technical Support.

Log Files

SAS Message Broker log files are located in **`/opt/sas/viya/config/var/log/rabbitmq-server/default`**.

SAS Cache Locator and Cache Server

Overview

SAS Cache Locator and SAS Cache Server provide a distributed cache technology to microservices in SAS Viya.

See Also

Tuning on page 647

Operate (Linux)

SAS Viya provides scripts in **/etc/init.d** that you use to stop, start, restart, and check the status of SAS Cache Locator and SAS Cache Server. The scripts are named `sas-viya-cachelocator-default` and `sas-viya-cacheserver-default`, respectively.

Syntax

How you run `sas-viya-cachelocator-default` and `sas-viya-cacheserver-default` depends on your operating system:

- Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

  sudo systemctl status | stop | start | restart sas-viya-cachelocator-default

  sudo systemctl status | stop | start | restart sas-viya-cacheserver-default

- Red Hat Enterprise Linux 6.x (or an equivalent distribution):

  sudo service sas-viya-cachelocator-default status | stop | start | restart

  sudo service sas-viya-cacheserver-default status | stop | start | restart

Usage Notes and Tips

- You must be logged on to the machine where the SAS Cache Locator and SAS Cache Server services reside. Also, you must have root-level privileges to run these scripts.
- There is a script with which you can manage and view the running state of all SAS Viya services. For more information, see "Start and Stop All Servers and Services" on page 462.
- On Linux systems that support systemd, use the `systemctl` command when running the individual service and server scripts. The `systemctl` command maintains a record of service status that the `service` command and a direct call does not use.

CAUTION! On Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x, do not mix System V init and systemd commands. Mixing the System V init (`service` command) with the systemd (`systemctl` command) causes several issues. The `systemctl` command knows nothing about a SAS Viya service started with the `service` command. If you start `sas-viya-cachelocator-default` on RHEL 7.x with the `service` command, and later attempt to shut down SAS Cache Locator using the `systemctl` command, the cache locator stops responding and does not shut down.

Examples

- To check status of SAS Cache Locator on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl status sas-viya-cachelocator-default
```

- To stop SAS Cache Server on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-cacheserver-default stop
```

- To start SAS Cache Locator on Red Hat Enterprise Linux 7.x (or an equivalent distribution) and SUSE Linux Enterprise Server 12.x:

```
sudo systemctl start sas-viya-cachelocator-default
```

- To restart CAS Cache Server on Red Hat Enterprise Linux 6.x (or an equivalent distribution):

```
sudo service sas-viya-cacheserver-default restart
```

Operate (Windows)

Using the Microsoft Management Console (MCC) Services snap-in, you can start, stop, and restart SAS Cache Locator and the SAS Cache Server.

Figure A.8 *SAS Cache Locator in the Services Snap-In*

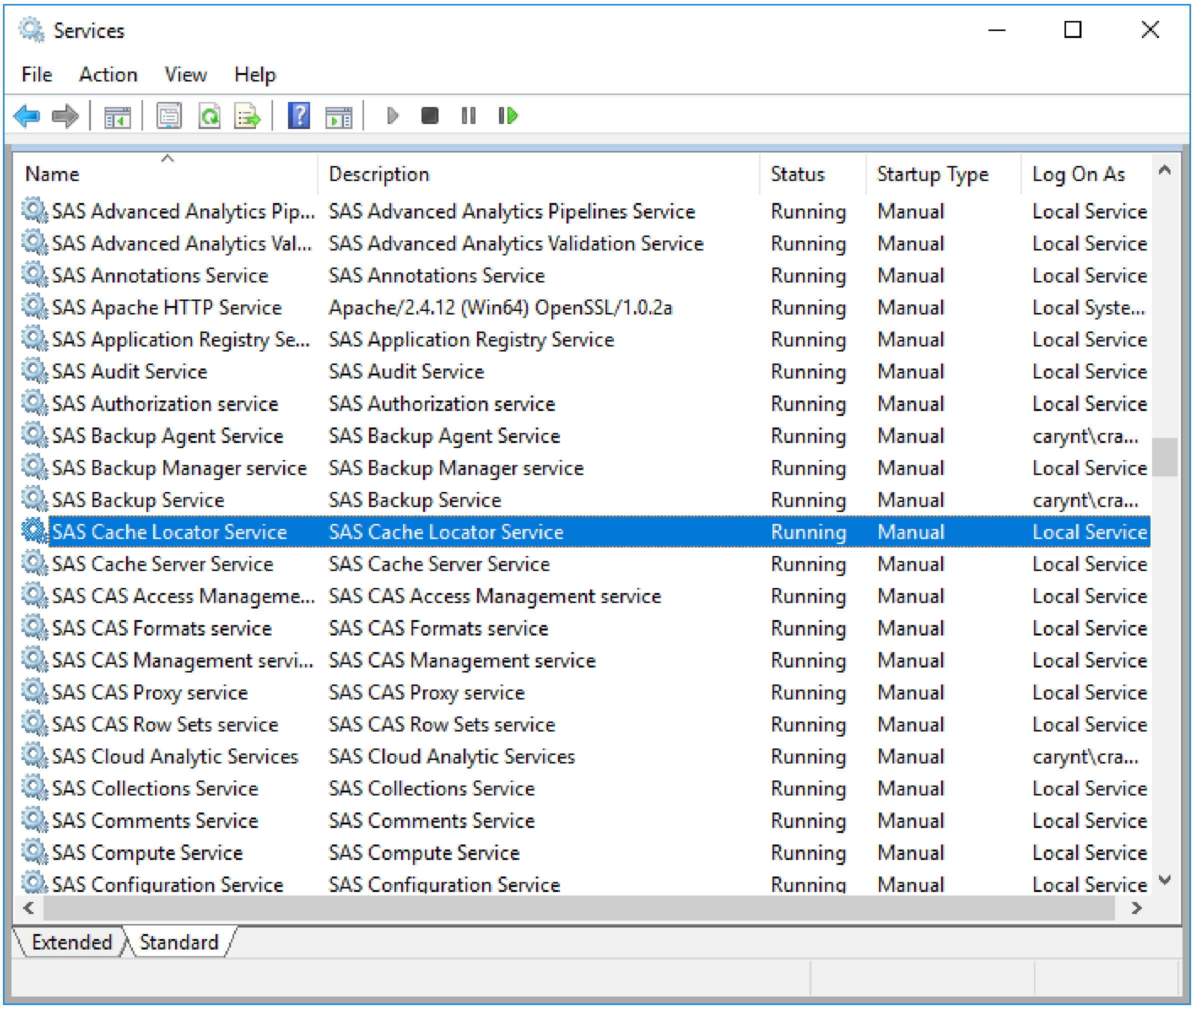

Name	Description	Status	Startup Type	Log On As
SAS Advanced Analytics Pip...	SAS Advanced Analytics Pipelines Service	Running	Manual	Local Service
SAS Advanced Analytics Val...	SAS Advanced Analytics Validation Service	Running	Manual	Local Service
SAS Annotations Service	SAS Annotations Service	Running	Manual	Local Service
SAS Apache HTTP Service	Apache/2.4.12 (Win64) OpenSSL/1.0.2a	Running	Manual	Local Syste...
SAS Application Registry Se...	SAS Application Registry Service	Running	Manual	Local Service
SAS Audit Service	SAS Audit Service	Running	Manual	Local Service
SAS Authorization service	SAS Authorization service	Running	Manual	Local Service
SAS Backup Agent Service	SAS Backup Agent Service	Running	Manual	carynt\cra...
SAS Backup Manager service	SAS Backup Manager service	Running	Manual	Local Service
SAS Backup Service	SAS Backup Service	Running	Manual	carynt\cra...
SAS Cache Locator Service	SAS Cache Locator Service	Running	Manual	Local Service
SAS Cache Server Service	SAS Cache Server Service	Running	Manual	Local Service
SAS CAS Access Manageme...	SAS CAS Access Management service	Running	Manual	Local Service
SAS CAS Formats service	SAS CAS Formats service	Running	Manual	Local Service
SAS CAS Management servi...	SAS CAS Management service	Running	Manual	Local Service
SAS CAS Proxy service	SAS CAS Proxy service	Running	Manual	Local Service
SAS CAS Row Sets service	SAS CAS Row Sets service	Running	Manual	Local Service
SAS Cloud Analytic Services	SAS Cloud Analytic Services	Running	Manual	carynt\cra...
SAS Collections Service	SAS Collections Service	Running	Manual	Local Service
SAS Comments Service	SAS Comments Service	Running	Manual	Local Service
SAS Compute Service	SAS Compute Service	Running	Manual	Local Service
SAS Configuration Service	SAS Configuration Service	Running	Manual	Local Service

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see “Start and Stop All Servers and Services”.

Concepts

SAS Cache Locator

SAS Cache Locator is a server that provides discovery information to SAS Viya microservices for the purpose of forming a distributed data cache. SAS Cache Locator is based on the open-source Apache Geode project.

SAS Cache Server

SAS Cache Server hosts long-lived data regions (a cache) and serves the contents to SAS Viya microservices. Like SAS Cache Locator, SAS Cache Server is based on the open-source Apache Geode project.

Configuration

SAS Cache Locator and SAS Cache Server embed the Apache Geode API within their respective SAS Viya microservices, cachelocator and cacheserver.

The cachelocator and cacheserver microservices enable the cache locator and the cache server to gain access to SAS Configuration Server (Consul) in order to dynamically register and to retrieve properties with the SAS Viya Configuration service. For more information, see "Non-Spring-Based Servers" on page 86.

When configuration changes are made to cachelocator and cacheserver, you must restart SAS Cache Locator and SAS Cache Server in order for their changes to take effect. For information about how to modify the configuration for cachelocator and cacheserver, see "Edit Configuration Instances" on page 75.

Log Files

Log files for SAS Cache Locator and SAS Cache Server are located in `/opt/sas/viya/config/var/log/cachelocator/default` and `/opt/sas/viya/config/var/log/cacheserver/default`.

Apache HTTP Server

Overview

SAS Viya uses Apache HTTP Server to serve static HTML content and to proxy client connections. A high-availability proxy environment is not installed by default, but is a supported configuration.

SAS Viya supports the following versions of Apache HTTP Server:

- Red Hat Linux 6.*x* uses Apache HTTP Server upstream v2.2.
- Red Hat Linux 7.*x* uses Apache HTTP Server upstream v2.4.
- SUSE Linux Enterprise Server 12.*x* uses Apache HTTP Server upstream v2.4.
- Windows uses Apache HTTP Server v2.4.34.

For more information, see "Apache httpd" in *SAS Viya for Linux: Deployment Guide*.

How To

Operate (Linux)

Note: You must be logged on to the machine where SAS HTTP Proxy Server resides, and you must have root-level privileges to run this script.

Note: For complete information about httpd arguments, see https://httpd.apache.org/docs/2.0/programs/httpd.html.

To the operate SAS HTTP Proxy Server on the appropriate operating system:

- On Red Hat Enterprise Linux 7.*x* (or an equivalent distribution):

```
sudo systemctl status | stop | start | restart httpd
```

- On SUSE Linux Enterprise Server 12.*x*:

  ```
  sudo systemctl status | stop | start | restart apache2
  ```

- On Red Hat Enterprise Linux 6.*x* (or an equivalent distribution):

  ```
  sudo service httpd status | stop | start | restart
  ```

Operate (Windows)

Using the Microsoft Management Console (MMC) Services snap-in, you can start, stop, and restart SAS HTTP Proxy Server.

Figure A.9 SAS Apache HTTP Server in the Services Snap-In

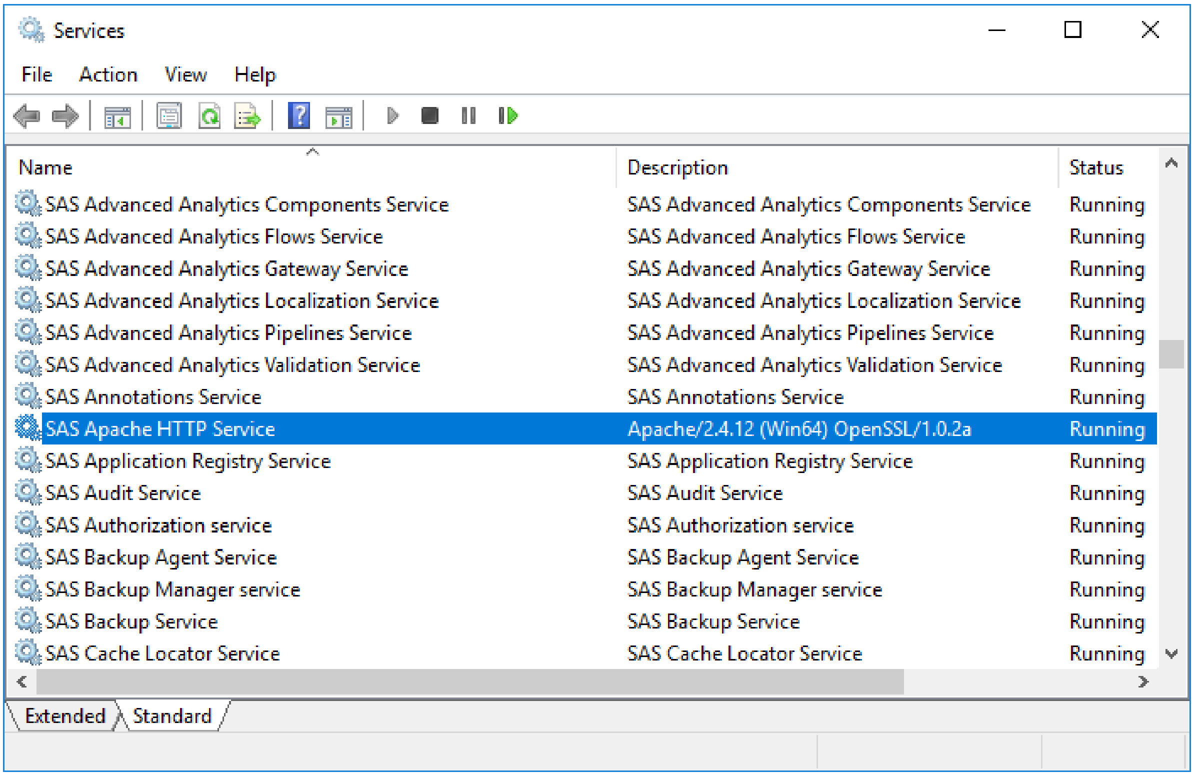

Name	Description	Status
SAS Advanced Analytics Components Service	SAS Advanced Analytics Components Service	Running
SAS Advanced Analytics Flows Service	SAS Advanced Analytics Flows Service	Running
SAS Advanced Analytics Gateway Service	SAS Advanced Analytics Gateway Service	Running
SAS Advanced Analytics Localization Service	SAS Advanced Analytics Localization Service	Running
SAS Advanced Analytics Pipelines Service	SAS Advanced Analytics Pipelines Service	Running
SAS Advanced Analytics Validation Service	SAS Advanced Analytics Validation Service	Running
SAS Annotations Service	SAS Annotations Service	Running
SAS Apache HTTP Service	Apache/2.4.12 (Win64) OpenSSL/1.0.2a	Running
SAS Application Registry Service	SAS Application Registry Service	Running
SAS Audit Service	SAS Audit Service	Running
SAS Authorization service	SAS Authorization service	Running
SAS Backup Agent Service	SAS Backup Agent Service	Running
SAS Backup Manager service	SAS Backup Manager service	Running
SAS Backup Service	SAS Backup Service	Running
SAS Cache Locator Service	SAS Cache Locator Service	Running

Because there is a particular sequence in which the servers and services must be started and stopped, the individual services are not configured to run automatically when the SAS Viya machine is booted.

Important: SAS Configuration Service (Consul), SAS Infrastructure Data Server (PostgreSQL), SAS HTTP Proxy Server (Apache HTTP Server), and SAS Message Broker (RabbitMQ) are dependencies for the other SAS Viya services. If you are operating one or more services individually, always start each of these four services first and stop them last.

Note: There is one service, SAS Services Manager, that you can use to start and stop all SAS Viya servers and services. SAS Services Manager recognizes the dependencies between services and starts and stops services in the correct sequence. For more information, see “Start and Stop All Servers and Services”.

Configure External Reverse Proxy

To configure SAS Viya for an external reverse proxy, including load balancers that act as a reverse proxy:

1 Add properties to ensure that applications that generate links to SAS Viya objects (such as SAS Visual Analytics reports) are aware of the external reverse proxy.

From the machine where the SAS Viya internal Apache HTTP Server resides, run the following commands to add the appropriate property values. Here is an example:

Note: Specify each of these commands on a single line. Multiple lines are used here to improve readability

```
source /opt/sas/viya/config/consul.conf

/opt/sas/viya/home/bin/sas-bootstrap-config kv write
config/viya/sas.httpproxy.external.hostname reverseproxy_hostname

/opt/sas/viya/home/bin/sas-bootstrap-config kv write
config/viya/sas.httpproxy.external.port reverseproxy_port
```

2 On the external reverse proxy, set the X-Forwarded-Proto and X-Forwarded-Port headers to the protocol and port that the client is using to connect to the external reverse proxy.

3 From the SAS Viya internal Apache HTTP server machine, comment out the X-Forwarded-Proto and X-Forwarded-Port lines in the petrichor.conf file as follows:

a Locate the petrichor.conf file according to the platform used:

UNIX	**`/etc/http/conf.d`**
SUSE Linux	**`/etc/apache2/conf.d`**
Windows	**`C:\ProgramData\SAS\Viya\etc\httpd\conf.d`**

b Make a backup copy of the petrichor.conf file.

c Locate the following lines in the petrichor.conf file:

```
# Default the X-Forwarded-Proto header to http
RequestHeader set X-Forwarded-Proto http

# Set the X-Forwarded-Proto header if request is over HTTPS
RewriteCond "%{HTTPS}" "on"
RewriteRule ^.*$ - [ENV=HTTPS:true]
RequestHeader set X-Forwarded-Proto https env=HTTPS
RequestHeader set X-Forwarded-Port 443 env=HTTPS
```

d Comment out these lines so that they appear as follows:

```
# Default the X-Forwarded-Proto header to http
#RequestHeader set X-Forwarded-Proto http

# Set the X-Forwarded-Proto header if request is over HTTPS
#RewriteCond "%{HTTPS}" "on"
#RewriteRule ^.*$ - [ENV=HTTPS:true]
#RequestHeader set X-Forwarded-Proto https env=HTTPS
#RequestHeader set X-Forwarded-Port 443 env=HTTPS
```

4 Start (or restart) the SAS Viya internal Apache HTTP Server.

Note:

- The external reverse proxy must use HTTPS.
- The external reverse proxy must forward requests through the SAS Viya internal Apache HTTP Server without changing the URL path.

- In environments with multiple SAS Viya internal Apache HTTP servers, you should configure the external reverse proxy to route requests to an active HTTPD instance. Round-robin routing or load-balanced routing is recommended.

Concepts

SAS Viya uses Apache HTTP Server as a web server. Apache HTTP Server serves static HTML content and proxies client communication.

A third-party load balancer is required in order to provide high availability for Apache HTTP Server. You can also install your own web server on a separate machine in order to proxy connections from the internet to Apache HTTP Server. For more information about making HTTP Server highly available, see "Apache httpd" in *SAS Viya for Linux: Deployment Guide*.

Log Files

Note: You must be logged on with root-level privileges to the machine where the service resides in order to view log files.

By default, Apache HTTP Server log files are located in the following directory, as appropriate for your operating system:

- On Red Hat Enterprise Linux and equivalent distributions:

  ```
  /var/log/httpd
  ```

- On SUSE Linux Enterprise Server 12.*x*:

  ```
  /var/log/apache2
  ```

31

Tuning

Tuning: Overview

In SAS Viya, you can tune your environment for performance and scalability. This document includes tuning methodologies and parameters for:

- Java Virtual Machine (JVM)

- Java Database Connectivity (JDBC) connection pool
- Lightweight Directory Access Protocol (LDAP) connection pool
- Apache HTTP Server
- SAS Infrastructure Data Server
- Operating system
- SAS servers

Performance requirements are usually identified in terms of transaction response time, number of transactions per second, throughput, resource utilization, total cost per transaction, availability, and more. Scalability often refers to the ability of a component to adapt readily to a greater or lesser intensity of use, volume, or demand, while meeting integral business objectives. The common objective of scaling a component or system is to increase the capacity for growth, increase the speed of the component, improve the efficiency, or shift or reduce the load on the component.

Tuning: Apache HTTP Server

Overview

You can improve the performance of the Apache HTTP Server by configuring other aspects of the web server. For example, to improve performance, rotate log files and configure the Multi-Processing Modules (MPMs).

For more information about MPMs, see http://httpd.apache.org/docs/2.4/mpm.html.

Recommendations

Linux

1 For sites with upward of 400 users, it is recommended that you enable the following Apache HTTP modules:

- Apache 2.2 and later: **worker**

  In **`/etc/sysconfig/httpd`**, uncomment the following line:

  ```
  HTTPD=/usr/sbin/httpd.worker
  ```

- Apache 2.4 and later: **mod_mpm_worker.so**

  In **`/etc/httpd/conf.modules.d/00-mpm.conf`**, *comment* the line ending in mod_mpm_prefork.so, and *uncomment* the line ending in mod_mpm_worker.so:

  ```
  #LoadModule mpm_prefork_module modules/mod_mpm_prefork.so
  LoadModule mpm_worker_module modules/mod_mpm_worker.so
  #LoadModule mpm_event_module modules/mod_mpm_event.so
  ```

2 Configure the Apache HTTP Server to use the worker MPM as follows:

- For Apache 2.2, modify the **`/etc/httpd/conf/httpd.conf`** file to adjust worker MPM settings. Add the ServerLimit setting and change the value for the other settings that are highlighted in the sample file below:

  ```
  # worker MPM
  # StartServers: initial number of server processes to start
  # MaxClients: maximum number of simultaneous client connections
  # MinSpareThreads: minimum number of worker threads which are kept spare
  # MaxSpareThreads: maximum number of worker threads which are kept spare
  ```

```
# ThreadsPerChild: constant number of worker threads in each server process
# MaxRequestsPerChild: maximum number of requests a server process serves
<IfModule worker.c>
   ServerLimit          32
   StartServers         10
   MaxClients           1024
   MinSpareThreads      25
   MaxSpareThreads      75
   ThreadsPerChild      32
   MaxRequestsPerChild  0
</IfModule>
```

- For Apache 2.4, add the following configuration block to the existing configuration file (it is recommended that you modify either the **/etc/httpd/conf.modules.d/00-mpm.conf** file or the **/etc/httpd/conf/httpd.conf** file):

```
<IfModule mpm_worker_module>
   ServerLimit          32
   StartServers         10
   MaxRequestWorkers    1024
   MinSpareThreads      25
   MaxSpareThreads      75
   ThreadsPerChild      32
   MaxConnectionsPerChild  0
</IfModule>
```

3 Identify a suitable log rotation strategy and modify the **/etc/httpd/conf/httpd.conf** file to configure the Apache rotatelogs tool to perform log rotation. For information about rotation strategies and configuration options, see http://httpd.apache.org/docs/2.4/programs/rotatelogs.html.

The following are sample **httpd.conf** file entries for configuring daily log rotation:

```
#ErrorLog logs/error_log
ErrorLog "|/usr/sbin/rotatelogs logs/error_log 86400"

#CustomLog logs/access_log combined
CustomLog "|/usr/sbin/rotatelogs logs/access_log 86400" combined
```

Windows

For sites with upward of 400 users, it is recommended that you apply the following tuning changes to the configuration of the WinNT MPM:

1 Edit the **conf\extra\httpd-mpm.conf** file

2 Add the ThreadLimit directive and increase the ThreadsPerChild within the mpm_winnt_module conditional block as highlighted in the sample file below:

```
# WinNT MPM
# ThreadsPerChild: constant number of worker threads in the server process
# MaxConnectionsPerChild: maximum number of connections a server process serves
<IfModule mpm_winnt_module>
    ThreadLimit             3072
    ThreadsPerChild         3072
    MaxConnectionsPerChild   0
</IfModule>
```

Tuning: Java Runtime Environment

Overview

Note: This information applies to Linux.

The goal of Java Runtime Environment (JRE) tuning is to improve performance in the services, particularly in the area of memory usage and garbage collection cycles. The goal is to also maximize the number of clients that the SAS web applications can support.

Recommendations

The default JRE tuning options that are applied for each service should be sufficient. However, you might need to limit how much the native memory usage grows for each Java process. To limit the growth, add the following lines to the ***Viya-installation-directory*/etc/sysconfig/sas-javaesntl/sas-java-services** file:

```
# Limit the number of "malloc arenas" to 1 (default behavior is to use (# of cores * 8))
export MALLOC_ARENA_MAX=1
```

Tuning: JDBC Connection Pool

Overview

Note: This information applies to Linux.

In Java Database Connectivity (JDBC) connection pooling, instead of creating connections every time they are requested, connections are reused. The JDBC connection pool is a collection of database connection objects that are available for reuse. It is maintained by a connection pooling module as a layer on top of the JDBC driver.

Recommendations

Configure Deployment Size

SAS Viya provides default JDBC connection pool settings for small, medium, and large deployments. By default, all services are configured to use the medium deployment settings.

To configure a different deployment size for a service, complete the following tasks:

1. From SAS Environment Manager, navigate to the **Definitions** view.
2. In the **Definitions** list, select **jvm**.
3. In the top right corner of the window, click ✳.
4. In the New jvm Configuration dialog box, complete the following tasks:
   a. Choose one or more services to which the new settings apply by clicking ↖ and selecting the services.

b. Click **OK**.

c. Click ✚.

d. In the **Name** field, specify `java_option_springdatasource_default`.

e. In the **Value** field, specify -`Dsas.deployment.springdatasource.defaults=size`, where *size* is small, medium, or large.

f. Click **Save**.

5. Click **Save**.

6. Restart all SAS Viya services.

See Also

- “Create Configuration Instances” on page 75
- “Start and Stop All Servers and Services” on page 462

Datasource Properties

By default, there are predefined property settings for small, medium, and large deployments. You can override these values. For example, you can set the Preferences Service to use the default property values for a small system, but override the default value of the `spring.datasource.tomcat.maxIdle` property by changing it from 2 to 3.

Override Default Property Values

To change the default property settings, complete the following tasks:

1. From SAS Environment Manager, navigate to the **Definitions** view.

2. In the **Definitions** list, select **spring**.

3. In the top right corner of the window, click ✳.

4. In the New spring Configuration dialog box, complete the following tasks:

   a. Choose one or more services to which the new settings apply by clicking ↘ and selecting the services.

   Note: Not all services use the default property settings. Instead, those services specify a scaling factor that enables them to have larger pool sizes, based on the deployment size that is specified in the `sas.deployment.springdatasource.defaults` property. For more information, see “Datasource Scaling Factor” on page 652.

   b. Click **OK**.

   c. Click ✚.

   d. In the **Name** field, specify a property from the Property Settings Table on page 652.

   e. In the **Value** field, specify the new size that you want to set the property.

   f. Click **Save**.

5. Click **Save**.

See Also

“Create Configuration Instances” on page 75

Default Datasource Property Settings

The following table provides the default property settings for the JDBC connection pool, based on the deployment size:

Table A.1 Property Settings Table

Property	Small Deployment	Medium Deployment	Large Deployment
datasource.tomcat.initialSize	2	2	2
datasource.tomcat.maxActive	6	10	20
datasource.tomcat.maxIdle	2	2	2
datasource.tomcat.minIdle	2	2	2

A service can also provide default files for small, medium, and large deployments in the service resource directory.

Datasource Scaling Factor

Not all services use the default property settings. Instead, those services specify a scaling factor that enables them to have larger pool sizes, based on the deployment size that is specified in the `sas.deployment.springdatasource.defaults` property. For example, the Authorization Service specifies a scaling factor of 10. Therefore, its maxActive value is 60 for small, 100 for medium, and 200 for large deployments. For a list of the default property values, see Table 31.94 on page 652

Configure the Scaling Factor

To define a scaling factor for a service, complete the following tasks:

1 From SAS Environment Manager, navigate to the **Definitions** view.

2 In the **Definitions** list, select **jvm**.

3 In the top right corner of the window, click ✱.

4 In the New jvm Configuration dialog box, complete the following tasks:

   a Choose one or more services to which the new settings apply by clicking ↖ and selecting the services.

   b Click **OK**.

   c Click +.

   d In the **Name** field, specify the `java_option_datasource_factor` property.

   e In the **Value** field, specify -`Dsas.datasource.custom.factor=multiplier`, where *multiplier* is the multiplier factor by which the property in the Property Settings Table on page 652 will be multiplied.

   f Click **Save**.

5 Click **Save**.

6 Restart all SAS Viya services.

See Also

- "Edit Configuration Instances" on page 75
- "Start and Stop All Servers and Services" on page 462

Scaling Factor Example

You can specify a multiplier factor for a service by using the `sas.datasource.custom.factor` property. The default value for a property is multiplied by the value that you specify. For example, for a medium deployment, the default value for the `spring.datasource.tomcat.maxActive` property is 10. If you set the multiplier factor to 5, the new maxActive value is 50. The factor must be greater than 0. The resulting maxActive value will be no less than `spring.datasource.tomcat.initialsize` and no more than 200.

Tuning: LDAP Connection Pool

Overview

Note: This information applies to Linux.

The Lightweight Directory Access Protocol (LDAP) service provider supports connection pooling. In LDAP connection pooling, the service provider maintains a pool of previously used connections. When a connection is closed or goes to garbage collection, it goes back to the pool to be used again.

By default, no configuration is required for the LDAP service provider to use connection pooling. However, configuration is needed to customize the setting for optimal performance.

Recommendations

Tune LDAP Connection Pool

1. In SAS Environment Manager, edit the **Identities service**.
2. Navigate to the **sas.identities.providers.ldap.connection** configuration instance and configure the following:

Table A.2 Properties and Values

Property	Value
pool.maxActive	30
pool.maxIdle	30

See Also

"Edit Configuration Instances" on page 75

Tuning: Operating System

Overview

There are a number of configuration changes and variables that you can set to tune SAS Viya for your performance and scalability needs. The following sections show how to configure the operating system settings that are relevant to SAS Viya post-deployment.

Linux

Recommendations

Tuning TCP/IP

For sites with upward of 400 users, it is recommended that you perform the following:

- Ensure that IPv6 is enabled.
- Permanently set the SAS recommended TCP/IP settings by using the following commands:

  ```
  /sbin/sysctl -w net.ipv4.tcp_fin_timeout=30
  /sbin/sysctl -w net.core.netdev_max_backlog=3000
  /sbin/sysctl -w net.core.somaxconn=3000
  /sbin/sysctl -w net.ipv4.tcp_keepalive_intvl=15
  /sbin/sysctl -w net.ipv4.tcp_keepalive_probes=5
  ```

Tuning for SAS Studio

The following options can be modified in the **/etc/sysctl.conf** file, when these conditions exist:

- For sites with upward of 40 concurrently logged-on users, who are running tasks that require rendering of graphs, the SEMMNI parameter should be increased to 4096.
- For sites with upward of 600 logged-on users, increase the PID_MAX parameter to 131072.

Windows

Recommendations

Update Windows Registry

The Windows registry must be updated. Microsoft recommends performing a system backup before editing the registry. To set the SAS recommended parameters, use the REGEDIT command as follows:

1. Access the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\Tcpip\Parameters** registry subkey. Add the DWORD value with a name of TcpTimedWaitDelay and a value of 30 (0xle).

2. Access the **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\services\AFD\Parameters** registry subkey. Add the following DWORD values:

Table A.3 AFD Service Parameters

Name	Recommended Value
EnableDynamicBacklog	1 decimal
MinimumDynamicBacklog	20 decimal
MaximumDynamicBacklog	1000 decimal
DynamicBacklogGrowthDelta	10 decimal

The recommended values specify the number of connections that you want to be available. These values request a minimum of 20 and a maximum of 1000 available connections. The number of available connections is increased by 10 each time there are fewer than the minimum number of available connections.

3 In addition, the user port range should be updated. From a command prompt, run the following commands, based on the version of your internet protocol:

```
netsh int ipvn set dynamicport tcp start=32768 num=32767
netsh int ipvn set dynamicport udp start=32768 num=32767
```

The value of *n* indicates the version of your internet protocol and is either 4 or 6.

4 Restart Windows.

Update System Configurations

The following list includes general recommendations for configuring Windows systems:

- Disable Windows indexing on any directories that are used by SAS software.
- Set Windows performance settings so that background processes are favored.
- Set the maximum power profile in the system BIOS for all systems, except Intel Sandy Bridge.
- Disable the C1E BIOS setting on Dell systems.

Tuning: SAS Cache Server

Overview

By default, the SAS Cache Server uses overflow to disk. An Apache Geode Resource Manager runs inside the SAS Cache Server, monitoring the current heap usage against the current maximum heap that is available. If the SAS Cache Server is configured with different values of initial heap size and maximum heap size, the Resource Manager may detect a premature critical heap situation. To avoid this, it is recommended that you set the initial heap size and maximum heap size to the same value.

For more information about Geode Resource Manager, see Managing Heap and Off-heap Memory.

Recommendations

If you increase the initial heap size option (-Xms) for the SAS Cache Server, increase the value so that it is equal to that of the maximum heap size option (-Xmx). These values can be increased if there is enough demand, but ideally they should be set to the same value to avoid the overhead of expansion and contraction. The default is set to -Xms256m and -Xmx256m. Depending on the usage patterns observed and the number of users, you might decide to increase these values. For example, -Xms1024m and -Xmx1024m are reasonable values for a multi-tenant deployment with a large number of users.

Configure Initial Heap Size

1 From SAS Environment Manager, navigate to the **All services** view.

2 In the **All services** list, select **Cache Server service**.

3 At the top of the content pane, enter `jvm` in the **Filter** field, and then click 🔍.

4 Click ✎.

5 In the Edit jvm Configuration window, click **+**.

6 In the Add Property window:

   a In the **Name** field, specify `java_option_xms`.

   b In the **Value** field, specify the new initial heap size as **`-Xms`*`size`***, where *size* is the number of bytes. Append *k* for kilobytes, *m* for megabytes, or *g* for gigabytes. For example, -Xms1024m.

   For more information, see Java Documentation.

7 Click **Save**.

See Also

- “Edit Configuration Instances” on page 75
- “Create Configuration Instances” on page 75

Configure Maximum Heap Size

1 In the Edit jvm Configuration window, update the **java_option_xmx** property with the new maximum heap size value.

   Important: For the SAS Cache Server, the maximum heap size should be the same as the value that you specified for the initial heap size.

2 Click **Save**.

Tuning: SAS Infrastructure Data Server

Overview

Note: This information does not apply to a programming-only deployment.

SAS Infrastructure Data Server provides a transactional store that is used to support SAS Viya. The server is configured automatically during deployment. However, to optimize its performance, it is recommended that you perform the tuning recommendations in this section.

Recommendations

Connection Settings

In SAS Environment Manager, edit the **SAS Infrastructure Data Server** service and modify the following properties to change the number of connections available to client:

Table A.4 Properties and Values

Configuration Instance	Property	Value
sas.dataserver.conf: common	max_connections For more information, see http://www.postgresql.org/docs/9.1/static/runtime-config-connection.html#RUNTIME-CONFIG-CONNECTION-SETTINGS	1027
sas.dataserver.conf: common	max_prepared_transactions For more information, see http://www.postgresql.org/docs/9.4/static/runtime-config-resource.html	1027
sas.dataserver.pool: common	num_init_children For more information, see http://www.pgpool.net/docs/pgpool-II-3.5.4/doc/pgpool-en.html#NUM_INIT_CHILDREN **Note:** This setting does not apply to Windows.	1024

Note: The `max_connections` value should be slightly higher than the `num_init_children` value to allow for direct connections outside pgpool for administrative use, such as backup and recovery.

See Also

"Edit Configuration Instances" on page 75

ulimit Settings

On Red Hat Linux deployments, there are recommended ulimit settings for the `sas` user. The settings reside in the **`/etc/security/limits.conf`** file.

```
sas soft nofile 150000
sas hard nofile 150000
sas soft nproc 100000
sas hard nproc 100000
sas soft stack 10240
sas hard stack 10240
```

Note: You might need to adjust additional Linux operating system settings in order to support these recommended ulimit settings.

Semaphore Settings

On Red Hat Linux and SUSE Linux deployments, there are recommended semaphore settings. The settings reside in the **/etc/sysctl.conf** file.

```
kernel.sem=512 32000 100 1024
 net.core.somaxconn=2048

for SEMMSL, SEMMNS, SEMOPM, and SEMMNI
```

For more information, including formulas and minimum values, see http://www.postgresql.org/docs/9.5/static/kernel-resources.html.

Note: Changing Linux semaphore settings requires a machine reboot.

Note: You might need to adjust additional Linux operating system settings in order to support these recommended semaphore settings. To optimize your PostgreSQL resources, you should also scale the server's working memory settings in accordance with “Special Considerations” on page 658.

Special Considerations

Specialized solutions or use cases might require further configuration tuning. If you need to experiment with the parameters for your optimized system performance, the most important parameters are:

shared_buffers
: specifies the amount of memory to be used for caching data. PostgreSQL also benefits from the file system cache, so shared_buffers should not be so large that they interfere with the file system cache. For a large database, set this parameter between 1 GB and up to 25% of the total system memory.

work_mem
: specifies the amount of memory to be used for sorts, hashing, and materialization, before writing to temporary disk files. Several running sessions can perform operations concurrently. Therefore, the total memory used might be many times the value of **work_mem**. Keep this in mind when choosing the value for this parameter. Set this parameter between 16 MB and 64 MB or more, for a specialized use case (for example, frequent very large sorts).

maintenance_work_mem
: specifies the maximum amount of memory to be used for vacuuming (reclaiming storage used by rows that are marked for deletion) and index builds. For a large database, set this parameter to 256 MB or more.

If your application can tolerate losing a transaction if the computer or storage crashes, you can set the synchronous_commit parameter to **Off** for faster updates.

Tuning: SAS Message Broker

Overview

SAS Message Broker, which is based on RabbitMQ, is an intermediary program that converts messages from the protocol of the sender of the message to the protocol of the receiver. The server is configured automatically during deployment. However, to optimize its performance, it is recommended that you perform the tuning recommendations in this section.

Recommendations

Modify Allocations

Memory Allocation

By default, SAS Message Broker is configured to use up to 40% of the physical RAM on the machine on which an instance runs. This value does not guarantee that more than 40% will be used, but it sets a threshold at which publishers are throttled (notified to slow down message sending). You must decide what percentage of memory to dedicate to the message broker.

For example, if your system has 250 GB and you want to dedicate 50 GB to SAS Message Broker, use the following calculation to begin throttling back at 40% of the dedicated memory:

```
(0.4 * 50 GB) / 250 GB = 0.08 ~ 0.10
```

In the above example, you start throttling back the message broker when it has consumed more than 10% of the available memory. It is difficult to determine the value that the memory threshold, **vm_memory_high_watermark**, should be set to on a system where SAS Message Broker is sharing resources with other services. When the threshold is reached, producers are blocked from sending additional messages until used memory falls below this threshold again. Alternatively, an absolute limit high watermark might be set. However, this value must be less than the amount of available RAM. Otherwise, the message broker will not start.

To set the memory threshold for SAS Message Broker, complete the following tasks:

1 On Linux, set the **vm_memory_high_watermark** parameter by editing one of the following files:

   - If your environment is enabled for Transport Layer Security (TLS), edit the **/opt/sas/*deploymentId*/config/etc/rabbitmq-server/rabbitmq.config.ssl** file.
   - If your environment is not enabled for TLS, edit the **/opt/sas/*deploymentId*/config/etc/rabbitmq-server/rabbitmq.config.tcp** file.

   On Windows, set the **vm_memory_high_watermark** parameter by editing the **C:\ProgramData\SAS\Viya\var\lib\rabbitmq-server\rabbitmq.conf** file.

2 Specify the following in the configuration file:

   ```
   vm_memory_high_watermark, percentRAM
   ```

   For more information, see Configuring the Memory Threshold.

Disk Space Allocation

By default, SAS Message Broker requires at least 50 MB of free disk space to operate. If this threshold is reached, SAS Message Broker slows down message sending and blocks connections. Therefore, it is recommended that you set the minimum free disk size to the amount of memory that is installed on the machine, if it is available. By configuring a large amount of free disk space, a constrained system is more likely to recover under heavy usage scenarios by providing adequate space for paging considerations. By default, paging of transient messages, which are written to the disk under high memory consumption, starts when the system gets halfway to the **vm_memory_high_watermark**.

To set the free disk size, complete the following tasks:

1 On Linux, set the **disk_free_limit** parameter by editing one of the following files:

   - If your environment is enabled for Transport Layer Security (TLS), edit the **/opt/sas/*deploymentId*/config/etc/rabbitmq-server/rabbitmq.config.ssl** file.

- If your environment is not enabled for TLS, edit the **`/opt/sas/`*`deploymentId`*`/config/etc/rabbitmq-server/rabbitmq.config.tcp`** file.

  On Windows, set the **disk_free_limit** parameter by editing the **`C:\ProgramData\SAS\Viya\var\lib\rabbitmq-server\rabbitmq.conf`** file.

2 Specify the following in the configuration file:

```
disk_free_limit, {mem_relative, 1.0}
```

For more information, see Configuring the Disk Free Space Limit.

Note: Using a disk space setting that is relative to the memory size assumes that the available disk space is greater than the amount of available memory.

Tuning: SAS Studio 5

Overview

SAS Studio is a development application for SAS that you access through your web browser. Heap size increases are recommended if you have more than 25 users accessing the SAS Studio 5 server.

Recommendations

It is recommended that you increase the initial heap size option (-Xms) and the maximum heap size option (-Xmx) for SAS Studio 5. The maximum heap size option also needs to be increased for SAS Compute Server.

Configure Initial Heap Size

1 From SAS Environment Manager, in the navigation bar, click 🔧 **Configuration**.

2 Navigate to the **All services** view.

3 In the **All services** list, select **SAS Studio Viya**.

4 At the top of the content pane, click **New Configuration**.

5 In the Select Definition dialog box, select **jvm**.

6 In the New jvm Configuration window, click **+**.

7 In the Add Property window:

   a In the **Name** field, specify **`java_option_xms`**.

   b In the **Value** field, specify the new initial heap size as **`-Xms4096m`**.

   For more information, see Java Documentation.

See Also

- "Edit Configuration Instances" on page 75
- "Create Configuration Instances" on page 75

Configure Maximum Heap Size for SAS Studio 5

1 In the New jvm Configuration window, click **+**.

2 In the Add Property window:

   a In the **Name** field, specify **`java_option_xmx`**.

   b In the **Value** field, specify the new maximum heap size as **`-Xmx4096m`**.

   c Click **Save**.

3 In the New jvm Configuration window, click **Save**.

4 Restart SAS Studio 5.

   On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-sasstudiov-default stop
sudo service sas-viya-sasstudiov-default start
```

   On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-sasstudiov-default
```

   On Windows, in Windows Services Manager, right-click **SAS Studio** and select **Restart**.

Configure Maximum Heap Size for SAS Compute Server

1 In the **All services** list, select **SAS Compute Service**.

2 At the top of the content pane, click **New Configuration**.

3 In the Select Definition dialog box, select **jvm**.

4 In the New jvm Configuration window, click **+**.

5 In the Add Property window:

   a In the **Name** field, specify **`java_option_xmx`**.

   b In the **Value** field, specify the new maximum heap size as **`-Xmx1024m`**.

   c Click **Save**.

6 In the New jvm Configuration window, click **Save**.

7 Restart the SAS Compute Service.

   On Red Hat Enterprise Linux 6.x, run the following commands:

```
sudo service sas-viya-compute-default stop
sudo service sas-viya-compute-default start
```

   On Red Hat Enterprise Linux 7.x and SUSE Enterprise Linux Server 12, run the following command:

```
sudo systemctl restart sas-viya-compute-default
```

   On Windows, in Windows Services Manager, right-click the **SAS Compute Service** and select **Restart**.

32

Using SAS Environment Manager

What Is SAS Environment Manager?

SAS Environment Manager is a web application for managing a SAS Viya environment. It includes a dashboard view, which provides a quick overall look of your environment’s health and status, as well as detailed views that enable you to examine and manage your environment in detail.

You can use the application to manage these areas of your environment (some of these functions might not be available to you, depending on your role and the products that are installed):

Data
CAS tables, caslibs, other data sources

Servers
configuration and information for CAS servers and launcher servers

User content
saved reports and data, favorites, and history

User information
users and groups from your directory service and SAS groups

License information
your SAS licenses and expiration dates.

System backups
backups and restores of system data

Configuration
configuration data for SAS Viya microservices

Contexts
values such as environment variables and port ranges that are used when launching a process

User-defined formats
user-defined data formats and format libraries

Logs
log messages from SAS applications and services.

Machines
information and metric data for the machines and services

Jobs
monitoring of current and past jobs and schedules for jobs

Domains
authentication domains (for storing a user ID and password), encryption domains (for storing an encryption key), and connection domains (for storing a user ID without a password)

Credentials
personal credentials for the authenticated user across authentication and connection domains

Mobile device access
lists that allow or prevent access to the system by specific mobile devices

Rules
access controls and rules that control who can access resources and content in your system

Quality Knowledge Bases
collections of files that store data and the logic that define data quality operations such as parsing, standardization, and matching (available only if SAS Data Quality is installed)

Publishing destinations
destinations for publishing decisions, models, and rule sets from SAS applications (available only if SAS Model Manager, SAS Decision Manager, or Model Studio is installed)

Tenants
information about tenants and status of tenant services (available only in a multi-tenant environment and only to provider administrators)

Note: If you are using only the SAS Viya programming interface, SAS Environment Manager is not deployed. See “Deployment Types” on page 1 for more information.

Accessing SAS Environment Manager

To access SAS Environment Manager, select **Manage Environment** (under **Administration**) in the **Applications** menu (☰).

Note: If you are using only the SAS Viya programming interface, SAS Environment Manager is not available. See "Deployment Types" on page 1 for more information.

When you log on to SAS Environment Manager, if you are a member of the SAS Administrators group, a prompt appears, asking you whether you want to opt in to your assumable groups. If you select **Yes**, your membership in the SAS Administrators group is in effect. See "Assumable Custom Groups" on page 364 for more information about assumable groups. The functions that are available to you in SAS Environment Manager depend on your group membership. See "Predefined Custom Groups" on page 364 for information about the functions that are available for groups.

If you are the tenant administrator in a multi-tenant system, not all functions are available. See "What is Available to a Tenant Administrator?" on page 674 for information.

To sign out of SAS, use the application bar. Click your name, and then click **Sign out**. When you click **Sign out**, you sign out of all SAS web applications.

Using the Dashboard

Overview

The Dashboard provides a quick view of the state of your system. It displays a set of tiles and reports, each of which summarizes an aspect of system status. The Dashboard is the default view when you first open SAS Environment Manager. You can return to the Dashboard from any page in SAS Environment Manager by clicking **Dashboard** from the navigation menu.

Click ⋮ then **Refresh** to refresh the data. Some tiles are also automatically refreshed, as noted in the tile's description on this page.

By default, the **Dashboard** displays these tiles:

- **Availability**
- **System Health**
- **Logged Issues**
- **Mobile Devices**

  If you are an administrator, you can also display system performance reports on the Dashboard by clicking **Show Reports**. This option is not visible if no reports are selected in the **My Dashboard Items** setting or if no user-created reports are pinned to the Dashboard. See "Personalizing Your Dashboard" on page 668 for more information.

Using the Dashboard Tiles

The Dashboard can include these tiles:

Availability

displays grids of color-coded boxes that correspond to the machines, services, and service instances in your environment. The colors reflect the status of each machine, service, and service instance. A cyan box indicates that the item is available, a yellow box indicates that it is partially available or in a warning state, and a red box indicates that it is unavailable.

The grids are updated every ten seconds.

Selecting a box in one of the grids highlights the corresponding boxes in the other two grids. The box that you selected is outlined with a solid line, and the associated boxes are outlined with a dashed line. These are the associations between the selected boxes:

- When you click on a box in the **Machines** grid, the services and service instances that are running on that machine are highlighted in the **Services** and **Service Instance** grids. If a machine is identified as unavailable, the service instances on that machine are identified as unavailable as well.
- When you click on a box in the **Services** grid, the machines where that service is running are highlighted in the **Machines** grid, and the instances of the service are highlighted in the **Service instances** grid.
- When you click on a box in the **Service instances** grid, the machines where the service instance is running are highlighted in the **Machines** grid, and the service is highlighted in the **Services** grid.

Note: To deselect a box, hold down the Ctrl key and click the box. You can also hold down the Ctrl key and press the spacebar.

Place your cursor over a box to view the name of the machine, service, or service instance.

Double-click on a box in the **Machines** grid to open the Machine Status dialog box, which lists the services that are running on that machine and their availability. Click **Machine Details** from the Machine Status dialog box to open the Machines page for the selected machine.

Click on a box in the **Service instances** grid to view the machine address and port where the instance is running.

Note: Place your cursor over or click on an instance of the postgres service to also identify whether the instance is a pgpool instance, whether it is a primary or standby data node, and whether the node has ssl enabled.

Use the **Filter** field to display only certain machines, services, and service instances. When you type characters in the **Filter** field and click 🔍 or press the **Enter** key, the boxes displayed in the **Availability** area change. The boxes that are displayed either match the filter that you specify or are associated with the boxes that are displayed. For example, entering "`laun`" in the **Filter** field might cause two **Services** boxes to display (for the Launcher service and the Launcher server), only the **Service instance** boxes associated with the displayed services, and only the **Machines** boxes associated with the displayed services.

System Health

displays graphs that give you a quick view of the state of the nodes (machines) in your SAS Viya cluster for a selected CAS server. The data displayed on the graphs reflects all of the work taking place on the nodes, not just CAS operations. If you are on a UNIX system, the tile displays the **Node Memory Usage** and the **Load Average** graphs. Use the buttons at the top of the tile to select the graph that you want to view. If you are on a Windows system, only the **Node Memory Usage** graph is displayed.

If your environment contains more than one CAS server, a menu above the graph enables you to select the server to view. When you log in to SAS Environment Manager, this tile attempts to connect to the default CAS server. If the default server cannot be found, the tile displays information for the first server to which it can connect. If it can connect to the default server, but the server does not respond within five seconds, the tile displays a message. You can then retry the server or choose another server.See "SAS Cloud Analytic Services: Overview " on page 470 for more information on CAS servers.

You specify the default server in the casManagement configuration property. See "Configuration Properties: How to Configure Services" on page 74 for more information.

The **Node Memory Usage** graph displays the host memory usage for each node in your cluster. This graph is displayed on both UNIX and Windows systems. Each bar represents a separate node. Bars for controller nodes use a different color than bars for worker nodes. The colors that are displayed depend on the theme that you use. Position your pointer over a bar on the graph to view the name of the node, its type, and its memory usage.

Note: If your environment contains both a primary and a secondary controller node, this graph displays information only for the controller that is currently active.

The **Load Average** graph displays the 1-minute load average over the past five minutes for the nodes in your cluster. This graph is displayed only on UNIX systems. The chart updates every ten seconds. Each node is represented by a separate line on the graph. The vertical scale of the graph changes depending on the largest value being displayed in the chart. Position your pointer on any of the lines on the graph to view the name of the node and the load average at the selected time.

The graphs update every ten seconds.

Logged Issues

displays a time series graph of the number of ERROR and FATAL level log messages captured by SAS Viya log files in the previous 30 minutes. Only the top five sources of ERROR and FATAL messages are included. If there have been no ERROR or FATAL messages in the past 30 minutes, a message is displayed in place of the graph. To view details about the messages or to filter the displayed messages, click ⋮ and then select **Open** to display the **Logs** page. This tile is updated when you refresh the Dashboard.

Custom Groups

displays the name and number of members for the top five custom groups (by number of members). Custom groups are created to control access to SAS Viya features. If you have a sufficient authorization level, you can use the **Users** page to manage custom groups. See "Manage Custom Groups" on page 350. This tile refreshes whenever the Dashboard is reloaded.

If you have the proper authorization, click on ⋮ and the click **Open** to display the **Users** page.

Mobile Devices

displays the type of mobile device access control in use and the number of successful and unsuccessful logon attempts. You can use the **Mobile Devices** page to manage mobile device access and view detailed information about access attempts. See "Mobile: How To" on page 270. This tile refreshes whenever the Dashboard is reloaded. If you do not have a sufficient authorization level, this tile does not appear.

If you have the proper authorization, click on ⋮ and then click **Open** to display the **Mobile Devices** page.

Reports

You can display two types of reports on the Dashboard.

User-selected reports are selected by right-clicking on the report in the Content page and selecting **Pin to dashboard**. References to the selected reports are stored in the folder **`Users/user_name/Application Data/SAS Environment Manager/Dashboard Items`**. Any user can pin reports to their dashboard.

System performance reports are provided by default in SAS Viya and are available only to SAS administrators. These reports are available:

Application Activity

performance and usage by application. See "Monitor Application Activity" on page 208 for more information about this report.

CAS Activity

CPU, memory usage, and system performance for CAS. See "Monitor CAS Activity" on page 209 for more information about this report.

Disk Space

disk usage history and forecast. See "Monitor Disk Space" on page 210 for more information about this report.

Infrastructure Data Server Tables

size and usage for SAS Infrastructure Data Server tables. See "Monitor SAS Infrastructure Data Server Tables" on page 211 for more information about this report.

Message Queue Activity
activity and traffic on the RabbitMQ message exchanges that are used by the operations infrastructure. See “Monitor Message Queue Activity” on page 212 for more information about this report.

System Activity
memory and CPU usage and network activity. See “Monitor System Activity” on page 212 for more information about this report.

User Activity
reports based on audit records. See “View Audit Record Reports and Tables” on page 284 for more information about this report.

By default, the reports are hidden. Click **Show Reports** to display thumbnails for the reports. If you display the reports and then log out of SAS Environment Manager, the reports are displayed when you log back in to SAS Environment Manager.

Select ⋮ in any report tile and then click on **Open** to display the full report in SAS Report Viewer.

Click › and ‹ to change the displayed reports.

The CAS tables that are used to create these reports are refreshed every five minutes. During a deployment, it might take longer than five minutes for data to appear in these reports. This delay is variable and depends on how quickly SAS Infrastructure Data Server, RabbitMQ, CAS, and authentication services are operational and able to respond.

Note: In SAS Environment Manager 3.2, some system reports used data from the CAS_SYSTEM table. In SAS Environment Manager 3.3, these reports now use data collected by the operations infrastructure, in order to provide a consistent view of SAS Viya metric data. The CAS_SYSTEM table will be removed after SAS Viya 3.4, so if you have created any reports that use data from this table, they must be rewritten.

Personalizing Your Dashboard

- To add or remove a tile, at the top of the window, select ***your_user_name*** ⇨ **Settings**. In the Settings dialog box, select **Dashboard**. Select the check boxes for the tiles that you want to display on the Dashboard.
- To remove a tile that is currently displayed on the Dashboard, click ⋮ and then click **Unpin**.
- To choose which system status reports to display on the Dashboard, at the top of the window select ***your_user_name*** ⇨ **Settings**. In the Settings dialog box, select **Public Dashboard Items**. Select the check boxes for the reports that you want to display on the Dashboard. These reports are available only to administrators.
- To add a system status report to those provided by default, place the report in the folder **`/Products/SAS Environment Manager/Dashboard Items`**. An administrator can then pin the report to the reports in their Dashboard.
- To add a report to your dashboard, navigate to the report in the **Content** area of SAS Environment Manager. Right-click on the report in the folder tree and select **Pin to dashboard** from the pop-up menu. The report is added to the report gallery and is copied to the folder /user/Application Data/SAS Environment Manager/Dashboard Items.
- To choose which of your reports to display on the **Dashboard**, at the top of the window select ***your_user_name*** ⇨ **Settings**. In the Settings dialog box, select **My Dashboard Items**. Select the check boxes for the reports that you want to display on the Dashboard.

Note: If you do not select any reports in the **My Dashboard Items** list or you do not pin any reports to the Dashboard, the report container and the options **Show Reports** and **Hide Reports** do not appear on the Dashboard.

SAS Environment Manager Functions

Accessing the Functions

To access a SAS Environment Manager page, select it from the navigation bar.

Depending on your organization’s environment and authorization policies, you might not have access to all pages.If you are not a member of the SAS Administrators group, or you do not opt in to the group, you can access only the **Dashboard**, **Data**, **Servers**, **Content**, **Jobs**, and **My Credentials** pages. **Quality Knowledge Bases** is also available if SAS Data Quality is installed. Your organization might also use other groups that could restrict your ability to access certain pages. See “Identity Management: Access to Functionality” on page 355 for more information.

Data

Select **Data** from the navigation menu to view and manage your data and data sources. You can select three views of your data.

The **Available** view displays information about all of the data tables that have been loaded into CAS memory. The view displays basic information for the selected table, such as the number of columns and rows, size, location, and dates of creation and last modification. The **Details** tab for a selected table displays detailed information about each variable in the table such as the type, raw length, formatted length, format used, and tags assigned. The **Sample Data** tab displays a selected number of rows of data (the default is 100). The **Profile** tab displays any selected profiles for the table and enables you to run a profile job for the table. From the Available view, you can also select the table for importing and view the table’s authorization.

The **Data Sources** view enables you to create a connection between a caslib and either a database server or a remote file system. You can navigate through the data source to locate a data table. When you select a data table, the **Details**, **Sample Data**, and **Profile** tabs are all available, as in the **Available** view.

The **Import** view enables you to create a connection between a caslib and either a local file, social media content, or Esri data. You can also create a connection between a caslib and a table or file on the **Data Sources** tab.

See “Understanding SAS Data Explorer” in *SAS Data Explorer: User’s Guide* for more information.

Servers

Select **Servers** to display information about the SAS Viya servers. The window displays basic information about the servers, such as the state, host, and port. While in the **Servers** view, you can also view detailed properties and system metrics for the server, including the settings and configuration values for the server and the users and groups that are superusers for the server. If you have the proper authorization, you can also assume the superuser access for the server, which enables you to edit the additional settings and configuration values.

See “SAS Cloud Analytic Services: How To (SAS Environment Manager)” on page 481 for more information.

Content

Select **Content** from the navigation menu to display folders that contain items that users have saved. When you open the Content page, you have access to your own data in the My Folder folder. If you have administrative access, you can also view the folders of other users. From this page, you can create, delete, move, and rename folders, create shortcuts, and manage the authorization for any folder or item that you select (if you have sufficient authorization). You can also export the reports in a folder to a package file and import the reports from a package file into a folder.

Each user’s folder contains several subfolders:

My Favorites
contains references to items identified as favorites, to enable quick access to often-used reports and data.

My Folder
contains saved items.

Application Data
contains items used by SAS Viya applications, such as items pinned to a user’s Dashboard in SAS Environment Manager.

My History
contains a list of the most recent items that you have accessed. You can select entries in this folder to quickly return to items that you have worked with recently.

See “Content Management: How To” on page 132 for more information.

Users

Select **Users** from the navigation menu to view information about users and groups, and to manage custom groups. The information displayed for users and groups comes from your organization’s directory service (such as LDAP or Microsoft Active Directory). Because this information is managed by your identity provider, it is displayed as read-only data in SAS Environment Manager.

You can also manage custom groups on the **Users** page. Custom groups enable you to manage special permissions for groups of users.

See “View User and Group Information” on page 350 for more information.

Licensed Products

Select **Licensed Products** from the navigation menu to view information about the licenses for your products. You can view a list of all your currently licensed products and see the expiration date, grace period, warning period, and maximum CPU count for each one. You can filter the list by any of the displayed criteria to make it easier to find products in the table.

See “Licensing: Overview ” on page 181 for more information.

Note: This page is not available for a tenant administrator.

Backup and Restore

Select **Backup and Restore** from the navigation menu to back up and restore your environment.

Configuration

Select **Configuration** from the navigation menu to manage the configuration settings for SAS Viya services. You can select from a list of basic services, all services, or definitions. When you select a service, the service’s configuration properties are displayed on the right side of the window. Click ✎ to change any of the displayed properties.

See “Introduction” on page 74 for more information.

Contexts

Select **Contexts** from the navigation menu to manage launcher contexts. A context is a collection of values such as environment variables and port ranges that are used when launching a SAS Viya instance. You can also specify values for a deployment, such as the deployment ID and the installation and configuration directories.

User-Defined Formats

Select $w.d **User-Defined Formats** to display information about all of the user-defined formats and format libraries that are available for the data. The window displays a list of available user-defined formats and format libraries. You can add and import new formats, as well as edit, copy, and delete existing formats. You can import formats from a SAS item store, although some steps are required outside of SAS Environment Manager. You can also create, delete, and change the search order of format libraries. This function is available only for administrators.

Logs

Select **Logs** from the navigation menu to view information about messages that have been written to the logs. You can view a chart of the number of log messages and a table of the detailed messages. By default, the chart and table reflect the messages logged during the previous 30 minutes, but you can select a different time range. You can also search for a specific message or filter the messages by level and source.

See “Logging: Overview” on page 191 for more information.

Note: This page is not available for a tenant administrator.

Machines

Select **Machines** from the navigation menu to monitor the machines in your environment and the service instances running on those machines. You can view this information:

- charts of the percent of CPU utilization and memory used for each machine
- status of predefined checks (such as disk or memory usage) for the selected machine
- service instances running on the selected machine, along with their current state
- properties for the selected server

See “Monitoring: How to (SAS Environment Manager)” on page 206 for more information.

Note: This page is not available for a tenant administrator.

Jobs

Select ⚙ **Jobs** from the navigation menu to monitor and schedule jobs. You can perform these tasks:

- View a table or a chart of jobs that are currently running and that have run in a specified time in the past.
- Filter the jobs to narrow the number of jobs displayed and change the time period for displaying jobs.
- Re-run jobs and delete jobs.
- Schedule jobs to run at a particular time or in response to a specific trigger.
- Run a scheduled job immediately.
- Create time-based triggers to control when scheduled jobs run.
- Unschedule, delete, and view the properties of scheduled jobs.

- Manage CAS tables.

See “Jobs: Overview” on page 169 for more information.

Domains

Select **Domains** from the navigation menu to manage domains used for authentication, encryption, and connection.

See “About the Domains Page” on page 338 for more information.

My Credentials

Select **My Credentials** from the navigation menu to create and manage credentials for authentication, encryption, and connection domains. See “Manage My Credentials” on page 346 for more information.

Mobile Devices

Select **Mobile Devices** from the navigation menu to manage how mobile devices access certain reports. You can use either a blacklist or a whitelist. If you use a whitelist, all mobile devices are blocked except for those listed in the whitelist. If you use a blacklist, all mobile devices are allowed except for those listed on the blacklist. The **Mobile Devices** page displays a table of recent access attempts by mobile devices, the devices listed in the blacklist, and the devices listed in the whitelist. The page indicates whether the blacklist or the whitelist is being enforced, and it enables you to select which list to use. You can also add devices to either list.

See “Mobile: How To” on page 270 for more information.

Rules

Select **Rules** from the navigation menu to manage access to specific locations and content.

See “General Authorization: How to (Rules Page)” on page 416 for more information.

Quality Knowledge Bases

Select **Quality Knowledge Bases** from the navigation menu to view and manage Quality Knowledge Bases. Quality Knowledge Bases are collections of files that store data and logic that define data quality operations such as parsing, standardization, and matching. This area is available only if you have licensed SAS Data Quality. See SAS Viya Administration: QKB Management for more information.

Publishing Destinations

Select **Publishing Destinations** from the navigation to manage destinations for publishing decisions, models, and rule sets using SAS Decision Manager, SAS Model Manager, and Model Studio. You can create, edit, delete, and view properties for destinations. This area is available only if SAS Model Manager, SAS Decision Manager, or Model Studio is installed and if you opt in to the SAS Administrators group when you sign in to SAS Environment Manager. See “Publishing Destinations: How To” in *SAS Viya Administration: Publishing Destinations* for more information.

Tenants

If you are the provider administrator of a multi-tenant environment, select **Tenants** from the navigation menu to view information about tenants and tenant services.

How To

Work with Information Displayed in Tables

When you are viewing information that is displayed in a table in SAS Environment Manager, use these tips to control how data is displayed:

- To sort a table, right-click on a column header and select **Sort**. You can sort the table by the contents of the column or add the column as a secondary sort criteria.
- To reorder the columns in a table, click on the column heading and drag the header to the new location.
- To prevent a column from being reordered, right-click on the column heading and select **Freeze**. The column is moved to the left of the table and cannot be reordered. To enable the column to be moved, right-click on the header and select **Unfreeze**.
- To select which columns are displayed, click on the **Options** icon ☰ on the right side of the table header and select **Manage columns** from the pop-up menu. The **Columns** window displays a list of hidden columns and displayed columns. Select the columns that you want to display and click **OK**.
- To reduce the number of items displayed in the table, use the **Filter** field. As you enter text in the field, the table changes to display only the items that contain text that matches the text that you enter. The table is filtered dynamically as you enter text. The text that you enter as filter text is not case sensitive.

Note: Not all of the tables in SAS Environment Manager use all of these features.

Manage Settings

Access Settings

To access the Settings window, select your user name in the upper right of the SAS Environment Manager window. Select **Settings** from the pop-up menu.

General

The **General** section includes settings that enable users to change the appearance of the web applications, enable warning and information messages to be displayed, and choose a profile picture. Here are the values:

- You can change the appearance of the web applications by using the **Theme** setting. The default theme is specified by the system administrator. The theme specifies the collection of colors, graphics, and fonts that appear in the application. You can choose from SAS themes or custom themes, if available.

  Select **Choose a theme**, and then select another theme from the drop-down list to change the look of the applications. The theme change takes affect after you close the Settings window.

  SAS themes:

  Illuminate
  This theme has a clean and uncomplicated color palette that is easy to use.

  Inspire
  This theme consists of vibrant and cohesive colors that shift the emphasis from the application to the content.

  High Contrast
  This theme presents a dark background with high-contrast foreground elements to meet the needs of users with low vision.

- If you want messages to display that you previously asked not to display, click **Reset Messages**. By default, all warning and information messages are displayed.
- You can select a profile picture to display as an avatar in the application bar, as well as other places within the application that use avatars. An avatar is the graphical representation of the user or the user's alter ego or character.

  Click **Choose Picture** and then select an image file to upload. The image file’s size can be up to 1 MB. The valid file types are BMP, GIF, JPEG, JPG, and PNG.

Region and Language

The **Region and Language** section includes values that enable users to specify the locale for regional formats and sorting, as well as for offline processes.

- The **Locale for regional formats and sorting** setting specifies the locale that is used for sorting data and formatting values such as dates, times, numbers, and currency. By default, the browser locale is used. Changes take affect after you sign out and sign back in.
- The **Locale for offline processes** setting specifies the locale that is used for offline jobs or background processes such as report distributions or notifications. By default, the locale of the Java Runtime Environment is used

Accessibility Settings

Several settings in the **Accessibility** section can assist people who rely on assistive technologies. Accessibility features are part of the global settings, which are applied to all SAS web applications. Accessibility features are not specific to SAS Environment Manager.

The following accessibility features are available:

- Select **Enable sounds** to hear audio indicators for events that occur within the user interface.
- Select **Enable visual effects** to show visual effects that indicate state changes. For example, when this setting is selected, you will see a subtle movement in the user interface if you delete an item.
- Select **Invert application colors** to make the user interface easier to see for users with sensitivity to certain bright colors (for example, a black-on-white display). You can also use the Ctrl+` (Ctrl+back quote) keyboard shortcut to invert the application colors.
- Select **Display tooltips when using the keyboard to navigate** to enable keyboard users to access tooltips. When this option is selected, putting keyboard focus on a control also displays the tooltip on the screen. By default, this option is not selected, so a mouse is required for you to see a tooltip. You can also select the location in the browser window to display the tooltip. By default, the tooltip is displayed in the bottom-right corner of the browser window.
- The focus indicator is an outline that indicates which user interface component is active. You can make the focus indicator easier to see by selecting **Customize the focus indicator settings** and adjusting the color, thickness, and opacity.

What is Available to a Tenant Administrator?

If you are the tenant administrator in a multi-tenant system, these functions are not available:

- Licensed Products
- Tenants
- Logs

- Machines

By default, the Dashboard for a tenant administrator contains these items:

- **System Health** tile
- **Mobile Devices** tile
- **Top 5 Custom Groups** tile
- Personal reports that are pinned to the Dashboard

Other Dashboard items that are listed in this document are not available to tenant administrators.

33

Command Line Interfaces

Command-Line Interface: Overview

Introduction

SAS Viya contains administrative command-line interfaces (CLIs). In SAS Viya, a CLI is a user interface to the SAS Viya REST services where you enter commands on a command line and receive a response from the system. You can use a CLI to interact directly with SAS Viya programmatically without a GUI.

How to Use This Document

This document describes administrative tasks that enable you to run the CLIs, gives information about how to use the integrated help to use the CLIs, and includes examples of how to use each CLI. Here is a reading strategy for performing these tasks:

- Set up your environment to run the CLIs.
- Familiarize yourself with the CLI syntax and the integrated help within the CLIs.
- See "Inventory" on page 678 for information about each CLI plug-in, and links to examples.

Inventory

The following administrative CLIs are available in SAS Viya:

Name	Scope and Examples
admin	Hosts other CLIs that run as plug-ins to this one. The top-level administrative command-line interface that is used to initialize, authenticate, and execute other plug-ins.
audit	Gets SAS audit information. See "CLI Examples: Audit" on page 698.
authorization	Gets general authorization information and manages rules. See "CLI Examples: General Authorization" on page 702.
backup	Manages backups. See "CLI Examples: Backup" on page 706.
cas	Manages CAS administration and authorization. See "CLI Examples: CAS Administration", "CLI Examples: CAS Authorization" on page 699, and "CLI Examples: Formats" on page 716. For information about cas environment variables, see CAS Administration: Details on page 710.

Name	Scope and Examples
compute	Manages the operations of the compute service. See “CLI Examples: Compute” on page 711.
configuration	Manages the operations of the configuration service. See “CLI Examples: Configuration” on page 711.
dagentsrv	Manages interactions with the SAS Data Agent server. See “Command-Line Interface” in *Cloud Data Exchange for SAS Viya: Administrator’s Guide*.
dataexplorer	Automates data management capabilities from the Data Explorer UI. See “CLI Examples: Data Explorer” on page 712.
devices	Manages mobile device blacklist and whitelist actions and information. See “CLI Examples: Device Management” on page 713.
folders	Gets and manages SAS folders. See “CLI Examples: Folders” on page 714.
fonts	Manages fonts that are provided by SAS as well as custom fonts that are registered in SAS Visual Analytics 8.3. See “CLI Examples: Fonts” on page 715.
identities	Gets identity information, and manages custom groups. See “CLI Examples: Identities” on page 720.
job	Manages the operations of the job flow scheduling service. See “CLI Examples: Job” on page 721.
launcher	Manages the operations of the launcher service. See CLI Examples: Launcher on page 721.
licenses	Manages SAS product license status and information. See “CLI Examples: Licensing” on page 722.
qkbs	Manages the SAS Quality Knowledge Bases (QKBs) on a Cloud Analytics Services (CAS) server. See “CLI Examples:Quality Knowledge Bases (QKBs)” on page 722.
reports	Manages SAS Visual Analytics 8.2 reports. See “CLI Examples: Reports” on page 724.
restore	Manages restore operations. See “CLI Examples: Restore” on page 727.
scoreexecution	Manages resources that are no longer used by the Score Execution service. See CLI Examples: Score Execution on page 727.

Name	Scope and Examples
tenant	Manages tenants in a multi-tenant deployment. See "CLI Examples: Tenant Administration" on page 727.
transfer	Promotes SAS content. See "CLI Examples: Transfer" on page 728.

Key Points

Here are key points for using the CLIs:

- SAS recommends that users with administrative privileges run the admin CLI in order to ensure optimal results. Users without administrative privileges are able to run the admin CLI, but the results might not always be as expected.
- To prepare to use the admin CLI plug-ins, see "Command-Line Interface: Preliminary Instructions" on page 681.
- Commands and subcommands are case-sensitive.
- You must precede the options of the commands with two hyphens (`--`) if the option is a word. If the option is a single letter, you can precede the option with two hyphens (`--`) or one hyphen (`-`). For example, you can use **`--help`** or **`-h`** for the Help global option.
- Linux special characters that are specified on the command line must be preceded (or escaped) with a backslash (\) so that they are not interpreted by the Linux shell. Linux special characters that are included in JSON template files do not need to be escaped.
- If a single parameter contains spaces (such as in long file names), you must enclose it in quotation marks in order to pass it as one item. Depending on what can be parsed by your operating system, other parameter values might need to be enclosed in quotation marks or escaped accordingly.
- Use the Help from within each CLI for information about the available commands, subcommands, and options. For the admin CLI plug-ins, see "Command-Line Interface: Syntax" on page 686.
- Timestamps that are returned from the sas-admin CLIs are in Universal Time Coordinated (UTC) format rather than the local time. By contrast, the timestamps returned from SAS Environment Manager are in local time.

  For example, in SAS Environment Manager, you might notice that a SAS Visual Analytics report has a modified date of April 26, 2018 09:23:47. This is the local time. However, the reports CLI will show that the modified date for the same report is 2018-04-26T13:23:47.314Z. This is UTC time.

Note: This document reflects sas-admin CLI functionality in the initial release of SAS Viya 3.4. The integrated Help supersedes this documentation and provides the most current information about any expanded or enhanced functionality.

About the Examples

The following points apply to all of the examples in this document:

- The examples assume that you have signed in to SAS Viya using the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.
- The examples explicitly specify all necessary options. In practice, you might find it more efficient and concise to use environment variables where available. Remember to clear values for any environment variables when appropriate.

- The examples generally were run in a Linux environment.
- The examples include line breaks within commands to enhance readability. Do not include line breaks when you submit a command.
- The examples generally include single quotation marks when quotation marks are required. Use the quotation marks that are appropriate for your platform.
- The examples generally assume that you are using the default profile rather than a named profile.

  Note: If you logged in using a named profile, you must do one of the following to use that profile:

  - set the **`SAS_CLI_PROFILE`** environment variable to the name of the profile. This will remain in affect until you log off from the environment.

    For example, suppose that you have set the **`SAS_CLI_PROFILE`** environment variable to **`Target1`**.

    Then, you can log on to the environment with the “Target1” profile with this command: `sas-admin auth login`

    If you are signing in to an environment that has been configured for Kerberos, use this command: `sas-admin auth kerberos`.

  - include the **`profile`** global option on each CLI command as follows:

    ```
    sas—admin --profile profile-name CLI-name CLI-commands
    ```

    Then, you can sign in to the environment with the “Target1”profile with this command: `sas-admin --profile Target1 auth login`

    If you are signing in to an environment that has been configured for Kerberos, use this command: `sas-admin --profile Target1 auth kerberos`

Command-Line Interface: Preliminary Instructions

Complete the following required preliminary tasks before you use a CLI.

Set the SSL_CERT_FILE Environment Variable

If your environment is enabled for Transport Layer Security (TLS), you must set the **`SSL_CERT_FILE`** environment variable to the path location of the trustedcerts.pem file (if using the SAS default truststore) or the path location of your site-signed certificate (if using an internal truststore).

Note: For CLI users on Linux who are running the CLIs directly on the SAS machine, you might be able to source the consul.conf file rather than set the SSL_CERT_FILE environment variable.

Set the SSL_CERT_FILE environment variable to the following value:

Truststore	Operating System	Path
SAS default truststore	Linux	**`/opt/sas/viya/config/etc/SASSecurityCertificateFramework/cacerts/trustedcerts.pem`**
	Windows	**`\Program Files\SAS\Viya\SASSecurityCertificateFramework\cacerts\trustedcerts.pem`**
internal truststore	Linux or Windows	*path_to_certificate*

Here is an example of how to set the environment variable in a Linux environment: `export SSL_CERT_FILE=/opt/sas/viya/config/etc/SASSecurityCertificateFramework/cacerts/trustedcerts.pem`

Download the CLI Executable

Overview

You can download the admin CLI directly from the SAS Support website, and install the plug-ins that you need. This is the recommended approach for running the admin CLI from client machines. The most current plug-ins are available with the downloaded version of the admin CLI.

Important: If you are installing (or upgrading to) the Windows version of SAS Viya 3.4 for the first time, you must download and install the latest version of the admin CLI to client machines.

All the plug-ins are available for installation with the downloaded admin CLI with the exception of the following plug-ins:

backup	Will be available as a plug-in for installation with the admin CLI in a future release.
restore	Will be available as a plug-in for installation with the admin CLI in a future release.
compute	Might be available as a plug-in for installation with the admin CLI in a future release.
tenant	Is not available as a plug-in for installation with the admin CLI.

Note: Alternatively, the admin CLI, along with the plug-ins, is installed on the SAS Viya server during deployment. If you are not downloading the admin CLI to run on a client machine or you want to run a plug-in that is not available with the downloaded admin CLI, you can run the admin CLI directly from this location on the SAS Viya server:

- Linux: **`/opt/sas/viya/home/bin`**
- Windows: **`\Program Files\SAS\Viya\bin`**

For SAS Viya 3.4, the dataexplorer CLI plug-in is only available for download. It is not installed in the above directory location on the SAS Viya server.

How To

Here are the steps to download the admin CLI and install the plug-ins:

Note: You must have a valid SAS profile to download the admin CLI.

1 Go to the Support / Downloads and Hot Fixes page at http://support.sas.com/downloads/package.htm?pid=2133, and download the appropriate file to your machine.

2 Prepare the files for your environment:

UNIX	Expand the file to a directory from which you plan to run the CLI:
	`tar -xvzf /`*`CLI-directory`*`/sas-admin-cli-1.1.2-download-linux.tgz`
	From the directory, make sure that execution permission is set:
	`chmod +x sas-admin`

Windows or Macintosh	Unzip the file and place it in a directory from which you plan to run the CLI.

CAUTION! Do not place the admin CLI download file in the following directory:

Table A.1 CLI Default Directory Prohibited from Download

Linux	**`/opt/sas/viya/home/bin`**
Windows	**`\Program Files\SAS\Viya\bin`**

3 Navigate to the directory location where you saved the CLI file. You must run the CLI commands directly from this directory, or you can add this directory to your system path to make the sas-admin CLI available to all CLI users.

4 Create a profile and sign in.

5 Install the CLI plug-ins by running these commands:

```
1 sas-admin plugins

2 sas-admin plugins list

3 sas-admin plugins list-repo-plugins

4 sas-admin plugins install --repo SAS plugin-name
```

1 Display the commands that are available to the admin CLI plug-ins command.

2 Display the admin CLI plug-ins that are currently installed.

3 List the plug-ins in the SAS repository that are available for installing.

4 Install a plug-in from the SAS repository.

Note: The plug-ins are installed in the home directory of the user who ran the **`sas-admin plugins`** command: ***home-directory*/.sas/admin-plugins**.

Run the Admin CLI Plug-In from a Shared Location

Important: SAS recommends that each admin CLI user should download a unique instance of the admin CLI and install a unique instance of a plug-in from the SAS repository. If your organization has a policy that enforces the download of only a single instance, it is still possible to make the CLI plug-ins available from a shared location as follows:

1 Navigate to the config.json file according to the operating system being used:

Table A.2 CLI Directory for config.json File

Linux:	***home-directory*/.sas/admin-plugins**
Windows:	***home-directory*\.sas\admin-plugins**
Macintosh:	***home-directory*/.sas/admin-plugins**

2 Open the config.json file, and for each CLI plug-in, find the line that contains the location, and modify the directory path to point to a shared location that contains the CLI admin plug-in directory. Here is an example of this section in the config.json file for the transfer CLI plug-in:

```
"transfer": {
      "location": "\shared-location\admin-plugins\\sas-transfer-cli-1.3.5-20180330.1522434840.exe",
      "description": "tool for promoting contents across environments",
      "version": "1.3.5"
```

3 Propagate this change to the config.json file for all users:***home-directory*/.sas/admin-plugins**

Create at Least One Profile

If you have not already created a profile for the environment that you want to use, complete the following steps:

1 If you downloaded the CLI executables separately, navigate to the directory on the machine that contains the CLIs. If you run CLIs directly from the SAS Viya server and you have not downloaded the CLI executables separately, from a command prompt on the SAS Viya server, navigate to the appropriate directory:

Table A.3 CLI Default Directory Prohibited from Download

Linux	**/opt/sas/viya/home/bin**
Windows	**\Program Files\SAS\Viya\bin**

2 At the command prompt, enter a command to initialize a new profile. Here are examples:

To create a default (unnamed) profile, enter:	`sas-admin profile init`
To create a profile called **prod**, enter:	`sas-admin --profile prod profile init`

You can use a named profile to access different environments using the same set of CLIs. See "Default Profile and Named Profile" on page 690 for information about why you might want to use a named profile.

Note: Running the **profile** global command creates a config.json file in this directory: ***home-directory*/.sas**. For more information, see "Overview" on page 689.

3 Respond to the subsequent prompts as follows:

`Service Endpoint`	Specify the URL for the SAS Viya environment. Use the following format: *communications-protocol*://*web-server-host-name*:*web-server-port* For example: `https://host.example.com`

`Output type`	Specify your preferred format for CLI output (`text`, `json`, or `fulljson`). **Note:** For more information about the output types, see "Output Type" on page 689.
`Enable ANSI colored output`	Specify whether to enable colored output (`y` or `n`).

4 Repeat steps 2 and 3 for any additional profiles that you want to create.

Use a Profile to Sign In

1 If you downloaded the CLI executables separately, navigate to the directory on the machine that contains the CLIs. If you run CLIs directly from the SAS Viya server and you have not downloaded the CLI executables separately, from a command prompt on the SAS Viya server, navigate to the appropriate directory:

Table A.4 CLI Default Directory Prohibited from Download

Linux:	**`/opt/sas/viya/home/bin`**
Windows:	**`\Program Files\SAS\Viya\bin`**

2 At the command prompt, enter a command to initiate the sign-in process. Here are examples:

To use your default (unnamed) profile (assuming that the **`SAS_CLI_PROFILE`** environment variable is not set), enter:	`sas-admin auth login`
To use a profile called **`prod`**, enter:	`sas-admin --profile prod auth login`

Note: Kerberos is supported for Linux and Windows environments. If you are signing in to an environment that has been configured for Kerberos, replace `auth login` with `auth kerberos` in the preceding examples.

3 At the subsequent prompts, enter your user ID and password. No prompts appear in environments that have been configured for Kerberos.

By default, your authentication remains active for 12 hours. You can use the **`auth logout`** command to sign out.

Note: When you run the **`auth login`** global command, a bearer token is written to the credentials.json file in this directory: ***`home-directory`*****`/.sas`**. For more information, see "Overview" on page 689.

Note: If you logged in using a named profile, you must do one of the following to use that profile:

- set the **`SAS_CLI_PROFILE`** environment variable to the name of the profile. This will remain in affect until you log off from the environment.

  For example, suppose that you have set the **`SAS_CLI_PROFILE`** environment variable to **`Target1`**.

  Then, you can log on to the environment with the "Target1" profile with this command: `sas-admin auth login`

If you are signing in to an environment that has been configured for Kerberos, use this command: `sas-admin auth kerberos`.

- include the **`profile`** global option on each CLI command as follows:

  ```
  sas—admin --profile profile-name CLI-name CLI-commands
  ```

  Then, you can sign in to the environment with the “Target1”profile with this command: `sas-admin --profile Target1 auth login`

  If you are signing in to an environment that has been configured for Kerberos, use this command: `sas-admin --profile Target1 auth kerberos`

See Also

“Command-Line Interface: Overview” on page 678

Command-Line Interface: Syntax

Structure

The basic structure of a command-line interface (CLI) command is:

sas-admin *interface-name [global options] command [command options] [subcommand] [subcommand options] [arguments]*

sas-admin
specifies the sas-admin CLI.

interface name
specifies the CLI plug-in.

[global options]
specifies options that are applicable to all CLIs.

command
specifies a command that is specific for the CLI that you are using.

[command options]
specifies options for the CLI-specific command that you are using.

[subcommand]
specifies a subcommand for the CLI command that you are using.

[subcommand options]
specifies options for the subcommand.

[arguments]
specifies arguments for options.

Here are some basic examples of issuing CLI commands:

Example: Change the output type that is used with the audit CLI to JSON.

```
sas-admin audit --output json
```

Example: Show more detailed information for the mobile device CLI **`blacklist list`** command.

```
sas-admin --verbose devices blacklist list
```

For information about the CLIs that are available, see “Inventory” on page 678.

Integrated Help

Global Commands

Use the integrated help within the CLI to learn about the available global commands, plug-ins, and global options. The global options apply to each CLI plug-in.

Example: List all the global commands, plug-ins, and global options for the admin CLI.

```
sas-admin help
```

Here is the output from this command in a Linux environment:

```
NAME:
   sas-admin - SAS Administrative Command Line Interface

USAGE:
   sas-admin [global options] command [command options] [arguments...]

VERSION:
   1.1.13

COMMANDS:
   authenticate, auth, authn    Handles authentication to the target environment.
   help, h            Shows a list of commands or help for one command.
   plugins            Manages plugins.
   profile, prof        Shows and updates options.

PLUGINS:
   audit
   authorization
   backup
   cas
   compute
   configuration
   devices
   folders
   fonts
   identities
   job
   launcher
   licenses
   qkbs
   reports
   restore
   scoreexecution
   tenant
   transfer

GLOBAL OPTIONS:
   --colors-enabled        Enables or disables ANSI colored output. [$SAS_CLI_COLOR]
   --help, -h            Shows help.
   --insecure, -k        Allows connections to TLS sites without validating the server certificates.
   --locale "en"        Specifies a locale to use. [$LC_ALL, $LANG]
   --log-file             Specifies the file to write log events to. [$SAS_LOG_FILE]
   --output             Specifies output format - text, json, fulljson. [$SAS_OUTPUT]
```

```
      --profile, -p "Default"    Specifies a named profile to use. [$SAS_CLI_PROFILE]
      --quiet, -q            Quiets spurious output, data only.
      --sas-endpoint         Sets the URL to the SAS services. [$SAS_SERVICES_ENDPOINT]
      --verbose            Shows detailed processing information and output.
      --version, -v        Prints the version.

   COPYRIGHT:
      (c) 2016-2018 SAS Institute Inc. All Rights Reserved.
```

Example: Show the version of the CLI.

```
sas-admin --version
```

Here is the output from this command in a Linux environment:

```
sas-admin version 1.1.13
```

CLI Plug-in

Use the integrated help within the CLI to learn about the available commands, subcommands, and options for each CLI plug-in. Use the same syntax for each CLI, substituting the CLI plug-in name or command that you are getting help for.

Example: List all the commands for the devices plug-in to the admin CLI.

```
sas-admin devices --help
```

Here is the output from this command in a Linux environment:

```
NAME:
   sas-devices

USAGE:
   sas-admin devices command [command options] [arguments...]

COMMANDS:
   authorized-devices    Manages the authorization of devices.
   blacklist        Manages the list of blacklisted devices.
   enforcement        Manages the policy enforcement of mobile devices.
   help, h        Shows a list of commands or help for one command.
   last-access        Manages the set of history records of the devices that use the mobile application.
   whitelist        Manages the list of whitelisted devices.
```

Example: List the subcommands of the **blacklist** command that is a part of the devices plug-in to the admin CLI.

```
sas-admin devices help blacklist
```

Here is the output from this command in a Linux environment:

```
NAME:
   sas-admin devices blacklist - Manages the list of blacklisted devices.

USAGE:
   sas-admin devices blacklist [arguments...]

COMMANDS:
   add          Adds a device to the blacklist.
   clear        Clears the blacklist of all device IDs that are present.
   delete       Deletes a device from the blacklist.
   list         Lists the devices in the blacklist.
```

Example: List the options of the **list** subcommand of the **blacklist** command that is a part of the devices plug-in to the admin CLI.

```
sas-admin devices blacklist help list
```

Here is the output from this command in a Linux environment:

```
NAME:
   sas-devices blacklist list - Lists the devices in the blacklist.

USAGE:
   sas-devices blacklist list [command options] [arguments...]

OPTIONS:
   --all    Returns all of the devices in the blacklist.
   --limit "10"    Specifies the maximum number of devices to return. The default value is 10.
   --start "0"    Specifies the 0-based offset of the first device to return. The default value is 0.
```

Command-Line Interfaces: Concepts

Output Type

You must specify an output type for your CLI when you create your profile. The output types for CLIs are as follows:

- text

  Specifies that the output from the CLI is in text format. This is the default format.

- JSON

  Specifies that the output from the CLI is in JSON format.

- fulljson

  Specifies that the output from the CLI is the entire JSON response. This option is useful when writing scripts in which you need access to the entire response in order to complete a task.

Global Command: Output

Use the **output** global command to specify the output format for a CLI command. Doing so overrides any output type that was specified when you created your profile or that was set in the **SAS_OUTPUT** environment variable.

Global Command: Profile

Overview

Use the **profile** global command to create the connection profile that defines your SAS Viya deployment. This process creates the following two files in the directory ***home-directory*/.sas**:

- config.json

Contains information about your SAS Viya deployment, including the name of the connection profile, the service endpoint, and the output type (**text**, **json**, or **fulljson**).

Note: If you downloaded the CLI executables, the config.json file is created in this directory:

Table A.5 CLI Directory for config.json File

Linux	***home-directory*/.sas/admin-plugins**
Windows	***home-directory*\\.sas\\admin-plugins**
Macintosh	***home-directory*/.sas/admin-plugins**

- credentials.json

  Will contain the authentication tokens for your session that are created after you issue the **auth login** global command.

Default Profile and Named Profile

You can create a default profile or a named profile. If you do not specify a named profile in the **sas-admin auth login** command, and the **SAS_CLI_PROFILE** environment variable is not set, then the default profile is used.

You can use a named profile if you need to work with two or more environments simultaneously from the same machine using the same set of CLIs. When you log on to an environment with a named profile, the associated token is stored for that specific profile. You can work in different environments at the same time by specifying different profile names in the commands. For example, suppose that you have different development, test, and production environments. You can create a separate, named profile for each environment to help distinguish the environment that you are connecting to.

You can also use a named profile to eliminate the requirement to create a profile with the correct settings every time you log on. Suppose that you know that you want to use the JSON output type and a certain endpoint. You can create a named profile with these settings. Then when you want to log on, you can specify this profile name and eliminate the need to create a new profile with these settings

Note: If you logged in using a named profile, you must do one of the following to use that profile:

- set the **SAS_CLI_PROFILE** environment variable to the name of the profile. This will remain in affect until you log off from the environment.

  For example, suppose that you have set the **SAS_CLI_PROFILE** environment variable to **Target1**.

  Then, you can log on to the environment with the “Target1” profile with this command: `sas-admin auth login`

  If you are signing in to an environment that has been configured for Kerberos, use this command: `sas-admin auth kerberos`.

- include the **profile** global option on each CLI command as follows:

  ```
  sas—admin --profile profile-name CLI-name CLI-commands
  ```

  Then, you can sign in to the environment with the “Target1”profile with this command: `sas-admin --profile Target1 auth login`

  If you are signing in to an environment that has been configured for Kerberos, use this command: `sas-admin --profile Target1 auth kerberos`

Examples

Here are typical examples of creating default and named profiles in an environment that has not been configured to use Kerberos:

Example: Create a default connection profile, specify the `text` output type, and specify the path to your SAS Viya deployment as the service endpoint.

Note: The `SAS_CLI_PROFILE` environment variable must not be set in order for this example to work.

```
sas-admin profile init
```

- At the prompt for **Service Endpoint**, enter the URL for the SAS Viya environment as follows:`https://`*`endpoint URL`*
- At the prompt for **Output type**, enter `text`.
- At the prompt for **Enable ANSI colored output**, enter `y` or `n`.

Here is the config.json file that was created. **“Default”** indicates that a default profile was created.

```
{
  "Default": {
    "ansi-colors-enabled": "false",
    "output": "text",
    "sas-endpoint": "https://endpoint URL"
  }
}
```

Example: In the same environment, create a connection profile that is named `Target1`, specify the `json` output type, and specify the path to your SAS Viya deployment as the service endpoint.

```
sas-admin --profile Target1 profile init
```

Here is the config.json file that was created. Notice that there are now two profiles: “Default” and “Target1”. Notice that the output type for the “Target1” profile is JSON.

```
{
  "Default": {
    "ansi-colors-enabled": "false",
    "output": "text",
    "sas-endpoint": "https://endpoint URL
  },
  "Target1": {
    "ansi-colors-enabled": "false",
    "output": "json",
    "sas-endpoint": "https://endpoint URL
```

See “Command-Line Interface: Preliminary Instructions” on page 681 for more information about creating profiles.

Example: In the same environment, log on using the `Target1` profile that you created in the previous example.

```
sas-admin --profile Target1 auth login
```

Global Command: Authenticate

Use the `authenticate` global command to log on or log off from the environment. This process stores a token in the credentials.json file.

Note: The token expires after 12 hours, so you might need to re-execute the command at a later time to reconnect to the environment.

To log on , run the **`auth login`** command. For syntax information, see “Structure” on page 686.

Note: Kerberos is supported for Windows and Linux environments. Sign in to environments that have been configured for Kerberos with this command: `auth kerberos`

Here are typical examples from a Linux environment that has not been configured for Kerberos:

- Use this command to log on with a default profile: `sas-admin auth login`. The **`SAS_CLI_PROFILE`** environment variable must not be set in order for this example to work.
- Use this command to log on with the profile named “Target1”: `sas-admin --profile Target1 auth login`.

  If the **`SAS_CLI_PROFILE`** environment variable is set to **`Target1`**, then you can log on to the “Target1” environment as follows: `sas-admin auth login`.

To log off, issue the **`auth logout`** command. For syntax information, see “Structure” on page 686. Here are typical examples:

- Use this command to log off from the SAS Viya environment that is specified in your default profile: `sas-admin auth logout`. The **`SAS_CLI_PROFILE`** environment variable must not be set in order for this example to work.
- Use this command to log off from a SAS Viya environment that is specified in the profile named “Target1”: `sas-admin --profile Target1 auth logout`.

  If the **`SAS_CLI_PROFILE`** environment variable is set to **`Target1`**, then you can log off from the “Target1” environment as follows: `sas-admin auth logout`.

Source Files: JSON Templates

Overview

For some CLI commands, you can include the command options in a separate JSON file by using the **`source-file`** option. This file is referred to as a template or source file. You might choose this option for the following reasons:

- to accelerate the process for entering CLI commands by including common options in a template
- to automatically create a template by piping output to a file

Details

- If a duplicate CLI option is specified on the command line and in the source file, the command-line option takes precedence over the option in the source file.
- Although some CLI options can be specified in the source file and on the command line, other CLI options can be specified only on the command line.

Create Caslibs from JSON Templates

Templates that are used for creating caslibs must be JSON files. You can generate sample JSON templates for creating caslibs with the cas CLI **`generate-cas-samples`** command.

Example: Generate a template for creating a caslib. In this example, a template is generated for creating a path caslib. You will generate the path.json sample template, and use it to create the new template.

1
```
sas-admin cas generate-cas-samples --output-location sample-template-path
```

2
```
sas-admin cas caslibs create --help
```

```
3 sas-admin cas caslibs create path --server serverA  --path path --name caslibA
 --source-file /sample-template-path/path.json
```

```
4 sas-admin --output json cas caslibs show-info --name caslibA > pathA.json
```

1 Generate sample templates. The path.json template is one of the files that are created. Specify the folder location for the samples templates in the **`output-location`** option.

2 List the types of caslibs that can be specified. You will use this value with the **`create`** command in the next step.

3 Create the new path caslib with the name caslibA. In the **`source-file`** option, specify the path.json sample template that you just created. In the next step, you will use the information about the new caslib to create a JSON template.

On the command line, specify any options for overriding the values in the sample template. In this example, the server and path were specified.

4 Redirect the JSON output for caslibA to a file named pathA.json, which you can then use as a template for creating new caslibs. Specify the use of JSON output with the **`output json`** global option.

Example: Create a caslib from a template. This example uses the template that you created in the previous example.

```
1 sas-admin cas caslibs create --help
```

```
2 sas-admin cas caslibs create path --server serverB  --path pathB --name caslibB --source-file pathA.json
```

```
3 sas-admin --output json cas caslibs show-info --name caslibB --server serverB
```

1 List the types of caslibs that can be specified. You will use this value with the **`create`** command in the next step.

2 Create a new caslib named caslibB by specifying the pathA.json file that you created in the previous example in the **`source-file`** option. On the command line, specify any options for overriding the values in the template. In this example, the server and path were specified.

3 Show information about caslibB in JSON format. Note that the server and path that were specified on the command line overwrote the server and path values that were in the template.

Create User-Defined Formats from JSON Templates

In templates that are used for adding a user-defined format, templates that include command options as well as the format ranges must be JSON files. You can generate sample JSON templates for creating user-defined formats with the cas CLI **`generate-cas-samples`** command.

Example: Create a JSON template for adding a user-defined format.

```
1 sas-admin cas generate-cas-samples --output-location sample-template-path
```

```
2 sas-admin cas format-libraries add-format json --format-library format_libraryA --server serverA
 --source-file  /sample-template-path/gender.json
```

```
3 sas-admin cas --output json format-libraries show-format-ranges --format 'en_us-$gender'
 --format-library format_libraryA --server serverA  > formatA_template.json
```

```
4 sas-admin cas format-libraries add-format json --format-library format_libraryB
 --server serverA --source-file formatA_template.json
```

1 Generate sample templates. The gender.json template is one of the files that are created. Specify the folder location for the samples templates in the **`output-location`** option.

2 Add a new format to a format library. In the **`source-file`** option, specify the gender.json sample template that you just created. In the next step, you will use the information about the new format to create a JSON template.

On the command line, specify any options for overriding the values in the sample template. In this example, the format library was specified. If the command is successful, the server returns a message that identifies the name of the format that you added. In this example, the `en_us-$gender` file is created. The `en_us-` prefix identifies the English locale.

3 Redirect the JSON output for the new format to a JSON file using the `format-libraries show-format-ranges` command. Specify the use of JSON output with the `output json` global option. The format name to use in the `format` option was returned from the server when you created the format. Since the gender format is a character format, you must precede it with a dollar sign ($) and enclose it in quotation marks.

You can use the JSON file as a template for creating new caslibs. Note that the format library that you specified when you created the new format is saved in the new JSON file.

4 Create a new format using the new JSON file as the `source-file` template. On the command line, specify any options for overriding the values in the template. In this example, a new format library was specified.

Create Policies from JSON Templates

Templates that are used for creating CAS server policies must be JSON files.

Important: For CAS to use resource management policies, the environment variable `CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT` must be set on the CAS server. For more information, see Environment Variables.

Global Caslib Policies

Example: Create a policy from a sample JSON template.

```
1 sas-admin cas generate-cas-samples --output-location sample-template-path

2 sas-admin  cas servers policies define global-caslibs
  --server serverA --source-file /sample-template-path/policies-examples/globalCaslibs.json
```

1 Generate sample templates. Specify the folder location for the sample templates in the `output-location` option. The policies-examples folder is created, and one of the files in the folder is globalCaslibs.json.

Edit the globalCaslibs.json file for your environment. You can add new policy options on page 544 or change values of existing policy options. Save the file.

Note: As an alternative to manually editing the JSON file, you can add additional options by using the `attributes` option.

2 Define a global caslib policy by specifying the edited globalCaslibs.json sample template file in the `source-file` option. The values in the globalCaslibs.json file are used to define the new policy. Information about the new policy is returned from the server after the policy is created. Make sure that the policy is correct.

Example: Create a JSON template for creating a global caslib policy.

```
1 sas-admin cas generate-cas-samples --output-location sample-template-path

2 sas-admin  cas servers policies define global-caslibs
  --server serverA --source-file /sample-template-path/policies-examples/globalCaslibs.json

3 sas-admin --output json cas servers policies show-info --policy globalCaslibs
 --server serverA > /path/global-caslib-template.json

4 sas-admin  cas servers policies define global-caslibs
 --attributes sessionTables:1000000000 --server serverA --source-file /path/global-caslib-template.json

5 sas-admin cas servers policies show-info --policy globalCaslibs --server serverA
```

1 Generate sample templates. Specify the folder location for the sample templates in the `output-location` option. The policies-examples folder is created, and one of the files in this folder is globalCaslibs.json.

Edit the globalCaslibs.json file for your environment. You can add new policy options on page 544 or change values of existing policy options. Save the file.

2. Define a global caslib policy by specifying the edited globalCaslibs.json template file in the **`source-file`** option. The values in the globalCaslibs.json file are used to define the new policy. Information about the new policy is returned from the server after the policy is created. You will use this new policy to create a JSON template.

3. To create the template, redirect the JSON output for the new policy to a JSON file using the **`servers policies show-info`** command. Specify the use of JSON output with the **`output json`** global option. The policy name to use in the **`policy`** option was returned from the server when you created the policy in the previous step. Open the new JSON file to verify that it is correct.

4. Create a new global caslib policy using the new JSON template in the **`source-file`** option. On the command line, specify any new attributes or attributes to override the values in the template. In this example, a session tables value was specified. The value that you assign to the session tables attribute must be specified in bytes.

5. Display information about the new global caslib policy to verify the presence of the desired attributes and the additional session tables attribute.

Priority Assignments Policies

Note: The groups to which you are assigning a priority-level must exist in LDAP or as a custom group in order for the policy to work.

Example: Create a policy from a sample JSON template.

```
1 sas-admin cas generate-cas-samples --output-location sample-template-path
```

```
2 sas-admin  cas servers policies define priority-assignments
  --server serverA --source-file /sample-template-path/policies-examples/priorityAssignments.json
```

1. Generate sample templates. Specify the folder location for the sample templates in the **`output-location`** option. The policies-examples folder is created, and one of the files inside this folder is priorityAssignments.json.

   Edit the priorityAssignments.json file for your environment. You can add new policy options on page 544 or change values of existing policy options. Save the file.

   Note: As an alternative to manually editing the JSON file, you can add additional options by using the **`attributes`** option in the next step.

2. Define a priority assignments policy by specifying the edited priorityAssignments.json sample template file in the **`source-file`** option. The values in the priorityAssignments.json file are used to define the new policy. Information about the new policy is returned from the server after the policy is created. Make sure that the policy is correct.

Example: Create a JSON template for creating a priority assignment policy.

```
1 sas-admin cas generate-cas-samples --output-location sample-template-path
```

```
2 sas-admin cas servers policies define priority-assignments
--server serverA --source-file /sample-template-path/policies-examples/priorityAssignments.json
```

```
3 sas-admin --output json cas servers policies show-info --policy priorityAssignments
 --server serverA > /path/priority-assignments-template.json
```

```
4 sas-admin cas servers policies define priority-assignments
--attributes userM:1 --server serverA --source-file /path/priority-assignments-template.json
```

```
5 sas-admin cas servers policies show-info --policy priorityAssignments --server serverA
```

1. Generate sample templates. Specify the folder location for the sample templates in the **`output-location`** option. The policies-examples folder is created, and one of the files in the folder is priorityAssignments.json.

Edit the priorityAssignments.json file for your environment. You can add new policy options on page 544 or change values of existing policy options. Save the file.

2 Define a priority assignments policy by specifying the edited priorityAssignments.json template file in the **source-file** option. The values in the priorityAssignments.json file are used to define the new policy. Information about the new policy is returned from the server after the policy is created. You will use this new policy to create a JSON template.

3 To create the template, redirect the JSON output for the new policy to a JSON file using the **servers policies show-info** command. Specify that JSON output is used with the **output json** global option. The policy name to use in the **policy** option was returned from the server when you created the policy in the previous step. Open the new JSON file to verify that it is correct.

4 Create a new priority assignments policy using the new JSON template in the **source-file** option. On the command line, specify the users and their priority levels with the **attributes** option. In this example, userM is added to the policy with priority 1. The server returns information about the new policy after it is defined. Verify that userM is now included in the policy as well as the groups that were already in the source file.

5 Display information about the new priority assignments policy to verify the presence of the desired attributes and the additional user.

Priority-Level Policies

Note:

Priority-level policies can be used to assign resources based on a user's group membership or by explicit assignment with the priority assignment policy. To facilitate assigning resources by group membership, define custom groups with the same names as the priority-level policy and assign users to the groups.

Example: Create a policy from a sample JSON template.

```
1 sas-admin cas generate-cas-samples --output-location sample-template-path

2 sas-admin  cas servers policies define priority-levels --server serverA --priority 3
--source-file /sample-template-path/policies-examples/cas-shared-default-priority-2.json
```

1 Generate sample templates. Specify the folder location for the sample templates in the **output-location** option. The policies-examples folder is created, and one of the files in the folder is cas-shared-default-priority-2.json.

Edit the cas-shared-default-priority-2.json file with values appropriate for your environment. You can add new policy options on page 544 or change values of existing policy options. Save the file.

2 Define a priority-level policy by specifying the edited cas-shared-default-priority-2.json template file in the **source-file** option. Use the **priority** option to specify the priority for the policy. In this example, the new policy is priority 3. Information about the new policy is returned from the server after the policy is created. Make sure that the policy is correct.

Example: Create a JSON template for creating a priority-level policy.

```
1 sas-admin cas generate-cas-samples --output-location sample-template-path

2 sas-admin cas servers policies define priority-levels --priority 3
--server serverA --source-file  /sample-template-path/policies-examples/cas-shared-default-priority-2.json

3 sas-admin --output json cas servers policies show-info --policy serverA-priority-3
 --server serverA > /path/priority-levels-template.json

4 sas-admin cas servers policies define priority-levels --priority 4 --cpu 30 --session-tables 2000000000
 --server serverA --source-file /path/priority-levels-template.json
```

1 Generate sample templates. Specify the folder location for the sample templates in the **output-location** option. The policies-examples folder is created, and one of the files in the folder is cas-shared-default-priority-2.json.

Edit the cas-shared-default-priority-2.json file with values appropriate for your environment. You can add new policy options on page 544 or change values of existing policy options. Save the file.

2 Use the edited cas-shared-default-priority-2.json template file in the **`source-file`** option to define a new policy. You will use this new policy to create a JSON template.

In this example, a policy for priority level 3 is specified with the **`priority`** option. Information about the new policy is returned by the server after the policy is created. The name of priority-level policies is generated as follows: *server*-**`priority`**-*priority*. Therefore, in this example, the policy is named serverA-priority-3.

3 To create the template, redirect the JSON output for the new policy to a JSON file using the **`servers policies show-info`** command. Specify the use of JSON output with the **`output json`** global option. Open the new JSON file to verify that the contents are correct.

4 Create a new policy using the new JSON template in the **`source-file`** option. This example a shows a policy for priority level 4.

Specify the appropriate values for priority level 4. This example specifies different values for the CPU and session-tables attributes with the **`cpu`** and **`session-tables`** options. The value that you assign to the **`session-tables`** attribute must be in bytes. The server returns information about the new policy after it is defined. Verify that the serverA-priority-4 policy was created with the correct values for CPU and session tables.

Note: To enable the policy resource settings by group membership, create a custom group that has the same name as the priority-level policy. For example, if the policy is named serverA-priority-4, you must create a custom group with the same name. Then, add the users to the group who will be granted this priority level. These users will automatically be granted the resource settings that are specified in the priority-level policy by being a member of this custom group.

See Also

"CAS Resource Management Policies" on page 544

Source Files: Data

Overview

For some CLI commands, you can pass data to the CLI by including it in a separate file by using the **`source-file`** option. This file is referred to as a template or source file. You might want to use this option to enter multiple data items at the same time rather than entering multiple CLI commands to enter data items.

Examples

Format Ranges

Templates that are used for adding the ranges for new user-defined formats must be CSV files or JSON files. You can generate sample templates for adding ranges to user-defined formats with the cas CLI **`generate-cas-samples`** command.

Here is an example of a formats source file in CSV format:

```
F,female
M,male
f,female
m,male
```

See Also

See "Creating Formats" on page 716 for examples of creating user-defined formats with source files.

CAS Server Paths Lists

Files that are used for adding paths to the paths list must be text files.

Here is an example of a paths list source file:

```
/opt/sas/viya/config/folderA
/opt/sas/viya/config/folderB
/opt/sas/viya/config/folderC
```

Example: Add multiple paths to the paths list for the specified CAS server using a template. This example assumes that you have a template text file that is formatted so that each path is on a separate line, as shown in the preceding source file.

```
1 sas-admin cas servers paths-list list --server serverA

2 sas-admin cas servers paths-list add-paths --server serverA --source-file /path_to_template/templateA.txt

3 sas-admin cas servers paths-list list --server serverA
```

1 Display the current paths list for serverA.

2 Add the paths in the template text file to the paths list for serverA.

3 Display the current paths list for serverA to confirm that the paths were added.

Devices Blacklist

Files that are used for adding mobile devices to the blacklist must be text files.

Here is an example of a mobile devices source file:

```
deviceID4
deviceID5
deviceID6
```

Example: Add multiple devices to the blacklist using a template. This example assumes that you have a template text file that is formatted with each device on a separate line, as shown in the preceding source file.

```
sas-admin devices blacklist add --source-file /path_to_template/templateA .txt
```

Command-Line Interface: Examples

CLI Examples: Audit

The following examples assume that you have already signed in to SAS Viya at the command line. See “Command-Line Interface: Preliminary Instructions” on page 681.

Examples

Example: List the records of audit entries for reports.

```
sas-admin audit list --application reports
```

Example: List the records of audit entries of type security and sort by user.

```
sas-admin audit list --sort-by user --type security
```

Example: List the records of audit entries of type security and state of success, and then write the results as CSV to an output file:

```
sas-admin audit list --state success --type security --csv /tmp/outputfile.text
```

See Also

- "Command-Line Interface: Overview" on page 678
- Auditing on page 283

CLI Examples: CAS Authorization

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Getting Access Information

Example: List tableA's direct access controls.

```
sas-admin cas tables list-controls --server serverA --caslib caslibA --table tableA
```

Example: List tableA's direct access controls and inherited settings.

```
sas-admin cas tables list-controls --server serverA --caslib caslibA --table tableA --list-type all
```

Example: Show effective (net) access to tableA for userA.

```
sas-admin cas tables list-controls --server serverA --caslib caslibA --table tableA
--control-type effective --user userA
```

Example: Show effective access to tableA for groupA and groupB.

```
sas-admin cas tables list-controls --server serverA --caslib caslibA --table tableA
--control-type effective --group 'groupA|groupB'
```

Example: Show the source of userA's access to tableA.

```
sas-admin cas tables list-controls --server serverA --caslib caslibA --table tableA
--control-type origin --user userA
```

Managing Access Controls

Example: Remove all direct access controls from tableA.

```
sas-admin cas tables clear-controls --server serverA --caslib caslibA --table tableA
```

Example: Set a simple row-level access control on the CARS table so that members of groupA can see only those rows where the value in the Make column is **`Ford`**.

```
sas-admin cas tables add-control --server serverA --caslib caslibA --table CARS --group groupA
--grant Select --where "make='Ford'"
```

Example: Set an identity-based, row-level access control on the Salary table. The reason is so that each authenticated user can see only those rows where the value in the User column is his or her own user ID.

```
sas-admin cas tables add-control --server serverA --caslib caslibA --table salary --group "*"
--grant Select --where "User='SUB::SAS.Userid'"
```

Example: Enable guests to read data in the Public caslib.

```
1 sas-admin cas caslibs add-control --server serverA --caslib Public --guest --grant ReadInfo
 --superuser

2 sas-admin cas caslibs add-control --server serverA --caslib Public --guest --grant Select
 --superuser
```

1 This example is applicable to only a deployment where guest access is enabled.

In the standard configuration, because only a privileged user can modify access to the Public caslib, the superuser option is specified here. Only a member of the Superuser role for the specified CAS server can obtain elevated privileges by specifying the superuser option.

2 The grants in this example support reading of data, but do not support just-in-time loading of data. Instead of granting LimitedPromote to guest, consider using a different technique for ensuring that data is loaded.

Because this example does not create and reuse a dedicated superuser session, you must specify the superuser option in each command where elevated privileges are needed.

Example: Enable groupA to read data (and perform just-in-time data loading) in a new caslib.

```
1 sas-admin cas caslibs add-control --server serverA --caslib caslibA --group groupA --grant ReadInfo

2 sas-admin cas caslibs add-control --server serverA --caslib caslibA --group groupA --grant Select

3 sas-admin cas caslibs add-control --server serverA --caslib caslibA --group groupA
 --grant LimitedPromote
```

Example: Make the same changes as in the preceding example, but use an access control transaction so that you can review your changes before you commit them to the server.

```
1 sas-admin cas sessions create --name mysess --server serverA --superuser

2 sas-admin cas transactions checkout --session-id XYZ --server serverA --caslib caslibA

3 sas-admin cas caslibs add-control --session-id XYZ --server serverA --caslib caslibA
--group groupA --grant ReadInfo

4 sas-admin cas caslibs add-control --session-id XYZ --server serverA --caslib caslibA
--group groupA --grant Select

5 sas-admin cas caslibs add-control --session-id XYZ --server serverA --caslib caslibA
--group groupA --grant LimitedPromote

6 sas-admin cas caslibs list-controls --session-id XYZ --server serverA --caslib caslibA

7 sas-admin cas transactions commit --session-id XYZ --server serverA

8 sas-admin cas sessions delete --session-id XYZ --server serverA
```

1 Start a session. If you are a member of the Superuser role for the associated CAS server, give the session Superuser status.

2 Check out the caslib into the session that you just started. Use the session ID that is returned from the preceding command. The session-id value `XYZ` is used here for simplicity. Checking out an object automatically starts a transaction.

3 Grant access within the transaction.

4 Grant access within the transaction.

5 Grant access within the transaction.

6 Review the results. Because you supply the session ID, the output reflects the uncommitted changes in your session.

7 After you review the output from the list-controls command, commit your changes.

8 If you are finished, it is a good practice to delete your session.

Example: Replace any direct access controls on tableA with access controls from an external JSON file. In this example, the replacement access controls are derived from tableB and are then applied to tableA.

```
1 sas-admin cas tables list-controls --server serverA --caslib caslibA --table tableB > ac.json

2 sas-admin cas tables replace-controls --server serverA --caslib caslibA --table tableA
 --source-file ac.json
```

```
3 sas-admin cas tables list-controls --server serverA --caslib caslibA --table tableA
```

1 This example writes the direct access controls for tableB to the file **`ac.json`** in the directory from which you are running your CLI.

Note: This example assumes that the profile that you are using specifies **`json`** as your default output type. Otherwise, you must use the global option --`output` to specify **`json`** as the output type for this command. That option must immediately follow `sas-admin`.

Note: You can reference an absolute path or a relative path.

2 Delete any direct access controls on tableA, and replace them with the access controls that you wrote to the ac.json file.

3 Review the new set of direct access controls on tableA.

Note: You can modify the output from tableB before you use it to replace direct access controls on tableA.

Details and Tips

Basics

- Throughout this topic, the term *access control* refers to an access control in the CAS authorization system. To manage access to content objects and functionality, see “CLI Examples: General Authorization” on page 702.
- You can add, delete, and replace only direct access controls.
- A request to delete a direct access control that does not exist does not generate an error.
- To modify access that a table inherits, set direct access controls on the parent caslib.
- You cannot modify access that a caslib inherits.
- Use of access control transactions is optional. You do not have to check out an object in order to modify its access controls.
- In the list-controls command, use the control-type and list-type options as follows:
  - Use the control-type option only if you want to obtain net access information (**`effective`**) or source information (**`origins`**).
  - The list-type options are not relevant if the control-type is **`effective`** or **`origin`**.
  - Use the list-type option only if you want to obtain inherited settings, in addition to direct access controls (**`all`**) or instead of direct access controls (**`inherited`**).
- You can obtain origins on page 403 information for only one identity at a time.
- The value **`serverA`** is used in the examples for simplicity. A more typical server name is **`cas-shared-default`**.
- See “Details” on page 710 for information about using environment variables with the CAS commands.

Principals

- To specify a particular identity (where supported), you must provide a user ID or a group ID rather than a name.

  CAUTION! The user ID and the group ID that you provide are not validated. Make sure the IDs that you provide are accurate.
- To specify multiple users or multiple groups (where supported), use the pipe character (|) as a delimiter and enclose the string in single quotation marks (for example: --`user 'userA|userB'`). You cannot specify both users and groups in a single request.

- The group * corresponds to Authenticated Users. To specify that principal, enter `--group '*'`.
- The Guest principal represents all users who connect as guests. To specify that principal, enter the option `--guest` and do not specify a value.

Permissions

- To specify a particular permission (where supported), use one of the following case-insensitive values: ReadInfo, Select, LimitedPromote, Promote, CreateTable, DropTable, DeleteSource, Insert, Update, Delete, AlterTable, AlterCaslib, or ManageAccess. You cannot specify multiple permissions in a single request.
- For information about the scope and purpose of each permission, see "CAS Authorization: Concepts" on page 394.

Fine-Grained Controls

- Row-level grants are always for the Select permission on a table. The syntax for row-level permission filters is the same as in other CAS authorization interfaces.
- You cannot set column-level permissions using this interface.

See Also

- "Command-Line Interface: Overview" on page 678
- CAS Authorization on page 383

CLI Examples: General Authorization

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Getting Access Information

Example: Show detailed properties of a specified rule.

```
sas-admin authorization show-rule --id d85144aa-79dc-4852-b949-645cc5ff8ffc --details
```

Example: Show effective access for a specified object URI. To return information about contributing rules, you must specify fulljson output in your profile. Specifying the **`output`** global option in the CLI command inline is insufficient.

```
sas-admin authorization explain --target-uri /SASHome/**
```

Example: List all the rules in the deployment.

```
sas-admin authorization list-rules
```

Managing Rules

Example: Give groupA Read access to reportA.

```
sas-admin authorization authorize --permissions Read --group groupA
--object-uri /reports/reports/33db163a-716e-4980-a5bc-6c42a0278c40
```

Example: Provide guest access to reportA.

```
sas-admin authorization authorize --permissions Read --guest
--object-uri /reports/reports/33db163a-716e-4980-a5bc-6c42a0278c40
```

Example: Grant Authenticated Users Read access to folderA and its child members.

```
sas-admin authorization authorize --permissions Read --authenticated-users
--object-uri /folders/folders/2414f911-d276-4357-8550-fcf03753c9e7/**
--container-uri /folders/folders/2414f911-d276-4357-8550-fcf03753c9e7
```

Example: Delete a rule.

```
1 sas-admin authorization show-rule --id d85144aa-79dc-4852-b949-645cc5ff8ffc --details

2 sas-admin authorization remove-rule --id d85144aa-79dc-4852-b949-645cc5ff8ffc
```

1 Review the rule’s properties so that you are certain you are deleting the correct rule.

2 Delete the rule.

Example: Change the principal in an existing rule so that the rule is assigned to Group B, which has **`groupB`** as its ID.

```
1 sas-admin authorization show-rule --id cd75a376-c5d4-4951-9e57-cf441610628c --details

2 sas-admin authorization update-rule --id cd75a376-c5d4-4951-9e57-cf441610628c --group groupB
```

1 Review the rule’s properties so that you are certain you are modifying the correct rule.

2 Modify the rule.

Example: Include the Update and Delete permissions in an existing rule that already grants the Read permission.

```
1 sas-admin authorization show-rule --id cd75a376-c5d4-4951-9e57-cf441610628c --details

2 sas-admin authorization update-rule --id cd75a376-c5d4-4951-9e57-cf441610628c
--grant --permissions Read,Update,Delete
```

1 Review the rule’s properties so that you are certain you are modifying the correct rule.

2 Modify the rule.

Example: Edit the description in an existing rule.

```
1 sas-admin authorization show-rule --id 0e8a6ce7-e51a-40cc-aeda-ee2a5efb53ca --details
2 sas-admin authorization update-rule --id 0e8a6ce7-e51a-40cc-aeda-ee2a5efb53ca
--description 'This is a revised description.'
```

1 Review the rule’s properties so that you are certain you are modifying the correct rule.

2 Modify the rule.

Managing Rules (Bulk Approach)

Adding Rules

Example: Add multiple rules, as specified in a referenced JSON file.

```
sas-admin authorization create-rules --file newrules.json
```

Note: The path to the input file is relative to the location of the CLI executable (sas-admin).

The input file is in JSON format. Here is an example that adds two rules:

```
[
   {
    "op": "add",
    "value": {
      "description": "Description for this rule.",
      "objectUri": "/folders/folders/156f833a-31ac-40f8-bf78-82e738daef36",
      "permissions": [
        "Secure"
      ],
      "principalType": "user",
      "principal": "userC",
      "type": "grant"
      }
   },
   {
    "op": "add",
    "value": {
      "description": "Description for this rule.",
      "objectUri": "/folders/folders/156f833a-31ac-40f8-bf78-82e738daef36",
      "permissions": [
       "Read",
       "Update",
       "Delete"
      ],
      "principalType": "group",
      "principal": "groupD",
      "type": "grant"
    }
   }
]
```

Because an ID for each rule is not specified in the preceding example, the rule ID is generated. For details about available parameters and values, see the Authorization REST API documentation on developer.sas.com.

Deleting Rules

Example: Delete multiple rules, as specified in a referenced JSON file.

```
   sas-admin authorization remove-rules --file oldrules.json
```

Note: The path to the input file is relative to the location of the CLI executable (sas-admin).

The only required content for the remove-rules command is the ID of each rule that you want to remove. Here is an example of the input file for a remove-rules command:

```
[
   {
    "op": "add",
    "value": {
      "id": "rule123456788"
    }
   },
   {
    "op": "add",
    "value": {
      "id": "rule123456789"
    }
   }
]
```

Notice that even though this file is used in a removal command, the value of the op parameter is **add**.

> **TIP** If you specify a rule ID in an input file that you use to create rules, you can reference the same file to remove the rules. The remove-rules command ignores values other than the rule ID.

Details and Tips

- Throughout this topic, the term *rule* refers to an authorization rule in the general authorization system. To manage access to CAS objects (such as caslibs and tables), see "CLI Examples: CAS Authorization" on page 699.
- To assign a rule to a principal type, use one of the following options:

Principal Type	Option
Guest	`--guest`
Authenticated Users	`--authenticated-users`
Everyone	`--everyone`

- To assign a rule to a particular identity, you must provide a user ID or a group ID, not a name. For example, to assign a rule to the SAS Administrators custom group, specify: `--group SASAdministrators`

  CAUTION! The user ID and the group ID that you provide are not validated. Make sure the IDs that you provide are accurate.
- You can obtain the ID for a user or group from the **Users** page in SAS Environment Manager.
- You can obtain the objectURI for a content object (such as a report) from the **Content** page in SAS Environment Manager. Select the object in the navigation pane. On the right, the **URI** field in the **Basic Properties** section contains the object URI. To target the object URI for a content object (such as a report) or a container (such as a folder), append a suffix. See "Rule Targets" on page 425.
- You can obtain the ID for a rule from the **Rules** page in SAS Environment Manager. Right-click a rule and select **Properties**. The last field in the Properties window contains the rule's ID.
- When you use the show-rule command, always specify that you want details to be returned. Some of the fields that can be essential to interpreting a rule are excluded from the default response. For example, a condition is not included in the default response.
- When you use the update-rule command, specify only the options for the rule properties that you want to modify. For any option that you specify in the update-rule command, provide the complete replacement value or values.
- When you use the explain command, the returned information indicates the effective (net) access of each relevant principal for all permissions.

  Note: In the output from the explain command, the `grant` and `prohibit` values indicate effective (net) access, not direct settings. For example, a `prohibit` value in the output from the explain command is usually caused by the lack of any relevant grant, rather than by the existence of a relevant Prohibit rule. See "Authorization Decisions" on page 401.
- Enabling or disabling guest access involves more than running the enable-guest-access or disable-guest-access command. See the guest access documentation.

See Also

- "Command-Line Interface: Overview" on page 678
- General Authorization on page 411

CLI Examples: Backup

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: List the first 50 backup jobs.

```
sas-admin backup list --limit 50
```

Example: Start a binary backup named backupA.

```
sas-admin backup start --slug backupA
```

See Also

- "Overview" on page 34
- "Command-Line Interface: Overview" on page 678

CLI Examples: CAS Administration

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Important: For Windows users of the CAS CLI who are using it for CAS administration and who are CAS administrators, CAS sessions are launched under the CAS service account. Also, any user who has credentials that are stored can use those credentials to run the CAS CLI for any purpose. See "External Credentials: How To" on page 338.

Generate CAS Samples

Example: Generate sample template files for creating caslibs, user-defined formats, and resource management policies.

```
sas-admin cas generate-cas-samples --output-location sample-location
```

Note: The **cas generate-cas-samples** command also generates an md5.txt file that contains checksums for each sample file. SAS Technical Support can use this file to validate the integrity of the sample files.

See Also

"Source Files: JSON Templates " on page 692

Enable Guest Access

Note: This information does not apply to Windows environments or to Linux environments that have been configured for Kerberos.

To enable guest access on the specified server, modify the direct access controls for the predefined caslibs on the server. To do this, use the controls that are defined in the specified source file. If you do not have a preexisting source file, you can generate one that contains the default access controls, make modifications to it, and use it as the source file.

Perform this action with elevated privileges if the user is able to assume the Superuser role. This is one of several required steps to enable guest access. For a complete set of instructions on how to enable guest access, see "Authentication: Guest Access" on page 327.

Example: Enable guest access for the predefined caslibs on the server by using a preexisting source file.

```
sas-admin cas facilitate-guest --source-file path-to-source-file --server serverA --superuser
```

Example: Enable guest access for the predefined caslibs on the server by creating a source file and using it to modify the direct access controls.

```
1 sas-admin cas generate-guest-controls --output-location /path/

2 sas-admin cas facilitate-guest --source-file path-to-controls-file --server serverA --superuser
```

1 Generate a source file from the default access controls. The generated source file is named **facilitate-guest-controls.txt**.

Make the desired modifications to access controls to the source file that you just generated.

2 Modify the direct access controls using the source file that you just modified.

Note: You can also remove guest access with the **cas remove-guest-controls** command. However, this command removes only the default set of direct access controls and not any other guest access controls that you might have applied. To view which direct access controls are removed, see the **facilitate-guest-controls.txt** file that is generated by the **cas generate-guest-controls** command.

Manage CAS Role Memberships

Example: List the administrative users on the specified CAS server.

```
sas-admin cas admin-users list --server serverA
```

Example: Add the user user1 to the Superuser role on the specified CAS server.

```
sas-admin cas admin-users add --user user1 --server serverA
```

Example: Remove the group with the ID group1 and name group1_name from the Superuser role on the specified CAS server. The action is performed without prompting the user for confirmation since the **force** option is used.

```
sas-admin cas admin-users remove --group group1 --server serverA --name group1_name --force
```

> **TIP** To remove a user or a group from the administrative users, you must specify the name of the user or the group as well as the identity of the user or the group. Use the name option to specify the name. Use the user option to specify the ID for a user. Use the group option to specify the ID for a group. In order to obtain the ID, use the **admin-users list** command.

Manage SAS Sessions

Example: List the sessions of which you are the owner on the specified CAS server.

```
sas-admin cas sessions list --server serverA
```

Example: Using elevated privileges, list 50 sessions for which the owner is user1 on the specified CAS server, and sort by **state** . You must be able to assume the Superuser role to run the command with elevated privileges.

```
sas-admin cas sessions list --server serverA --superuser --limit 50 --owner user1 --sort-by state
```

Example: Using elevated privileges, list all sessions for which the name contains the string dataExplorer on the specified CAS server. You must be able to assume the Superuser role to run the command with elevated privileges.

```
sas-admin cas sessions list --server serverA --superuser --all --name-contains dataExplorer
```

Example: Using elevated privileges, show additional information about the session with ID 12345 on the specified CAS server. You must be able to assume the Superuser role to run the command with elevated privileges.

```
sas-admin cas sessions show-info --superuser --session-id 12345  --server serverA
```

Manage Tables

Note: These examples explicitly specify the **server** and **caslib** required options. However, using environment variables might be more efficient for these options. For more information about the environment variables, see "Details" on page 710.

Example: For the specified caslib and CAS server, list all tables with names that contain the string visual, sort by **state**, and return a maximum number of 50 tables.

```
1 sas-admin cas help tables

2 sas-admin cas tables help list

3 sas-admin cas tables list --caslib caslibA --server serverA --name-contains visual
--sort-by state --limit 50
```

1. Review the Help for the cas tables command.
2. Review the Help for the list subcommand of the cas tables command.
3. Issue the command to list the tables for the specified caslib and CAS server with the appropriate subcommand and options.

Example: Load the given table for the specified caslib and CAS server.

```
sas-admin cas tables load --table airlines --server serverA --caslib caslibA
```

Example: Unload the given table for the specified caslib and CAS server.

```
sas-admin cas tables unload --table airlines --server serverA --caslib caslibA
```

Example: Show information about the given table for the specified caslib and CAS server.

```
sas-admin cas tables show-info --table airlines --server serverA --caslib caslibA
```

Manage Caslibs

Important: The caslibs command prompts for the required options for the different types of caslibs that you can create. For information about the available options for each type of caslib, see the **dataSource** option for the **addCaslib** action here: addCaslib Action. Be aware that there are more required options for the cas caslibs CLI command than there are for the addCaslib action set.

Note: These examples explicitly specify the **server** required option. However, using an environment variable might be more efficient for this option. For more information about the environment variables, see "Details" on page 710.

Example: Generate sample template files for creating caslibs.

```
sas-admin cas generate-cas-samples --output-location /sample-template-path
```

Example: On the specified server, create a caslib that is based on a file path.

```
sas-admin cas caslibs create path --name caslibA --path /tmp/dept --server serverA
```

Example: On the specified server, create a caslib that is based on an instance of an Impala server.

```
sas-admin cas caslibs create impala --authentication-domain domain  --server serverA
--name caslibA --schema schema --impala-server Impala-server
```

Example: On the specified server, create a caslib that is based on an instance of a Postgres server.

```
sas-admin cas caslibs create postgres --authentication-domain domain --server serverA
 --name caslibA --postgres-server Postgres-server --postgres-database Postgres-database
```

Example: On the specified server, list the first 20 global caslibs. You must be able to assume the Superuser role to run the command with elevated privileges.

```
sas-admin cas caslibs list --server serverA --scope global --limit 20 --superuser
```

Example: On the specified server, list the global caslibs starting at caslib number 21. You must be able to assume the Superuser role to run the command with elevated privileges.

```
sas-admin cas caslibs list --server serverA --scope global --start 21 --superuser
```

Example: Delete the given caslib from the specified server.

```
sas-admin cas caslibs delete --server serverA --name caslibA
```

See Also

"Create Caslibs from JSON Templates" on page 692

Manage Servers

Note: These examples explicitly specify the **server** required option when applicable. However, using an environment variable might be more efficient than this option. For more information about the environment variables, see "Details" on page 710.

Example: List all the CAS servers.

```
sas-admin cas servers list --all
```

Example: Adds multiple paths to the paths list on the specified server so that a caslib that is based on the paths can be created. The active list in this example is the whitelist. You can specify multiple paths in the file that is referenced in the **source-file** option. See caslib paths list on page 500 for more information.

```
sas-admin cas servers paths-list add-paths --server serverA --source-file /path-to-source-file/source.txt
```

Example: On the specified server, list the identities who can create and delete session and global caslibs.

```
sas-admin cas servers privileges list --server serverA
```

Example: On the specified server, grant the ability to manage global caslibs to a group.

Note: A reference to the group * corresponds to Authenticated Users. To specify the Authenticated Users group, enter --`group '*'`

```
sas-admin cas servers privileges modify --server serverA --grant --global --group group-name
```

Example: Show detailed information about the specified server. Show the dates and times in the local time zone.

```
sas-admin cas servers show-info --server serverA --use-local-datetime
```

Example: Create a new resource management policy for a global caslib. You can also create a global caslib policy from a JSON template on page 694.

```
sas-admin cas servers policies define global-caslibs
--attributes cpu:30,Public:1000000000,HPS:4000000000 --server serverA
```

Example: Create a new resource management policy for priority assignments. You can also create a priority assignment policy from a JSON template on page 695.

```
sas-admin cas servers policies define priority-assignments
 --attributes userA:4,userB:3 --server serverA
```

Example: Create a new resource management policy for priority level 5 on serverA. This creates a policy with the name **serverA-priority-5**. You can also create a priority level policy from a JSON template on page 696.

```
sas-admin cas servers policies define priority-levels --cpu 10 --global-casuser 500000000
--global-casuser-hdfs 100000000 --session-tables 2500000000 --priority 5  --server serverA
```

Example: Create a new resource management policy for priority level 4 on serverA. Specify a source file that contains values for priority level 3, but override them with the values for priority 4 by including them on the command line. This example overrides the value for CPU in the source file, by specifying the **`cpu`** option on the command line. This example assumes that you have a preexisting JSON file for a priority level 3 policy named priority3-levels.json. You can also create a priority level policy from a JSON template on page 696.

```
sas-admin cas servers policies define priority-levels --priority 4 --server serverA --cpu 15
 --source-file /path/priority3-levels.json
```

Note:

- For CAS to use resource management policies, the environment variable **`CAS_ENABLE_CONSUL_RESOURCE_MANAGEMENT`** must be set on the CAS server. For more information, see Environment Variables.
- When you create a resource management policy for priority levels, the policy is named ***`CAS-server-`*****`priority-`*****`priority-level`***.
- When you create a resource management policy for priority levels and you use the **`source-file`** option to specify values, you must include the **`priority`** option on the command line.
- CPU sharing is not supported on Windows. As a result, when creating a resource management policy for priority levels on Windows, the **`cpu`** option is allowed, but will not be honored.
- The values that you assign to the **`global-casuser`**, **`global-casuser-hdfs`**, and **`local-tables`** attributes must be specified in bytes.
- Options that are specified the command line take precedence over the options in the source file.

See Also

- “Create Policies from JSON Templates” on page 694
- “CAS Resource Management Policies” on page 544

Details

- When running the CAS CLI on a UNIX machine or a Macintosh machine, with the CAS server running on a Windows machine, a Windows pathname that contains backslashes must be enclosed in single quotation marks. Here is an example: **`'\\tmp\sas'`**.
- The CAS CLI supports the following environment variables:
  - SAS_CLI_DEFAULT_CAS_SERVER
  - SAS_CLI_DEFAULT_CASLIB
  - SAS_CLI_DEFAULT_CAS_SESSION

  You can assign values to the environment variables that you want to remain in effect throughout your session. If the CAS CLI command requires the **`server`**, **`caslib`**, or **`session-id`** options, and the environment variables are set, then you can omit the required options from the CAS CLI command.

  For example, suppose the following:

  - **SAS_CLI_DEFAULT_CAS_SERVER** is set to **`serverA`**.
  - **SAS_CLI_DEFAULT_CASLIB** is set to **`caslibA`**.

  You can then run this command without specifying the required **server** and **caslib** options: `./sas-admin cas tables show-info --table airlines`.

Note: Some commands do not support the use of some of the environment variables. For example, the CAS CLI ignores the CAS environment variables for the following commands:

- caslib remove-control
- caslib delete
- tables remove-control
- sessions delete

You must explicitly specify all required options when using these commands.

See Also

- "Command-Line Interface: Overview" on page 678
- SAS Cloud Analytic Services on page 469
- Identity Management on page 349
- Data Administration on page 141

CLI Examples: Compute

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: List the compute contexts.

```
sas-admin compute contexts list
```

Example: Validate the compute context session with the specified ID.

```
sas-admin compute contexts validate --id 389fee7a-e164-4e45-b836-a301638e9945
```

Example: Delete the compute context session with the specified name.

```
sas-admin compute contexts delete --name "SAS Job Execution compute context"
```

Example: List the launcher contexts.

```
sas-admin compute launchers list
```

Example: Delete the launcher context with the specified ID.

```
sas-admin compute launchers delete --id 8fbdd5f8-a2ee-42a5-a228-8737a0cf778f
```

Example: List the compute sessions.

```
sas-admin compute sessions list
```

See Also

"Command-Line Interface: Overview" on page 678

CLI Examples: Configuration

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: List all the configurations that exist in the Configuration service.

```
sas-admin configuration configurations list
```

Example: Download the configurations with the specified definition name (**`spring`** in this example) from the SASLogon service, and write the output to a file.

```
sas-admin configuration configurations download --definition-name spring --service SASLogon
--target path-to-output-file
```

Example: List the expectation objects in the Configuration service.

```
sas-admin configuration expectations list
```

Example: Show the expectation with the specified ID.

```
sas-admin configuration expectations show --id 185a046e-4e88-4e29-86cf-61d04b9abd07
```

Details

- You can delete a configuration with the configuration CLI.

  CAUTION! Do not use the **`delete`** subcommand of the configurations command unless you are sure that you want to delete your configuration.

See Also

"Command-Line Interface: Overview" on page 678

CLI Examples: Data Explorer

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: List the data sources.

```
sas-admin dataexplorer data-sources list
```

Example: List the connections for the specified data source.

```
sas-admin dataexplorer data-sources connections list --source dsName
```

Example: List the tables for the specified data source and connection.

```
sas-admin dataexplorer data-sources tables list --source dsName --connection connection
```

Example: List the rows for the specified table for the specified data source and connection.

```
sas-admin dataexplorer data-sources rows list --source dsName --connection connection --table tableName
```

Example: Delete a table from the specified data source and connection.

```
sas-admin dataexplorer data-sources tables delete --source dsName --connection connection --table tableName
```

Example: Create a new table from the specified data source and connection.

```
sas-admin dataexplorer data-sources tables create --target-data-source dsName
--target-connection connection --target-table-name tableName --source-data-source dsName
--source-connection connection --source-table-path path-to-table
```

Example: Create a new table from a local file.

```
sas-admin dataexplorer data-sources tables create --target-data-source dsName
--target-connection connection --target-table-name tableName --local-file path-to-file
```

Example: Create a new table from a table URI.

```
sas-admin dataexplorer data-sources tables create --target-data-source dsName
--target-connection connection --target-table-name tableName
--source-table-uri uri-for-table
```

Example: Query for the specified job.

```
sas-admin dataexplorer jobs query --job-id jobID
```

See Also

- "Command-Line Interface: Overview" on page 678
- SAS Data Explorer: User's Guide

CLI Examples: Device Management

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: Determine whether the device with ID device1 is authorized for use in the environment.

```
sas-admin devices authorized-devices validate --device-id device1
```

Example: Show whether the blacklist or whitelist is being enforced.

```
sas-admin devices enforcement status
```

Example: Add the device with ID device1 to the whitelist.

```
sas-admin devices whitelist add --device-id device1
```

Example: List the devices that are enabled in the whitelist.

```
sas-admin devices whitelist list
```

Example: Add the device with ID device1 to the blacklist.

```
sas-admin devices blacklist add --device-id device1
```

Example: Remove the device with ID device1 from the blacklist.

```
sas-admin devices blacklist delete --device-id device1
```

Example: From a list of the devices of type iPhone that have connected or attempted to connect to the server, add a specific iPhone to the blacklist.

```
1 sas-admin devices last-access list --device-type iPhone

2 sas-admin devices blacklist add --device-id device-id
```

1. List the last-access attempts of all devices of type iPhone.
2. Add a device to the blacklist using the device ID that was identified in the previous step.

Example: List in fulljson output the last access attempts to the server for all devices.

```
sas-admin --output fulljson devices last-access list
```

See Also

- "Command-Line Interface: Overview" on page 678
- Mobile on page 269

CLI Examples: Folders

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: Add subfolderA as a child folder to folderA. In this example, the path to folderA is **/parent-folder/folderA**.

```
sas-admin folders create --name subfolderA --parent-path /parent-folder/folderA
```

Example: Update the name of folderA to departmentA. In this example, the path to folderA is **/parent-folder/folderA**.

```
sas-admin folders update --path /parent-folder/folderA --name departmentA
```

Example: Delete the departmentA folder including any non-folder members. In this example, the path to departmentA is **/parent-folder/departmentA**.

```
sas-admin folders delete --path /parent-folder/departmentA --recursive
```

Example: Move folderA to folderB. In this example, the path to folderA is **/parent-folder/folderA**, and the path to folderB is **/parent-folder/folderB**.

```
3 sas-admin folders move --path /parent-folder/folderA --parent-path /parent-folder/folderB

4 sas-admin folders list-members --path /parent-folder/folderB

5 sas-admin folders list-members --path /parent-folder/folderB --tree
```

1 Move folderA to folderB.

2 List the members of folderB to confirm that folderA is a member.

3 Display the members of folderB in a visual tree structure using the **tree** option.

Example: List the members of folderA using the ID of folderA.

```
sas-admin folders list-members --id folderA-ID
```

Example: Create folderA. Add subfolderA as a child folder to folderA using the ID of folderA.

```
1 sas-admin folders create --name folderA

2 sas-admin folders create --name subfolderA --parent-id parent-folder-ID
```

1 Create folderA and note the ID.

2 Issue the command to create subfolderA and specify the ID of folderA as the parent folder.

Details

- Many of the folders commands require the ID of a folder as an argument. The ID of a folder is displayed when you create the folder. The ID of folders is also displayed when you list folders.
- When you delete a folder that contains non-folder content (such as reports), the CLI displays a message that you cannot delete the folder, because it is not empty. To delete all contents of folders including non-folder content, you must use the **recursive** option of the **delete** subcommand.

See Also

- "Command-Line Interface: Overview" on page 678
- Content Management on page 131

CLI Examples: Fonts

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: Add and register the open-source OpenSans-Bold.tff Google font.

```
sas-admin fonts add --uri https://server:port/fonts/OpenSans-Bold.tff
```

Note: In a Windows environment, to add a font with a file uniform resource indicator (URI), copy the file to the Windows server and then add the font to the environment with this command: `sas-admin fonts add --uri file://hostname/path-to-font`

Example: List the currently registered fonts in the system.

```
sas-admin fonts list
```

Example: Show information about the Arial Symbol font.

```
1 sas-admin fonts list --name "Arial Symbol"

2 sas-admin fonts show-info --id font-ID-of-Arial-Symbol-font
```

1 Identify the ID of the Arial Symbol font.

2 Show information about the Arial Symbol font using the specified font ID.

Example: Delete fontA from the system.

```
1 sas-admin fonts list --name fontA

2 sas-admin fonts delete --file-id file-id of fontA
```

1 Identify the file ID of fontA.

2 Delete fontA from the system using the specified file ID.

Details

- The fonts CLI can be used to add web open font format (WOFF), TrueType (TTF), and TrueType Collection (TTC) fonts to the Fonts service. Once added, the fonts are available for the following purposes:
  - □ To render content in web browsers with web open font format (WOFF) and TrueType (TTF) fonts.

    TrueType Collection (TTC) fonts are not displayed in web browsers. Therefore, users cannot select text that uses TrueType Collection (TTC) fonts if they are included in a SAS web application such as SAS Visual Analytics.
  - □ To render contents for printing with TrueType (TTF) and TrueType Collection (TTC) fonts.

    Web open font format (WOFF) fonts are not supported for printing PDF or SVG output. Therefore, WOFF fonts must be paired with a matching TTF or TTC font for print support.

  Note: The following SAS system fonts are WOFF only, and are not supported for printing:

  - Noto Sans
  - Noto Sans JP

  - Noto Sans KR
  - Noto Sans SC
  - Noto Sans TC
  - Noto Sans Thai

  Given the TTC and WOFF font limitations, you might need to upload a combination of font formats in order to satisfy both of the preceding purposes.

- A font might not be available from a web server that is accessible from the middle tier without authentication. If so, you can upload the font to the SAS Viya server and reference it with the file:/// URI scheme.
- In multi-tenant environments, the fonts are shared across all the tenants. Maintenance is supported only on the provider tenant.
- You are responsible for obtaining licensing of any fonts that are registered with the system. In a multi-tenancy configuration, this includes licensing the fonts for use by all tenants.

See Also

"Command-Line Interface: Overview" on page 678

CLI Examples: Formats

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Creating Format Libraries

Note: By default, the tables for newly created format libraries are stored in the **`FORMATS`** library. You can specify a different caslib to contain the tables with the **`format-libraries create --caslib`** option.

Example: Create a format library from a SAS catalog.

```
sas-admin cas format-libraries create --format-library fmtlib1 --server serverA
 --source-path /path-to-catalog/catalog-name.sas7bcat --su
```

Example: Create a format library from a CAS table.

```
sas-admin cas format-libraries create --format-library fmtlib1 --server serverA
 --source-path /path-to-table/table-name.sashdat --su
```

Example: Create a format library from a CAS item store.

```
sas-admin cas format-libraries create --format-library fmtlib1 --server serverA
 --source-path /path-to-table/item-store-name --su
```

Creating Formats

Note: In the following examples, a dollar sign ($) is used to indicate a character format. SAS character format names must be quoted.

Example: Generate sample templates for creating user-defined formats.

```
sas-admin cas generate-cas-samples --output-location /sample-template-path
```

Example: Create a user-defined format from a JSON template. For more information, see "Source Files: JSON Templates " on page 692. This example assumes that you already have a JSON template file, as follows:

```
{
    "libraryName": "userformats1",
    "locale": "en_US",
    "name": "$ynm",
    "ranges": [
        "1=yes",
        "2=no",
        "3=maybe"
    ],
    "version": 1
}
```

```
1 sas-admin cas format-libraries add-format json --server serverA --source-file /path-to-template/ynm.json
```

```
2 sas-admin cas format-libraries show-format-ranges --server serverA
 --format '$ynm' --format-library userformats1
```

1 Create the ynm format from the JSON template. Since the **`format`** and **`format-library`** options are not included on the command, the library and name values in the template are used.

2 Display information about the ynm format to verify that the userformats1 library and ynm name from the template were used.

Example: Create a user-defined format from a JSON template, but specify a different format and format library on the command line. For more information, see “Source Files: JSON Templates ” on page 692. This example uses the JSON template from the previous example.

```
1 sas-admin cas format-libraries add-format help json
```

```
2 sas-admin cas format-libraries add-format json --server serverA
 --format '$ynmA' --format-library userformats2 --source-file /path-to-template/ynm.json
```

```
3 sas-admin cas format-libraries show-format-ranges --server serverA
 --format '$ynmA' --format-library userformats2
```

1 Review the options that are available for the **`json`** subcommand. Since you are specifying a format and format library that is different from those that were specified in the JSON template, you will use the **`format`** and **`server`** options on the command line.

2 Create the ynmA format in the userformats2 library using the JSON template. Specify the ynmA format and userformats2 format library on the command line with the **`format`** and **`format-library`** options. These options override the options that are specified in the JSON template.

3 Display information about the format to verify that the userformats2 library and ynmA name from the command line overwrote the values in the template.

For more information, see “Create User-Defined Formats from JSON Templates ” on page 693.

Example: Create a user-defined format and specify the ranges from a CSV template. For more information, see “Source Files: Data” on page 697. This example assumes that you already have a CSV template file, as follows:

```
F,female
M,male
f,female
m,male
```

```
1 sas-admin cas format-libraries add-format help csv
```

```
2 sas-admin cas format-libraries add-format csv --server serverA --format '$gender'
 --format-library userformats1 --source-file /path-to-template/gender.csv
```

```
3 sas-admin cas format-libraries show-format-ranges --server serverA --format '$gender'
 --format-library userformats1
```

1 Review the options that are available for the **`csv`** subcommand.

2 Create the gender format in the userformats1 library using the CSV template to specify the ranges. Since the gender format is a character format, you must precede it with a dollar sign ($) and enclose it in quotation marks. The other command options must be specified on the command line.

3 Display information about the format.

Example: Create the same user-defined format for different locales. Use the **show-format-ranges** option with the formats that you created. This example uses the CSV file from the previous example.

```
1 sas-admin cas format-libraries add-format csv --server serverA --format '$gender'
 --format-library userformats1 --source-file /path-to-template/gender.csv

2 sas-admin cas format-libraries add-format csv --server serverA --format '$gender'
 --format-library userformats1 --source-file /path-to-template/gender.csv --format-locale ja_JP

3 sas-admin cas format-libraries add-format csv --server serverA --format '$gender'
 --format-library userformats1 --source-file /path-to-template/gender.csv --format-locale ko_KR

4 sas-admin cas format-libraries show-formats --server serverA --format-library userformats1

5 sas-admin cas format-libraries show-format-ranges--server serverA
--format-library userformats1 --format '$gender'

6 sas-admin cas format-libraries show-format-ranges --server serverA
--format-library userformats1 --format '$gender' --ignore-locale

7 sas-admin cas format-libraries show-format-ranges --server serverA
--format-library userformats1 --format 'ja_JP-$gender'
```

1 Create the gender format in the userformats1 library with no specified locale. Since the gender format is a character format, you must precede it with a dollar sign ($). Since no locale is specified, the format will be created with no locale.

2 Create another gender format in the userformats1 library, but specify the Japanese locale.

3 Create another gender format in the userformats1 library, but specify the Korean locale.

4 List the SAS formats in the userformats1 library to verify that the formats were created. You should see the following formats: $gender, ja_jp-$gender, and ko_kr-$gender.

5 Display information about the $gender format that was created with no specified locale. The locale defaults to the operating system of the CAS server. If the CAS server is physically located in Japan and the default locale has not been changed, the locale will default to Japanese. If the CAS server is physically located in Korea and the default locale has not been changed, the locale will default to Korean. If the CAS server is physically located in the United States and the default locale has not been changed, the locale will default to English.

6 Display information about the $gender format that was created with no specified locale, but specify the **ignore-locale** option. You will see that no locale is displayed.

7 Display information about the ja_JP-$gender format that was created with the Japanese locale. The Japanese locale is displayed, and takes precedence over the locale of the operating system of the CAS server.

Example: Create a user-defined format. Then save the format library that contains the new format as a SASHDAT file (table) in the data source that is associated with the active caslib. You can load a format library from the table in future sessions. This example assumes that you already have a CSV template file, as follows:

```
low - <2.7,"Very economical"
2.7 - <4.1,"Small"
4.1 - <5.5,"Medium"
5.5 - <6.9,"Large"
6.9 - high,"Very large"
```

```
1 sas-admin cas caslibs create path --name caslibA --server serverA --path /path-to-caslib
```

```
2 sas-admin cas format-libraries add-format csv --format '$enginesize' --format-library Casformats
 --server serverA --source-file /path-to-source-file/enginesize.csv

3 sas-admin cas format-libraries export --caslib caslibA --format-library Casformats
 --server serverA --table enginefmt --su

4 cas tables list --caslib caslibA
```

1 Create a path caslib called caslibA to contain the new format. The SASHDAT file that is created will be placed in the location that you specify in the **`--path`** option.

2 Create a new character format called enginesize by importing its ranges from a CSV file. Add the new format to the Casformats library.

3 Export the Casformats library that contains the enginesize format to the table enginefmt in caslibA. The **`--table`** option specifies the name of the table and the results in a file named enginefmt.sashdat. The file is placed in the location that you specified in the **`--path`** option when you created caslibA. You must specify the superuser option in commands where elevated privileges are needed.

4 List the tables in caslibA to verify that you see the enginefmt table.

Details and Tips

- You can create a SAS format library from an existing SAS catalog, CAS table, or item store file. The extension of the file that you specify in the **`source-path`** option of the **`cas format-libraries create`** command identifies the source of the SAS format library as follows:

Extension	Source of the Format Library
.sas7bcat	SAS catalog
.sashdat	CAS table
no extension or any other extension	CAS item store file

- If you run the show-format-ranges command for a format, the results are displayed according to these rules:
  - If you request information about a non-locale format (such as **`$charfmt`**) and you do not specify the **`ignore-locale`** option, the default locale format will be returned. If the default system locale is en_US and the **`ignore-locale`** option was not specified, the returned format for **`$charfmt`** is en_US-$charfmt.
  - If you request information about a non-locale format (such as **`$charfmt`**), you must specify the **`ignore-locale`** option in order to return the actual format with no locale. If the **`ignore-locale`** option was specified, the returned format for **`$charfmt`** is $charfmt.
  - If you request information about a format that has a non-existent locale, the format for the default system locale is returned. Suppose that you request information about en_bogus-$charfmt, which is a format that does not exist. If the default system locale is en_US, the returned format is **`en_US-$charfmt`**.
  - If you request information about an invalid format, an error is returned. Suppose that you request information about en-US-$charfmt. The hyphen in the locale name (en-US) should have been an underscore (en_US). No format is returned, and an error is displayed.
- When you create a format library, the search order defaults to none. If you want to specify how the new format library appears in the SAS format search order, you must do the following:
  - specify one of the following values for the **`search-order`** option in the **`format-libraries create`** command: append, prepend, or replace.

  - specify the **`superuser`** option in the **`format-libraries create`** command.

See Also

See “Overview” on page 142 for more information about user-defined formats.

CLI Examples: Identities

The following examples assume that you have already signed in to SAS Viya at the command line. See “Command-Line Interface: Preliminary Instructions” on page 681.

Examples

Example: Add user1 to the group that has the ID 4444.

```
sas-admin identities add-member --user-member-id  user1 --group-id 4444
```

Example: Create a group with the name Salesgroup, the group ID 8888, and the description of “Custom sales group”.

```
sas-admin identities create-group --id 8888 --name Salesgroup --description "Custom sales group"
```

Example: Remove user1 from the sales and marketing groups.

```
1 sas-admin identities list-memberships --user-id user1

2 sas-admin identities remove-member --group-id sales-group --user-member-id user1

3 sas-admin identities remove-member --group-id marketing-group --user-member-id user1
```

1. Verify the groups that user1 belongs to.
2. Remove user1 from the group that has the group ID sales-group.
3. Remove user1 from the group that has the group ID marketing-group.

Example: Show details about the group that has the group ID ABC.

```
sas-admin identities show-group --id ABC
```

Details

- If a group is created with no name, the specified ID will be used for the name.
- The following identities commands list only 50 items at a time by default:
  - list-groups
  - list-members
  - list-memberships

  To list more than 50 items, you can use the --`limit` option.

  Example:

  ```
  sas-admin identities list-members --group-id ABCD --limit 100
  ```

See Also

- “Command-Line Interface: Overview” on page 678
- “Identity Management: Overview” on page 349

CLI Examples: Job

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: List the job flows.

```
sas-admin job flows list
```

Example: Generate a template for a flow and write the output to a file.

```
sas-admin job flows generate-template --file-out /tmp/output.txt
```

Example: List job flow scheduling service objects.

```
sas-admin job schedulers list
```

Details

Here is an example of a template file that is generated from a flow that you created:

```
{
    "name": "Replace with name of the flow",
    "description": "(Optional) Replace with description of the flow",
    "triggerType": "Replace with trigger type, select one of: runnow, manual, event",
    "triggerCondition": "Replace with either any or all",
    "flowProperties": {},
    "defaultJobProperties": {}
}
```

See Also

"Command-Line Interface: Overview" on page 678

CLI Examples: Launcher

The following examples assume that you have already signed in to SAS Viya using the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: List all launcher contexts.

```
sas-admin launcher contexts list --all
```

Example: Show detailed information about the SAS Launcher context with the specified ID.

```
sas-admin launcher contexts show --id context-id
```

Example: List the SAS Launcher options set mappings.

```
sas-admin launcher options-set-mappings list --all
```

Example: Show detailed information about the SAS Launcher options that are set with the specified ID.

```
sas-admin launcher options-sets show --id options-set-ID
```

See Also

- "Command-Line Interface: Overview" on page 678
- "SAS Launcher Server and Launcher Service" on page 566

CLI Examples: Licensing

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: List site information, which is the site name, the site number, the operating system name, the release number, the server date, the grace period, and the warning period.

```
sas-admin licenses site-info list
```

Example: List all products in the system whose name contains "Visual Analytics".

```
sas-admin licenses products list --name-contains "Visual Analytics"
```

Example: List all products in the system that are expired.

```
sas-admin licenses products list --expired
```

Example: List the products in the system that have these identifiers: 827, 921, and 985.

```
sas-admin licenses products list --product-ids 827,921,985
```

Example: List the number of products with a current license that are deployed.

```
sas-admin licenses count --current
```

Example: List the Data-Connector products whose licenses are covered within the grace period.

```
sas-admin licenses data-connectors list --grace
```

See Also

- "Command-Line Interface: Overview" on page 678
- SAS Licensing on page 181

CLI Examples:Quality Knowledge Bases (QKBs)

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Examples

Example: List the details about all environments.

```
sas-admin qkbs environments list --all
```

Example: List all the QKBs that are registered in the system.

```
sas-admin qkbs qkbs list --all
```

Example: Delete a specified QKB.

```
sas-admin qkbs qkbs delete --context serverA --environment CAS --qkb qkbName
```

Example: Delete a specified import job.

```
sas-admin qkbs jobs delete --id QKB-job-ID
```

Example: List the contexts of a particular environment type. Currently, **CAS** is the only supported environment type.

```
1 sas-admin qkbs contexts help list

2 sas-admin qkbs contexts list --environment CAS -all

3 sas-admin qkbs contexts list --environment CAS --superuser-role TRUE --state running
```

1 Display the options that you can use to filter the contexts list.

2 Notice the values that are returned for type, state, and Superuser role. These are examples of valid values that you can use for these options when filtering the search results. For example, you will see examples of valid values for **type** and **state**. You will also see that the valid values for **Superuser role** are **TRUE** or **FALSE**.

3 Search for the QKBs using appropriate values for type, state, and Superuser role.

Example: Submit a job to import a QKB to the CAS server. Enter a separate command to check the status of the job to determine whether it ran successfully.

```
1 sas-admin qkbs qkbs import --context serverA --environment CAS --destination qkbname
   --source /path-to-qarc/demoqkb.qarc

2 sas-admin qkbs jobs show-info --id QKB-job-ID
```

1 Submit a job to import a QKB to a CAS environment using the serverA context. Specify the path and filename of the source QKB archive file and the name of the QKB after the import is complete. When the job is successfully submitted, you will receive a message that includes the ID of the import job. Running the command in this manner submits the job, but does not determine whether the job ran successfully.

2 Display detailed information about the import job that you just ran to determine whether it ran successfully. You must specify the ID that you obtained in the previous step.

Example: Submit a job to import a QKB to the CAS server. Use the **poll** option to show whether the QKB was successfully registered and whether the import job ran successfully.

```
sas-admin qkbs qkbs import --context serverA --environment CAS --destination qkbname
 --source /path-to-qarc/demoqkb.qarc --poll
```

Example: Check the status of a long running QKB job from another Linux shell. This example assumes that you have the ID of the import job for which you want to check the status. The ID should have been returned after you submitted the job.

```
1 ssh userID@server

2 sas-admin qkbs jobs show-info --id QKB-job-ID
```

1 Open a new shell or terminal on the CAS server that your QKB job is running on.

2 Display information about the import job that you just ran. You must specify the ID of the import job that you received when you submitted the job.

Details

- Due to the size of the imported QKBs, you cannot use the **verbose** option. If you attempt to use the **verbose** option, an error results.
- Because QKB files can be large, the processing time for an import job might be lengthy. To check the status of import jobs, you can open another shell and run the **qkbs jobs show-info** command.
- If an import job fails while it is being registered, you must repeat the import process.

- To import or delete a QKB, you must be logged in as a user who has the ability to assume the Superuser role for the CAS server where the QKB is being imported to or deleted from. There is no separate **superuser** or **su** option.

See Also

- "Command-Line Interface: Overview" on page 678
- SAS Viya Administration: QKB Management

CLI Examples: Reports

The following examples assume that you have already signed in to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Single Command Examples

Example: Show information about the report that has the ID a85235e7-fad1-4f8a-9ad9-ea0d576619e1.

```
sas-admin reports show-info  --id a85235e7-fad1-4f8a-9ad9-ea0d576619e1
```

Example: List the detailed output of all reports that were created after 2017-05-23.

```
sas-admin reports list --created-after 2017-05-23 --details
```

Example: List the reports in the SAS Viya system that were modified by user1.

```
sas-admin reports list --modified-by user1
```

Example: List a maximum of 50 reports that are sorted by ID and in descending order.

```
sas-admin reports list --details --sort-by ~id --limit 50
```

Example: Delete the report that has the ID a85235e7-fad1-4f8a-9ad9-ea0d576619e1.

```
sas-admin reports delete --id a85235e7-fad1-4f8a-9ad9-ea0d576619e1
```

Multiple Command Examples

Note: Some of these tasks must be performed by a user with administrative privileges.

Example: Update the theme that is used by a SAS Visual Analytics report. This example assumes that you already have a SAS Visual Analytics report that is using the Marine theme (the default).

```
1 sas-admin reports themes list

2 sas-admin reports themes show-info --theme-id theme_ID

3 sas-admin reports list --name reportA

4 sas-admin reports themes update --report-id report_ID --theme-id theme_ID
```

1 List the themes that are available for SAS Visual Analytics reports, and record the ID of the theme that you want to use.

2 Show detailed information about the identified theme for updating the report. You use the ID of the theme that was identified in the previous step.

3 List information about reportA that you want to update, and record the ID of the report.

4 Update the theme that is used in reportA to the theme that you identified in step 1. You use the ID of the report that you identified in the previous step.

Note: You can update a report's theme only to a theme of type **system** or **custom**. Themes of types **legacy** and **retired** are allowed in existing reports, but cannot be used to update a theme.

Example: Export the translation worksheets for reports that need to be translated, translate the worksheets, and then import the translated worksheets into the report. Some of these tasks must be performed by a user with administrative privileges whereas the actual translation might be performed by another user.

You need to determine what reports need to be translated and what languages that the reports need to be translated into. You also need to identify the base language that the reports were created in. The translation worksheets that are exported will contain strings to translate, and these strings are written in the base language that the report was created in. You must save the translation worksheets using the UTF-8 character encoding.

For this example, you learn that you need to translate the reports that were created after 01MAY2018 and whose names contain the text “VAN”. You also learn that the reports need to be translated into French, Spanish, and Japanese. You determine that the base language that the reports were written in is English. See “Details” on page 726 for additional information about translation worksheets.

```
1 sas-admin reports list --created-after 2018-05-01 --limit "100" --name-contains VAN

2 sas-admin reports translations export --report-id report_ID
--output-location /output-path-for-worksheet --report-locale fr-FR

3 sas-admin reports translations export --report-id report_ID
--output-location /output-path-for-worksheet --report-locale es-ES

4 sas-admin reports translations export --report-id report_ID
--output-location /output-path-for-worksheet --report-locale ja-JP

5 ######## LOCALE FOR THIS TRANSLATION WORKSHEET ########
fr-FR
############# BEGIN TRANSLATABLE STRINGS ##############
property.label = translated value
property1.label = translated value
property2.label = translated value

6 sas-admin reports translations import --source-file /output-path-for-worksheet/report-name_fr-FR.reports

7 sas-admin reports translations import --source-file /output-path-for-worksheet/report-name_es-ES.reports

8 sas-admin reports translations import --source-file /output-path-for-worksheet/report-name_ja-JP.reports
```

1 List the reports that were created after 01MAY2018 and whose names contain the text “VAN”. Use the **`limit`** option to specify the maximum number of reports to list. The default value is 20. To make sure that you list all the reports that were created after 01MAY2018, specify a value of 100. If you receive a message at the end of the report list that indicates that more reports are available, list the reports again with a higher value for the **`limit`** option.

After running the code to list the reports, you find that there are four reports that you need to translate. In other words, there are four reports that were created after 01MAY2018 and whose names contain the text “VAN”. Note the report ID of each report. You will need the report ID to export the translation worksheets.

If you do not have administrative privileges, you might be able to obtain the report ID from SAS Environment Manager.

2 Export the translation worksheet for the first of the four reports in French (fr-FR),

3 Export the translation worksheet for the first of the four reports in Spanish (es-ES).

4 Export the translation worksheet for the first of the four reports in Japanese (ja-JP).

Repeats steps 2, 3, and 4 for the second, third, and fourth reports. Be sure to update the report ID of the report in the command when you move to the second, third, and fourth report. You will issue the command a total of 12 times in this example. (There are four reports, and each report is translated into three languages). Afterward, you should have 12 translation worksheets in the location that is specified in the output location option.

Note: If another user is performing the translation, then provide them with the translation worksheets. They can proceed with step 5.

5 Here is an example of the lines that exist in each translation worksheet. The line following **`LOCALE FOR THIS TRANSLATION WORKSHEET`** indicates the language of the translation worksheet.

In the French translation worksheets (*report_name*_fr-FR), locate the line that contains **`BEGIN TRANSLATABLE STRINGS`**. In the following lines, edit the values to the right of the equal sign (=) and provide the appropriate values, in French. Save the file.

In the Spanish translation worksheets (*report_name*_es-ES), locate the line that contains **`BEGIN TRANSLATABLE STRINGS`**. In the following lines, edit the values to the right of the equal sign (=) and provide the appropriate values, in Spanish. Save the file.

In the Japanese translation worksheets (*report_name*_ja-JP), locate the line that contains **`BEGIN TRANSLATABLE STRINGS`**. In the following lines, edit the values to the right of the equal sign (=) and provide the appropriate values, in Japanese. Save the file.

Note: If the translation was performed by another user, then this user must provide the translated worksheets to the user who is performing the import.

6 Import the French translation worksheet (*report_name*_fr-FR).

7 Import the Spanish translation worksheet (*report_name*_es-ES).

8 Import the Japanese translation worksheet (*report_name*_ja-JP).

Details

- The **`translations export`** command appends the locale value that is specified as the report locale to the name of the translation report. For example, if you export a translation worksheet for reportA and specify the US English locale as the value of the **`report-locale`** option, the translation worksheet will be saved in a file with the name:

  ```
  reportA_en-US.reports
  ```

  The **`translations export`** command includes a locale section in the translation worksheet that includes the locale name that you specified as the report locale. Here is an example of the locale section for the locale en-US:

  ```
  ######## LOCALE FOR THIS TRANSLATION WORKSHEET ########
  en-US
  ```

  Note: Before importing a translation worksheet, you must make sure that the language in the locale section of the worksheet matches the language in the name of the translation worksheet. The name of the file that you are importing must be appended with a locale value.

- The character encoding of the files that contain the translation worksheets must be UTF-8.
- The **`clear-results-cache`** option clears the results cache of report data for SAS Visual Analytics reports in SAS Viya. In multi-tenancy environments, this command can be run only on the provider tenant, but it clears the results cache across all tenants.
- The reports CLI lists only 20 reports at a time by default. To list more than 20 reports, you can use the --`limit` option. To list 50 reports, enter --`limit 50`.
- To specify what report number to start the list with, you can use the --`start` option. Suppose that you have listed the first 20 reports, and you want to list the next 20 reports, starting with report number 21, enter --`start 21`.

See Also

“Command-Line Interface: Overview” on page 678

CLI Examples: Restore

The following examples assume that you have already signed in to SAS Viya at the command line. See “Command-Line Interface: Preliminary Instructions” on page 681.

Examples

Example: Show the history of restore operations.

```
sas-admin restore list
```

Example: Start a restore operation of a specified backup, and specify that the name of the restore operation is restoreA.

```
sas-admin restore start --backup-name backup-ID --slug restoreA
```

See Also

- “Overview” on page 34
- “Command-Line Interface: Overview” on page 678

CLI Examples: Scoreexecution

The following examples assume that you have already signed in to SAS Viya using the command line. See “Command-Line Interface: Preliminary Instructions” on page 681.

Examples

Example: Display detailed information about the unused resources.

```
sas-admin scoreexecution --output fulljson list-hanging-resources
```

Example: Display information about unused resources in a table format.

```
sas-admin scoreexecution --output text list-hanging-resources
```

Example: Display only the URIs of the unused resources.

```
sas-admin scoreexecution --output text --quiet list-hanging-resources
```

Example: Write the URIs of the unused resources to a file named uris.txt.

```
sas-admin scoreexecution list-hanging-resources --file uris.txt
```

Example: Delete the unused resources listed in the file uris.txt.

```
sas-admin scoreexecution remove-hanging-resources --file uris.txt
```

See Also

- “Command-Line Interface: Overview” on page 678
- SAS Decision Manager: User’s Guide

CLI Examples: Tenant Administration

The following examples assume that you have already signed in to SAS Viya at the command line. See “Command-Line Interface: Preliminary Instructions” on page 681.

Examples

Note: The **create**, **delete**, **onboard**, and **offboard** commands are currently reserved for use only in Ansible playbooks. Do not use these commands at an operating system prompt.

Here are typical examples of commands that can be used with the CLI:

Example: List the tenants.

```
sas-admin tenant list
```

Example: Show the properties for the specified tenant.

```
sas-admin tenant show --id tenantID
```

Example: Check the health of the specified tenant.

```
sas-admin tenant health-check --id tenantID
```

Example: Enable the specified tenant.

```
sas-admin tenant enable --id tenantID
```

Example: Disable the specified tenant.

```
sas-admin tenant disable --id tenantID
```

Details

- The Help for the ID of a tenant states that the string must match this pattern: ^[a-z]+[a-z0-9]*

  This pattern means that the ID of a tenant must start with a lowercase letter, followed by any number of lowercase letters. Use of numbers is optional.

- When you enable a tenant so that users can sign in to it, the access policy of the tenant is changed from **providerTenantUsersOnly** to **allUsers**.

  When you disable a tenant so that users can no longer sign in to it, the access policy of the tenant is changed from **allUsers** to **providerTenantUsersOnly**.

  For more information about the access policy of a tenant, see “Edit the Properties of a Tenant” in *SAS Viya Administration: Multi-tenancy*.

See Also

- “Command-Line Interface: Overview” on page 678
- SAS Viya Administration: Multi-tenancy

CLI Examples: Transfer

The following examples assume that you have already signed in to SAS Viya at the command line. See “Command-Line Interface: Preliminary Instructions” on page 681.

Examples

Example: See the commands and subcommands that are available for the transfer import command.

```
sas-admin transfer import --help
```

Example: Export a new transfer package from an object that has the resource URI "/reports/reports/faa7f5f2-0822-4ca0-9f92-23bda3e02738", and name the package “Export Report”.

```
sas-admin transfer export --name "Export Report"
```

```
--resource-uri "/reports/reports/faa7f5f2-0822-4ca0-9f92-23bda3e02738"
```

Example: Get the mapping information for the transfer package that is named “Export Report” from the previous example.

```
1 sas-admin transfer list --name "Export Report"

2 sas-admin transfer get-mapping --id transfer-package-ID
```

1 Locate the ID for the transfer package that you want to export.

2 Issue the command to retrieve the mapping information for the “Export Report” transfer package that has the specified ID.

Example: Upload a package source environment to the target environment, and write the mapping file to the specified location.

```
1 sas-admin --profile source transfer download --id transfer-package-ID --file /tmp/MyPackage.json

2 sas-admin --profile target transfer upload --file /tmp/MyPackage.json --mapping /tmp/map.txt
```

1 Download the package to your local machine and store it in a package file that is named MyPackage.json.

2 Upload the MyPackage.json file to the target environment, and write the mapping file to the file map.txt.

Details

You can specify information about the export or import operation that you want to perform using the Transfer service REST API standards, as follows:

- export

  You can specify information about the export operation that you want to perform using the request option of the transfer export command. The option accepts JSON input of type ExportRequest. The content of the input can be contained in a quoted string or in a file. The filename must begin with the at sign (@). The filename can be specified in either of two forms:

  ```
  @filename.txt

  @/path/.filename.txt
  ```

  The content of this option makes up the POST request through which the export package is sent.

  Table A.6 HTTP POST Export Request Members

Name	Type	Description
version	integer	The schema version number of the JSON media type. This is version 1.
name	string	The name of the export job that is used to export objects from a source system to a transfer package. This is also the name of the transfer package that is being created.
description	string	A short description of the export job.
items	list	The list of URIs to include in the transfer package.

Here is a sample JSON file of type ExportRequest. The name of the file is export.json.

```
{
  "version": 1,
  "name": "My reports",
```

```
  "description": "Export of all my reports",
  "items": [
    "/reports/reports/4d083692-3c9a-4f2c-945f-e96fad972036",
    "/folders/folders/d4f1533a-229d-4d45-a5c9-b8a21fbc1e39"
  ]
}
```

You can use either of the following syntax options of the command to export the information:

- `sas-admin transfer export --request @/path-to-file/export.json`

  Note: When you include the information in a file, the first character following the request option must be the at sign (@). Therefore, if a pathname is used, it must start with the at sign (@).

- UNIX: `sas-admin transfer export --request '{ "version": 1, "name": "My reports", "description": "Export of all my reports", "items": [ "/reports/reports/4d083692-3c9a-4f2c-945f-e96fad972036", "/folders/folders/d4f1533a-229d-4d45-a5c9-b8a21fbc1e39" ] }'`

  Windows: `sas-admin transfer export --request "{ \"version\": 1, \"name\": \"My reports\", \"description\": \"Export of all my reports\", \"items\": [ \"/reports/reports/4d083692-3c9a-4f2c-945f-e96fad972036\", \"/folders/folders/d4f1533a-229d-4d45-a5c9-b8a21fbc1e39\" ] }"`

  Note: When running the transfer CLI in a Windows environment, single quotation marks are not allowed. You must enclose the **`request`** option data in double quotation marks, and then escape the embedded double quotation marks with a backslash (\).

- import

  You can specify information about the import operation that you want to perform using the request option of the transfer import command. The option accepts JSON input of type ImportRequest. The content can be contained in a quoted string or in a file. The filename must begin with an at sign (@). The filename can be specified in either of two forms:

  @*filename*.txt

  @/*path*/.*filename*.txt

  The content of this option makes up the POST request through which the import package is sent.

  Table A.7 HTTP POST Import Request Members

Name	Type	Description
version	integer	The schema version number of the JSON media type. This is version 1.
name	string	The name of the import job that is used to import objects from a transfer package to a target system.
description	string	A short description of the import job.
packageUri	string	The package to import.

Here is an example of a JSON file of type ImportRequest that is named import.json:

```
{
      "version": 1,
      "name": "My reports",
      "description" : "import all my reports ",
        "packageUri": "/transfer/packages/a2ef940e-14ac-4960-9a9a-7689702b06f0"
    }
```

You can use either of the following syntax options of the command to import the information:

- □ `sas-admin transfer import --request @/path-to-file/import.json`

  Note: When you include the information in a file, the first character following the request option must be the at sign (@). Therefore, if a pathname is used, it must start with the at sign (@).

- □ UNIX:`sas-admin transfer import --request '{ "version": 1, "name": "My reports", "description" : "import all my reports ", "packageUri": "/transfer/packages/a2ef940e-14ac-4960-9a9a-7689702b06f0" }'`

  Windows:`sas-admin transfer import --request "{ \"version\": 1, \"name\": \"My reports\", \"description\" : \"import all my reports \", \"packageUri\": \"/transfer/packages/a2ef940e-14ac-4960-9a9a-7689702b06f0\" }"`

  Note: When running the transfer CLI in a Windows environment, single quotation marks are not allowed. You must enclose the **`request`** option data in double quotation marks, and then escape the embedded double quotation marks with a backslash (\).

See Also

- "Promotion: How to Import (Command-Line Interface)" in *SAS Viya Administration: Promotion (Import and Export)*
- "Command-Line Interface: Overview" on page 678

Command-Line Interface: Troubleshooting

Message: `token expired and refresh token is not set`

Explanation: You are not currently authenticated to SAS Viya at the command line. See "Command-Line Interface: Preliminary Instructions" on page 681.

Message: `flag provided but not defined`

Explanation: You might have specified a global option in the wrong location. See "Command-Line Interface: Syntax" on page 686.