



THE  
POWER  
TO KNOW.

# **SAS<sup>®</sup> Federation Server 4.1**

## **Administrator's Guide**

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2014. *SAS® Federation Server 4.1: Administrator's Guide*. Cary, NC: SAS Institute Inc.

**SAS® Federation Server 4.1: Administrator's Guide**

Copyright © 2014, SAS Institute Inc., Cary, NC, USA

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

**For a hard-copy book:** No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

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

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

**U.S. Government Restricted Rights Notice:** Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions set forth in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

May 2014

SAS provides a complete selection of books and electronic products to help customers use SAS® software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit [support.sas.com/bookstore](http://support.sas.com/bookstore) or call 1-800-727-3228.

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

---

# Contents

<i>What's New in SAS Federation Server 4.1</i> . . . . .	vii
<b>Chapter 1 • Overview</b> . . . . .	<b>1</b>
Introduction . . . . .	1
Services Provided by SAS Federation Server . . . . .	3
Components of SAS Federation Server . . . . .	5
Supported Data Sources . . . . .	6
<b>Chapter 2 • Upgrading SAS Federation Server</b> . . . . .	<b>9</b>
Overview . . . . .	9
Cross-Platform Considerations . . . . .	9
Migrating to SAS Federation Server 4.1 . . . . .	10
Promoting Content to SAS Federation Server 4.1 . . . . .	11
ODBC Drivers . . . . .	12
<b>Chapter 3 • Configuring SAS Federation Server</b> . . . . .	<b>13</b>
Overview . . . . .	14
Configure Temporary Storage for SAS Utility Files . . . . .	14
Configuring the Windows Environment . . . . .	15
Configuring the UNIX Environment . . . . .	17
Configuring Server Accounts . . . . .	26
SAS Federation Server Configuration Reference . . . . .	32
Configuring SAS Federation Server Manager . . . . .	41
<b>Chapter 4 • SAS Federation Server Administration</b> . . . . .	<b>45</b>
Overview . . . . .	46
SQL Scripting for SAS Federation Server Administration . . . . .	46
The SAS Federation Server Database . . . . .	50
SAS Federation Server Resource Cache . . . . .	51
Connection Pooling . . . . .	54
SQL Logging . . . . .	56
Server Logging Configuration . . . . .	65
Utilities for SAS Federation Server . . . . .	70
Uninstalling SAS Federation Server . . . . .	72
<b>Chapter 5 • SAS Federation Server Security</b> . . . . .	<b>75</b>
Overview . . . . .	76
Authentication . . . . .	76
Authorization . . . . .	76
Permissions . . . . .	78
Object Privilege Inheritance . . . . .	81
Row-Level Security . . . . .	86
Data Masking . . . . .	92
Server Encryption . . . . .	96
<b>Chapter 6 • Working with Federated Data</b> . . . . .	<b>99</b>
Overview of Data Federation . . . . .	99
Configuring Data Source Access . . . . .	100
Federated SQL Views . . . . .	107

Data Caching .....	112
Memory Data Store .....	117
Understanding Data Federation and Best Practices .....	119
<b>Chapter 7 • Driver Reference for SAS Federation Server .....</b>	<b>123</b>
Database Functionality and Driver Performance .....	124
DB2 Driver Reference .....	124
FedSQL Driver Reference .....	129
Greenplum Driver Reference .....	132
MDS Driver Reference .....	136
Netezza Driver Reference .....	139
ODBC Driver Reference .....	143
Oracle Reference .....	151
PostgreSQL Driver Reference .....	156
SAP Reference .....	161
SAP HANA Driver Reference .....	175
SAS Data Set Reference .....	182
SASHDAT Driver Reference .....	186
Teradata Reference .....	189
<b>Appendix 1 • Administration DDL Statements Reference .....</b>	<b>195</b>
ALTER SERVER Statement .....	195
CREATE DATA SERVICE Statement .....	199
DROP DATA SERVICE Statement .....	202
ALTER DATA SERVICE Statement .....	203
CREATE CATALOG Statement .....	205
DROP CATALOG Statement .....	206
ALTER CATALOG Statement .....	207
CREATE SCHEMA Statement .....	207
DROP SCHEMA Statement .....	209
ALTER SCHEMA Statement .....	209
CREATE DSN Statement .....	211
DROP DSN Statement .....	214
ALTER DSN Statement .....	214
CREATE CACHE Statement .....	215
REFRESH CACHE Statement .....	217
ALTER CACHE Statement .....	217
DROP CACHE Statement .....	218
PURGE CACHE Statement .....	218
DROP AUTHID Statement .....	218
GENERIC OPTIONS Syntax .....	219
ALTER GENERIC OPTIONS Syntax .....	219
GRANT and DENY Statements .....	220
REVOKE Statement .....	221
<b>Appendix 2 • Information Views .....</b>	<b>223</b>
Visibility Rules for Information Views .....	224
AUTHORIZATION_IDENTIFIERS View .....	226
CACHES View .....	227
CATALOGS View .....	228
CATALOG_PRIVILEGES View .....	228
COLUMNS View .....	229
COLUMN_PRIVILEGES View .....	230
CONFIG_CATALOGS View .....	230
CONFIG_DATA_SERVICES View .....	231
CONFIG_DSNS View .....	231

CONFIG_OBJECTS View .....	232
CONFIG_SCHEMAS View .....	233
DATA_SERVICES View .....	233
DATA_SOURCE_NAMES View .....	234
DS_PRIVILEGES View .....	235
DSN_CONTENT View .....	236
DSN_LINEAGE View .....	237
DSN_PRIVILEGES and EFFECTIVE_DSN_PRIVILEGES Views .....	237
IDENTITY View .....	239
MESSAGES View .....	239
OBJECTS View .....	240
OBJECT_PRIVILEGES View .....	240
PRIVILEGES and EFFECTIVE_PRIVILEGES Views .....	241
SCHEMAS View .....	243
SCHEMA_PRIVILEGES View .....	243
X_COLUMN_PRIVILEGES/X_EFFECTIVE_COLUMN_PRIVILEGES View ...	244
X_OBJECT_PRIVILEGES/X_EFFECTIVE_OBJECT_PRIVILEGES View .....	246
<b>Glossary .....</b>	<b>249</b>



# What's New in SAS Federation Server 4.1

---

## Overview

The following new features are available in SAS Federation Server 4.1:

- There are four new SAS Federation Server drivers:
  - SAS Federation Server Driver for SAP HANA
  - SAS Federation Server Driver for Netezza
  - SAS Federation Server Driver for PostgreSQL
  - SAS Federation Server SASHDAT Data Source Driver
- data masking function to protect sensitive data
- SQL scripting for SAS Federation Server administration
- API support for interacting with and managing SAS Federation Server using a REST interface

---

## New Drivers for SAS Federation Server

### ***SAS Federation Server Driver for SAP HANA***

The SAS Federation Server Driver for SAP HANA enables Read and Write access to SAP HANA data sources. The driver supports both native SQL and FedSQL dialects. The SAS Federation Server Driver for SAP HANA is also available as a data service in SAS Federation Server Manager. For additional information, see “[SAP HANA Driver Reference](#)” on page 175.

### ***SAS Federation Server Driver for Netezza***

The SAS Federation Server Driver for Netezza enables SAS Federation Server to create, read, update, and drop Netezza tables. The driver supports both native SQL and FedSQL dialects. The SAS Federation Server Driver for Netezza is also available as a data service in SAS Federation Server Manager. For additional information, see “[Netezza Driver Reference](#)” on page 139.

## **SAS Federation Server Driver for PostgreSQL**

The SAS Federation Server Driver for PostgreSQL enables SAS Federation Server to create, read, update, and drop PostgreSQL tables. The driver supports both native SQL and FedSQL dialects. The SAS Federation Server Driver for PostgreSQL is also available as a data service in SAS Federation Server Manager. For additional information, see [“PostgreSQL Driver Reference” on page 156](#).

## **SASHDAT Data Source Driver**

The SAS Federation Server Driver for SASHDAT is a write-only driver designed for use with Hadoop on SAS LASR Analytic Server. SAS LASR Analytic server integrates Hadoop by storing SAS data in the Hadoop Distributed File system (HDFS). Using the SAS Federation Server Driver for SASHDAT, you can create files in the SASHDAT format in HDFS. These files can later be quickly loaded into the SAS LASR Analytic server for processing. The SAS Federation Server Driver for SASHDAT is also available as a data service in SAS Federation Server Manager. For additional information, see [“SASHDAT Driver Reference” on page 186](#).

---

## **Data Masking**

SAS Federation Server delivers a new data masking function as a method of ‘de-identifying’ sensitive information within data sources. The purpose of data masking is to protect the original data by using a functional substitute in situations where the audience is not privileged to access the original data. The primary focus of data masking is to protect data that is classified as personally identifiable information (PII). For additional information, see [“Data Masking” on page 92](#).

---

## **SQL Scripting for SAS Federation Server**

SAS Federation Server provides SQL language scripting capabilities to handle administrative needs for start-up and shutdown events. SQL scripts, executed at start-up or shutdown, manage auditing or related event notifications. Administrators can also manage content of in-memory MDS tables to a persistent data store when the server is stopped, as shown in the [“Example SQL Script” on page 48](#).

---

## **REST Interface for SAS Federation Server**

SAS Federation Server delivers a new API that provides a simple method to access data on SAS Federation Server. The API is implemented as a REST interface that provides direct interaction with SAS Federation Server, including metadata and data queries and SQL submissions. The SQL can include SAS Federation Server administration DDL or queries, DML, and DDL for back-end data sources. With the REST interface, you can perform the following tasks:

- Discover servers and data sources that are accessible through SAS Federation Server.
- Submit ad hoc queries for execution.
- Operate on query results.

The API is supported through a Java mid-tier service that communicates with SAS Federation Server through the JDBC driver. This service is delivered as a single WAR file that is deployed as part of the SAS Federation Server Manager installation. The WAR file does not depend on the SAS middle tier or use SAS WIP services.

The API provides examples of all media formats used, with in-line examples for collection objects of known media types. The collection usage indicates the media type that is used. Templated resource values are based on names of resources within the data architecture of SAS Federation Server. Resources might change during operations as data is manipulated. (for example, dropping or adding a table). The API returns a data element by name as it exists at the time of a request, so this information should be considered as requests are processed. This information should also be considered when designing a consuming application. For additional details, see the documentation delivered with the REST interface.

**x** *What's New in SAS Federation Server 4.1*

# Chapter 1

## Overview

---

<b>Introduction</b> .....	<b>1</b>
About SAS Federation Server .....	1
<b>Services Provided by SAS Federation Server</b> .....	<b>3</b>
Data Access Technology .....	3
Threaded Services .....	3
Multi-User Services .....	3
Performance .....	4
Data Storage Support .....	4
Standards-Based Interface for SQL .....	5
Security .....	5
<b>Components of SAS Federation Server</b> .....	<b>5</b>
Introduction .....	5
Federation Server Drivers .....	5
Language Driver .....	5
FedSQL .....	6
<b>Supported Data Sources</b> .....	<b>6</b>
Overview .....	6
SAS Data Set .....	7
SAP .....	7
Third-Party Relational Databases .....	7

---

## Introduction

### *About SAS Federation Server*

SAS Federation Server is a data server that provides scalable, threaded, multi-user, and standards-based data access technology in order to process and seamlessly integrate data from multiple data services. The server acts as a hub that provides clients with data by accessing, managing, and sharing SAS data as well as several popular relational databases.

SAS Federation Server enables powerful querying capabilities, as well as improved data source management. With SAS Federation Server, you can efficiently unite data from many sources, without moving or copying the data.

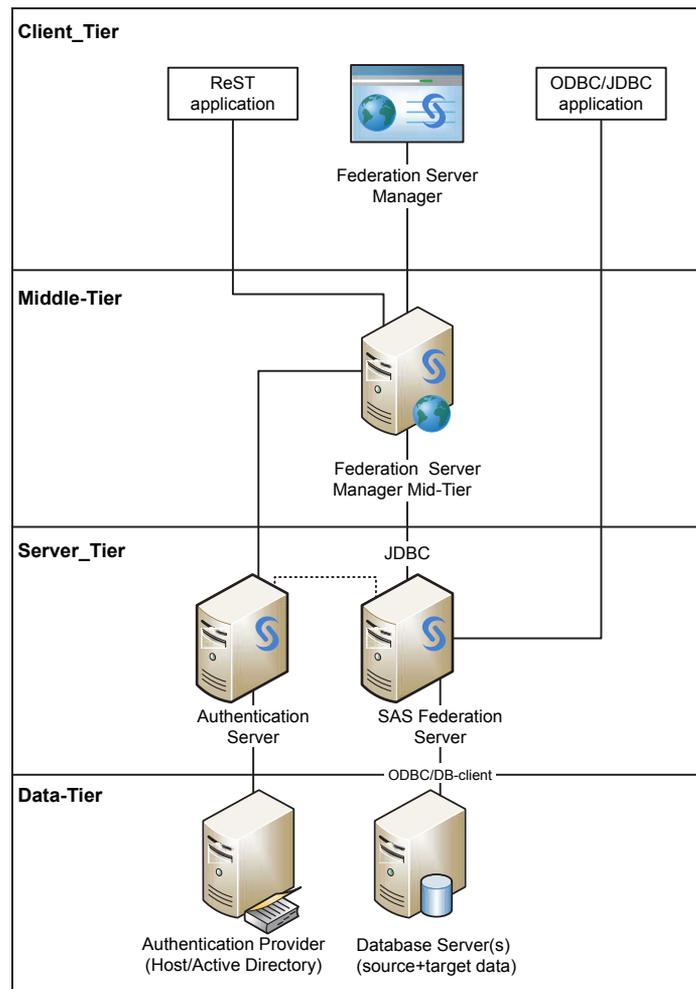
SAS Federation Server provides the following data capabilities:

- A central location for setup and maintenance of database connections.

- Access to popular database systems including DB2, Netezza, Oracle, SAP, SQL Server, PostgreSQL, Teradata, and Greenplum.
- ODBC and native drivers to connect to select data sources.
- Threaded data access technology that enhances enterprise intelligence and analytical processes.
- Multi-user services that enable multiple clients to access the same data concurrently.
- Ability to reference data from disparate data sources with a single query, known as data federation. also includes its own SQL syntax, Federated Query Language (FedSQL), to provide consistent functionality – independent of the underlying data source.
- A data abstraction layer, providing the ability to present a consistent data model throughout the organization. This is accomplished through the use of FedSQL views.
- Data access control with user permissions and data source security.

The following figure illustrates the architecture of SAS Federation Server:

**Figure 1.1** SAS Federation Server Architecture



---

## Services Provided by SAS Federation Server

### ***Data Access Technology***

The data access technology that is provided by SAS Federation Server consists of a set of run time components that provide a scalable, threaded, multi-user, and standards-based way to process and seamlessly integrate data from multiple data sources. The components provide the data access services that are required by business intelligence and analytical processes.

### ***Threaded Services***

Threads are an integral part of a high-performance, scalable system, and they are one of the main features of SAS Federation Server data access technology. Most threaded functionality can be further boosted in an environment in which multiple processors work in parallel. However, performance boosts can also be obtained with multi-threaded processes on a single processor machine.

A threaded service is a method of processing that divides a large job into several smaller jobs that can be executed in parallel. Threaded services control and execute requests by using multiple threads to increase data throughput. A thread is a single path of execution of a process in a single CPU. A thread can also be thought of as a basic unit of program execution in a thread-enabled operating environment. In a symmetric multiprocessing (SMP) environment, which uses multiple CPUs, multiple threads can be spawned and processed simultaneously. Regardless of whether there is one CPU or many, each thread is an independent flow of control that is scheduled by the operating system.

SAS® Federation Server provides threaded services that execute multiple user requests in parallel. Here are examples:

- Each connection to the SAS® Federation Server is managed on a separate thread. This enables multiple users to execute requests in parallel and reduces the probability of a user request being blocked while other user requests are processed.
- Complex requests (or large individual requests) are separated into units of work that are then executed in parallel. For example, filtering operations that require scanning large tables can be processed in parallel, and operations such as sorting can be processed by dividing the result set into subsets, sorting each subset in parallel, and then merging the sorted subsets into the final result set.
- Threading is also used to return result sets on multiple threads. For example, the FedSQL processor can request result sets from disparate data sources on separate threads. By reading data simultaneously, the FedSQL processor can acquire the data faster and expedite results to the client.

In addition to threaded services, some data services provide threaded I/O, which further enhance performance.

### ***Multi-User Services***

Multi-user services enable multiple clients to access the same data concurrently. If the data source supports this capability, SAS Federation Server enables two or more clients to write to the same table at the same time without destroying or losing updates. This process is referred to as concurrent update access.

SAS Federation Server uses Integrated Object Model (IOM) technology. IOM technology is a set of object-based interfaces to features or services. The technology enables application developers to use industry-standard programming languages, programming tools, and communication protocols to develop client programs that access these services on IOM servers.

A multi-user environment automatically ensures data protection during concurrent updates. The data services support concurrent updates by locking the data that is being updated and releasing the lock when updates are complete. This prevents loss of data or loss of updates that are due to simultaneous updates.

## **Performance**

SAS Federation Server integrates both user scalability and processing scalability to provide increased performance.

SAS Federation Server supports the following performance capabilities:

- In a multi-user environment, the server automatically scales to the number of concurrent users.
- The server provides rapid access to large amounts of data.
- The server is available under many 64-bit operating environments, which enables the server to scale in-memory processes.
- The server provides application-based, high-performance data reading by supporting a variety of cursor types, multi-row fetch capabilities, and positioned update of result sets.

## **Data Storage Support**

SAS Federation Server provides access to several types of data, which allows you to work with multiple data sources as if they were a single resource, regardless of where the information is stored. SAS Federation Server supports SAS® data sets, SAP®, and third-party relational databases that include:

- IBM® DB2, Netezza
- Oracle®
- Microsoft® SQL Server®
- Teradata®
- Greenplum®
- PostgreSQL®

By supporting several data sources, SAS Federation Server gives you the flexibility to configure data storage based on specific needs. You choose the type of data storage that is most appropriate for the particular needs of an application, based on functionality that is provided by each data source.

For more information about the supported data sources and the functionality that they provide, see the [SAS Federation Server Driver Reference on page 124](#).

## **Standards-Based Interface for SQL**

SAS Federation Server provides a standards-based interface for the SQL, which defines the data access model for the server. That is, an application creates, requests, and manipulates data by submitting SQL statements.

An application can submit SQL statements by using JDBC and ODBC drivers. The SQL is interpreted by the FedSQL processor, which supports a standard dialect across all back-end data sources. For more information, see [“Components of SAS Federation Server” on page 5](#).

## **Security**

SAS Federation Server security services ensure that both the server and its data are protected against unauthorized access. SAS Federation Server supports configurable authorization processes and other security features, including encryption. In addition, SAS Federation Server provides the ability to control access to SAS data sets that are placed under exclusive control of the server.

For more information, see [“Authentication” on page 76](#).

---

# **Components of SAS Federation Server**

## **Introduction**

SAS Federation Server consists of a set of components that provide the functionality that is required by data integration, business intelligence, and analytic processing. For example, SAS Federation Server provides several types of drivers and two interfaces that you use to connect to the server. The following topics describe each component.

## **Federation Server Drivers**

A SAS Federation Server driver interacts with a data source to read and write proprietary file formats. Each supported data source has a driver that communicates with the data service in its own language to resolve data access requests, and to manage physical files and database tables. For example, to process an SAP table, an application uses the Driver for SAP, and to process an Oracle table, an application uses the Driver for Oracle.

A Federation Server driver provides connectivity to and from the data source, submits SQL statements to the data source, and sends data to and from the data source. That is, a Federation Server driver receives an SQL expression as input and returns the result as output. Each Federation Server driver supports the database functionality of the underlying data source.

For more information about data sources and SAS Federation Server drivers, see [“Overview” on page 100](#).

## **Language Driver**

Language drivers implement the SAS® Federation Server languages by processing a request and sending the parsed query to the appropriate Federation Server driver that

satisfies the request and returns the result. The multi-threaded languages provide a powerful way to create and query data.

## FedSQL

FedSQL is the implementation of SQL that SAS Federation Server uses to access relational data. FedSQL is designed to be ANSI SQL:1999 core compliant with some extensions.

For applications, FedSQL provides a common SQL syntax across all data sources. That is, FedSQL is a vendor-neutral SQL dialect that accesses data from various data sources without requiring the application to submit queries in the SQL dialect that is native to the data source. In addition, a single FedSQL query can target data in several data sources and can return a single result set. When possible, FedSQL queries are optimized with multi-threaded algorithms to resolve large-scale operations.

The FedSQL language driver parses FedSQL requests, and then sends the parsed query to the appropriate Federation Server driver to determine whether the functionality can be handled by the data source.

For the SAS Federation Server Driver for Base SAS, the parsed FedSQL request is interpreted 'as is' because FedSQL is the native SQL dialect for the BASE data service.

For the relational database drivers, if the data source supports the functionality, then the FedSQL request is translated to the data source's native SQL dialect. By enabling as much functionality as possible to be handled by the data source, performance is enhanced. However, if a data source does not support the requested functionality, then the FedSQL language driver attempts to compensate by completing the balance of the operation.

For complete FedSQL statement reference, see the *SAS FedSQL Reference Guide*.

---

## Supported Data Sources

### Overview

SAS Federation Server supports disparate data sources by providing software in the form of Federation Server drivers, which access the physical data that an application processes. Federation Server drivers provide access to third-party relational databases, such as DB2®, Greenplum®, ODBC data sources, Oracle®, and Teradata®, by connecting to a remote server process.

SAS Federation Server supports the following data sources: SAS data set, SAP®, and several third-party relational databases.

The following table lists the data sources with their supported data service types:

Data Sources	Data Service Types	Default Drivers	Supported Drivers
SAS data sets	BASE	BASE	BASE
DB2	DB2UNXPC	DB2	ODBC, DB2
Greenplum	GREENPLUM	Greenplum	ODBC, Greenplum

Data Sources	Data Service Types	Default Drivers	Supported Drivers
Netezza	NETEZZA	Netezza	ODBC, Netezza
Oracle	ORACLE	ORACLE	ODBC, ORACLE
PostgreSQL	POSTGRESQL	PostgreSQL	ODBC, PostgreSQL
SAP	SAP	SAP	SAP
SAP HANA	SAPHANA	SAP HANA	ODBC, SAP HANA
SQL Server	SQLSERVER or QLSVR	ODBC	ODBC
Teradata	TERADATA	TERADATA	ODBC, TERADATA
ODBC_FED	Native Catalogs	ODBC	ODBC
ODBC	Logical Catalogs	ODBC	ODBC
Memory Data Store	MDS	MDS	MDS

### SAS Data Set

The SAS data set is the Base SAS proprietary file format for SAS software, which contains data values that are organized as a table of observations (rows) and variables (columns). The supported file format is the same as SAS data sets that are created by the BASE engine for Version 7 and later. A supported SAS data set uses the extension `.sas7bdat`.

The SAS Federation Server Driver for Base SAS provides Read and Update access to legacy SAS data sets. In addition, the driver creates SAS data sets that can be accessed by both SAS Federation Server and Base SAS software. The driver supports standard Base SAS storage functionality such as indexing, general integrity constraints, and SAS formats and informats. For more information about supported functionality and compatibility guidelines, See [“SAS Data Set Reference” on page 182](#) .

### SAP

For information about supported functionality and compatibility guidelines, See [“SAP Reference ” on page 161](#) .

### Third-Party Relational Databases

SAS Federation Server can access data in several third-party relational databases. The relational database drivers read, update, and create tables for those third-party relational databases on behalf of the Federation Server client. Each driver supports most of the FedSQL functionality. The Federation Server drivers support native database functionality by using the SQL dialect that is implemented by the third-party databases. For details about supported functionality and compatibility guidelines, see the specific data source reference:

- “DB2 Driver Reference” on page 124
- “Greenplum Driver Reference” on page 132
- “Netezza Driver Reference” on page 139
- “ODBC Driver Reference” on page 143
- “Oracle Reference” on page 151
- “PostgreSQL Driver Reference” on page 156
- “Teradata Reference” on page 189

## Chapter 2

# Upgrading SAS Federation Server

---

<b>Overview</b> .....	<b>9</b>
<b>Cross-Platform Considerations</b> .....	<b>9</b>
Overview .....	9
Prepare the Database for the Destination Platform .....	10
<b>Migrating to SAS Federation Server 4.1</b> .....	<b>10</b>
Overview .....	10
Invoking the Migration Utility on Windows .....	10
Invoking the Migration Utility on UNIX .....	11
<b>Promoting Content to SAS Federation Server 4.1</b> .....	<b>11</b>
<b>ODBC Drivers</b> .....	<b>12</b>
Overview .....	12
Modifying ODBC for UNIX .....	12

---

## Overview

To upgrade to SAS Federation Server 4.1, you must perform certain tasks to update your databases and configuration files. There are two options for upgrading to SAS Federation Server 4.1:

- **Migration:** Migration is needed only when upgrading from DataFlux Federation Server 2.x to SAS Federation Server 4.1. In this case, use the Migration tasks to update your databases to be compatible with SAS Federation Server 4.1.
- **Promotion:** Promotion is the movement of selected content from Federation Server 3.x to SAS Federation Server 4.1 on the same version platform (for example, from 3.1.x to 4.1).

---

## Cross-Platform Considerations

### Overview

When migrating or promoting SAS Federation Server to a different platform (for example, migrating from Windows to Linux), you must prepare the SYSCAT database for the destination platform. If applicable, you can also prepare the SQL\_LOG database

for the destination platform. The preparation must be done before performing the migration or promotion tasks that are presented in this chapter. If you are not changing platforms, you can skip this topic.

### **Prepare the Database for the Destination Platform**

To prepare the SYSCAT database for the destination platform:

1. Perform a backup of the 2.x or 3.x SYSCAT.TDB file as outlined in “[Database Backup and Restore](#)” on page 71. Perform the backup from the 2.x or 3.x installation path. This is an example of the backup command for UNIX:

```
install-path/bin/dfsutil backup -db syscat /tmp/syscat.tbk
```

2. Move the SYSCAT backup file to the destination machine, where SAS Federation Server 4.1 is installed.
3. Working in the installation area for 4.1, perform a restore on the backup file. See “[Database Backup and Restore](#)” on page 71 for additional details. This is an example of the restore command for UNIX:

```
install-path/bin/dfsutil restore -db syscat /tmp/syscat.tbk
```

*Note:* The restored file is placed in the `/var` directory of the 4.1 installation path, and is named **SYSCAT.TDB** in Windows, and **syscat.tdb** in UNIX. The location of the restored file is where the `cfg.TRANPATH` entity is pointing in `dfs_entities.dtd`. The default directory is usually `/var` but might be configured for a different path.

4. To perform a migration, use the 4.1 installation location when prompted for the old 2.x installation location. If you are performing a promotion, you are finished because the database file has been copied to the new location.
5. (Optional) Repeat the preceding steps for the SQL\_LOG transactional database, if needed. However, the SQL\_LOG database file does not require migration. The restore places the database in the correct directory of the 4.1 installation area.

---

## **Migrating to SAS Federation Server 4.1**

### **Overview**

SAS Federation Server 4.1 installation contains a migration utility that enables you to migrate DataFlux Federation Server 2.1.x. The Windows installation contains a shortcut labeled **Migrate 2.1.x Database**, and the UNIX installer uses the **migrate** option with **dfsadmin**. Each option requires that you point the migration utility to the 2.1.x installation directory.

### **Invoking the Migration Utility on Windows**

The migration process involves copying data from your old installation into the current installation and upgrading it from the 2.x format to the 4.1 format. When running the migration utility, you will be asked for the location of the 2.1.x installation as shown in the following task.

1. After installing SAS Federation Server, run the **Migrate 2.1.x Database** tool, which launches a mini-installer.
2. Click **Next** in the Setup dialog box.
3. At Choose Old Install Location, navigate to the root directory for Federation Server 2.1.x (for example, *C:\Program Files (x86)\DataFlux\FederationServer\2.1*).
4. Click **Next** and **Start** to begin the migration.
5. When the migration completes, click **Next** and **Finish**.

*Note:* Migration results are written to the `dfs_log_213_310.xml` file located in the `\etc\migrate` directory.

### Invoking the Migration Utility on UNIX

The UNIX utility, `dfsadmin`, contains a migration tool to migrate a Federation Server system database (`syscat.tdb`) from 2.1.x to 4.1. Use the following syntax to invoke migration:

```
./bin/dfsadmin migrate <2.1.x_installation_directory>
```

The migration tool reads from the 2.1.x federation server configuration file to determine where the 2.1.x system database (`syscat.tdb`) file is located and then migrates it to the location for release 4.1. Migration results are written to a log file, `dfs_log_213_310.xml`, located in `/etc/migrate`.

---

## Promoting Content to SAS Federation Server 4.1

To promote SAS Federation Server databases:

1. Stop the 3.x servers using **Windows Services Manager**, the **Start** menu, or the `dfsadmin` utility for UNIX.
2. Copy the 3.1 or 3.2 SYSCAT.TDB (Windows) or `syscat.tdb` (UNIX) and the SQL\_LOG.TDB files to the 4.1 installation location (for example, `<SAS Home>\SASFederationServer\4.1\fedserver\var`).
3. Start the 4.1 servers using **Windows Services Manager**, the **Start** menu, or the `dfsadmin` utility for UNIX.

**TIP** If any custom configuration changes were made to the existing configuration files, update the new configuration files with these configurations. For more information, see [SAS Federation Server Configuration Reference on page 32](#)

*Note:* When promotion is performed as an ‘Upgrade in Place’ (UIP), the port number for the prior release (3.x) is reassigned to release 4.1. As a result, release 3.x is disabled and the associated service or daemon will no longer run.

---

## ODBC Drivers

### Overview

SAS Federation Server 4.1 is installed with release 7.10 of the Wire Protocol ODBC drivers. The installation directory for Windows has changed to `C:\Program Files\DataFlux\ODBC\7.1`. All of the existing ODBC DSNs should use the new 7.10 Wire Protocol ODBC drivers, and modifications are not necessary. If you are running ODBC on UNIX, you must make some minor modifications as presented in the next topic.

### Modifying ODBC for UNIX

To use the new 7.1 ODBC drivers, take one of the following actions:

- Configure the existing `odbc.ini` file to point to the new drivers in `<4.1_SASHome>/SASFederationServer/4.1/fedserver/lib/fx*27.so`. For example, change the `Driver=` value to:  

```
Driver=<4.1_SASHome>/SASFederationServer/4.1/fedserver/lib/FXora27.so.
```
- Use the `dfdbconf` tool to configure the data sources to use the new drivers. You should already have `ODBCINI` and `ODBCINST` set from a previous installation. If using `dfdbconf`, you must export the `ODBCINI` environment variable. The environment variable for `ODBCINST` is set when you execute `dfsadmin` for SAS Federation Server 4.1. Here are export examples:

```
export ODBCINI=<4.1_SASHome>/SASFederationServer/3.x/fedserver/etc/odbc.ini
export ODBCINST=<4.1_SASHome>/SASFederationServer/3.x/fedserver/etc/odbcinst.ini
```

For more information, see “Configuring ODBC Connections” on page 24.

## Chapter 3

# Configuring SAS Federation Server

<b>Overview</b> .....	<b>14</b>
<b>Configure Temporary Storage for SAS Utility Files</b> .....	<b>14</b>
The Directory Location .....	14
UNIX .....	14
Windows .....	14
<b>Configuring the Windows Environment</b> .....	<b>15</b>
Overview .....	15
Federation Server Directory Permissions .....	15
Configuring ODBC Connections .....	16
Starting and Stopping the Windows Service .....	16
<b>Configuring the UNIX Environment</b> .....	<b>17</b>
Overview .....	17
UNIX File System and Directory Permissions .....	17
Setting Environment Variables .....	18
Configuring ODBC Connections .....	24
Starting and Stopping the Server Daemon .....	25
<b>Configuring Server Accounts</b> .....	<b>26</b>
Overview .....	26
The System User Account .....	26
Administrator Account .....	27
The Trusted User Account .....	27
Shared Logins .....	28
Best Practices for Setting Shared Logins .....	29
<b>SAS Federation Server Configuration Reference</b> .....	<b>32</b>
Locale Support .....	32
Key Configuration Files .....	33
About the Server Configuration Files .....	34
About Option Names and Option Sets .....	34
Required Configuration Options .....	35
Configuration Options .....	36
<b>Configuring SAS Federation Server Manager</b> .....	<b>41</b>
Start the Configuration .....	41
Update JVM .....	42
Federation Server Manager Service .....	42

---

## Overview

This chapter focuses on post-installation configuration of SAS Federation Server for Windows and UNIX environments. Also included is a configuration reference that covers all of the possible configuration options for SAS Federation Server. SAS Federation Server Manager configuration is also covered in this chapter. SAS Federation Server Manager is the graphical interface for administration of SAS Federation Server.

---

## Configure Temporary Storage for SAS Utility Files

### The Directory Location

SAS Federation Server and other SAS applications create temporary utility files that are written to the default `/tmp` or `/temp` directory that is set in your environment. It is recommended that you specify a location for these files to ensure that there is enough space for processes, such as threaded applications, to create utility files. Use one of the procedures below to configure an environment variable to accommodate the utility directory and files. The name of the SAS utility directory is determined by the following:

```
SAS_util<serial><pid>_<node>
```

- `<serial>` is a unique 4-to 6-digit hexadecimal serial number that distinguishes each directory from the other directories that are created by the same process.
- `<pid>` is the process ID number, which is represented as an 8 digit hexadecimal number.
- `<node>` is the name of the host, or machine on which the process is running.

### UNIX

In UNIX, set the location for utility files using `TKUTILLOC` in an export statement:

```
export TKUTILLOC=~/directory_1/dfs
```

The utility directory and files are created in the specified directory. When SAS Federation Server is started, you should see a directory similar to `SAS_util000100000EF0_devmachine` that contains `*.ut1` files

### Windows

In Windows, use the Control Panel to set the `TKUTILLOC` environment variable:

1. From the Control Panel, select **System and Security** and then select **System**.
2. Select **Advanced system settings** to open the System Properties window.
3. Click on environment variables.
4. Under system variables click **New** and set **TKUTILLOC** as the variable name with the path to the directory that will store utility files. Utility files contain a `.ut1` extension.

5. Click **OK** and start the Federation Server.

When SAS Federation Server is started, you should see a directory similar to `SAS_util000100000EF0_devmachine` that contains `*.utl` files.

*Note:* If you set the directory to a location that does not exist, TKUTILLOC does not create the directory and reverts to the default temporary directory. For example, in Windows, `C:\Users\user_1\AppData\Local\Temp`.

---

## Configuring the Windows Environment

### Overview

This section outlines the necessary configuration procedures and server tasks that you must complete following installation of SAS Federation Server in a Windows environment.

### Federation Server Directory Permissions

The recommended directory permissions for SAS Federation Server installed on a Windows® platform are listed in the following table:

Directories	Users	Default Permissions
<code>SASHome</code>	Installer, Administrator	Full Control
<code>SASHome\fedserver \server instance</code>	Process user	Read and Execute, List Folder Contents
<code>SASHome\fedserver \server instance\var</code>	Installer, Administrator	Full Control
	Process user	Read, Write, List Folder Contents
	The user who backs up SAS Federation Server; Backup Administrator	Read, List Folder Contents
TranPath as specified in the server configuration file	Installer, Administrator	Full Control
	Process user	Read, Write, List Folder Contents
	The user who backs up SAS Federation Server; Backup Administrator	Read, List Folder Contents

*Note:* All other users have no access.

## Configuring ODBC Connections

This section describes how to configure ODBC connections for SAS Federation Server in a Windows environment.

To access a database through ODBC with SAS Federation Server, a driver for ODBC for the specific database must be used. SAS Federation Server comes with the ODBC drivers for several databases. The database and database connection must also be configured as an ODBC data source.

To access a database through a vendor supplied client, the client must be installed and configured according to the vendor documentation.

To add an ODBC data source, use the ODBC Data Source Administrator in Microsoft Windows. To set up a new ODBC connection:

1. Click **Start** ⇒ **Control Panel**.
2. Double-click **Administrative Tools** ⇒ **Data Sources (ODBC)**.  
*Note:* In Windows 7, the view of the Control Panel can vary. If you do not see Administrative Tools when you open the Control Panel, click System and Security to access Administrative Tools, Data Sources (ODBC).
3. Click **Add**.
4. In the ODBC Data Source Administrator window, select the driver that is appropriate for your data source.
5. In the ODBC Driver Setup dialog box, enter the Data Source Name, Description, and database-specific configuration options. These values are required, and can be obtained from your database administrator.

Use the vendor supplied client configuration utility for connections that are not ODBC.

## Starting and Stopping the Windows Service

The SAS® Federation Server runs as a Windows service that is accessible through the Control Panel or Management Console. To access the service:

1. Click **Start** ⇒ **Settings** ⇒ **Control Panel**
2. Double-click **Administrative Tools** ⇒ **Computer Management**.
3. Expand the **Services and Applications** folder.
4. Click **Services**.
5. Click **SAS Federation Server**.
6. Click either **Stop the service** or **Restart the service**.

*Note:* Version number differences between the database schema and the Federation Server itself can terminate the start process. Major version numbers must match. The minor version number of the schema can be 1 less than the minor version number of the server.

### *Modifying the Service Log On*

At installation, SAS Federation Server service is configured to start using the local system account. Because this account can have some restrictions, such as accessing

network drives, it is suggested that you modify the service log on account to an account that has the appropriate privileges to run SAS Federation Server.

To modify the SAS Federation Server service log on:

1. Select **Control Panel** ⇒ **Administrative Tools**.
2. Double-click **Services**, and select the **SAS Federation Server service**.
3. Click the **Log On** tab, select **This account**, and enter Account and Password credentials for a user with administrative privileges.

---

## Configuring the UNIX Environment

### Overview

This chapter outlines the necessary configuration procedures and server tasks that you must complete following installation of SAS Federation Server in a UNIX environment.

### UNIX File System and Directory Permissions

The recommended file permissions for Federation Server installed on a UNIX platform are listed in the following table:

Directories	Users	Default Permissions
<b>SASHome</b>	Installer, Administrator	Read, Write, Execute
<b>SASHome/SAS Federation Server</b>	Process user	Read, Execute
<b>SASHome/SAS Federation Server/server instance/var</b>	Installer, Administrator	Read, Write, Execute
	Process user	Read, Write, Execute
	The user who backs up SAS Federation Server; Backup Administrator	Read, Execute
<b>TranPath</b> as specified in the server configuration file, <b>dfs_serv_common.xml</b>	Installer, Administrator	Read, Write, Execute
	Process user	Read, Write, Execute
	The user who backs up SAS Federation Server; Backup Administrator	Read, Execute

*Note:* All other users have no access.

## Setting Environment Variables

### Overview

Before configuring SAS Federation Server driver software, set environment variables as outlined in the following sections.

### Set the LANG Environment Variable

If using BASE data sets with SAS Federation Server, the **LANG** environment variable must be set before bringing up the server. This environment variable is needed for the VALIEDATEFMT table.

Most UNIX or Linux systems use the **LANG** environment variable to specify the desired locale and this variable is often already set in your environment. Locale names vary among different UNIX or Linux operating systems, so use a value that is supported by your version of UNIX or Linux.

- Invoke the `locale` command to show your current locale.
- Use `locale -a` to display a list of all the locales that are currently installed on the machine.

For more information about setting locale environment variables, consult the documentation for your operating system.

### Setting Environment Variables for DBMS-Based Drivers

Before configuring your Federation Server, you should determine the following information about your DBMS:

- The version or release of the DBMS client shared libraries installed on your operating system. This is important due to potential incompatibilities between DBMS versions or releases.
- The location of the DBMS client shared libraries. This is important so that the correct client libraries can be loaded.

*Note:* The steps outlined in this chapter assume that the ODBC drivers were installed during installation of SAS Federation Server.

Refer to the following sections for detailed instructions on configuring your environment to interface with your Federation Server driver software:

- [“SAP Reference” on page 161](#)
- [“DB2 Driver Reference” on page 124](#)
- [“ODBC Driver Reference” on page 143](#)
- [“Oracle Reference” on page 151](#)
- [“Teradata Reference” on page 189](#)

### SAS Federation Server Driver for SAP

SAP® software requires extensive configuration before it can be used. For more information, see [“Installing and Configuring the SAS Federation Server Driver for SAP” on page 165](#).

### SAS Federation Server Driver for DB2

The SAS Federation Server Driver for DB2 uses shared libraries that are referenced as shared objects in UNIX. You must add the location of the shared libraries to one of the system environment variables, and, if necessary, indicate the DB2 version that you have installed at your site. Before setting the environment variables as shown in the examples below, you must also set the following environment variables:

- The INSTHOME environment variable must be set to your DB2 home directory.
- The DB2DIR environment variable should also be set to the value of INSTHOME.
- The DB2INSTANCE environment variable should be set to the DB2 instance configured by the administrator.

#### AIX

```
Bourne Shell    $ LIBPATH=$INSTHOME/lib:$LIBPATH
                $ export LIBPATH
```

```
C Shell        $ setenv LIBPATH $INSTHOME/lib:$LIBPATH
```

#### HP-UX and HP-UX for the Itanium Processor Family Architecture

```
Bourne Shell    $ SHLIB_PATH=$INSTHOME/lib:$SHLIB_PATH
                $ export SHLIB_PATH
```

```
C Shell        $ setenv SHLIB_PATH $INSTHOME/lib:$SHLIB_PATH
```

#### Linux for Intel Architecture, Linux for x64, Solaris, and Solaris for x64

```
Bourne Shell    $ LD_LIBRARY_PATH=$INSTHOME/lib:$LD_LIBRARY_PATH
                $ export LD_LIBRARY_PATH
```

```
C Shell        $ setenv LD_LIBRARY_PATH $INSTHOME/lib:$LD_LIBRARY_PATH
```

### SAS Federation Sever Driver for Greenplum

To use ODBC with Greenplum, you must set the **ODBCINI** environment variable to the **odbc.ini** file located in the Federation Server installation path:

```
export ODBCINI=$installpath/fedserver/etc/odbc.ini
```

When you run **dfsadmin**, the **ODBCINST** environment variable is set to the **odbcinst.ini** file located in the federation server installation path. Here is an example:

```
export ODBCINST=$installpath/fedserver/etc/odbc.ini
```

### SAS Federation Server Driver for Netezza

The Netezza ODBC drivers are ODBC API-compliant shared libraries that are referenced as shared objects in UNIX. You must include the full path to the shared libraries in the shared library path as shown below so that the ODBC drivers can be loaded dynamically at run time.

**AIX**

```
Bourne Shell  $ LIBPATH=$ODBCHOME/lib64:$LIBPATH
               $ export LIBPATH
```

```
C Shell      $ setenv LIBPATH $ODBCHOME/lib64:$(LIBPATH)
```

**HP-UX for the Itanium Processor Family Architecture**

```
Bourne Shell  $ SHLIB_PATH=$ODBCHOME/lib64:$SHLIB_PATH
               $ export SHLIB_PATH
```

```
C Shell      $ setenv SHLIB_PATH $ODBCHOME/lib64:$(SHLIB_PATH)
```

**Linux for Intel Architecture, Linux for x64, Solaris, and Solaris for x64**

```
Bourne Shell  $ LD_LIBRARY_PATH=$ODBCHOME/lib64:$LD_LIBRARY_PATH
               $ export LD_LIBRARY_PATH
```

```
C Shell      $ setenv LD_LIBRARY_PATH $ODBCHOME/lib64:$(LD_LIBRARY_PATH)
```

**SAS Federation Server Driver for ODBC**

To configure ODBC data sources, you might have to edit the `.odbc.ini` file in your home directory. Some ODBC Driver vendors allow system administrators to maintain a centralized copy by setting the environment variable `ODBCINI`. Please refer to your vendor documentation to find specific information. Additional information for ODBC can be found in “[Configuring ODBC Connections](#)” on page 16.

The Drivers for ODBC are ODBC API-compliant shared libraries, referred to as shared objects in UNIX. You must add the location of the shared libraries to one of the system environment variables so that drivers for ODBC are loaded dynamically at run time. You must also set the `ODBCHOME` environment variable to your ODBC home directory before setting the environment variables as shown in the following examples.

**Linux for Intel Architecture and Linux for x64**

```
Bourne Shell  $ LD_LIBRARY_PATH=$ODBCHOME/lib:$LD_LIBRARY_PATH
               $ export LD_LIBRARY_PATH
```

```
C Shell      $ setenv LD_LIBRARY_PATH
               $ODBCHOME/lib:$LD_LIBRARY_PATH
```

**Solaris and Solaris for x64**

```
Bourne Shell  $ LD_LIBRARY_PATH=$ODBCHOME/lib:$LD_LIBRARY_PATH
               $ export LD_LIBRARY_PATH
```

### Solaris and Solaris for x64

```
C Shell      $ setenv LD_LIBRARY_PATH
              $ODBCHOME/lib:${LD_LIBRARY_PATH}
```

### AIX

```
Bourne Shell $ LIBPATH=$ODBCHOME/lib:$LIBPATH
              $ export LIBPATH
```

```
C Shell      $ setenv LIBPATH
              $ODBCHOME/lib:${LIBPATH}
```

### HP-UX and HP-UX for the Itanium Processor Family Architecture

```
Bourne Shell $ SHLIB_PATH=$ODBCHOME/lib:$SHLIB_PATH
              $ export SHLIB_PATH
```

```
C Shell      $ setenv SHLIB_PATH
              $ODBCHOME/lib:${SHLIB_PATH}
```

### SAS Federation Server Driver for Oracle

You can connect to any Oracle server from SAS Federation Server using the Oracle driver, client versions 10g or 11g. If you are using Oracle client version 11g, run these additional commands to successfully connect with Oracle from SAS Federation Server:

```
cd
  /installation_directory/fedserver/lib
cp tkeora11.so tkeora.so
```

*Note:* For HP64, the file extensions are .sl.

To use the SAS Federation Server Driver for Oracle, you must set the `ORACLE_HOME` environment variable. In addition, you must make sure that the shared library path variable (the name of this variable is operating system dependent) points to the location of the Oracle shared libraries. This is required since the SAS Federation Server Driver for Oracle executable uses Oracle shared libraries and needs to know where they are located at your site.

The following are examples for the various operating systems:

### AIX

```
Bourne Shell $ LIBPATH=$ORACLE_HOME/lib:$LIBPATH
              $ export LIBPATH
```

```
C Shell      $ setenv
              LIBPATH=$ORACLE_HOME/lib:$LIBPATH
```

**HP-UX and HP-UX for the Itanium Processor Family Architecture**

```
Bourne Shell  $
               SHLIB_PATH=$ORACLE_HOME/lib:$SHLIB_PATH
               $ export SHLIB_PATH
```

```
C Shell      $ setenv SHLIB_PATH
               $ORACLE_HOME/lib:$SHLIB_PATH
```

**Linux for Intel Architecture, Linux for Itanium-based Systems, Solaris, and Solaris for x64**

```
Bourne Shell  $
               LD_LIBRARY_PATH=$ORACLE_HOME/lib:$LD_LIBRARY_PATH
               $ export LD_LIBRARY_PATH
```

```
C Shell      $ setenv LD_LIBRARY_PATH
               $ORACLE_HOME/lib:$LD_LIBRARY_PATH
```

**SAS Federation Server Driver for SAP HANA**

The SAS Federation Server Driver for SAP HANA uses an ODBC interface to access SAP HANA. The SAS Federation Server Driver for SAP HANA requires the 64-bit ODBC driver for SAP HANA. The SAP HANA client includes the ODBC driver.

These are the prerequisites for configuration of the SAS Federation Server Driver for SAP HANA:

- You have downloaded the SAP HANA client software from SAP Service Marketplace and installed and configured the ODBC driver.
- For more information about how to obtain the software, see the SAP HANA Master Guide on [http://help.sap.com/hana\\_appliance/](http://help.sap.com/hana_appliance/).
- For information how to install and configure the ODBC driver refer to the SAP HANA Client Installation Guide on [http://help.sap.com/hana\\_appliance/](http://help.sap.com/hana_appliance/).
- You must include the full path to the shared library in the shared library path as shown below so that the ODBC drivers can be loaded dynamically at run time.

**AIX**

```
Bourne Shell  $ LIBPATH=/usr/sap/hdbclient:$LIBPATH
               $ export LIBPATH
```

```
C Shell      $ setenv LIBPATH /usr/sap/hdbclient:$LIBPATH
```

**HP-UX for the Itanium Processor Family**

```
Bourne Shell  $ SHLIB_PATH=/usr/sap/hdbclient:$SHLIB_PATH
               $ export SHLIB_PATH
```

**HP-UX for the Itanium Processor Family**

```
C Shell      $ setenv SHLIB_PATH /usr/sap/hdbclient:$SHLIB_PATH
```

**Linux for Intel Architecture, Linux for x64, Solaris, and Solaris for x64**

```
Bourne Shell $ LD_LIBRARY_PATH=/usr/sap/hdbclient:$LD_LIBRARY_PATH
              $ export LD_LIBRARY_PATH
```

```
C Shell      $ setenv LD_LIBRARY_PATH /usr/sap/hdbclient:$LD_LIBRARY_PATH
```

The SAS Federation Server Driver for SAP HANA can use data sources defined in the `odbc.ini` file to identify the SAP HANA server. The general format of the `odbc.ini` is shown below:

```
[ODBC Data Source]
SERVERNODE=hana_host:hana_port
```

For example:

```
[SAPHHANADSN]
SERVERNODE=hanasrv1.mycompany.com:30015
```

Set the `ODBCINI` environment variable to the location and name of your `odbc.ini`:

```
Bourne Shell  ODBCINI=path-to/odbc.inieport ODBCINI
```

```
C Shell      setenv ODBCINI path-to/odbc.ini
```

**SAS Federation Server Driver for Teradata**

The SAS Federation Server Driver for Teradata uses shared libraries, referred to in UNIX as shared objects. You must add the location of the shared libraries to one of the system environment variables.

**AIX**

```
Bourne Shell $ LIBPATH=TERADATA-CLIENT-LOCATION:$LIBPATH
              $ export LIBPATH
```

```
C Shell      $ setenv LIBPATH TERADATA-CLIENT-LOCATION:$LIBPATH
```

**HP-UX**

```
Bourne Shell $ SHLIB_PATH=TERADATA-CLIENT-LOCATION:$SHLIB_PATH
              $ export SHLIB_PATH
```

```
C Shell      $ setenv SHLIB_PATH TERADATA-CLIENT-LOCATION:$SHLIB_PATH
```

**HP-UX for the Itanium Processor Family**

```
Bourne Shell    $ SHLIB_PATH=TERADATA-CLIENT-LOCATION:$SHLIB_PATH
                $ export SHLIB_PATH
                $ LD_PRELOAD=/usr/lib/hpux64/libpthread.so.1
                $ export LD_PRELOAD
```

```
C Shell        $ setenv SHLIB_PATH TERADATA-CLIENT-LOCATION:$SHLIB_PATH
                $ setenv LD_PRELOAD /usr/lib/hpux64/libpthread.so.1
```

**Linux for Intel Architecture, Linux for x64, Solaris, and Solaris for x64**

```
Bourne Shell    $ LD_LIBRARY_PATH=TERADATA-CLIENT-LOCATION:$LD_LIBRARY_PATH
                $ export LD_LIBRARY_PATH
```

```
C Shell        $ setenv LD_LIBRARY_PATH TERADATA-CLIENT-LOCATION:$LD_LIBRARY_PATH
```

**TTU 8.2 and HP-UX**

HP-UX users with TTU 8.2 must create two symbolic links from the `/usr/lib/pa20_64` directory with the following commands:

```
$ ln -s /usr/lib/pa20_64/libicudatatd.sl libicudatatd.sl.34
$ ln -s /usr/lib/pa20_64/libicuuctd.sl libicuuctd.sl.34
```

**Configuring ODBC Connections****Overview**

To access a database through ODBC with SAS Federation Server, an ODBC driver for the specific database must be used. SAS Federation Server is delivered with branded ODBC drivers for several databases. You can opt to use one of these ODBC drivers or you can use a third-party driver. The database must also be configured as an ODBC data source when using an ODBC driver.

**Configuring ODBC Connections Using the SAS Federation Server Drivers**

SAS Federation Server includes an ODBC configuration tool used to configure the ODBC data sources that are installed with SAS Federation Server.

**ODBC Configuration Tool**

Use the ODBC configuration tool, `dfdbconf`, to add new data sources to your ODBC configuration.

To add a new data source:

1. From the root directory of SAS Federation Server installation, run: `./bin/dfdbconf`.
2. Select **A** to add a data source.

3. Select a template for the new data source by choosing a number from the list of available drivers.
4. Set parameters for the driver as you are prompted to do so. The new data source is added to the `odbc.ini` file.

**TIP** You can also use `dfdbconf` to modify an existing data source if it is no longer needed.

### Testing the Connection

Once you have added all of your data sources, use the interactive ODBC Viewer application, `dfdbview`, to test your connection. For example, if you added a data source called `my_oracle`, run `./bin/dfdbview my_oracle` (from the installation root) to test the connection. You might be prompted for a user name and password. If the connection succeeds, you will see a prompt from which you can enter SQL commands and query the database. If the connection fails, `dfdbview` displays error messages describing one or more reasons for the failure.

### Configuring ODBC Connections Using Third Party ODBC Drivers

- To access a database through ODBC with SAS® Federation Server, an ODBC driver for the specific database must be used. The database must also be configured as an ODBC data source when using an ODBC driver.
- To access a database through a vendor supplied client, the client must be installed and configured according to the vendor documentation.
- Verify the connection with a third-party client tool before attempting connection through SAS Federation Server.

### unixODBC Driver Manager

unixODBC is an open-source product that implements the ODBC API. If unixODBC is required and not already installed, visit <http://www.unixODBC.org> and download the required software.

The following configurations are required when using the unixODBC driver manager with SAS Federation Server:

1. Include unixODBC in the `PATH` and `LD_LIBRARY_PATH`.
2. The `odbcini` and `odbcinst.ini` installed with SAS Federation Server are for use with the ODBC driver manager, also installed with SAS Federation Server. Since the third-party ODBC driver will likely use the unixODBC driver manager, you will need to update the `odbcini` and `odbcinst` files used by unixODBC and update the OBCDINI environment variable accordingly.
3. `DM_UNICODE=utf-16` is required in the advanced options of any data service that is used with the driver manager.

Use the vendor supplied client configuration utility for non-ODBC connections. For more information about configuring third-party databases, See “[Setting Environment Variables for DBMS-Based Drivers](#)” on page 18 .

## Starting and Stopping the Server Daemon

Start and stop the SAS Federation Server daemon using the `dfsadmin` application, located in the `bin` directory of the installation root. This application is run using the

command syntax: `./bin/dfsadmin yourcommand` where `yourcommand` is one of the following:

To start SAS Federation Server: `./bin/dfsadmin start`

To stop SAS Federation Server: `./bin/dfsadmin stop`

Additional information for the `dfsadmin` application can be found in the “[dfsadmin Utility](#)” on page 70 topic.

---

## Configuring Server Accounts

### Overview

SAS Federation Server uses the following accounts for administration, authentication and data authorization:

- “[The System User Account](#)” on page 26  
The system user account is the most privileged account for SAS Federation Server.
- “[Administrator Account](#)” on page 27  
An administrator account is a user account created in Authentication Server, and then granted ADMINISTER privilege on SAS Federation Server.
- “[The Trusted User Account](#)” on page 27  
The trusted user account establishes a trust relationship between the SAS Federation Server and the Authentication Server.
- “[Shared Logins](#)” on page 28  
Shared logins help manage user connections from SAS Federation Server to the data. With shared logins, users can access data without needing individual logins to access the data source.

### The System User Account

#### Introduction

The SYSTEM user is a privileged account which means that it carries more privileges than an administrator account. There is nothing on SAS Federation Server that the system account cannot do because the account has implicit privileges to all user and data objects. The operating system process that invokes and runs SAS Federation Server is automatically a system user. By default, the account used to install SAS Federation Server is registered as a system user account. Other user accounts can also be registered as system users in the `dfs_serv_common.xml` configuration file.

The system user should identify users who will be administrators of SAS Federation Server, and make them administrators by granting them the ADMINISTER privilege on the server object. Like SYSTEM users, administrators are unconditionally and implicitly granted all privileges on SAS Federation Server. However, if these users are revoked their ADMINISTER privilege, then they become standard users that have privileges granted to and denied from them. A SYSTEM user can never be denied privileges.

If a Data Source Name (DSN) is created by either the system user or an administrator, the DSN is created using the AS ADMINISTRATOR clause, which means that the

ADMINISTRATOR role owns the DSN, not the individual creating it. Therefore, if the administrator user is later removed from the system, the DSN will not be deleted with the user.

Use the system user account to define one or more administrators for SAS Federation Server. As a best practice, all configuration and administration should be performed by the administrator.

### **Configuring the SystemUsers OptionSet**

The SystemUsers option set defines the system user accounts for the SAS Federation Server. When system users grant or deny privileges to others, the grantor is reflected in the system tables as the system user ID. Each system user should be a domain-qualified user name.

```
<SystemUsers>
  <Option name="Account">domain\uid1</Option>
  <Option name="Account">domain\uid2</Option>
</SystemUsers>
```

### **Altering the System User Account**

You can add or remove system users by updating configurations in the SystemUsers option set explained above.

## **Administrator Account**

### **Overview**

An administrator account is a user account created in Authentication Server, and then granted ADMINISTER privilege on the SAS Federation Server. Only system users can grant the ADMINISTER privilege.

Administrators have implicit privileges to perform every other action including the following:

- create and drop data source names (DSN)
- grant and deny privileges to other accounts
- create and drop data services, catalogs, and schemas

### **Setting the ADMINISTER Privilege**

To define a user as an administrator for SAS Federation Server, grant the ADMINISTER privilege to their account using the following syntax:

```
GRANT serverpriv ON servername TO "user-ID"
GRANT administer ON FedServer1 TO "user1"
```

For further details, reference the [GRANT and DENY DDL statements on page 220](#).

## **The Trusted User Account**

### **Overview**

The trusted user account is set up in the Authentication Server and configured in SAS Federation Server using the [ALTER SERVER statement on page 195](#) or SAS Federation Server Manager. Use this account to establish a trust relationship between the

SAS Federation Server and the Authentication Server. The trust relationship is required for certain features to work such as definer's rights views.

A few notes about the Trusted User account:

- A trusted user is a user ID that has to be able to authenticate using the authentication method that is deployed for the installation.
- As stated above, this account is used only by SAS Federation Server to connect to Authentication Server in certain scenarios like definer's rights views. You would never log on using this account.
- The trusted user should not be a system or administrative user on Authentication Server or SAS Federation Server.

### **Using DDL to Configure the Trusted User**

Use the [ALTER SERVER statement on page 195](#) to define the trusted user.

```
ALTER SERVER {OPTIONS(xset TRUSTED_USER_UID 'user-ID',
xset TRUSTED_USER_PWD 'password')}
```

*Note:* User IDs and passwords must be added or dropped in pairs.

### **Configure the Trusted User in Federation Server Manager**

You can also configure a trusted user account using SAS Federation Server Manager.

To define the trusted user:

1. Locate the SAS Federation Server object in the tree and select the **Action Menu** in the upper left corner.
2. Using the action menu, select **Properties** to display the server properties dialog box.
3. Click the **Security** tab in the Federation Server **Properties** dialog box.
4. Set the Trusted user ID and password and click **OK**.

## **Shared Logins**

### **Overview**

Using Shared Logins, users can access data when they do not or should not need to know the credentials used to access this data. Shared Logins help manage user connections from SAS Federation Server and outside of the database. Only one set of credentials is used to access data for groups of people who will never know the credentials that are being used. Authorization management occurs and is managed at a layer above the database. Using shared logins is optional. If using shared logins to connect to data sources, they must be configured in SAS Federation Server before they will function properly in a DSN.

### **Shared Login Manager**

When using shared logins to authenticate users to a data source, users do not need to know the credentials that they are using because the shared login manager is retrieving credentials for the user that is logged on and providing the credentials to SAS Federation Server so that the server can connect the user to the database through the appropriate data service or DSN.

- The shared login account is created in Authentication Server which includes the login to be shared and its domain, consumers of the login, and a login key.

- The shared login manager and login key is configured in SAS Federation Server using administrative DDL or SAS Federation Server Manager.

The shared login manager login identifies the credentials used to connect to the Authentication Server and retrieve the shared login. It is required that the shared login manager be a manager of that shared login in Authentication Server. The shared login is not directly readable by consumers of the shared login. It can be read only by managers or the owner. The shared login manager's password is encrypted.

The shared login key identifies which shared logins created in Authentication Server are available to a SAS Federation Server instance. The key defined in SAS Federation Server must match the key that is part of the shared login definition in the Authentication Server.

- After the shared login manager account is configured, the account information can be referenced from DSNs in SAS Federation Server. The DSN is configured with a credentials search order of SHARED.

Configure shared logins in SAS Federation Server using the **ALTER SERVER** DDL statement. You can also configure a shared login through SAS Federation Server Manager, using the **Security** dialog box located in the properties of a Federation Server object.

### **Shared Login Manager Configuration Using DDL**

Use the [ALTER SERVER statement on page 195](#) to set shared logins and to configure the shared login key. For example,

```
alter server {options xset sharedloginkey 'DefaultKey'}
alter server {OPTIONS(xset SHAREDLOGINMANAGER 'manager user id',
xset SHAREDLOGINPASSWORD 'manager password')}
```

*Note:* User IDs and passwords must be added or dropped in pairs.

### **Shared Login Configuration Using Federation Server Manager**

1. Locate the Federation Server object in the tree and select the **Action Menu** in the upper left corner.
2. From the action menu, select **Properties** to display the server properties dialog box.
3. Select the **Security** tab and configure the shared login with the following information:
  - Shared login manager user ID
  - Shared login manager password
  - Shared login key

After shared logins are configured, they can be referenced to establish connections to data sources. Reference the [“Configuring Data Source Access” on page 100](#) topic for additional information.

## **Best Practices for Setting Shared Logins**

### **Overview**

Shared logins are created in the Authentication Server when there is a need to share logins among multiple users. Logins must be owned by users, not groups. Therefore, shared logins are the only mechanism to provide this functionality.

Shared logins consist mainly of the login and domain to share, and the consumers who use that login. The consumers will typically list one or more groups, which is a best practice. A conflict can arise when a particular user is in a consumer group (directly or indirectly) of multiple principal maps for the same domain. The following scenarios explain how shared login conflicts are resolved.

### **Scenario 1: Application Users**

In the following scenario, an application exists which requires the use of a particular set of database credentials to access its protected data. For example, an HR application has data content stored in Oracle and DB2.

To manage credentials:

1. Identify all the users of the HR application. The users might have different roles or data access privileges, but they all need to access the data. These users (or subgroups) will all be placed in the group HR\_USERS.
2. Create a shared login for each domain. In this case, the administrator would create an HR\_ORACLE and HR\_DB2 shared login. For both shared logins, the administrator would specify the HR\_USERS group as a consumer member of the shared login. Each shared login would contain the appropriate principal and domain for the database.
3. Specify the GROUP option to qualify the users with the shared login, either in the DSN itself, or on the connection string specifying the DSN. In this case, the GROUP option would be HR\_USERS.
4. Set authorizations on different users and groups to control which set of users can perform which operations, for example SELECT versus UPDATE versus DELETE. All of the users and groups should be members of the HR\_USERS group.

At connect time, the HR\_USERS group is used to identify the correct shared login for each underlying database connection. If the connecting user were a consuming member of another shared login, the GROUP value would properly identify which one to use.

### **Algorithm When Using the GROUP Option**

Shared logins are initially considered candidates for outbound credentials selection if the domain and shared login key match. If the domain is empty, shared logins for any domain initially qualify. This also applies to the shared login key, which is configured as part of SAS Federation Server.

If the GROUP connection string option is specified (which is derived from the consumer group in the DSN configuration), then only maps where the group is a direct or indirect consumer will be considered a candidate for outbound credentials selection. The basic algorithm selects a map based on the closeness of the specified group to the map.

### **Candidate Map Processing**

1. If the user is not a direct or indirect member of the shared login consumer group, the map is not a candidate; otherwise,
2. If the GROUP is not a direct or indirect consumer of the map, the map is not a candidate; otherwise,
3. The distance from the GROUP to the map is computed by following the group member-of relationship up to the group that is the direct consumer. The candidate map is retained only if the distance is less than or equal to the current minimum distance to the map. The current minimum distance is updated.

After all candidate maps have been processed,

1. if exactly one has been retained, use the associated credentials; or
2. if two or more were retained, check the closest two and use the credentials associated with the closest of the two (error if the distances are the same unresolved conflict); or
3. if no maps have been retained, use no credentials.

### **Scenario 2: Organized Consuming Users**

In the following scenario, the administrator has organized the users based on company organization or another classification.

The administrator wants to use this relationship so users qualify for a particular shared login, for example:

- The administrator wants to grant access to Oracle account EXECUTIVE\_USER to his most privileged users, identified by the MARKETING\_EXECUTIVE group.
- The administrator wants to grant access to Oracle account MARKETING\_USER to members of the marketing division in the company, identified by the MARKETING group.
- The administrator wants to grant access to Oracle account STANDARD\_USER to all other known users in the system, identified by the USERS group.
- The administrator created groups which reflect the company's organizational chart. The MARKETING group reflects all members of the marketing organization, with the MARKETING\_EXECUTIVE group included as a member of the MARKETING group.

The administrator creates shared logins for the EXECUTIVE\_USER, MARKETING\_USER and STANDARD\_USER Oracle accounts. Next, he assigns MARKETING\_EXECUTIVE, MARKETING and USERS consuming groups, respectively, to these shared logins.

Then, the shared login chosen would be as follows:

- For members of the MARKETING\_EXECUTIVE group, they would be closest to the shared login identified by that group, even though they were likewise members of the MARKETING and USERS groups. Therefore, this set of users would consume the EXECUTIVE\_USER Oracle account.
- For members of the MARKETING group, they would be closest to the Shared Login identified by that group, even though they were likewise members of the USERS group. Therefore, this set of users would consume the MARKETING\_USER Oracle account.
- All other known users would qualify only for the shared login identified by the USERS group. Therefore, this set of users would consume the STANDARD\_USER Oracle account.

In this scenario, the administrator would not use the GROUP option, since the option accepts only a single value, and no single value works for all users. The administrator would omit the GROUP option and allow a closeness algorithm to identify which shared login to use.

### **Algorithm When No GROUP Option Is Specified**

If the GROUP connection string option is not specified, then all maps where the user is a direct or indirect consumer will be considered a candidate for outbound credentials selection. The basic algorithm selects a map based on the closeness of the specified user

to the map, or USERS if the user is not a consumer, and finally PUBLIC if USERS is not either. If the user is PUBLIC, selection is done based on PUBLIC alone.

### **Candidate Map Processing for a User**

1. If the user is not a direct or indirect consumer of the map, the map is not a candidate; otherwise,
2. The distance from the user to the map is computed by following the group member-of relationship up to the group that is the direct consumer. The candidate map is retained if the distance is less than or equal to the current minimum distance to the map. The current minimum distance is updated.

After all candidate maps have been processed,

1. if exactly one has been retained, use the associated credentials; or
2. if two or more were retained, check the closest two and use the credentials associated with the closest of the two (error if the distances are the same unresolved conflict); or
3. if maps have not been retained, use no credentials.

### **Candidate Map Processing for USERS**

- If USERS is a direct consumer of the map, then the candidate map is retained.
- If a candidate map has been retained already, return an error (unresolved conflict).
- After all candidate maps have been processed, if exactly one has been retained, return OK and the associated credentials; or
- Continue on to Candidate Map Processing for PUBLIC.

### **Candidate Map Processing for PUBLIC**

- If USERS is a direct consumer of the map, then the candidate map is retained.
- If a candidate map has been retained already, return an error (unresolved conflict).
- After all candidate maps have been processed, if exactly one has been retained, return OK and the associated credentials; or
- Return OK but empty credentials if no candidate maps retained.

### **Path Length Computation Details**

If USERS or PUBLIC is a member of another group, and that group is a map consumer, direct or indirect, the path length is not incremented when traversing from the user to the map. For the purposes of map selection, this effectively makes placing either of these two groups in another group a shorthand way to place all users in that group.

## **SAS Federation Server Configuration Reference**

### **Locale Support**

At this time, SAS Federation Server supports the English, United States of America (en\_US) locale. The following table outlines the character representations for output (display) format. There are no deviations from these formats:

Character Type	Format
Number	dddd.d.fffffffffff
Date	yyyy-mm-dd
Time	hh:mm:ss
Timestamp	yyyy-mm-dd hh:mm:ss [.fffffffffff]

*Note:* Configure database drivers and clients to match this behavior to ensure that conversions are handled correctly.

## Key Configuration Files

The following table lists the key configuration files for SAS Federation Server:

File or Script Name	Description
<code>dfs_serv.xml</code> , <code>dfs_serv_common.xml</code>	These are the core configuration files for SAS Federation Server. They specify the system users, the location of the internal database, and other key configuration settings necessary for proper functionality of SAS Federation Server. These configuration files are located in the <code>/etc</code> directory of the Federation Server installation path. Detailed configuration information is presented in the <a href="#">Configuration Options on page 36</a> .
<code>dfs_entities.dtd</code>	The <code>dfs_entities.dtd</code> file contains the values that were supplied during the install process. These values are referenced by other configuration files such as <code>dfs_serv.xml</code> and <code>dfs_serv_common.xml</code> files.
<code>dfs_log.xml</code>	This is the logging facility configuration file for SAS Federation Server. It specifies logging options for SAS Federation Server from information-only to debug and trace. This file is installed in the <code>/etc</code> directory of the installation path. For more information, see <a href="#">Server Logging Configuration on page 65</a> .
<code>dfs_log_SQL_Logging.xml</code>	This is the configuration file that is used to facilitate SQL Logging. This file is located in the <code>/etc</code> directory of the installation path. For additional information, see <a href="#">“SQL Logging” on page 56</a> .
<code>dfs_pwencode.xml</code>	This is the logging file used when running the tool to encrypt a password. The logs are located in the <code>/var</code> directory of the Federation Server installation path. See <a href="#">Utilities for SAS Federation Server on page 70</a> for additional information about password encryption.

## About the Server Configuration Files

SAS Federation Server uses the `dfs_entities.dtd` file to store values that are supplied during installation. These values are used by the other configuration files, `dfs_serv_common.xml`, `dfs_serv.xml`, `dfs_log4sas.xml`, and `dfs_log_sql_logging.xml`. For example, a port number supplied during installation is recorded in `dfs_entities.dtd` and the port number option in `dfs_serv.xml` points to the definition in the `.dtd` file:

---

<code>dfs_serv.xml</code>	<code>&lt;Option name="Port"&gt;&amp;cfg.port;&lt;/Option&gt;</code>
<code>dfs_entities.dtd</code>	<code>&lt;!ENTITY cfg.Port "2171"&gt;</code>

---

## About Option Names and Option Sets

### Overview

The `dfs_serv_common.xml` and `dfs_serv.xml` configuration files consist of a combination of option names and option sets that are explained below.

### Option Names

Option names specify a **name=value** pair as configuration file options. They can stand alone in the configuration file or are contained within an option set. Here are the different types of option name configurations that appear in `dfs_serv.xml`:

This specifies a simple name=value pair as a configuration option:

```
<Option name="XXX">yyy</Option>
```

Example: `<Option name="Port">21030</Option>`

This represents a name=value pair where the value is a space-delimited string of values:

```
<Option name="XXX">yyy zzz aaa</Option>
```

Example: `<Option name="AdminLoginManagementPolicy">ADD  
REMOVE UPDATE</Option>`

### Option Sets

An option set consists of various options (option names). Option names that belong in an OptionSet will not be assessed correctly if they are placed outside of the OptionSet. For example, the SystemUsers option set requires at least one option name, Account, be defined in order to function properly:

```
<SystemUsers>
  <Option name="Account">domain\uid1</Option>
  <Option name="Account">domain\uid2</Option>
</SystemUsers>
```

### Child Option Sets

Here is an example of an option set within an option set which is called a child option set. The License option set shown in the following example contains two options to define: Provider and Location (where the license provider and the location of the license file is specified):

```

<OptionSet name="License">
  <OptionSet
name="Primary">
    <Option name="Provider">SAS</Option>
    <Option name="Location">path_to_license_file</Option>
  </OptionSet>
</OptionSet>

```

For a complete list of configuration options for SAS Federation Server, refer to [“Configuration Options” on page 36](#).

## Required Configuration Options

### Overview

The following configurations are required for SAS Federation Server operations.

### Security Provider

The SecurityProvider option set provides information about the Federation Server security provider, including the threaded kernel extension name and other information specific to the security provider. Configuration of the extension and database options are required.

```

<OptionSet name="SecurityProvider">
  <Option name="extension">extension_name</Option>
  <Option name="Database">database_name</Option>
</OptionSet>

```

### Authentication Server

The AuthenticationServer option set defines the location of the Authentication Server that serves as the back end for SAS Federation Server. The Authentication Server is required to authenticate SAS Federation Server users and store user information such as logins and group memberships.

The option, IOM\_URI contains connection information for the Authentication Server and uses the following format:

```

iom://<machine>:<port>;Bridge;CLSID=2D1BCDBF-F900-4CA9-85F6-95ECDBAF2122
<OptionSet name="AuthenticationServer">
  <OptionSet name="PrimaryServer">
    <Option name="URI">IOM-URI</Option>
  </OptionSet>
</OptionSet>

```

### License

The License Option Set provides information about the type of licensing issued for SAS Federation Server. Within the Primary option set, specify the license provider and the path to the license file.

```

OptionSet name="License">
  <OptionSet name="Primary">
    <Option name="Provider">SAS</Option>
    <Option name="Location">C:\Federation Server\server1\etc\license</Option>
  </OptionSet>

```

The Primary option set is a required configuration for operation of SAS Federation Server. You cannot start the server without a valid license.

## Configuration Options

### Overview

The following sections reflect the options that are available in the system configuration files, `dfs_serv.xml` and `dfs_serv_common.xml`. A `dfs_entities.dtd` can exist in specific configurations. The `dfs_entities.dtd` file contains the values supplied during installation of SAS Federation Server. These values are referenced by the system configuration files.

### AppendEnv OptionSet

#### Syntax

```
<OptionSet name="AppendEnv">
  <Option name="FIREBIRD">drive:\install_loc\firebird</Option>
</OptionSet>
```

#### Description

The AppendEnv option set locates the specified OS environment variable and appends the specified option to the environment variable's current value. If the environment variable does not exist, it is created and set to the specified value. The AppendEnv option does not add a delimiter between the existing and appended environment variable values. Therefore, if a delimiter is needed, it should be included at the beginning of the specified value.

**TIP** Each AppendEnv, PrependEnv and SetEnv option is processed entirely in the order in which they appear in the configuration file.

### SetEnv OptionSet

#### Syntax

```
<OptionSet name="SetEnv">
  <Option name="FIREBIRD">[drive]:\install_dir\lib\fbembed</Option>
  <Option name="FIREBIRD_LOG">[drive]:\install_dir\var\log</Option>
  <Option name="FIREBIRD_TMP">[drive]:\FDS_Tmp</Option>
</OptionSet>
```

#### Description

This option set appears in `dfs_serv_common.xml`. The SetEnv option sets the OS environment variables to specific values. If the environment variable does not exist, it will be created and set to the option value. If the environment variable does exist, the value will be updated to the option value.

Set FIREBIRD\_TMP as an environment option should the default database directory run out of space. Once the default directory has no available space, the engine switches to the directory specified in FIREBIRD\_TMP.

**TIP** Each AppendEnv, PrependEnv and SetEnv option is processed entirely in the order in which they appear in the configuration file.

### Memory Size Option

#### Syntax

```
<Option name="MemSize">nnnnn [(K|k|M|m|T|t) [(B|b)]]</Option>,
```

#### Example

```
<Option name="MemSize">1G</Option>
```

**Description**

The MemSize system option specifies the total amount of memory available for each SAS Federation Server session. If a setting is not specified, all system memory is available for use by SAS Federation Server. However, SAS Federation Server will use only as much memory as it needs to complete a process. Setting a value that is too low will result in out-of-memory conditions.

**Transactional Data Store Options****Syntax**

```
<Option name="FIREBIRD">drive:\install_dir\lib\fbembed</Option>
<Option name="FIREBIRD_LOG">drive:\install_dir\var\log</Option>
```

**Description**

The FIREBIRD environment variable specifies the location of the Transactional Data Store installation files.

The FIREBIRD\_LOG environment variable specifies the location of the log files for Transactional Data Store. The configuration file generated during installation sets the FIREBIRD\_LOG option to the var\log directory of the installation path. If FIREBIRD\_LOG is not set, the federation server will default to one of two locations:

- **TranPath:** If the TranPath environment variable is set, FIREBIRD\_LOG is set to the TranPath value.
- **ContentRoot:** If TranPath is not set, FIREBIRD\_LOG is set to the ContentRoot value as defined in the configuration file.

**TranPath Option****Syntax**

```
<Option name="TranPath">directory</Option>
```

**Description**

The TranPath option identifies the location in which to store the Federation Server Database files.

By default, the Federation Server Database files are stored in ContentRoot as defined in the configuration file. Federation Server Database system files cannot be stored on remote file systems. If ContentRoot is a network file system or share, use the TranPath option to redirect the Federation Server Database files to a local directory on the machine where SAS Federation Server is installed.

**Rules and Dependencies**

Here are the rules and dependencies for the TranPath option:

- A relative directory specified in the ContentRoot tag is resolved against the server's working directory.
- A relative directory specified in the TranPath tag is resolved against the directory specified in the ContentRoot tag, or the working directory of the server if none is specified.
- If the TranPath option is omitted, it defaults to the directory specified in the ContentRoot tag, or the working directory of the server if none is specified.
- The TranPath directory is used to resolve the server's system catalog database name specified in the provider-specific SecurityProvider or Database tag for the Transactional Data Store security provider (**tkescfb**) provider.

**PrependEnv OptionSet**

## Syntax

```
<OptionSet name="PrependEnv">
  <Option name="FIREBIRD">drive:\install_loc\firebird</Option>
</OptionSet>
```

## Description

This option appears in `dfs_serv_common.xml`. The PrependEnv option will find the indicated OS environment variable and prepend the option value to the OS environment variable value. If the environment variable does not exist, it will be created and set to the option value. The PrependEnv option will not add a delimiter of any sort between the existing and new environment variable value. If a semicolon (;) is needed, then the option value should include it at the end.

**TIP** Each AppendEnv, PrependEnv and SetEnv option is processed entirely in the order in which they appear in the configuration file.

**SystemUsers OptionSet**

## Syntax

```
<SystemUsers>
  <Option name="Account">domain\uid1</Option>
  <Option name="Account">domain\uid2</Option>
</SystemUsers>
```

## Description

This option appears in `dfs_serv_common.xml` and defines the system user account(s) that are given all privileges to SAS Federation Server including all user and data objects. This privilege cannot be revoked or denied. When system users grant or deny privileges to others, the grantor is reflected in the system tables as the SYSTEM user ID. Each system user should be a domain-qualified user name.

**SecurityProvider OptionSet**

## Syntax

```
<OptionSet name="SecurityProvider">
  <Option name="extension">extension_name</Option>
  <Option name="Database">database path</Option>
</OptionSet>
```

## Description

This is a required configuration for SAS Federation Server. The security provider option set provides information about the server's security provider, including the threaded kernel extension name and other information specific to the security provider.

**FunctionDispatchManager Option**

## Syntax

```
<Option name="FunctionDispatchManager">tktsfd</Option>
```

## Description

This option appears in `dfs_serv_common.xml`. The Function Dispatch Manager tells FedSQL to load an extension that implements SQL functions, including row-level security. This option should always be set to `tktsfd`.

**Database Option**

## Syntax

```
<Option name="Database">syscat.tdb</Option>
```

**Description**

Identifies the name of the database (syscat.tdb) to be used for the Transactional Data Store security provider (**tkescfb**).

The database name is appended to the TranPath value as defined elsewhere in the configuration file. The default TranPath value for Windows is **Drive:\Program Files\SASHome\SASFederationServer\version\var**, so by default, syscat.tdb will reside in that location. The default TranPath value for UNIX is **<install\_dir>/SASHome/SASFederationServer/version/var**, so by default, syscat.tdb will reside in that location.

**CAUTION:**

**The syscat database must reside on a local file system as indicated by the default installation paths above.** If the **var** directory is pointed to a remote file system such as a network file share or storage area network (SAN), the syscat database cannot be created, resulting in multiple application errors.

**ContentRoot Option****Syntax**

```
<Option name="ContentRoot">content_root_path</Option>
```

**Description**

This option, configured in **dfs\_serv\_common.xml**, defines the content root for SAS Federation Server. The content root is used to resolve all relative pathnames specified in SAS Federation Server configuration, such as a schema path. It is recommended that the value for ContentRoot be set to an absolute, fully qualified path.

**Rules**

If the ContentRoot option is not set, files will be written to the install directory.

- Content root is absolute or relative to the install directory.
- TRACEFILEPATH is absolute or relative to content root.
- TRACEFILE names are resolved against the TRACEFILEPATH path. Paths that do not match are rejected.
- PRIMARYPATH paths in schema configuration options are absolute or relative to content root.
- SCHEMA=(PRIMARYPATH) connection string options are resolved against PRIMARYPATH schema configuration path.

**Port Option****Syntax**

```
<Option name="Port">port_number</Option>
```

**Description**

Indicates the port on which SAS® Federation Server will start.

**Authentication Server OptionSet****Syntax**

```
<OptionSet name="AuthenticationServer">
  <OptionSet name="PrimaryServer">
    <Option name="URI">IOM-URI</Option>
  <OptionSet>
</OptionSet>
```

**Description**

A required configuration for SAS® Federation Server, this option defines the location of the Authentication Server that serves as the back end for the .

IOM\_URI contains connection information for the Authentication Server and uses the following format:

```
iom://<machine>:<port>;Bridge;CLSID=2D1BCDBF-F900-4CA9-85F6-95ECDBAF2122
```

For example, IOM\_URI can be used to connect to an Authentication Server that is running on port number 24140 and machine name **myhost**. The CLSID for an Authentication Server should always be 2D1BCDBF-F900-4CA9-85F6-95ECDBAF2122:

```
iom://yourserver.yourdomain.com:21030;Bridge;CLSID=2D1BCDBF-F900-4CA9-85F6-95ECDBAF2122
```

**License OptionSet****Syntax**

```
<OptionSet name="License">
  <OptionSet name="Primary">
    <Option name="Provider">SAS</Option>
    <Option name="Location">path_to_license_file</Option>
  </OptionSet>
</OptionSet>
```

**Description**

This option set, configured in `dfs_serv_common.xml`, provides information about the type of licensing issued for the . The Primary option set is a required configuration. Within the Primary option set, specify the license provider and the path to the license file. The location option points to a setinit file.

**ObjectServerParms Option****Syntax**

```
<Option name="ObjectServerParms">object_server_parameters</Option>
```

**Description**

Identifies IOM object server parameters. Values for `object_server_parameters` include:

- **Clientencryptionlevel = (EVERYTHING | CREDENTIALS | NONE)**— Specifies the client encryption level to use. Valid values include:
  - **NONE**—Nothing is encrypted.
  - **CREDENTIALS**— Login credentials are encrypted. Note: these are the login credentials used to authenticate to the SAS® Federation Server, and NOT logins used as outbound credentials to connect to third-party databases. It also does NOT include credentials passed in administration DDL, such as CREATE ACCOUNT REGISTRATION or CREATE USER.
  - **EVERYTHING**— All client/server communications are encrypted. Setting a value of EVERYTHING can affect server performance.

**NetworkEncryptAlgorithm Option****Syntax**

```
<Option name="NetworkEncryptAlgorithm">algorithm |
("algorithm1", "algorithm2", ...)</Option>
```

**Description**

algorithm | ("algorithm1", "algorithm2", ...) — Specifies the algorithm or algorithms that can be used for encrypting data that is transferred between a client and a server

across a network. When you specify two or more encryption algorithms, use a space or a comma to separate them, and enclose the algorithms in parentheses. If more than one algorithm is specified, the client session negotiates the first specified algorithm with the server session. If the client session does not support that algorithm, the second algorithm is negotiated, and so on.

This is set with the CLIENTENCRYPTIONLEVEL OBJECTSERVERPARMS option. Possible values include: SASProprietary which is the default, and AES if implementing DFSecure.

For more information about Advanced Encryption Standard (AES), see the *Authentication Server Administrator's Guide*.

### **SelectStarExpansion Option**

#### Syntax

```
<Option name="SelectStarExpansion">ALL</Option>
```

#### Description

This option modifies the behavior of the SELECT \* expansion for table columns. The configuration options are ALL or VISIBLE. If set to ALL, the SAS® Federation Server attempts to expand SELECT \* to all of the physical columns in the table and fails if the user does not have the SELECT privilege to one or more columns. If set at VISIBLE, which is the default value, SAS® Federation Server traverses the visible path, expanding the SELECT \* privilege to those columns for which the user has the SELECT privilege.

### **Deadlock Protection Option**

#### Syntax

```
<Option name="SystemDBCTimeOut">milliseconds</Option>
```

#### Description

A deadlock is sometimes caused by competing resources on the server resulting in perpetual wait time for the tasks to complete. This option controls the wait time in milliseconds before timing out a lock attempt that is causing the deadlock. If a connection cannot be acquired within the specified time limit, the request fails and the deadlock connection is aborted, allowing the remaining connection to run to completion. The default is <=0 which waits 'forever'. This option is configured in `dfs_serv_common.xml`.

---

## Configuring SAS Federation Server Manager

### **Start the Configuration**

SAS Federation Server Manager starts a configuration dialog box when you first login after installation. Use the following web address to start Federation Server Manager configuration again: `http://host:port/fsmanager/web/config`

For example: `http://FedSvrManager:21077/fsmanager/web/config`

#### **CAUTION:**

**Before configuring SAS Federation Server Manager, the Federation Server service or daemon must be started.**

1. After launching the configuration URL, you are redirected to a login page. Log in to SAS Federation Server Manager using the administrator user and password that was set during installation.

**TIP** If you do not remember what the admin credentials are, look in `fsmanager.cfg` located in the `\etc` directory of the installation path. The password is case-sensitive.

2. After logging in, you are presented with an **Initial Setup** dialog box. Specify the Authentication Server host name and port number and click **Save**.
3. At the third and final configuration dialog box, you can edit Authentication Server information, configure databases and edit the fetch size for data fetched between SAS Federation Server and the web application server.
4. When you are finished with configuration, click **Log Off** in the upper right corner. You will be redirected to the main login page for SAS Federation Server Manager.

## Update JVM

SAS Federation Server Manager uses Jetty, an opened source Java application web server. It is best practice to update the following JVM settings to suit your environment:

*Note:* When making changes to `fsmanager.cfg` it is best to work with a text editor that can handle the Unicode character set. Some plain text editors might convert characters which can corrupt the file.

-Xms (initial heap size)

-Xmx (max heap size)

-XX MaxPermSize

-Xss (stack size)

To update JVM settings for the Jetty Web Server:

1. Open the `fsmanager.cfg` file located in `SASHome\FSManager\version\etc`.
2. Use the `java_args` option to specify JVM settings.

```
# java_args = <arg1> [<arg2>...]
# Additional arguments for the Java executable used to run the web application.
#
# example: java_args = -Xss1024k -Xmx512m
```

3. Save the `fsmanager.cfg` file.
4. Run the **Refresh Service Options** script located in the `FSManager` directory of the installation root. The JVM settings that are specified in the `fsmanager.cfg` file will be applied to Jetty.

## Federation Server Manager Service

### Introduction

SAS Federation Server Manager runs as a Microsoft Windows service called Federation Server Manager. You can start and stop the service using the Microsoft Management Console (MMC) or through **Administrative Tools** in the **Control Panel**.

In UNIX, SAS Federation Server Manager runs as a daemon administered from a command line using `fsmadmin` to start and stop the daemon.

### **The Federation Server Manager Service in Windows**

Start and stop the service using the MMC or the **Control Panel** ⇒ **Administrative Tools**.

1. Click **Start** ⇒ **Control Panel**.
2. Double-click **Administrative Tools** ⇒ **Computer Management**.
3. Expand the **Services and Applications** folder.
4. Click **Services**.
5. Click **Federation Server Manager** *<instance>*.
6. Click either **Stop the service** or **Restart the service**.

### **Modifying the Windows Service Log On**

When SAS Federation Server Manager is installed, the Federation Server Manager service is configured to start with the local system account. Because this account can have restrictions, such as accessing network drives, you should modify the service properties so the service log on uses an account with the necessary privileges to access the required network drives and files. For security reasons, you should assign administrative privileges only if necessary.

To modify the SAS Federation Server Manager log on:

1. Select **Control Panel** ⇒ **Administrative Tools**.
2. Double-click **Services**, and select the **Federation Server Manager service**.
3. Click the **Log On** tab and select **This account**. Enter the account and password credentials for a user with administrative privileges.
4. Restart the **Federation Server Manager service**.

### **The Federation Server Manager Daemon in UNIX**

Start and stop the daemon using the `fsmadmin` application included in the installation. This application can be run using the command-line command: `./bin/fsmadmin command` from the installation root directory, where *command* is one of the following options:

Command	Description
start	Starts the SAS Federation Server Manager daemon. For example: <code>./bin/fsmadmin start</code>
stop	Stops the SAS Federation Server Manager daemon. For example: <code>./bin/fsmadmin stop</code>
status	Checks the status of the SAS Federation Server Manager daemon.
help	Displays help information for the <code>fsmadmin</code> command.
version	Displays the version information for <code>fsmadmin</code> .



## Chapter 4

# SAS Federation Server Administration

<b>Overview</b> .....	<b>46</b>
<b>SQL Scripting for SAS Federation Server Administration</b> .....	<b>46</b>
Overview .....	46
SQL Scripting Phases .....	46
About the Configuration for SQL Scripting .....	46
Example SQL Script .....	48
<b>The SAS Federation Server Database</b> .....	<b>50</b>
Overview .....	50
Working with the SAS Federation Server Database .....	50
<b>SAS Federation Server Resource Cache</b> .....	<b>51</b>
Overview .....	51
Managing Named Server Caches .....	51
Managing Cache Configuration Properties .....	52
Cache Properties .....	53
<b>Connection Pooling</b> .....	<b>54</b>
Overview .....	54
Options .....	55
<b>SQL Logging</b> .....	<b>56</b>
Overview .....	56
Configuring SQL_Logging .....	56
Configuring a 3rd Party DBMS for SQL Logging .....	57
ARM Transactions .....	57
The SQL_LOG Data Service and DSN .....	59
The EVENTS Table .....	59
SQL Logging Performance Tuning .....	65
<b>Server Logging Configuration</b> .....	<b>65</b>
Introduction .....	65
Initial Logging Configuration .....	66
Logging Thresholds .....	68
Modifying the Server Logging Configuration .....	69
Trace Log .....	69
<b>Utilities for SAS Federation Server</b> .....	<b>70</b>
Introduction .....	70
UNIX Utilities .....	70
Windows Utilities .....	71
<b>Uninstalling SAS Federation Server</b> .....	<b>72</b>

---

## Overview

SAS Federation Server can be managed with the use of various DDL statements and SQL commands or through the SAS Federation Server Manager user interface. SAS Federation Server Manager provides dialog boxes and wizards that guide you through the completion of a task, without requiring that you know the FedSQL commands required to perform the task. If you are not familiar with FedSQL and federation, you might want to use the SAS Federation Server Manager, which is intuitive and easy to use. If you are already familiar with SQL or FedSQL, you might prefer to use the Administration DDL statements presented in [Appendix 1 on page 195](#). Most of the functions performed with administration DDL, can be accomplished using SAS Federation Server Manager.

---

## SQL Scripting for SAS Federation Server Administration

### Overview

SAS Federation Server provides SQL language scripting capabilities to handle administrative needs for start-up and shutdown events. Administrators can write and execute scripts to manage auditing or related event notifications. To run a script, add the name of the script to the main configuration file, `dfs_serv.xml` using an external entity defined in the DTD section. Additional configuration information, including an example, appears at the end of this topic.

### SQL Scripting Phases

SQL scripts execute in one of two phases: **Startup.Epilog** and **Stop.Prolog**.

#### **Startup.Epilog**

**Startup.Epilog** is executed on start-up of SAS Federation Server before listening for connections and after administration DDL can be executed. These scripts can connect to any configured data service and execute SQL that is run as the process user.

#### **Stop.Prolog**

**Stop.Prolog** is executed on SAS Federation Server shutdown after quiescing or dropping client connections. These scripts can connect to any configured data service and execute SQL that is run as the process user.

### About the Configuration for SQL Scripting

#### **XML Format**

SQL scripts are specified in an XML configuration file as an SQL node, an **OptionSet** element, and a name attribute of SQL:

```
<OptionSet name="SQL">  
...
```

```
</OptionSet>
```

SQL nodes can be nested with the outer-most node, which is the script. Nested nodes correspond to SQL commands that run within the script. Additional XML elements provide execution context for the SQL nodes.

### **Elements of SQL Scripting**

The following `Option` and `OptionSet` elements are valid within an SQL node:

```
<OptionSet name="SQL">...</OptionSet>
```

Specifies an SQL node containing a nested SQL script composed of other elements.

```
<Option name="name">SQL node name</Option>
```

Specifies the name of the SQL node. This option is used for logging context.

```
<Option name="SQLErrorMediationAction">Error mediation action</Option>
```

Specifies the required action that is needed to mediate errors during SQL command execution:

**STOP**

Specifies that the script terminate execution when it encounters an error that satisfies the current SQL state and match mode criteria. STOP is the default mediation action.

**CONTINUE**

Specifies that the script continue executing the next SQL node regardless of encountering an error that satisfies the current SQL state and match mode criteria. This is used when a set of SQL commands should run without regard to the success of those commands previously executed within the same SQL node.

Error mediation for the SQL node is inherited from the nesting SQL node, if it exists.

```
<Option name="SQLErrorMediationSSMatchMode">SS match mode</Option>
```

Specifies the SQL state match mode used to identify specific SQL states requiring mediation action:

**EXCLUDE**

Action taken when the SQL state does not match one of the states specified within the set of scoped states of the current SQL node.

**INCLUDE**

Action taken when the SQL state matches one of the states specified within the set of scoped states of the current SQL node.

**IGNORE**

Specifies that SQL states within scope are ignored. This causes the mediation action to be honored based on the success or failure of the command regardless of the SQL state. IGNORE is the default match mode.

The SQL state match mode of the SQL node, which is inherited from the nesting SQL node, if it exists.

```
<Option name="SQLState">SQL state expression</Option>
```

Specifies an SQL state or prefix that is used to identify specific SQL states requiring error mediation. An SQL state can be specified as a full 5-character mnemonic such as **HY001**, or as a prefix matching any SQL state starting with a specified prefix. In addition, a leading **+** character (concatenation operator) can be prepended to the mnemonic to add the SQL state to the current set of constraining SQL states. Without the leading concatenation operator, the specified SQL states replaces any SQL states that were previously specified.

The SQL states for the SQL node are inherited from the nesting SQL node if any.

```
<Option name="ConnectionString">connection-string</Option>
```

Specifies the connection string that is used to access the data source for which SQL commands are submitted for execution. The connection string for the SQL node is DSN=ADMIN.

```
<Option name="Command">SQL command</Option>
```

Specifies the SQL command to execute. The status and SQL state of the execution is processed according to the current error mitigation configured in the containing SQL node lineage.

Nested parameterized SQL commands can be specified where the inner command is executed once for each row materialized in the outer command's result set.

Parameters are specified using @n syntax where n is a single parameter number corresponding to the n'th column of the outer command's result set. Note that commands can be nested only once. Here is an example of the nested command:

```
<OptionSet name="SQL">
...
<Option name="Command">outer query command</OptionSet>
<OptionSet name="SQL">
...
<Option name="Command">parameterized inner command 1</Option>
</OptionSet>
...
<OptionSet name="SQL">
...
<Option name="Command">parameterized inner command n</Option>
</OptionSet>
</OptionSet>
```

```
<Option name="Condition">SQL Boolean scalar result query</Option>
```

Specifies a query that resolves to a scalar Boolean result. Nonconforming commands will fail, causing script execution to end. If the command result is 1, all sibling SQL nodes are executed. Otherwise, they are skipped. Only the first condition is processed within an SQL node parent.

As a precursor to executing a set of SQL commands, a condition can be used to check for the existence of a table or row within a table or a value within a row. A condition on the outermost level SQL node will effectively make the entire script's execution dependent on the result of the specified query.

The following example condition returns 1 when table T has at least one row matching the **WHERE** clause, which is not shown. If the query returns 1, sibling SQL nodes contained in the parent SQL node are executed:

```
<Option name="Condition"> select cast(case when count(*) > 0 then 1
else 0 end as integer) from T where ...</Option>
```

### Example SQL Script

This example script copies the content of in-memory MDS tables to a persistent data store when the server is stopped. You can use an inverse script to load the tables back into the MDS service in the **Startup.Epilog** phase. Error mediation prevents **CREATE TABLE** commands from stopping script execution when the table already exists in the data store. Errors with SQL states beginning with **42S** are excluded from

the stop action. Also, if the C\_STORE\_MDS catalog requires credentials to connect, you can supply those in the connection string.

```
<?xml version="1.0" encoding="utf-8" ?>
<OptionSet name="SQL">
  <!--
    Default phase, SQL error mediation control
  -->
  <Option name="name">MDS to STORE_MDS store</Option>
  <Option name="SQLErrorMediationAction">STOP</Option>
  <Option name="SQLErrorMediationSSMatchMode">EXCLUDE</Option>

  <!--
    MDS Store script
  -->
  <OptionSet name="SQL">
    <Option name="Phase">Stop.Prolog</Option>
    <Option name="SQLState">42S</Option>
    <Option name="ConnectionString">
      driver=FEDSQL;conopts=((security=NO;catalog=C_STORE_MDS);
                            (security=NO;catalog=C_MDS))
    </Option>
  <!--
    Result set generator command: Enumerate all MDS tables...
  -->
  <Option name="Command">
    select TABLE_SCHEM,
           TABLE_NAME
    from DICTIONARY.TABLES
      where TABLE_CAT='C_MDS' and
            TABLE_TYPE='TABLE'
  </Option>
  <!--
    Result set iterator command:
    Store MDS catalog tables in C_STORE_MDS "mirror" catalog
  -->
  <OptionSet name="SQL">
    <Option name="Command">
      create table C_STORE_MDS.S."@2" as
        select * from C_MDS."@1"."@2" where 1=0
    </Option>
  </OptionSet>
  <OptionSet name="SQL">
    <Option name="Command">
      delete from C_STORE_MDS.S."@2"
    </Option>
  </OptionSet>

  <OptionSet name="SQL">
    <Option name="Command">
      insert into C_STORE_MDS.S."@2"
        select * from C_MDS."@1"."@2"
    </Option>
  </OptionSet>
</OptionSet>
</OptionSet>
</OptionSet>
```

This example script was saved as `store_mds.xml`. After saving the script, edit the `dfs_serv.xml` configuration file as highlighted in the following example:

```
<?xml version="1.0" encoding="UTF-8"?>

<DOCTYPE Config [
<ENTITY % entities SYSTEM "dfs_entities.dtd">
%entities;
<ENTITY MDS_SCRIPT SYSTEM "store_mds.xml">
]>

<Config name="TSConfig">

<!-- Common server options -->
&SERVER_COMMON;

<!-- Run the MDS script -->
&MDS_SCRIPT;

</Config>
```

For additional information about the `dfs_serv.xml` configuration file, see [“SAS Federation Server Configuration Reference” on page 32](#).

## The SAS Federation Server Database

### Overview

The SAS Federation Server database is a transactional database, or system catalog (SYSCAT) that contains configuration metadata. Configuration metadata includes the list of created data services, DSNs, privileges, and other information generated as a result of configuring SAS Federation Server. This information is stored in a database because the metadata must be in a consistent state which requires the use of ACID transactions (atomicity, consistency, isolation and durability).

The system catalog contains information about the configuration of SAS Federation Server. This information can be returned to the user through queries against [information views on page 224](#).

### Working with the SAS Federation Server Database

#### Creation of the Database

The SAS Federation Server database is created once when SAS Federation Server is initially invoked. At that time, a set of system tables is created to hold various objects that are created as the server is configured. For example, when a data service is created, the system tables are updated to hold the definition of the new data service. Each time a change is made to the server configuration, the system tables in the database are modified. The database can be backed up at any time to capture and preserve a particular server configuration. The default location of the database is `<installation root>\SASHome\SASFederationServer\version\var` and is called `SYSCAT.TDB`. The name and location are contained in the `dfs_entities.dtd` configuration file.

### **Changing the Database Location**

If the location of the SAS Federation Server database changes, the following configurations require updates:

#### **dfs\_serv\_common.xml/dfs\_entities.dtd**

Using `dfs_entities.dtd`, change the value of the `TranPath` option to point to the new location of the SYSCAT transactional database. This configuration change effects the `dfs_serv_common.xml` configuration file.

#### **dfs\_log\_SQL\_Logging.xml/dfs\_entities.dtd**

Using `dfs_entities.dtd`, change the value for the `FIREBIRD_LOCK` environment variable to point to the new location of the SYSCAT transactional database. This configuration change updates the location of the database lock files for both the `SQL_LOG` and SYSCAT transactional databases. When updating the `FIREBIRD_LOCK` environment variable, use an absolute path only.

#### **dfs\_log.xml**

Change the value of the `DFS_DBAPPENDER_DB` entity to point to the new location. This changes the location of the `SQL_LOG` transactional database.

### **Backing Up SAS Federation Server Database**

It is a best practice to periodically back up the SAS Federation Server database, including the system catalog, SYSCAT.TDB. As mentioned previously, you can backup the database at any time to capture and preserve a particular server configuration. SAS Federation Server can remain active during backup. There is no need to shut the server down. For additional information, refer to *Database Backup and Restore* in “[Utilities for SAS Federation Server](#)” on page 70.

## **SAS Federation Server Resource Cache**

### **Overview**

Data that is used frequently can be cached from Authentication Server and retained on SAS Federation Server until the cached information is refreshed or purged. This data cache can help improve server performance at both ends by reducing the number of calls needed from SAS Federation Server to Authentication Server.

### **Managing Named Server Caches**

SAS Federation Server maintains several internal resource caches, all of which are designed to improve the performance of potentially expensive operations. An administrative user can manage common cache properties by name by using the `ALTER SERVER` DDL statement. Among the cached resources are user and group identity information. This information is required in authorization enforcement and multi-tiered authentication, privilege information, and result sets generated from the execution of definer's rights views.

SAS Federation Server can cache resources that are related to authentication, reducing roundtrips to Authentication Server. Several of these configurable caches are periodically repopulated as SAS Federation Server captures information from Authentication Server during the authentication process. The cache names are prefixed with `AS`. By default, resources related to Authentication Server are not cached.

SAS Federation Server can also cache privilege information, reducing internal queries to various system tables related to privileges, thereby improving the rendering of authorization enforcement decisions. The authorization cache is periodically updated as SAS Federation Server performs authorization enforcement and processes DDL such as GRANT, DENY, REVOKE, and various DROP commands. The authorization cache is named **Authorization** and is configured at maximum level by default.

SAS Federation Server can cache result sets of definer's rights views, improving query execution and data access performance. For information about enabling caching, see [Managing Cache Configuration Properties on page 52](#).

The following cache namespace table describes the information cached under each name.

Cache Name	Description
AS	All Authentication Server (AS) cached resources
AS.Name	Name to identifier mappings
AS.Name.Subjects	User name to Authentication Server identifier cache
AS.Name.Groups	Group name to Authentication Server identifier cache
AS.Subject	Per user cache resources
AS.Subject.Groups	User group memberships cache
AS.Subject.Principals	User owned principals cache
AS.List	Directory listings
AS.List.Subjects	User listings cache
AS.List.Groups	Group listings cache
Authorization	Privileges cache
ResultSet	Result Sets
ResultSet.View	View result sets cache

*Note:* SAS Federation Server Manager does not display these values. To view them, use the SQL Console to select from the Information Views (for example, ).

```
SELECT *
FROM CONFIG_DATA_SERVICES WHERE DATA_SERVICE_NAME=
'__SERVER__'
```

### Managing Cache Configuration Properties

Common cache management operations are handled using the ALTER SERVER command with CACHE list-valued options. This CACHE list-valued option is keyed by the NAME option (similar to the CONOPTS list-valued option, keyed by DRIVER). Values of the NAME option must be one of the names listed in the preceding table.

This statement resets, drops, or adds individual properties of the named cache:

```
ALTER SERVER {OPTIONS( cache-option-list [,cache-option-list ...] )}
cache-option-list ::= CACHE( NAME cache-name , cache-properties )
```

This statement drops properties currently persisted with the named cache and reverts their run-time settings to defaults:

```
cache-option-list ::= DROP CACHE( NAME cache-name )
```

This statement resets or adds properties of the named cache as a complete set, replacing any existing properties:

```
cache-option-list ::= SET|ADD|XSET CACHE( NAME cache-name, cache-properties )
```

The **NAME** option is required and specifies the name of the cache to be managed. Properties of the cache are replaced or created within the sublist. Normal generic SQL options syntax applies to the cache option and the associated suboptions outlined in Cache Properties.

## Cache Properties

### **TIMEOUT** *timeout*

All caches support the **TIMEOUT** option. The value for **TIMEOUT** specifies the length of time, in seconds, that a resource can be cached before being considered stale and marked for on-demand refresh. When a resource becomes stale, it is typically refreshed and reached on its next access. Here are the default **TIMEOUT** values associated with each of the caches:

<b>NAME</b>	<b>Default TIMEOUT Value</b>
ResultSet	1800 (30 minutes)
ResultSet.View	1800 (30 minutes)
Authorization	-1 (infinite)
All others	0 (not applicable — not cached)

The **TIMEOUT** property can be restored to a default several ways once it is explicitly configured. In the following scenario, the configured **TIMEOUT** values for result set caching are as follows:

```
ResultSet          = 3600 (1 hours)
ResultSet.View     = 3600
```

The following statement overrides both of these **TIMEOUT** values:

```
ALTER SERVER {options cache(name ResultSet, xset timeout 300)}
```

The statement sets the timeout of **ResultSet** to 300 seconds explicitly and also sets all children (for example, **ResultSet.View** to 300 seconds. Note that the statement only persists the new **TIMEOUT** value for the cache, **ResultSet**, but changes the current value for all the children, **ResultSet.View**, as well. This allows top-down run-time management of **TIMEOUT** values while preserving the configured defaults of child names.

To reset **TIMEOUT** to the original default value, issue the **TIMEOUT** option with no value:

```
ALTER SERVER {options cache(name ResultSet.View, xset timeout)}
```

Cache properties are inherited from the parent namespace when the cache configuration is dropped altogether:

```
ALTER SERVER {options drop cache(name ResultSet.View)}
```

Afterward, the **ResultSet.View** cache inherits the **TIMEOUT** value from the parent namespace, **ResultSet**, which is 300 seconds.

#### **PURGE** | **FLUSH**

Specifies that the named cache should be refreshed. Associated resources are reacquired and cached on next access and can be flushed immediately. This option is not persisted and using it does not affect existing properties that have already been configured for the named cache. All caches support the **FLUSH** option.

#### **LEVEL** *level*

Controls the caching granularity of the named cache. This property applies to the Authorization cache only. Valid values are as follows:

<b>ALL / OBJECT</b>	Cache privileges for columns, tables and all higher level secure objects. This is the default privilege caching level.
<b>CONTAINER</b>	Cache privileges for schemas and all higher level secure objects.
<b>NONE / OFF</b>	Used to turn off all privilege caching.

---

## Connection Pooling

### Overview

Connection pooling is a reserve of database connections that are maintained in SAS Federation Server so that the connections can be reused as future requests to the database are required. Opening and maintaining a database connection for each user, especially requests made dynamically, is costly and resource intensive. With connection pooling, connections are created and placed into the pool to be used over again so that a new connection to a back-end data source does not have to be reestablished. This practice reduces the amount of time it takes to establish a connection to a database. If all the pooled connections are in use, and the pool is large enough to hold a new connection, then a new connection is made and added to the pool.

If connection pooling is enabled, the client connects to a data source as usual. If there is an existing database connection in the connection pool that meets the client's requirements (for example, a connection to the desired database using the applicable credentials), then that connection will be used by the client. Otherwise, a new connection is created.

When the client disconnects, the server evaluates whether to keep the underlying connection in the connection pool, or whether to free it. If the connection will not be pooled, then the connection is freed when the client frees its connection handle.

## Options

The following options control connection pooling on the server. The options are controlled by the [ALTER SERVER DDL statement on page 195](#).

### Enable Connection Pooling

```
CONNECTION_POOLING
[N|O] | F [ALSE] | OFF | 0 | Y [ES] | T [RUE] | ON | 1
```

This option controls whether connection pooling is enabled or disabled for the server. If connection pooling is switched on, connections to databases are not disconnected immediately when the client requests to disconnect from the database. The connections are put into a pool of connections that can be reused by subsequent requests to connect to the same database with the same attributes and credentials. Connections used for Memory Data Store cannot be pooled.

To disable connection pooling, use **DROP\_CONNECTION\_POOLING**. Configuring **ENABLE\_CONNECTION\_POOLING** using a value of 0 (zero) is invalid and does not disable connection pooling.

### Connection Pool Timeout

```
CONNECTION_POOL_TIMEOUT
seconds
```

This option identifies the time in seconds an unused connection stays in the connection pool. The default is 60 seconds. If the time is exceeded, the connection is removed from the pool and the connection is closed. A value of -1 indicates that the connection never times out and can stay in the pool indefinitely. These connections are freed when the server is stopped. To disable the timeout, use **DROP\_CONNECTION\_POOL\_TIMEOUT**. Configuring **CONNECTION\_POOL\_TIMEOUT** using a value of 0 (zero) does not disable the timeout.

### Maximum Unused Connections

```
CONNECTION_POOL_MAXSIZE maxsize
```

This option identifies the maximum number of unused connections in the connection pool. The default is 50. If the maximum number of connections is reached and a new connection is added to the connection pool, the oldest connection is removed from the pool and that connection is closed. If this option is set at 0, the default of 50 is used. Connections used for Memory Data Store cannot be pooled.

### Drop Connection Pooling

```
DROP CONNECTION_POOLING
```

This option drops and also disables the connection pooling option.

### Drop Connection Pool Timeout Option

```
DROP CONNECTION_POOL_TIMEOUT
```

This option removes the connection pool timeout value. If connection pooling is enabled and connection pool timeout option has been dropped, it will use the default timeout value, 60.

### Drop Connection Pool Maxsize Option

```
DROP CONNECTION_POOL_MAXSIZE
```

This option removes connection pool maxsize value. If connection pooling is enabled and connection pool maxsize option has been dropped, it will use the default maxsize value, 50.

---

## SQL Logging

### Overview

SQL Logging is the ability to view SQL statements submitted to SAS Federation Server, combined with other information such as the user who submitted the SQL and a breakdown of the SQL into prepare, execute and cursor phases. Metrics are also available, including elapsed time, number of rows fetched, and size of data moved (fetched or inserted). SQL Logging provides critical information for all server activity, so it can easily be determined who is accessing the system, when they were connected, and what work they performed.

Errors and informational messages that occur when writing SQL Log records to the EVENTS table are recorded in the server's log and appear with a prefix of **DBAppender<SQL\_LOG>:Append**. See the “[Server Logging Configuration](#)” on page 65 topic for more information about the logging facility for SAS Federation Server.

### Configuring SQL\_Logging

#### Introduction

Configuration of SQL Logging is facilitated using `dfs_log_SQL_Logging.xml` located in the `/etc` directory of the Federation Server installation path.

#### Enable SQL Logging

To enable SQL Logging, open `dfs_log_SQL_Logging.xml` for editing and set the top-level logger to TRACE. There is no need to change any of the configuration parameters. The logging level value should always be set at TRACE so that all information is captured.

Following top-level logging are transactions that you can use to control SQL Logging detail. Remove the comment marks from each line to display logging details:

#### Enable SQL Logging in SAS Federation Server Manager

SAS Federation Server always starts with SQL Logging set at the default level in the configuration file, but can be enabled or disabled dynamically for a server session. This is done through the **SQL Log** tab of the Federation Server Properties dialog box. When SQL Logging is modified from the **SQL Log** tab, it is active for the server session only. When the server is restarted, the level of logging reverts to the value specified in the SQL logging configuration file.

*Note:* Configuring SQL Logging in SAS Federation Server Manager activates logging for the server session only.

To enable SQL Logging in SAS Federation Server Manager:

1. Select a Federation Server in the tree.
2. Open the action menu in the upper left corner and select **Properties**.
3. The Federation Server Properties window appears. Click the **SQL Log** tab.
4. Click to select **On. Log SESSION** and select the events to record.

## Configuring a 3rd Party DBMS for SQL Logging

If you are using a database other than the SQL\_LOG database that installs with SAS Federation Server, set up your database according to vendor specifications. SAS Federation Server is shipped with example configuration files for a few select databases such as Oracle and DB2. These files are located in the `/etc` directory of the Federation Server installation path. After setting up your database and specifying the domain, configure a connection to the driver, and modify the data types to suit your environment as outlined below.

### Change the Domain

Using the ALTER DATA SERVICE command, change the domain for SQL\_LOG to the domain that was created for the data service. Here is an example:

```
ALTER DATA SERVICE SQL_LOG domain ORAI
```

### Modify the Connection String

Modify the connection string to include the values for your SQL Logging database. Here is an example,

```
param name="connectionString"
value="CATALOG=*;DEFAULT_CATALOG=catalog_name;DRIVER=driver;
CONOPTS=(DSN={yourdsn};UID='your user id';PWD='your password')"
```

For a list of possible connection options see the [“Database Functionality and Driver Performance” on page 124](#) in this guide.

### Create the Table / Modify Data Types

At the end of the CREATE TABLE statement, update `'IN <name of your table space>'` to reflect the value for your setup.

Modify the columns that appear in the SQL Logging configuration file to map to data types supported by the new database. A complete list of SQL Logging columns and data types appears in [“The EVENTS Table” on page 59](#).

```
CREATE TABLE &DFS_DBAPPENDER_TABLE;
  (ARM_NAMESPACE VARCHAR(256),
  TRAN_TIMESTAMP TIMESTAMP,
  APP_HANDLE VARCHAR(36),
  APP_ID VARCHAR(36),
  APP_NAME VARCHAR(50),...
  ) IN <name of your table space>"/>
```

*Note:* For additional setup information, refer to the sample configuration files located in the `/etc` directory of the Federation Server installation path.

## ARM Transactions

The table below shows the ARM transactions that are captured in SAS Federation Server. When SQL Logging is enabled, information in each of the transactions is captured. Following the table is a brief explanation of each of the transactions.

Name	Type	Associated Namespace
SESSION	Session transaction	Perf.ARM.FederationServer.Session.Transaction.SESSION
DBC	Database connection	Perf.ARM.SQLServices.Connection.Transaction.DBC

Name	Type	Associated Namespace
DBTRAN	Database transaction	Perf.ARM.SQLServices.Connection.Transaction.DBTRAN
SQL	SQL statement	Perf.ARM.SQLServices.Statement.Transaction.SQL
Prepare	SQL Statement	Perf.ARM.SQLServices.Statement.Prepare
Execute	SQL Statement	Perf.ARM.SQLServices.Statement.Execute
CURSOR	SQL Statement	Perf.ARM.SQLServices.Statement.Transaction.CURSOR
Fetch Scroll	SQL Statement	Perf.ARM.SQLServices.Statement.FetchScroll
SetPos	SQL Statement	Perf.ARM.SQLServices.Statement.SetPos
BulkOps	SQL Statement	Perf.ARM.SQLServices.Statement.BulkOperations
Fetch	SQL Statement	Perf.ARM.SQLServices.Statement.Fetch

- **SESSION:** (Session Transaction) A session transaction starts when a user initiates a server session.
- **DBC:** (Database Connection) A database connection transaction is a child object of the SESSION transaction. A database connection begins when a user connects to a data source and ends when the user disconnects from the data source.
- **DBTRAN:** (RDBMS Transaction) DBTRAN is the actual database transaction. It is a child object of the DBC transaction. A DBTRAN transaction begins with an established driver connection, or when a previous transaction is committed or rolled back, such that a new one begins. DBTRAN records are written to the log only if AUTOCOMMIT is set to OFF. The DBTRAN transaction stops when AUTOCOMMIT is set to ON or when a COMMIT or ROLLBACK command is issued. SQL statements can span DBTRAN transaction boundaries.
- **SQL:** (SQL Statement) SQL is a logical transaction. It encapsulates a series of activities related to one SQL statement. It is a child object of a DBC transaction. An SQL transaction starts when a user issues an SQL statement. Regardless of the statement type (DQL, DML or DDL) the SQL transaction stops when the statement is either closed, or the call to Prepare ends. Subsequent executions of the same statement are recorded under the same SQL transaction, even if the statement is a DQL and the result set associated with it has been is closed.
- **Prepare:** The prepare transaction measures the Prepare phase of an SQL statement. It is a child object of an SQL transaction. The Prepare transaction starts when a user Prepares an SQL statement and stops when the call to prepare returns.
- **Execute:** The execute transaction measures the Execute phase of an SQL statement. It is a child object of the SQL transaction. The EXEC transaction starts when a user executes an SQL statement and stops when the call to execute returns.
- **CURSOR:** CURSOR is a logical transaction. CURSOR is a child object of an SQL transaction and it encapsulates all operations executed in a cursor, including reading, positioning and updates. The CURSOR transaction starts when the Execute transaction finishes. It stops when the cursor is closed. All operations on the same result set belong to the same CURSOR transaction.

- **Fetch/Fetch Scroll:** The FETCH transaction is a child object of the CURSOR transaction. The FETCH transaction has an Execute transaction as its predecessor. It is started when a user issues the first fetch on a result set using Fetch or Fetch Scroll. It stops when the call to Fetch or Fetch Scroll returns.
- **SetPos:** The SetPos transaction is a child object to a CURSOR transaction. The SetPos transaction has an execute transaction as its predecessor. It is started when a user issues a SetPos call and stops when the call returns.
- **BulkOperations :** The BulkOperations transaction is a child object of a CURSOR transaction. The BulkOperations transaction has an Execute transaction as its predecessor. It is started when a user issues a call to BulkOperations and stops when the call returns.

### **The SQL\_LOG Data Service and DSN**

Whether SQL Logging is enabled or disabled on SAS Federation Server, the default configuration will automatically create an SQL\_LOG data service and DSN. When the data service is created, the server also creates an SQL\_LOG transactional database and creates an EVENTS table within the database. This table contains data captured for specific activity in the server, such as information about SQL statements submitted by connected users. The data service, DSN, database and table are always created so that they are available, even if the server is not initially invoked with SQL Logging enabled. As noted above, SQL Logging can be enabled or disabled dynamically at any time, so the server ensures that the SQL Logging constructs are always created when the server starts.

When connected to the SQL\_LOG data service and DSN:

- Catalog functions are restricted so that they only return the SQL\_LOG table. Therefore, only the SQL\_LOG catalog, schema, and table are visible.
- Privileges will restrict access to anything other than the SQL\_LOG table, therefore, only the SQL\_LOG catalog, schema and table are visible.
- The administrator can assign CONNECT privilege on the SQL\_LOG DSN. Federation Server SQL Authorization Enforcement is enabled by default.
- The administrator can assign CONNECT, SELECT, and DELETE privileges only.
- You cannot create new tables or insert into tables through the SQL\_LOG data service.

The valid privileges are CONNECT, SELECT and DELETE.

- The SQL\_LOG data service will allow SELECT or DELETE privileges for the active SQL Logging table and associated columns.
- CONNECT is valid on the SQL\_LOG data service, catalog, and DSN.

Use GRANT, REVOKE, or DENY to set privileges for the SQL\_LOG data service, DSN, catalog, schema, table, and associated columns.

### **The EVENTS Table**

#### **About the EVENTS Table**

The EVENTS table resides in the SQL\_LOG database that is created with the SQL\_LOG data service as a result of enabling SQL Logging. The EVENTS table contains transactional data records written from SAS Federation Server.

Data captured and stored in the EVENTS table includes the number of bytes inserted from an SQL transaction. This does not include any literal data sent to SAS Federation Server. Only data sent through bound memory buffers, such as parameter data, is included. Also, the number of bytes inserted reflects the amount of data stored in the bound memory locations. It does not reflect the size of the data on disk.

### **Managing the EVENTS Table**

The EVENTS table must be managed manually. With SQL Logging enabled, data records are continuously written to the EVENTS table. Therefore, the table increases in size when the server is active and processing requests. SAS Federation Server maintains logging data indefinitely until the table is managed or purged. You can use the SQL console window to issue commands to manage the EVENTS table. By connecting with the SQL\_LOG DSN which has FedSQL dialect enabled, use any SELECT or DELETE statement that is supported by the FedSQL language to manage the table.

#### **CAUTION:**

**Do not change the name of the EVENTS table.**

### **Determine the Size of the EVENTS Table**

Because the size of the EVENTS table increases as activity is logged, you should check the size of the table periodically to determine whether records need to be archived or deleted. Use the following statement to calculate the amount of rows in the EVENTS table:

```
SELECT COUNT(*) FROM EVENTS
```

### **Archiving Data in the EVENTS Table**

As the EVENTS table grows in size, you can move data to another table for archive purposes. This is accomplished by creating a federated DSN to the SQL\_LOG DSN and another data source to use for storing the archived data.

- Use the following statement to move data to a new table:

```
CREATE TABLE
<archive_catalog>.<archive_schema>.<archive_table> AS SELECT *
FROM SQL_LOG.SQL_LOG.EVENTS WHERE <where_condition>
```

- To determine what records are needed, use the WHERE condition with dates:

```
WHERE TRAN_TIMESTAMP > TIMESTAMP '2012-01-25 01:00:00'
```

- To move data to an existing table, use the INSERT statement:

```
INSERT INTO
<archive_catalog>.<archive_schema>.<archive_table> SELECT *
FROM SQL_LOG.SQL_LOG.EVENTS WHERE <where_condition>
```

### **Deleting Data in the EVENTS Table**

Use the DELETE statement to remove outdated records that are no longer needed or have been archived. Here is an example:

```
DELETE FROM SQL_LOG.SQL_LOG.EVENTS
WHERE TRAN_TIMESTAMP < TIMESTAMP '2011-04-11 12:00:00'
```

## Columns and Data Types

The following table presents the columns and associated data types that reside in the EVENTS table.

**Table 4.1** Column and Data Types of the EVENTS Table

Column	Data Type	Description
APP_HANDLE	VARCHAR(36)	A unique ID that is associated with an application instance.
APP_ID	VARCHAR(36)	The application ID.
APP_NAME	VARCHAR(50)	The application name.
ARM_NAMESPACE	VARCHAR(256)	The logger name (namespace) of the logging event.
AUTHORIZATION_ID	VARCHAR(128)	UUID for the user from the Authentication Server.
AUTHORIZATION_NAME	VARCHAR(128)	The user name from the Authentication Server.
BYTES_FETCHED	BIGINT	The size of data read in bytes.
BYTES_INSERTED	BIGINT	The size of data inserted in bytes.
CACHE_VIEW_CATALOG	VARCHAR(256)	The cache view catalog name.
CACHE_VIEW_NAME	VARCHAR(256)	The cache view name.
CACHE_VIEW_SCHEMA	VARCHAR(256)	The cache view schema name.
CATALOG_NAME	VARCHAR(256)	The catalog name.
CLIENT_CORRELATOR	VARCHAR(56)	The transaction's client correlator (base64 encoded).
CONNECTION_DRIVER	VARCHAR(25)	The driver that was used for the connection. For example, FEDSQL, ORACLE, TERADATA, ODBC, MYSQL, DB2, and others.
CONNECTION_NAME	VARCHAR(256)	The expanded connection string.
CONNECTION_TRAN_HANDLE	VARCHAR(36)	The DBC transaction under which the current transaction is assigned to.
CURR_SAS_TIMEOFDAY	DOUBLE PRECISION	The current time-of-day for the ARM event.
CURR_SYSTEM_CPU_TIME	DOUBLE PRECISION	The process current system CPU time for the ARM event.

Column	Data Type	Description
CURR_USER_CPU_TIME	DOUBLE PRECISION	The process current user CPU time for the ARM event.
CURRENT_CORRELATOR	VARCHAR(56)	The transaction's correlator (base64 encoded).
CURSOR_TRAN_HANDLE	VARCHAR(36)	The CURSOR transaction under which the current transaction is assigned to.
DBTRAN_STATE	VARCHAR(15)	The state of the current transaction, such as OPEN, CLOSED.COMMIT, CLOSED.ROLLBACK.
DBTRAN_TRAN_HANDLE	VARCHAR(36)	The UUID that points to the driver's DBTRAN transaction handle, which is not its parent (the CONNECTION handle is parent) since the SQL can span multiple DBMS transactions.
EVENT_SEQUENCE	BIGINT	A unique value associated with the ARM record. Values increase for each record, usually with an increment of 1 (database-specific).
EXEC_PARAM_DATA	VARCHAR(1024)	The XML format for encoding parameter array data.
GROUP_NAME	VARCHAR(128)	The group name of the application instances.
IO_COUNT	BIGINT	The total number of process disk, tape or related input and output operations for the transaction event.
IP_ADDRESS	VARCHAR(48)	The IP address of the client.
LOGIN_ID	VARCHAR(128)	The current user ID that is associated with the transaction.
MEM_CURRENT	BIGINT	The current process memory utilization for the transaction event.
MEM_HIGH	BIGINT	The highest amount of process memory used for the transaction event.
OBJECT_NAME	VARCHAR(256)	The object name used in the SQL statement.
OBJECT_TYPE	VARCHAR(60)	The type of object that was accessed.
PARENT_CORRELATOR	VARCHAR(56)	The transaction's parent correlator (base64 encoded).
PREDECESSOR_TRAN_HANDLE	VARCHAR(36)	The driver's Execute transaction handle. This ties the Fetch transaction to an execution and its metrics.

Column	Data Type	Description
ROWS_DELETED	BIGINT	The number of rows deleted.
ROWS_FETCHED	BIGINT	The number of rows read.
ROWS_INSERTED	BIGINT	The number of rows inserted.
ROWS_UPDATED	BIGINT	The number of rows updated.
SCHEMA_NAME	VARCHAR(256)	The name of the schema being accessed.
SERVER_MESSAGE	VARCHAR(500)	The message associated with the event.
SESSION_TRAN_HANDLE	VARCHAR(36)	The SESSION transaction that the current transaction is assigned to.
SOURCE_FILE_NAME	VARCHAR(128)	The filename where the logging request was issued.
SQL_DIALECT	VARCHAR(15)	The dialect that is being used: FEDSQL or NATIVE.
SQL_TRAN_HANDLE	VARCHAR(36)	The SQL transaction that the current transaction is assigned to.
STATEMENT_ID	BIGINT	The SQL statement hash. The value derived from the SQL statement content.
STATEMENT_NAME	VARCHAR(256)	The SQL statement name.
STATEMENT_PLAN	VARCHAR(15000)	Column valued for SQL statement types only. <i>Note:</i> The plan value can truncate if the character limit is exceeded. For example, Oracle has a VARCHAR limit of 4000 while SQL Server is 8000.
STATEMENT_STATE	VARCHAR(15)	The state of the SQL statement, such as S0, S1 and S2.
STATEMENT_TEXT	VARCHAR(15000)	The text of the SQL statement. <i>Note:</i> The plan value can truncate if the character limit is exceeded. For example, Oracle's VARCHAR limit is 4000 while SQL Server is 8000.
STATEMENT_TYPE	VARCHAR(15)	The type of SQL statement, such as DQL, DQL.Metadata (catalog methods), DML, or DDL. Empty if unknown.
THREAD_CURRENT	BIGINT	The current process thread count for the transaction event.

Column	Data Type	Description
THREAD_HIGH	BIGINT	The process highest thread count for the transaction event.
TRAN_CLASS_ID	VARCHAR(36)	The UUID of transaction class.
TRAN_HANDLE	VARCHAR(36)	The UUID of transaction instance.
TRAN_NAME	VARCHAR(50)	SESSION, DBC, DBTRAN, SQL, Execute, Fetch, CURSOR. For additional information see <a href="#">“ARM Transactions”</a> on page 57.
TRAN_START_SAS_TIMEOFDAY	DOUBLE PRECISION	The time-of-day value for the current transaction start event.
TRAN_STATE	VARCHAR(15)	The state of the transaction: START, STOP, UPDATE, BLOCK, UNBLOCK, DISCARD.
TRAN_STATUS	VARCHAR(15)	The transaction status: UNKNOWN, ABORTED, GOOD, FAILED, STOP.
TRAN_STOP_SAS_TIMEOFDAY	DOUBLE PRECISION	The time-of-day value for the current transaction stop event.
TRAN_TIMESTAMP	TIMESTAMP	The current timestamp of the transaction.
TRANRESP_SYS_CPU_TIME	DOUBLE PRECISION	The calculated system CPU time for the duration of the transaction.
TRANRESP_TIME	DOUBLE PRECISION	The calculated elapsed time for the duration of the transaction.
TRANRESP_USER_CPU_TIME	DOUBLE PRECISION	The calculated user CPU time for the duration of the transaction.
TRANSTART_SYS_CPU_TIME	DOUBLE PRECISION	The process system CPU time for the current transaction start event.
TRANSTART_USER_CPU_TIME	DOUBLE PRECISION	The process user CPU time for the current transaction start event.
TRANSTOP_SYS_CPU_TIME	DOUBLE PRECISION	The process system CPU time for the current transaction stop event.
TRANSTOP_USER_CPU	DOUBLE PRECISION	The process user CPU time for the current transaction stop event.

## SQL Logging Performance Tuning

### Introduction

You can configure SQL Logging so that it will have minimum impact on SAS Federation Server performance, especially during periods of heavy processing.

### Using Maximum Buffered Events (*MaxBufferedEvents*)

Maximum buffered events tells the server how many events to buffer before writing them to the EVENTS database. You can set this option using the `MaxBufferedEvents` option in the SQL Logging configuration file. The default setting for `MaxBufferedEvents` is 100 meaning that 100 events are buffered before writing them to the EVENTS table. Setting `MaxBufferedEvents` to a higher number might consume server resources but will show improved performance during periods of heavy processing. A sustainable high value for `MaxBufferedEvents` is approximately 500. Here is the syntax used in the SQL Logging configuration file:

```
<param name="maxBufferedEvents" value="100" />
```

### Configure Indexing on the EVENTS Table

When using the default configuration of the TRAN data store, writes to the SQL Logging database can sometimes affect overall performance, especially during periods of heavy processing. This is caused by the indexes that are present on the EVENTS table. During these high activity periods, you can drop indexes onto the EVENTS table by updating the SQL\_LOG data service as shown here:

```
alter service SQL_LOG {options AUTOINDEX off}
```

To activate indexing on the EVENTS table:

```
alter service SQL_LOG {options AUTOINDEX on}
```

Because indexes are required to process SQL queries in SAS Federation Server Manager, it is not advisable to view SQL Logging information using SAS Federation Server Manager during periods when indexes are not active. You can re-enable indexes after processing activity decreases on SAS Federation Server. At that time, you should be able to view SQL Logging information in SAS Federation Server Manager.

An alternative to dropping indexes is to configure SQL Logging to use an external, or third party, database. For more information see [“Configuring a 3rd Party DBMS for SQL Logging” on page 57](#).

## Server Logging Configuration

### Introduction

The logging facility is a flexible, configurable framework used for collecting, categorizing, and filtering events that are generated by processes and writing events to a variety of output devices. The `dfs_log.xml` configuration file controls the destination, contents, and formats of the logging facility log for SAS Federation Server. The configuration file specifies options for loggers, appenders, levels, filters, and the layout of log messages. You can change logging levels dynamically without stopping the server.

## Initial Logging Configuration

The default logging facility configuration for SAS Federation Server includes a definition for the RollingFileAppender. The appender routes events to a rolling log file.

The rolling log file is configured as follows:

- A new log is created when the date changes and when a new server process is started.
- Events are written by using a layout that includes the current date, current time, logging level, process ID, the user identity that is associated with the event, and a message.
- The name of the rolling log file follows the following convention:  

```
dfs_%d_%S{pid}.log
```

 where `%d` is the date and `%S{pid}` is the process ID number (PID) for SAS Federation Server.
- The rolling log files are placed in the `/var/log` directory.
- When a new rolling log file is created, a heading is written to the file. The heading identifies the server's host machine, operating system, and server start-up command.

The following table lists the loggers that reference the RollingFileAppender:

Logger Name	Logging Level	Description
Admin	Info	processes log events that are relevant to system administrators or computer operators.
App	Info	processes log events that are related to specific applications. For example, metadata servers, OLAP servers, stored process servers, and workspace servers use loggers that are named <code>App.class.interface.method</code> to record method calls that are issued to the server.
App.Server	Info	Server top-level object run-time and interface events.
	Debug	Method call and return events.
	Trace	Method parameters.
App.Session	Info	Session object run-time and interface events.
	Debug	Method call and return events.
	Trace	Method parameters.
App.Connection	Info	Connection object run-time and interface events
	Debug	Method call and return events.
	Trace	Method parameters.

Logger Name	Logging Level	Description
App.Statement	Info	Statement object run-time and interface events
	Debug	Method call and return events.
	Trace	Method parameters.
App.Program	Info	General application independent events including errors and warnings from arbitrary services or the OS
Audit	Info	Processes log events to be used for auditing. These events include updates to public metadata objects, user access to SAS libraries, accepted and rejected user authentication requests, and administration of users, groups, and access controls.
Audit Authentication	Info	Authentication provider events.
Audit Table	Info	Federation server specific events.
Audit Table Connection	Info	Audit events related to server connections, including connection pooling and dynamic connections.
Audit.Table.Security	Info	Federation Server authorization events.
Audit.Table.Security.Provider	Info	Detailed authorization services run-time events relating to user identity management and access control logic and enforcement decisions.
Logging	Error	SAS logging facility configuration and run-time events.
Logging.Appender	Error	Appender-specific configuration and run-time events.
Logging.Appender.DB	Error	DB Appender-specific events (used in SQL Logging).
Cradle	Info	General server process framework services, start-up and termination events.
IOM	Info	Processes log events for servers that use the Integrated Object Model (IOM). The IOM interface provides access to Foundation SAS features such as the SAS language, SAS libraries, the server file system, results content, and formatting services. IOM servers include metadata servers, OLAP servers, stored process servers, and workspace servers.
IOM.Proxy	Info	Server to server outcall events.
	Debug	Method call and return events.
	Trace	Method parameters.
IOM.PE	Warn	Communication protocol engine events.

Logger Name	Logging Level	Description
Perf	Error	Processes log events that are related to system performance.
Perf.ARM	Error	Application Response Measurement performance events.
Perf.ARM.IOM.Session	Error	Session API performance events.
Perf.ARM.IOM.Environment	Error	Environment API performance events.
Perf.ARM.IOM.Connection	Error	Connection API performance events.
Perf.ARM.IOM.Statement	Error	Statement API performance events.
Perf.ARM.FederationServer	Warn	Federation server API independent performance events.
Perf.ARM.SQLServices	Warn	Local SQL services performance events.
<root>	Error	All events produced from SAS® Federation Server.

## Logging Thresholds

The SAS logging facility provides six thresholds: TRACE, DEBUG, INFO, WARN, ERROR, and FATAL. Thresholds are used to ignore log events that are lower than a particular level, or to filter messages so that only a single message level is logged.

When a log event occurs, up to three levels of filtering can take place:

1. filtering log events by comparing the log event level to the log event's logger level.
2. filtering log events by comparing the log event level to the appender's threshold.
3. filtering log events by comparing the log event level to the threshold that is specified in the filter definition, which is a part of the appender configuration.

In the first two cases, if the log event level is lower than the logger or appender threshold, the logging facility ignores the log event. Otherwise, processing of the log event continues.

In the third case, the log event level is compared to the filter threshold. If there is a match, the log event can be either accepted or denied. If there is no match, the filtering process continues to the next filter in the filtering policy.

The logging levels, from the lowest to the highest, are as follows:

Level	Description
TRACE	Produces the most detailed information about your application. This level is primarily used by SAS Technical Support or development.
DEBUG	Produces detailed information that you use to debug your application. This level is primarily used by SAS Technical Support or development.

Level	Description
INFO	Provides information that highlights the progress of an application.
WARN	Provides messages that identify potentially harmful situations.
ERROR	Provides messages that indicate that errors have occurred. The application might continue to run.
FATAL	Provides messages that indicate that severe errors have occurred. These errors will probably cause the application to end.

*Note:* The logging level must be enclosed in quotation marks.

By default appenders do not have a threshold but a threshold can be configured. When set, all log events that have a level lower than the threshold are ignored by the appender.

### **Modifying the Server Logging Configuration**

You can modify the logging facility configuration for SAS Federation Server by modifying the `dfs_log.xml` file located in the `/etc` directory of the installation path. Before modifying the file, be sure to make a backup copy.

Here are some examples of changes that you might want to make:

- configure `RollingFileAppender` to use a different log filename, to roll over log files more or less frequently, or to roll over log files based on file size rather than date.
- specify additional appenders.
- use filters to limit the events that are written to an appender.
- configure a different message layout for an appender.

### **Trace Log**

By tracing each internal API routine that is called by the application, a trace log records transactions that can be used for debugging connection and processing issues. For example, you can request information that traces the FedSQL statements that are submitted to a data service.

*Note:* By default, tracing is not activated for server logging. You should not activate tracing unless you are instructed to do so by Technical Support.

Tracing can be activated by using the following methods:

- connection string options
- server start-up options
- data service connection arguments
- system options

When you activate tracing, you also specify the physical location where the transaction records are saved. Because SAS Federation Server supports one root file trace log directory and multiple subdirectories, you can group trace logs if necessary.

---

## Utilities for SAS Federation Server

### Introduction

SAS Federation Server contains several utilities that assist with management of your server environment. Utilities include database migration, password encryption and system database backup and restore. Utilities are available for the UNIX and Windows operating systems.

### UNIX Utilities

#### **dfsadmin Utility**

SAS Federation Server for UNIX contains the **dfsadmin** utility, located in the **bin** directory of the installation root. Run any of the commands in **dfsadmin** using the following syntax: `./bin/dfsadmin yourcommand` where *yourcommand* is one of the following commands:

Command	Description
<b>start</b>	Starts SAS Federation Server. Example: <code>./bin/dfsadmin start</code>
<b>stop</b>	Stops SAS Federation Server. Example: <code>./bin/dfsadmin stop</code>
<b>status</b>	Checks the run status of SAS Federation Server.
<b>migrate</b>	Migrate the Federation Server database from an older version options: <code>&lt;old_install_dir&gt;</code>
<b>crypt</b>	Print the encrypted version of a password options: <code>[&lt;password&gt;]</code> For additional information see Password Encryption.
<b>help</b>	Displays command usage and options.
<b>version</b>	Displays version information.

#### **Password Encryption**

The **dfsadmin** utility contains a **crypt** command from where you can key the password on the command line in plain text, or let it prompt you so that the password is not echoed to terminal. The password is encrypted using the encryption type installed on the server.

Here is an example:

```
./bin/dfsadmin crypt my_passwd
```

Key a password with the **crypt** command:

```
[admin@reason fedserver]$ ./bin/dfsadmin crypt my_passwd
{SAS002}1950043249132FE80CEC515733C508D5
```

Invoke the **crypt** command and let it prompt for a password:

```
[admin@reason fedserver]$ ./bin/dfsadmin crypt
Enter password:
Confirm password:
Encrypted value: {SAS002}1950043249132FE80CEC515733C508D5
```

*Note:* Enclose the encrypted password in single quotation marks when using it in a connection string.

### **Database Backup and Restore**

It is recommended that you backup the Federation Server databases periodically, especially the system catalog, SYSCAT.tdb. A utility named **dfsutil** is provided with SAS Federation Server installation and is located in **/bin** of the installation directory. With **dfsutil**, dynamic database backups can be performed without disruption to server operations. You can back up the system database (SYSCAT.tdb) and other databases such as the SQL Logging database (SQL\_Log). However, the SQL Logging database cannot be backed up with **dfsutil** if it was changed to a third-party data store. You can also restore databases using **dfsutil** as long as the database was backed up using the same utility.

*Note:* If you backup a database with **dfsutil**, then you are required to restore that database using **dfsutil**.

The procedures below describe how to use **dfsutil** to perform a backup or restore for a SAS Federation Server database. Note that you can run **dfsutil** while the server is running.

### **Backup**

To backup a database, use the **-db** parameter with the **dfsutil** command. Note that you can run this command while the server is running. For example:

```
install-path/bin/dfsutil
backup -db syscat
/path/to/backup
install-path/bin/dfsutil
backup -db sql_log
/path/to/backup
```

### **Restore**

To restore a database, use the **-db** parameter with the **dfsutil** command. For example:

```
install-path/bin/dfsutil
restore -db syscat
/path/to/backup
install-path/bin/dfsutil
restore -db sql_log
/path/to/backup
```

## **Windows Utilities**

### **Password Encryption**

Use the **dfs\_crypt** utility to encrypt passwords. **dfs\_crypt**, is located in the **/bin** directory of the installation path for SAS Federation Server.

- Navigate to **[drive:]\Program Files\SASHome\FedServer\server1\bin>** and enter: **dfs\_crypt password**

- An encrypted password is returned: `{SAS002}A984D9084D59A9EA`

*Note:* Enclose the encrypted password in single quotation marks when using it in a connection string.

### **Database Backup and Restore**

It is recommended that you backup the Federation Server databases periodically, especially the system catalog, SYSCAT.tdb. A utility named `dfsutil` is provided with SAS Federation Server installation and is located in `/bin` of the installation directory. With `dfsutil`, dynamic database backups can be performed without disruption to server operations. You can back up the system database (SYSCAT.tdb) and other databases such as the SQL Logging database (SQL\_Log). However, the SQL Logging database cannot be backed up with `dfsutil` if it was changed to a third-party data source. You can also restore databases using `dfsutil` as long as the database was backed up using the same utility.

*Note:* If you backup a database with `dfsutil`, then you are required to restore that database using `dfsutil`.

The procedures below describe how to use `dfsutil` to perform a backup or restore for a SAS Federation Server database. Note that you can run `dfsutil` while the server is running.

#### **Backup**

To backup a database, use the `-db` parameter with the `dfsutil` command. Note that you can run this command while the server is running. For example:

```
install-path\bin\dfsutil
backup -db syscat
\path\to\backup
install-path\bin\dfsutil
backup -db sql_log
\path\to\backup
```

#### **Restore**

To restore a database, use the `-db` parameter with the `dfsutil` command. For example:

```
install-path\bin\dfsutil
restore -db syscat
\path\to\backup
install-path\bin\dfsutil
restore -db sql_log
\path\to\backup
```

---

## **Uninstalling SAS Federation Server**

The uninstallation program for SAS Federation Server removes all directories and configuration files from the installation root. This includes the Federation Server Database which contains the system catalog. Before uninstalling SAS Federation Server, you should:

1. Install and configure a new instance of SAS Federation Server.
2. Migrate the federation server database.

3. Move or copy unique files, including configuration, from the current server to the new server instance.

*Note:* To preserve your existing configuration files, copy them from the `/etc` directory before performing the uninstallation.

The following items are deleted when you uninstall SAS Federation Server:

- The transactional database named `SYSCAT.TDB`, which is located in the `/var` directory of the installation path by default. `SYSCAT.TDB` contains the system tables for SAS Federation Server.
- The SQL Logging database named `SQL_LOG.TDB` which is also located in the `/var` directory of the installation path by default. `SQL_LOG.TDB` contains monitoring data for SQL Logging.
- Configuration files especially if modifications were made after installation. The configuration files are located in the `/etc` directory and the `/etc/migrate` directory of the installation path. The configuration filenames are listed in the table below.
- Log files that reside in the `/var/log` directory of the installation path.

**Table 4.2** Federation Server Configuration Files

Configuration Files in the <code>/etc</code> Directory:	Configuration Files in the <code>/etc/migrate</code> Directory:
<code>dfs_log.xml</code>	<code>dfs_log_213_310.xml</code>
<code>dfs_log_SQL_Logging.xml</code>	<code>dfs_serv_213_310.xml</code>
<code>dfs_serv.xml</code>	<code>dfs_serv_213_310_sql.xml</code>
<code>dfs_serv_common.xml</code>	
<code>dfs_entities.dtd</code>	



## Chapter 5

## SAS Federation Server Security

---

<b>Overview</b> .....	<b>76</b>
<b>Authentication</b> .....	<b>76</b>
<b>Authorization</b> .....	<b>76</b>
Overview .....	76
Data Source Authorization .....	76
SAS Federation Server Authorization .....	77
<b>Permissions</b> .....	<b>78</b>
Overview .....	78
Permission Types .....	78
Administration Permissions .....	80
SQL Permissions .....	80
<b>Object Privilege Inheritance</b> .....	<b>81</b>
About Object Privilege Inheritance .....	81
Object Privilege Summary .....	82
User and Group Privileges .....	83
Maintaining Security Definitions for Tables, Views, or Columns .....	84
Granting Privileges .....	84
Privilege Caching .....	85
<b>Row-Level Security</b> .....	<b>86</b>
Introduction .....	86
Row-Level Security Privilege Assembly Rules .....	87
The RLS Library and Library Reference .....	88
<b>Data Masking</b> .....	<b>92</b>
Overview .....	92
About the SYSCAT.DM.MASK Function .....	92
Data Masking Rule Types and Arguments .....	93
<b>Server Encryption</b> .....	<b>96</b>
Introduction .....	96
SAS Proprietary Encryption .....	96
DataFlux Secure .....	97
Password Encryption .....	97

---

## Overview

Properly configured security for SAS Federation Server ensures that both the server and its data are secure. Data is protected against unauthorized access, and can be guaranteed secure transmission lines for transferring data. The security features are flexible with respect to the types and amount of security, and both security setup and maintenance are easily managed.

---

## Authentication

SAS Federation Server must be configured to use an Authentication Server. When a user connects to SAS Federation Server to access data, the user's authenticating credentials are passed to the Authentication Server for validation. Once the credentials are validated, the Authentication Server will identify the user based on the submitted credentials. SAS Federation Server can then make requests to the Authentication Server for information about the user, including logins and group membership. If a user is authenticated but cannot be identified in the Authentication Server, that user becomes a member of the PUBLIC group. All users that are identified in the Authentication Server are members of the USERS group.

Only individual user objects and shared logins own outbound accounts. Groups, including USERS and PUBLIC, cannot own accounts, including shared logins. If a user is SYSTEM or ADMINISTRATOR, their personal credentials are used to authenticate. SYSTEM and ADMINISTRATOR accounts are discussed further in [“Configuring Server Accounts” on page 26](#).

It is assumed that the Authentication Server has been installed, configured and populated with user and group information before running SAS Federation Server. For more information, see the *Authentication Server Administrator's Guide*.

---

## Authorization

### Overview

Authorization determines what privileges a user or group object contains in order to gain access to resources and associated data sources. Because SAS Federation Server requires access to underlying data sources, authorization can happen at two distinct locations:

- Data Source Authorization
- Federation Server Authorization

### Data Source Authorization

An example of data source authorization would be an Oracle database, which provides its own layer of security for its data. Data source authorization cannot be bypassed by SAS Federation Server. For example, if an Oracle administrator denies privileges to a user on table T1, then that user will always be denied access to table T1, no matter what

privileges are set in SAS Federation Server. SAS Federation Server authorizations are more restrictive on the underlying data.

## **SAS Federation Server Authorization**

### **About Authorization**

SAS Federation Server authorizations are applied to all administration DDL. Most administration DDL is performed by an administrator only (defined as a user who has the ADMINISTER privilege on SAS Federation Server), but some commands, such as CREATE CACHE, have specific privileges which can be assigned to users and groups.

In the case where a user is connected to data sources providing customer data, SAS Federation Server authorizations are applied over the underlying data source. SQL statements submitted to the server are first parsed and then evaluated against the privileges defined in SAS Federation Server. If the action is not permitted from SAS Federation Server, an error is returned to the user, and no SQL is sent to the underlying data source. If the action is permitted, the SQL statement is evaluated, and the FedSQL processor determines what SQL should be sent to the underlying data sources. In summary, if the underlying data source does not permit the SQL action, then an error is returned. Otherwise, the SQL action is performed and results sent back to the user.

For example, an administrator can configure the server so that a particular user cannot access table T1 even if the underlying data source allows it. So SAS Federation Server authorizations can be used to restrict the type of activity that an administrator wants to allow on the server.

SAS Federation Server authorizations are also very powerful when used in conjunction with shared logins. Shared logins allow many users to be mapped to the same single login for an underlying data source. This allows for easy back-end data source user management, since each user of SAS Federation Server does not require an individual login. However, this alone would mean that all users of that shared login would have the same privileges to the accessible data. However, SAS Federation Server authorization can be used to restrict individual access to data, no matter what the shared login is allowed to access in the underlying data source.

As with other system metadata, SAS Federation Server authorization process uses an internal database to store security definitions for users, groups and objects. Privileges can be set on individual users, or on groups, which affect all members of the group. By default, no users (except those defined as system users) are granted any specific privileges on any objects in SAS Federation Server, and the lack of any privilege anywhere results in a DENY from the server's authorization subsystem. The administrator must specifically grant privileges before a user can perform any actions through SAS Federation Server.

Refer to Technical Support, Object Level Security for best practices related to authorization.

### **Case Sensitivity**

When security definitions are stored in SAS Federation Server, they are stored as entered in the GRANT or DENY SQL syntax. Using the following statement as an example:

```
GRANT SELECT ON TABLE "MyTable" TO BOB
```

SAS Federation Server creates a system table object for table "MyTable". When performing privilege comparisons at run time (for example, when the user Bob issues a SELECT statement against table "MyTable"), the server will perform case sensitive or insensitive comparisons, depending on the data source. Case sensitivity from the data

source is registered during the CREATE DATA SERVICE DDL statement, and the system automatically determines the correct setting. However, the administrator can specify case sensitivity as well via the CASE\_SENSITIVITY option in the CREATE DATA SERVICE DDL. It is highly recommended that the server default to the correct value. Specifying an incorrect value for case sensitivity can result in incorrect privilege determination.

Case sensitivity settings only apply to relations and columns. Data services, catalogs, schemas and DSNs are always treated in a case-insensitive manner. Also, user and group identifiers are treated as case insensitive.

---

## Permissions

### Overview

At times the terms permissions and privileges are used interchangeably. To clarify, permissions represent a specific ability while privileges are a combination of applying the permission to a user and an object. For example, a user's privileges include all the permissions that were granted on various objects in SAS Federation Server.

### Permission Types

Federation Server permissions are divided into two general types:

- Administration permissions
- SQL permissions

The following table provides a description for each of these permissions.

**Table 5.1** SAS Federation Server Permissions

Permission	Description
ADMINISTER	Controls the ability to configure the server. This privilege can be set on the server only. Users who are granted ADMINISTER privilege are automatically granted all other privileges.
ALTER TABLE	Controls the ability to add or drop columns in a table or create or drop indexes with the ALTER TABLE statement. The authorization is set on the server, data service, catalog, and schema objects.
CONNECT	Controls the ability of the user to connect, using either the DSN or data service. The authorization is set on the server, data service, and DSN.
CREATE TABLE	Controls the ability to create new tables or views with the CREATE TABLE statement. The authorization is set on the server, data service, catalog, and schema objects.
DELETE	Controls the ability to delete data with the DELETE statement or analogous method call. The authorization is set on the server, data service, catalog, and schema objects.
DROP TABLE	Controls the ability to remove tables with the DROP TABLE statement. The authorization is set on the server, data service, catalog, and schema objects.

Permission	Description
INSERT	Controls the ability to add data with the INSERT statement or analogous method call. The authorization is set on the server, data service, catalog, schema, table, and column objects.
REFERENCES	Controls the ability to create a foreign key reference to an existing table. The authorization is set on the server, data service, catalog, schema, table, and column objects.
SELECT	Controls the ability to retrieve data with the SELECT statement. The authorization is set on the server, data service, catalog, schema, table, and column objects.
TRACE	Controls the ability of the user to enable tracing and create trace files. This privilege can be set on the server only.
UPDATE	Controls the ability to modify data with the UPDATE statement or analogous method call. The authorization is set on the server, data service, catalog, schema, table, and column objects.
CREATE DSN	Controls the ability of the user to create a DSN. The authorization can be set on the server or data service.
ALTER VIEW	Controls the ability to alter a definer's rights or invoker's rights view. The authorization can be set on the server, data service, catalog or schema objects. <sup>1,2,3</sup>
CREATE VIEW	Controls the ability to create a definer's rights or invoker's rights view. The authorization can be set on the server, data service, catalog or schema objects. <sup>1</sup>
DROP VIEW	Controls the ability to drop a definer's rights or invoker's rights view. The authorization can be set on the server, data service, catalog or schema objects. <sup>2</sup>
CREATE CACHE	Controls the ability to create, drop and refresh a cache from a definer's rights view. The authorization can be set on the server, data service, catalog or schema objects, and on individual definer's rights views.
ALTER CACHE	Controls the ability to enable, disable and refresh a cache. The authorization can be set on the server, data service, catalog or schema objects, and on individual definer's rights views.
CREATE TABLESPACE	Allows a user to create tables in a schema that they do not own to implement a data cache operation. CREATE TABLESPACE applies to data cache operations. The authorization can be set on the server, data service, catalog or schema objects.
ALTER	Controls the ability to alter a view or a table. The authorization can be set on the table or view object. <sup>1,2</sup>
DROP	Controls the ability to drop a view or a table. The authorization can be set on the table or view object. <sup>2</sup>

Permission	Description
	<p><sup>1</sup> Because definer's rights views effectively allow a user to impersonate the schema owner, the creation of definer's rights views requires special consideration. It can be desirable to grant CREATE VIEW privilege to a user, but only for the intention of creating invoker's rights views. Likewise, if ALTER on a view allows switching invoker's rights views to definer's rights views, then that privilege is quite powerful as well. Only the schema owner can create definer's rights views, along with the system user and administrator, who have all privileges. In addition, only the schema owner can issue an ALTER VIEW command to change an invoker's rights view or a definer's rights view. Therefore, a grant ALTER VIEW or ALTER on a view only allows a user to change the view name.</p> <p><sup>2</sup> If the ALTER or DROP permission is explicitly set on the object, that privilege definition is honored. If the privilege is not set, the following applies: The proper ALTER or DROP privilege is searched on the container object. If the Alter privilege is on a table, ALTER TABLE is searched. If the privilege is on a view, the search target is ALTER VIEW. This applies to the DROP privilege as well.</p> <p><sup>3</sup> ALTER VIEW is required to execute the DESCRIBE VIEW FedSQL statement. See the <i>SAS FedSQL Reference Guide</i> for additional information.</p>

### Administration Permissions

The administration permissions are:

- ADMINISTER
- CONNECT
- CREATE DSN
- CREATE CACHE
- ALTER CACHE
- CREATE TABLESPACE
- TRACE

Administration permissions are always enforced, regardless of the 'Federation Server SQL Authorization Enforcement' setting for any DSN.

### SQL Permissions

The SQL privileges are:

- SELECT
- UPDATE
- INSERT
- DELETE
- REFERENCES
- ALTER TABLE
- ALTER VIEW
- DROP TABLE
- DROP VIEW
- CREATE TABLE
- CREATE VIEW

The enforcement of these privileges to users or groups is controlled through the 'Federation Server SQL Authorization Enforcement' option in the DSN. When that setting is disabled, a user connecting with that DSN can perform any SQL action on the data, regardless of what is defined in SAS Federation Server for that user. When the setting is enabled, SAS Federation Server privilege definitions are enforced for that user on that connection. The standard setting for most users will be to enable the SQL authorization enforcement. It will typically only be disabled on DSNs where certain privileged users are granted CONNECT privilege.

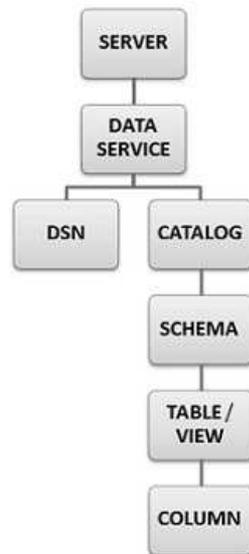
---

## Object Privilege Inheritance

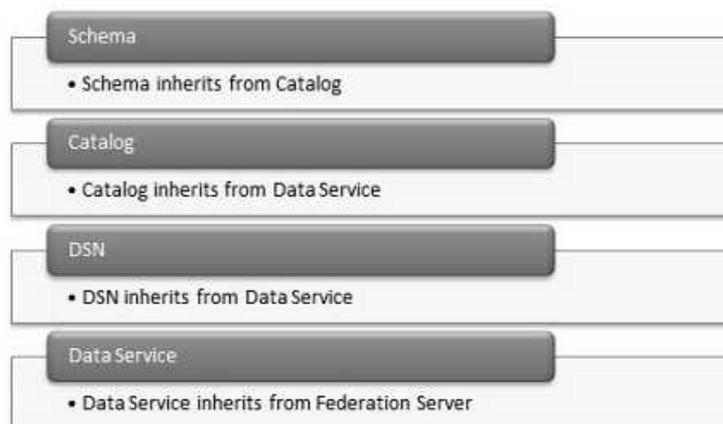
### About Object Privilege Inheritance

SAS Federation Server contains an inherent hierarchy of objects, in the following order:

**Figure 5.1** SAS Federation Server Privilege Inheritance



Where privileges on the server are inherited by the data service, privileges on the data service are inherited by the DSN and catalog, and so on. This inheritance hierarchy allows an administrator to set general security rules on higher level objects, and then only set exceptions on the more specific (subordinate) objects.

**Figure 5.2** SAS Federation Server Container Object Inheritance

For example, there is a group called SALES\_GROUP whose members are allowed to select most objects in the SALES\_DATA data service. An administrator can assign SELECT privilege on the SALES\_DATA data service to the SALES\_GROUP. The SELECT privilege is inherited on all the catalogs of the SALES\_DATA data service, and all the schemas of those catalogs, and all the related tables and views. There is a stipulation that the SALES\_GROUP is not allowed to see any data in a single catalog of the SALES\_DATA data service called EXECUTIVE\_DATA. An administrator could then deny SELECT privilege to the SALES\_GROUP on that particular catalog. Members of the SALES\_GROUP would then not be able to select any data from the EXECUTIVE\_DATA catalog or any of its schemas. An administrator can elect to grant all privileges on the EXECUTIVE\_DATA catalog to the EXECUTIVE\_GROUP. An administrator can also deny SELECT privilege to any member of the EXECUTIVE\_GROUP on any subordinate object of the EXECUTIVE\_DATA catalog. In this way, general authorizations are defined on higher level objects, and then exceptions are set on subordinate objects. This minimizes the number of privileges that an administrator must establish, and thereby reduces administration overhead. For any object in the hierarchy of SAS Federation Server, an administrator can request information about privileges held by any user or group, including where in the hierarchy the privilege was set and who the grantee of the privilege is, which can be a group that the user is a member of, directly or indirectly.

### Object Privilege Summary

The following table summarizes SAS Federation Server objects with their associated privileges. If privileges are inherited, the field is marked as 'Yes'. The fields are blank if there is no privilege inheritance for the object. Inheritance runs from right to left.

Object / Privilege	Column	Row	Table/ View	Schema	Catalog	DSN	Data Service	Federation Server
SELECT	Yes	Yes	Yes	Yes	Yes		Yes	Yes
UPDATE	Yes		Yes	Yes	Yes		Yes	Yes
INSERT	Yes		Yes	Yes	Yes		Yes	Yes
REFERENCES	Yes		Yes	Yes	Yes		Yes	Yes

Object / Privilege	Column	Row	Table/ View	Schema	Catalog	DSN	Data Service	Federation Server
DELETE	Yes		Yes	Yes	Yes		Yes	Yes
ALTER TABLE/ VIEW TABLE			Yes	Yes	Yes		Yes	Yes
DROP TABLE/ VIEW			Yes	Yes	Yes		Yes	Yes
CREATE TABLE/ VIEW				Yes	Yes		Yes	Yes
CREATE/ ALTERCACHE			Yes	Yes	Yes		Yes	Yes
CREATE TABLESPACE				Yes	Yes		Yes	Yes
ADMINISTER								Yes
TRACE								Yes
CREATE DSN <sup>1</sup>							Yes	Yes
CONNECT <sup>2</sup>						Yes	Yes	Yes

*Note:* <sup>1</sup>Special DSN inheritance applies to CREATE DSN where privileges are checked for inheritance on the data service first and then on the server. <sup>2</sup>For the CONNECT privilege, privileges are first checked on the DSN, and then checked for inheritance on the data service and finally, on the server.

### User and Group Privileges

Privileges can be assigned to a user or a group on any object. Group privileges are granted and denied in the same manner as individual users. Users who are members of a group inherit the privileges from the group unless explicitly denied in the individual user account. Also, privileges set on a group that is “closer” to a user take precedence over privileges set on groups that are more distant.

For example, an administrator grants SELECT privilege on table T1 to group G1, and that is the only privilege on that table. If Bob is a direct member of group G1, then Bob has privileges to select from T1. The administrator then denies SELECT privilege on table T1 to group G10, where group G1 is a direct member of group G10. User Bob continues to have privileges to select from T1 because group G1 gives him the SELECT privilege. Group G1 is closer in relation to Bob than group G10—of which Bob is also a member—but indirectly through membership in group G1. If the administrator now denies SELECT privilege on table T1 to group G2, where Bob is also a direct member of group G2, there is a conflict in privileges. Group G1 indicates that Bob is granted SELECT privilege., However, group G2 indicates that Bob is denied SELECT privilege. In these cases where there is a conflict, the user is denied the privilege. In this example, Bob does not have SELECT privilege on table T1.

## Maintaining Security Definitions for Tables, Views, or Columns

### Overview

If an administrator adds security definitions to a table, view, or column and then later drops the table, view, or column, the security definitions are removed. There might be times when a batch job drops the object and then re-creates it. In this case, the administrator must reestablish security on the object.

### Tables

As a suggested best practice, use the DELETE command to remove all of the rows from the table, but leave the table definition intact so that the security definitions remain with the table. Note that indexes created in the data store remain on the table as well. This practice offers an advantage if the table is merely being repopulated.

### Views and Columns

There is no way to alter a view definition without dropping the view and re-creating it. In this case, security for the view needs to be reestablished. The same applies for column security. Column security must be reestablished for columns that are dropped and added at a later time.

## Granting Privileges

### Overview

A system user has no restrictions on SAS Federation Server and can grant any privilege to any user. A SAS Federation Server administrator can grant any privilege on the server to any user except the ADMINISTER privilege itself. In other words, an administrator cannot create other administrators. Only the system user can do this.

Besides system users and administrators, the only other user that can grant a privilege to other users is a DSN owner. This user can grant the CONNECT privilege (and only that privilege) to other users.

### Grantor Precedence

The order of precedence between the three groups of users who can grant privileges is as follows:

- system user
- administrator
- DSN owner

Privileges granted from a system user cannot be overridden by administrators, and privileges granted from an administrator cannot be overridden by a DSN owner.

Here are some examples:

- A system user denies SELECT privilege to user John on the server object. An administrator cannot grant SELECT privilege to user John on the server object. An administrator can grant SELECT privilege to user John on a different object, such as a data service, catalog, or schema.

- An administrator grants CONNECT privilege to group SALES\_GROUP on DSN SALES\_DATA. The DSN owner cannot deny CONNECT privilege to group SALES\_GROUP on the same DSN.

### **Privilege Determination Summary**

To summarize, privileges can be determined by group membership and by object hierarchy. The precise algorithm is described here:

- System and administrator users are not denied access to any objects.
- For other users, a specific privilege is first checked at the current object in the hierarchy, in the following order:
  1. On the current object, evidence of the privilege is first searched in the identification of the specific user. If a GRANT or DENY determination is made, the status is returned for the privilege, and privilege lookup stops.
  2. User privilege search continues on the current object under the first level of group membership. If any group indicates DENY, the user is denied access. If all groups indicate GRANT, then the user is granted access. If a privilege determination cannot be made, the next level of group membership is checked.
  3. On the current object, the privilege is searched for under the next level of group membership, per the same rules as the previous item. If no specific determination is made, repeat for all levels of group membership.
  4. If no determination is made on the current object, then privilege determination goes to the next higher level in the object hierarchy. The privilege search algorithm repeats as described, first under the identification of the specific user, and then group membership.
  5. If all objects in the hierarchy are searched and no privilege determination is made, then a DENY is returned for the privilege type for the use.

## **Privilege Caching**

### **Overview**

Privilege caching entails capturing and reusing privileges enforced on a statement or request submitted to SAS Federation Server. This improves performance by reducing queries against internal system tables.

Privileges are cached on demand. Each time privileges on a securable for a given grantee are checked, the cache is examined initially to see whether the privilege has already been cached. If so, enforcement cost is reduced by limiting or eliminating queries against security-related system tables. Once uncached privileges are retrieved from system table queries, they are cached for subsequent use, stabilizing the cache based on actual access patterns.

### **Configuring Privilege Caching**

Use the [ALTER SERVER DDL statement on page 195](#) to cache privileges. For example:

```
alter server {options xset CACHE(name 'Authorization',
level <x level>, purge )}
```

- The default level is 2.

- The purge option frees the current cache. If privilege caching is enabled, the system dynamically builds the cache again.
- Use the purge option to trim memory usage in a long-running server instance.

### Level Control

0	Privilege caching is disabled. All levels are purged.
1	Schema, catalog, server, and service-level privileges are cached. Level 2 is purged.
2	Column, row, table, schema, catalog, server, and service-level privileges are cached.

## Row-Level Security

### Introduction

Row-level security (RLS) for SAS Federation Server provides additional security on tables and views by restricting data access on a row-by-row basis. Row-level security allows selection of specific rows for a given set of users and groups. Because row-level security applies only to rows that can be selected, its implementation is a function of the SELECT privilege. The SELECT privilege can be granted without restriction, or with a predicate applied. When a predicate is applied, you are implementing row-level security because the predicate now restricts what rows are returned.

For example, an administrator might choose to grant the SELECT privilege to USER1 on table T1. In this case, USER1 is allowed to see all the data in the table. However, an administrator might allow USER1 to select only from rows where a column or set of columns meet certain criteria, such as for the northeast sales region. In this case, the GRANT statement is:

```
GRANT SELECT ON TABLE CATALOG.SCHEMA.T1 TO SALES
WHERE SALES_REGION = 'NORTHEAST'
```

When members of the SALES group select from table T1, the predicate is automatically attached to the table. When BOB, a member of the SALES group, issues the statement, **SELECT \* FROM T1**, the query is effectively transformed to **SELECT \* FROM T1 WHERE SALES\_REGION = 'NORTHEAST'**.

Reference [Administration DDL on page 195](#) for syntax details about the GRANT statement.

When the predicate is automatically attached to the table, it must contain a valid WHERE clause. The syntax can include sub-queries and references to other tables. However, any external data referenced in the predicate must be available on the user's connection and the user must have the SELECT privilege to access the data.

If the data is coming from a different data source, then:

- The user's DSN must scope to that data source.
- The user must have SELECT privilege on the referenced data.

For example, a user can select rows only from a table in which the user's name is listed in the "USER NAME" column of the table. To apply this rule to all members of the SALES group, an administrator issues the following GRANT statement:

```
GRANT SELECT ON TABLE CATALOG.SCHEMA.T1 TO SALES WHERE
SYSCAT.RLS.CURRENT_USER() = \"USER NAME\"
```

When the user BOB in the SALES group selects from the table, **SELECT \* FROM T1**, the query is transformed to:

```
SELECT * FROM T1 WHERE 'BOB' = "USER NAME"
```

The [RLS Library and Library Reference on page 88](#) contains details about callable functions for row-level security.

## Row-Level Security Privilege Assembly Rules

### Overview

The SELECT privilege for a user can be derived through group memberships. A user can be a member of multiple groups, each granting SELECT privilege with attached row-level security. One exception is the schema owner. Since a schema owner effectively has all privileges granted, group membership is not traversed for privileges at the schema level.

RLS predicates are assimilated at each level in the securable hierarchy, starting with the table or view object and progressing to the server object.

The procedure is repeated over each object in the authorization hierarchy:

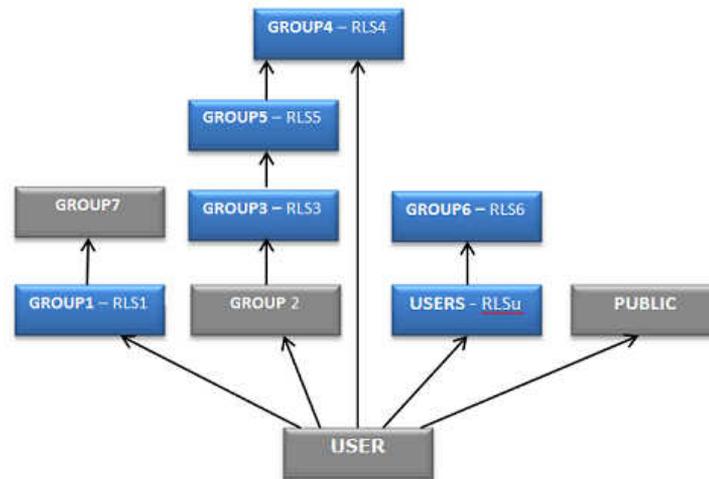
- the current user. If the current user has an RLS condition applied, no other RLS conditions are considered.
- any groups of which the current user is a member.
- followed by USERS and PUBLIC only if no other RLS conditions were discovered.

If an RLS predicate is discovered at a specific level, it is combined with the other RLS predicates at that level only, using an OR operator, and the process stops. The procedure produces predicates representing the summation of RLS privileges closest to the user. This approach gives preference to organizational security policies closer to the user over those more distant. An unconditional GRANT or DENY at the same level as an RLS predicate, will be honored and all RLS predicates will be ignored.

### Example: RLS Predicate Union

In this example USER is a direct member of the user-defined groups GROUP1, GROUP2 and GROUP4. USER is also a member of the two built in groups, USERS and PUBLIC. The RLS query returned for USER is (RLS1 OR RLS4 OR RLS3).

*Note:* Blue nodes contain a conditional privilege; gray nodes contain NO privilege.

**Figure 5.3** RLS Predicate Union

The first predicate, RLS1, is encountered at level 1 in the graph, so the remaining RLS predicates are captured at that level for the current graph, which does not include the secondary graphs, USERS and PUBLIC. The GROUP1 node is marked as visited, and its parent (“member of”) associations will not be navigated since it contributed a predicate.

After marking it as visited, the procedure skips GROUP2 since it has no assigned privileges and proceeds to GROUP4 where the node is marked as visited and RLS4 is captured and combined to the query using an OR operator. There are no more nodes at level 1, so the procedure continues with the parent associations of GROUP2.

GROUP3 is marked as visited and predicate RLS3 is captured and combined to the query using an OR operator. The privilege has been satisfied at this point because an RLS query is available. The two graphs starting at USERS and PUBLIC are not searched.

The resulting RLS query is: **RLS1 OR RLS4 OR RLS3**.

## The RLS Library and Library Reference

### Overview

Use the RLS library to look up functions that reference your data source. The RLS library resides in the SYSCAT.RLS catalog in the Federation Server Database. The general syntax is **SELECT SYSCAT.RLS.RLS\_function('name')**. For example:

```
SELECT SYSCAT.RLS.group_id('GROUP1')
```

### RLS Functions

In addition to row-level security, RLS Library user functions can be used with other SAS Federation Server tasks such as FedSQL views and queries by including an RLS function in your SELECT statement: **SYSCAT.RLS.RLS\_function**. Following are the callable functions for row-level security.

#### auth\_id

Returns the authentication provider identifier for the specified user or group name.

#### current\_user

Returns the name of the current user.

**current\_id**

Returns an opaque authentication provider specific user identifier.

**domain**

Returns the name of the domain in which the current user is authenticated.

**group\_id**

Returns the authentication provider group identifier for the specified group name.

**login**

Returns the domain-qualified user id that is used to authenticate the current user.

**userid**

Returns the SANs domain user id that is used to authenticate the current user. Note that the userid function is similar to the login function, but it is not domain-qualified.

**ip\_addr**

Returns the client IP address of the current user's session.

**is\_admin**

Returns TRUE or FALSE indicating if the current user is an administrator.

**is\_process\_user**

Returns TRUE or FALSE indicating if the current user is the process user.

**member\_of**

Returns TRUE or FALSE indicating if the current user is a member of the specified group.

**groups**

Returns a single group name or result set identifying the group memberships of the current user.

The following example uses RLS functions **current\_user** and **member\_of** to qualify users for SELECT on specific rows in the HR.EMPLOYEES table.

```
grant SELECT on HR.EMPLOYEES to USERS
  where (syscat.rls.current_user() in ("Name", "MgrName"))
  or syscat.rls.member_of("Payroll", 'DEEP')
```

Authenticated users must meet the following conditions:

1. The name of the current user matches that of the “Name” column (the row pertains to the current user).
2. or the name of the current user matches that of the “MgrName” column (the row pertains to the manager of the current user).
3. or the current user is a direct or indirect (DEEP) member of the “Payroll” group.

**RLS Library Reference**

Following are the data formats for the row-level security user functions described above.

**WVARCHAR(n) auth\_id(WVARCHAR(n) [[authorization]])**

Returns an authentication identifier as defined by the authentication provider, as a result of passing input for user name.

**WVARCHAR(n) current\_user()**

Returns the name of the current user. This is the authorization identifier of the currently authenticated user, rather than the login used.

**WVARCHAR(n) current\_id()**

Returns a user identifier as defined by the authentication provider. Typically, this is a static identifier by which the current user is known. Applications can associate this

identifier with an internal organization user identifier such as an employee number or account number.

**WVARCHAR(n) domain()**

The name of the domain in which the current user is authenticated.

**WVARCHAR(n) group\_id(WVARCHAR(n) [[authorization]])**

Returns a group identifier as defined by the authentication provider, as a result of passing input for group name

**WVARCHAR(n) login()**

The login used to authenticate the current user.

**WVARCHAR(n) userid( BITupn )**

The domain qualified user ID. If the upn parameter is TRUE, the format of the returned user ID is user@domain. Otherwise, the format is domain\user on Windows systems and just userid on all other systems. The userid function returns the authenticated user ID as specified by the authentication service. The authentication service can reside on a different host

**WVARCHAR(n) ip\_addr()**

Returns the client IP address of the current user's session.

**BIT is\_admin()**

Returns TRUE or FALSE if the current user is or is not an administrator.

**BIT is\_process\_user()**

Returns TRUE or FALSE if the current user is, or is not the process user.

**BIT member\_of( WVARCHAR(n) group [, WVARCHAR(n) options] )**

Returns TRUE or FALSE if the current user is, or is not a member of the specified group. Can assert direct or indirect membership. The group parameter is a group name by default and a group identifier if the 'ID' or 'DEEP' option is present in the options string. The options string is a blank or comma separated string consisting of one or more of 'ID' and 'DEEP' option keywords. The current user must be a direct member of the specified group unless the 'DEEP' keyword is specified. 'DEEP' checks for both direct and indirect group membership. Direct membership is tested by default.

**TABLE(WVARCHAR(n) group) groups( WVARCHAR(n) [[authorization] WVARCHAR(n) [, options]]**

Returns a single group name or identifier column result set containing the current user's group memberships. The available options are 'ID' or 'DEEP'. Can be restricted to direct memberships only. The authorization parameter is a user or group name by default and a user or group identifier if the 'ID' option is present in the options string. The options string is a blank or comma separated string consisting of one or more of 'ID' and 'DEEP' option keywords. A deep group membership listing is returned if the 'DEEP' keyword is specified, the default being a shallow listing.

*Note:* A trusted user must be set in order for the GROUPS function to return a result set if you pass in a user other than the current user. If you pass in the current user or a group as the first argument to GROUPS trusted user is not required.

**Using the 'ID' and 'DEEP' Options**

This topic outlines the behavior of the 'ID' and 'DEEP' options when used with the **member\_of** and **groups** RLS functions. Consider the following queries:

```
SYSCAT.RLS.GROUPS(user_name_or_id [, options])
```

Returns a group membership list for **user\_name\_or\_id**.

**SYSCAT.RLS.MEMBER\_OF**(group\_name\_or\_id [, options])  
Returns TRUE if the current user is a member of the group specified in group\_name\_or\_id.

The following behavior applies:

- **user\_name\_or\_id** returns a string literal indicating the user name or user ID. If **options** contains 'ID', the argument is treated as a user ID ('6C6C9AD1E2646F0469DD6A3D1874D167') rather than a user name ('USER1').
- **group\_name\_or\_ID** returns a string literal indicating the group name or group ID. If **options** contains 'ID', the argument is treated as a group ID ('78319AD1E2646F0469DD6A3D1874A2F7') rather than a group name ('GROUP1').
- **options** returns a string literal containing 'GROUP', 'ID', or both ('GROUP, ID'). If multiple options are specified, they can be separated in the string by a blank in single quotation marks (' ') or comma in single quotation marks (','). If **options** contains 'ID' the argument is treated as an ID rather than a name. If **options** contains 'DEEP', group membership is checked for direct and indirect membership.

For example, consider that USER1 is a member of GROUP1 and GROUP1 is a member of GROUP2, and USER1 runs the following queries:

```
Select SYSCAT.RLS.GROUP_ID('GROUP1')
Returns 'BEA892C6D4A40464C8A144D89FFE6463'
```

```
Select SYSCAT.RLS.GROUP_ID('GROUP2')
Returns '45C562900C7333C49B1706B38DBA75B0'
```

```
Select SYSCAT.RLS.CURRENT_ID()
Returns '5790EE3F6A24A7349AA2254600793411' for USER1
```

```
Select SYSCAT.RLS.MEMBER_OF('GROUP1')
Returns TRUE for USER1.
```

```
Select SYSCAT.RLS.MEMBER_OF('GROUP2')
Returns FALSE for USER1.
```

```
Select SYSCAT.RLS.MEMBER_OF('GROUP2', 'DEEP')
Returns TRUE for USER1.
```

```
Select
SYSCAT.RLS.MEMBER_OF('BEA892C6D4A40464C8A144D89FFE6463', 'ID')
Returns FALSE for USER1.
```

```
Select
SYSCAT.RLS.MEMBER_OF('45C562900C7333C49B1706B38DBA75B0', 'ID')
Returns TRUE for USER1.
```

```
Select
SYSCAT.RLS.MEMBER_OF('BEA892C6D4A40464C8A144D89FFE6463',
'DEEP, ID')
Returns TRUE for USER1.
```

The following queries use the **GROUPS** function which returns a result set:

```
Select * from SYSCAT.RLS.GROUPS('USER1')
Returns a result set showing direct group membership for USER1:
```

```
"GROUP"
'USERS'
'GROUP1'
'PUBLIC'
```

```
Select * from SYSCAT.RLS.GROUPS('USER1', 'DEEP')
```

Using DEEP in the SELECT statement returns a result set showing direct and indirect group membership for USER1:

```
"GROUP"
  'USERS'
  'GROUP1'
  'GROUP2'
  'PUBLIC'
```

---

## Data Masking

### Overview

Data masking is a method of hiding sensitive information within data sources. The purpose of data masking is to protect the original data by using a functional substitute in situations where the audience is not privileged to access the original data. The primary focus of data masking is to protect data that is classified as sensitive, or as personally identifiable information (PII).

Data masking with SAS Federation Server is implemented using a set of rules (rule types) and arguments that are run with a system function, `SYSCAT.DM.MASK`. Data masking rules consist of individual rule types that define the specific masking action or algorithm to apply to the data. The rules in effect are as follows:

- ENCRYPT and DECRYPT
- HASH

These rule types are valid for use in FedSQL queries and are applied over literal values or individual columns to hide personally identifiable information.

### About the `SYSCAT.DM.MASK` Function

The `SYSCAT.DM.MASK` function accepts defaults configured as package options in addition to the various arguments associated with each rule type. For example, the `KEY` argument for the `ENCRYPT` rule type defaults to the value configured as the `ENCRYPT_KEY` package option. To configure default masking parameters as package options, use the `ALTER SERVER` DDL command. The following example sets a default encryption key used by the `ENCRYPT` and `HASH` rule types:

```
ALTER SERVER {options PACKAGE(name 'DM',
    ENCRYPT_KEY '212e8ba6b7f84796a87a985d54277f2f')}
```

Configured parameters revert to a static default value if they are dropped.

*Note:* Column data encrypted with the static default key cannot be reversed with the `DECRYPT` rule type.

Use the `SYSCAT.DM.MASK` function with the specified rule types and arguments to mask a value containing PII. The rule type argument must be a string constant.

Argument names are not case sensitive. Here is the syntax:

```
SYSCAT.DM.MASK( 'rule-type', value [, rule-arguments])
```

Here is the syntax for the rule arguments:

```
[, 'rule-arg-name1', 'rule-arg-value1',
```

```
[, 'rule-arg-name2', 'rule-arg-value2', ...] )
```

**Table 5.2** Data Masking Arguments and Descriptions

Argument	Description
<i>rule-type</i>	Name of the rule type — ENCRYPT, DECRYPT, or HASH.
<i>value</i>	Data value that requires masking (for example, a column reference or other value expression).
<i>rule-arg-name1</i>	Name of the first masking rule argument.
<i>rule-arg-value1</i>	Value of the first masking rule argument.
<i>rule-arg-name2, ...</i>	Name of the second and subsequent position masking rule arguments.
<i>rule-arg-value2</i>	Value of the second and subsequent position masking rule arguments.

Here is an example of data masking that uses the **HASH** rule-type. This rule masks the 'LastName' column in the HR.EMPLOYEES table using an HMAC-MD5 hash and aliases the result as LN:

```
select SYSCAT.DM.MASK('HASH', "LastName",
                    'alg', 'MD5',
                    'key', 'abc123!' ) as "LN"
from HR.EMPLOYEES
```

## Data Masking Rule Types and Arguments

### ENCRYPT / DECRYPT

Here is the syntax for ENCRYPT and DECRYPT:

```
SYSCAT.DM.MASK( 'ENCRYPT', value [, rule-arguments ] ),
SYSCAT.DM.MASK( 'DECRYPT', value [, rule-arguments ] )
```

The ENCRYPT rule encrypts a single value using symmetric key encryption. Encrypted values cannot be decrypted if no KEY argument is specified and the ENCRYPT\_KEY package configuration is not configured.

The DECRYPT rule decrypts a previously encrypted value using symmetric key encryption. The DECRYPT rule returns NULL if a KEY argument is not specified and the ENCRYPT\_KEY package configuration is not configured.

Use HASH if uniqueness or reversibility (decryption) is not a requirement.

Unique	Results are unique if the input values are unique.
Deterministic	Results are deterministic if DETERMINISTIC YES is specified.
Reversible	Possible if key is poor quality or divulged.

## Arguments

### ALG

Argument type: Required. Case-insensitive string constant.

Specifies the algorithm name which is one of the following:

AES/FIPS*	AES/FIPS: RSA licensed FIPS compliant AES encryption.
AES*	RSA-licensed AES encryption.
SAS002	SAS Proprietary.
BASE64	Base 64 encoding.
SAS004	Alias for AES/FIPS.
SAS003	Alias for AES.
SAS001	Alias for BASE64.

*Note:* \*AES/FIPS and AES require installation of SAS/SECURE software.

### KEY

Argument type: Case-insensitive string constant.

Symmetric encryption/decryption key. The key defaults to the **ENCRYPT\_KEY** package configuration parameter or, for **ENCRYPT** only, the static default when none is configured.

Exception: Not valid with the SAS002 algorithm, which does not accept an encryption key.

### DETERMINISTIC

Argument type: Case-insensitive Boolean string constant or 1 or 0.

Use with **ENCRYPT** only.

Boolean string constant values must contain one of {YES, TRUE, ON, 1, NO, FALSE, OFF, 0 }.

Use this option when deterministic output is required. Controls whether encrypted value is DETERMINISTIC. The default value is FALSE.

Exception: Not valid with the SAS002 algorithm, which does not support a custom key.

### EXPAND\_PREC

Argument type: Case-insensitive Boolean string constant or 1 or 0.

Use with **ENCRYPT** only.

Boolean string constant values must contain one of {YES, TRUE, ON, 1, NO, FALSE, OFF, 0 }.

This option causes the precision of the output value to accommodate the possibility of data bloat when using the specified encryption or encoding method.

**EXPAND\_PREC** is active by default. If the encrypted value does not fit in the column, this option returns an empty string for **VARCHAR** or **NVARCHAR**, or all-blank for **CHAR** or **NCHAR**.

- an empty string for **VARCHAR** or **NVARCHAR**.
- a completely padded blank for **CHAR** or **NCHAR**.

### CASE

Argument type: Case-insensitive string constant.

Use with **ENCRYPT** only.

Controls case in output. Does not apply to the BASE64 algorithm, which produces upper and lower cased characters.

U[PPER] Use upper case hexits A-F.

L[OWER] Use lower case hexits a-f.

## STRIP

Argument type: Case-insensitive string constant.

Use with **ENCRYPT** only.

Specifies whether to strip trailing whitespace from the input value prior to encryption. By default, values are not stripped.

Valid values are BLANKS, UNICODESP, UNICODESPACE, ANY, ALL, WS. BLANKS specifies to strip ASCII blank (0x20) characters only. UNICODESP and UNICODESPACE specify to strip Unicode whitespace. ANY, ALL, and WS specify to strip Unicode whitespace as well as certain format control characters.

## HASH

Syntax

```
SYSCAT.DM.MASK( 'HASH', value [, rule-arguments ] ),
SYSCAT.DM.MASK( 'DECRYPT', value [, rule-arguments ] )
```

The HASH rule hashes a single value into a fixed-length hash digest or HMAC string and is not reversible.

Unique	Results may be unique if the input values are unique.
Deterministic	Yes
Reversible	No

Algorithm	Value Type	Return Type
MD5	CHAR or VARCHAR	CHAR(32)
	WCHAR or WVARCHAR	WCHAR(32)
	other	BINARY(16)
SHA256	CHAR or VARCHAR	CHAR(64)
	WCHAR or WVARCHAR	WCHAR(64)
	other	BINARY(32)

## Arguments

### ALG

Argument type: Case-insensitive string constant.

Required. Specifies the algorithm name:

MD5 Robert Rivest's 128-bit algorithm (1991).

SHA256\* NSA'S 256-bit algorithm (2001). SHA256 is the default.

*Note:* \*SHA256 requires installation of SAS/SECURE software.

### CASE

Argument type: Case-insensitive string constant.

Controls case in output. Does not apply to the BASE64 algorithm, which produces upper and lower cased characters.

U[PPER] Use upper case hexits A-F.

L[OWER] Use lower case hexits a-f.

### KEY

Argument type: Case-insensitive string constant.

By specifying a KEY argument, or defaulting to the ENCRYPT\_KEY parameter configured in the package, a 'hash message authentication code' (HMAC) that complies with RFC 2104 (<http://tools.ietf.org/pdf/rfc2104>) is computed. Otherwise, the specified raw digest is computed directly.

### STRIP

Argument type: Case-insensitive string constant.

Use with **ENCRYPT** only.

Specifies whether to strip trailing whitespace from the input value prior to encryption. By default, values are not stripped.

Valid values are BLANKS, UNICODESP, UNICODESPACE, ANY, ALL, WS. BLANKS specifies to strip ASCII blank (0x20) characters only. UNICODESP and UNICODESPACE specify to strip Unicode whitespace. ANY, ALL, and WS specify to strip Unicode whitespace as well as certain format control characters.

## Server Encryption

### Introduction

SAS Federation Server supports two methods of encryption strength: SAS Proprietary Encryption and DataFlux Secure.

### SAS Proprietary Encryption

SAS Proprietary Encryption is a fixed encoding algorithm that is included with SAS Federation Server. It requires no additional product licenses and is the default encryption method if DataFlux Secure is not installed. The SAS Proprietary Encryption algorithm is strong enough to protect your data from casual viewing. SAS Proprietary Encryption provides a medium level of security. SAS/SECURE and SSL provide a high level of security.

## DataFlux Secure

### Overview

DataFlux Secure is an add-on product that provides industry encryption capabilities in addition to the SAS Proprietary Encryption algorithm. DataFlux Secure requires additional licensing and it must be installed on each server that will use encryption. DataFlux Secure provides encryption of data in transit. It does not provide authentication or authorization capabilities.

The AES – 256-bit keys encryption algorithm is used by SAS Proprietary Encryption and DataFlux Secure.

### Specifying the Encryption Method

SAS Proprietary Encryption (SASProprietary) is the default encryption for SAS Federation Server. You can also decide how much data is encrypted in communication between a client and SAS Federation Server. This is specified by setting the **CLIENTENCRYPTIONLEVEL** using the **ObjectServerParms** option in **dfs\_serv.xml**. See the [SAS Federation Server Configuration Reference on page 40](#) for more information on this option.

## Password Encryption

SAS Federation Server provides a utility to encrypt user passwords from plain-text format. The method of encryption depends on the encryption method in use for SAS Federation Server. See “[Utilities for SAS Federation Server](#)” on [page 70](#) for additional information about password encryption.



## Chapter 6

# Working with Federated Data

---

<b>Overview of Data Federation</b> .....	<b>99</b>
<b>Configuring Data Source Access</b> .....	<b>100</b>
Overview .....	100
Data Services .....	100
DSNs .....	102
Catalogs and Schemas .....	105
<b>Federated SQL Views</b> .....	<b>107</b>
Overview .....	107
Federated SQL Views as Data Abstraction .....	107
Invoker and Definer's Rights Views .....	108
Requirements for Definer's Rights Views .....	108
Required Ownership for Federated SQL Views .....	110
Creating Federated SQL Views .....	110
Dynamic Connections .....	111
<b>Data Caching</b> .....	<b>112</b>
Overview .....	112
Views and Caching .....	113
Requirements for Cached Views .....	113
Working with Cached Views .....	114
Disabling and Enabling Cached Views .....	116
<b>Memory Data Store</b> .....	<b>117</b>
Overview .....	117
The MDS Data Service .....	117
The MDS Database .....	118
<b>Understanding Data Federation and Best Practices</b> .....	<b>119</b>
Overview .....	119
Data Model .....	119
Data Security .....	120

---

## Overview of Data Federation

By supporting several data sources, SAS Federation Server provides the flexibility to configure data storage based on specific needs. You choose the type of data storage that is most appropriate for the particular needs of an application, based on functionality that is provided by each data source. The first step in working with federated data is to configure access to your data sources.

Before creating federated (FedSQL) data views and caching data, make sure that your trusted user and shared logins are configured and ready for use, and have been configured and are ready for use. For additional information, see [“Configuring Server Accounts” on page 26](#).

For a comprehensive view of data federation and best practices, refer to [“Understanding Data Federation and Best Practices” on page 119](#) at the end of this chapter.

---

## Configuring Data Source Access

### Overview

To access data, the administrator must create and configure data services for SAS Federation Server. Data services contain connection information and driver specifics to connect with data sources such as Oracle® or Base SAS® data sets.

Relational databases provide authorization that limits the operations that can be performed on their data. As noted previously in [“Data Source Authorization” on page 76](#), SAS Federation Server respects authorizations that are defined and enforced on a third-party database.

### Data Services

#### Overview

Data services contain information that identifies the location of data tables. If a data source does not support native catalogs, SAS Federation Server enables you to define a logical catalog name to use as an SQL identifier. This allows unique identification of each data source when performing heterogeneous operations.

Data services that require logins must be associated with a domain in the Authentication Server. When users connect to the data service through a data source name (DSN), the domain name is used to retrieve user credentials associated with that data service. The credentials are then passed along to the back-end database. User credentials are stored in the Authentication Server.

Data services can also contain optional information to control SAS Federation Server driver behavior, such as locking semantics and tracing. Data services form the foundation for connectivity to a data source and you can assign privileges that control user access to the data.

#### Creating a Data Service

You can use administration DDL or SAS Federation Server Manager to create a data service. A new DSN is automatically generated each time a new data service is created. This is a standard DSN given the same name as the data service. If the data service is renamed, the DSN name remains unchanged.

By default, a BASE data service is created the first time that SAS Federation Server is started. Only one BASE data service can exist in a SAS Federation Server installation, and it cannot be modified or deleted.

#### Native Catalog Support

When creating a data service for a data source that supports native catalogs, and using the REGISTER option, the server attempts to connect to the database to acquire a list of

catalogs. Credentials are required to secure the connection. If the connection cannot be made, creation of the data service fails. The same requirement for pre-registered credentials applies when creating a data service to ODBC with native catalog support, or for any data services that support native catalogs.

### **Identifier Case Sensitivity**

When creating an ODBC data service, the server must query the data source to acquire its identifier case sensitivity property. The identifier case sensitivity property is used to create security entries in the server's system tables and is stored with the data service.

### **Database Login Prerequisite**

Due to the requirement for a database connection described above, the following database login prerequisite applies to SQLSERVER, ODBC and ODBC\_FED data services. These actions must be completed before creating the data service, and must be accomplished through the Authentication Server Manager.

1. A database login must be registered in the domain that will be associated with the new data service.
2. The domain must be registered in Authentication Server for it to be accessible when creating the data service using administration DDL or SAS Federation Server Manager.

### **Creating a Data Service with Administration DDL**

Here is an example of the DDL statement, CREATE DATA SERVICE. See [CREATE DATA SERVICE DDL on page 199](#) for details and a complete list of options. You can also alter and drop data services using DDL.

```
CREATE [DATA] SERVICE data-service
TYPE data-service-type
[CATALOG [NAME] catalog-name]
[DOMAIN [NAME] domain-name]
[REGISTER [( catalog-name1 [,catalog-name2 ...]) | ALL]
[register-options]]
[data-service-options]
```

Use CATALOG to register a catalog name for data sources that do not support native catalogs.

Use REGISTER to register the native catalog names for those data sources that support native catalogs, such as SQL Server. If an identical catalog name is encountered, a warning message is issued and the catalog is not registered. In this case use the [CREATE CATALOG DDL on page 205](#) statement to provide a mapped name for the native name.

### **Driver Search Order**

If specifying more than one driver when creating a data service, the first driver listed in the statement is deemed the default driver if a driver is not specified in a connection string. For example, if two drivers are specified when creating a data service, the first driver is used as the default driver if **driver=** is not specified in a connection string. If a driver is specified in a connection string using **conopts-configuration-list** when creating the data service, it is used when connecting to a data source.

## DSNs

### Overview

DSNs are resources that provide connection information for data sources accessed through SAS Federation Server. The administrator assigns permissions to determine what users can connect and how they connect. For example, to connect to a data source, a user must be granted **CONNECT** permission on SAS Federation Server, a specific data service, or a specific DSN.

A DSN references a specific data source to which it will connect and defines how SQL security is enforced. It can be configured so that SAS Federation Server enforces SQL privileges defined for the data service. The **CREATE DSN** permission is required to create a DSN. You can configure DSNs using Administration DDL statements or by using SAS Federation Server Manager.

### DSN Types

#### Standard DSN

A standard DSN is a single-service DSN created for a particular data service and is parented to that data service. The scope is limited to one data service and contains connection information, such as server name, port, path or other connection options specific to a data service.

#### Federated DSN

A federated DSN is a collection of one or more DSNs. Unlike the standard DSN which is parented to a data service, the federated DSN is parented to the federation server itself, even if it only contains DSNs from a single data service. Federated DSNs can contain other federated DSNs. Since federated DSNs are typically used to provide access to data from multiple, disparate data sources, the FedSQL dialect is required.

#### System DSNs

These system DSNs are created during installation of SAS Federation Server:

- The **ADMIN DSN** is created automatically at server start up and is used for the purpose of sending administration SQL to the server. The ADMIN DSN is also used to query [Information Views on page 224](#). The SAS Federation Server Manager automatically connects with the ADMIN DSN to display information such as the registered list of data services and DSNs. Any user expected to use SAS Federation Server Manager to accomplish tasks, such as creating views and caching views, will require **CONNECT** permission to the ADMIN DSN. SAS Federation Server automatically checks user privileges when administration SQL is submitted. Users can submit administration SQL for which they have privileges, such as selecting against Information Views. Some administration SQL can be executed by the server administrator only. See the *SAS FedSQL Reference Guide* for details.
- A **SQL\_LOG DSN** and data service are created when SQL Logging is enabled on SAS Federation Server. At that time the server creates an **EVENTS** table for the purpose of capturing server activity that reflects information about SQL statements submitted by connected users. Additional information can be found in [“SQL Logging” on page 56](#).

The **CONNECT** permission is not assigned to a DSN by default, and must be granted by the administrator or by the DSN owner, to users or groups that are connecting to data sources.

**DSN Permissions**

DSN permissions are assigned using GRANT, REVOKE, or DENY DDL statements.

The permissions for standard and federated DSNs are:

- CREATE DSN
- ALTER or DROP DSN
- CONNECT

**CREATE DSN**

To create a DSN, one of the following conditions must be met:

- The user is a system user.
- The user is an administrator of the server.
- The user has the CREATE DSN permission on the server. Note that this is the only way that a user who is not an administrator can create a federated DSN.
- The user has the CREATE DSN permission on the data service, and the user is creating a standard DSN.

**ALTER/DROP DSN**

To alter or drop a DSN, one of the following conditions must be met:

- The user is a system user.
- The user is an administrator of the server.
- The user is the owner of the DSN.

**CONNECT**

A user must have CONNECT permission to establish connection to a DSN. This permission is effective from the user object, inherited through the hierarchy, or acquired through group permissions. For a standard DSN, the CONNECT permission must be on (in order of inheritance):

- The DSN,
- The parent data service of the DSN, or
- SAS Federation Server.

For a federated DSN, the CONNECT permission must be on (in order of inheritance):

- The DSN, or
- SAS Federation Server.

Permissions granted on a federated DSN override any permissions that exist for child DSNs that are contained within the federated DSN. If a user has CONNECT permission on a federated DSN, permissions on any of the child DSNs contained within (standard or federated) are ignored, even if the user is explicitly denied CONNECT on any of the child DSNs

For additional information about permission assignment, see the topic on [Permissions on page 78](#).

**Enabling Federation Server SQL Authorization Enforcement**

When Federation Server SQL Authorization Enforcement is enabled, the FedSQL driver is also required, and the SQL dialect is automatically set to FedSQL. With FedSQL an additional layer of object-level security is enabled for the connection and SQL statements are secured before processing them. If Federation Server SQL Authorization

Enforcement is disabled, object-level security is bypassed and a user is granted all privileges regardless of what the user has been granted or denied. If Federation Server SQL Authorization Enforcement is disabled, an administrator can choose either FedSQL dialect or data source (native) dialect. For example, if you are connected to Oracle, then native dialect would be SQL supported by Oracle. The SQL dialect for Base data services is always FedSQL.

Security is enabled by default for all new DSNs. However, if you need to enable SAS Federation Server security on a DSN, use DDL options with CREATE DSN and set SECURITY to YES.

Here is an example DDL statement that enables SAS Federation Server security:

```
CREATE DSN "DSN1" UNDER BASE
DESCRIPTION 'creating DSN1' NOPROMPT
'DRIVER=BASE;CATALOG="catalog1_BASE";SCHEMA=(name="schema1_BASE")' {OPTIONS
(SEcurity YES)}
```

### **Credentials Search Order (CSO) for DSN Connections**

Connections made with a DSN use a credentials search order (CSO) as specified in the DSN configuration. By default, the credentials search order is PERSONAL, SHARED. Other valid values are SHARED, (PERSONAL, SHARED) and (SHARED, PERSONAL).

At connection request, SAS Federation Server attempts to select a user ID and password for each data service connection based on the associated domain:

- PERSONAL means the server attempts to select credentials directly owned by the user.
- SHARED means the server attempts to select credentials from a shared login of which the user is a consumer. Credentials are extracted on behalf of the user using the shared login manager's identity. Shared Logins must be configured in SAS Federation Server and defined in Authentication Server for this option to function properly. See [Shared Login Manager Configuration on page 29](#) for more information.
  - If a DSN is configured as CSO(SHARED) and a shared login is not found for any of the DSN's connections, the connection will fail immediately.
  - If the credentials search order is not configured on the DSN or is not CSO(SHARED), the connection is still attempted. If credentials are specified on the connection string, those will be used first. If credentials are not supplied, the server attempts to find shared logins for the user. If shared login credentials are not found, the server attempts to use personal credentials. If personal credentials are not found, the connection fails.

Also, if GROUP=*groupname* is specified with the DSN or supplied in the connection string, shared login selection is limited to those candidates in which the specified group is a consumer. The user must be a direct or indirect member of the group, or no shared login will be selected. The GROUP= option does not have any effect on personal login extraction.

### **DSN Login Credentials**

If data services require credentials, a DSN can be configured to specify how database logins are retrieved. The DSN can be configured to use the personal credentials of the user, or retrieve the login from a shared login. If you are using a shared login, you can specify a consumer group from the DSN. This is required only to identify what shared

login to use if multiple shared logins are available in the same domain for connecting users.

When using SAS Federation Server Manager, an administrator can specify personal credentials or a shared login to the underlying databases for the purpose of managing data services. SAS Federation Server Manager connects to a data service behind the scenes and data services use a credential search order of PERSONAL, SHARED (CSO=PERSONAL,SHARED). Therefore, if an administrator has both a personal and a shared login, the personal login will be chosen. If an administrator does not have a personal login, but has multiple shared logins available, the connection might be disallowed. See “[Best Practices for Setting Shared Logins](#)” on page 29 for rules on how shared logins are selected.

### **Creating a DSN with Administration DDL**

Using administration DDL, you can create standard and federated DSNs with various configuration options. For a complete list of configuration options, refer to the [CREATE DSN DDL statement on page 211](#).

#### **Standard DSN**

Here is the syntax for creating a standard DSN under a data service:

```
CREATE DSN dsn-name UNDER data-service
        create-dsn-options [ AS ADMINISTRATOR ]
```

*Note:* If a DSN is created by a user other than the system user or administrator, the DSN is owned by the individual user. If that user is later removed from the system, DSN ownership should be transferred to another user.

#### **Federated DSN**

Federated DSNs are objects of SAS Federation Server. Therefore, they are not created under a data service. Here is the syntax for creating a federated DSN:

```
CREATE DSN dsn-name
        create-dsn-options
        ADD "(" dsn-name ["," ...] ")"
```

## **Catalogs and Schemas**

### **Overview**

The terms catalog and schema are defined as ANSI SQL standards and refer to the organization of data in a relational database. That is, data is contained in tables, tables are grouped into schemas, and schemas are grouped into catalogs. Catalog and schema names can be used in SQL statements to qualify table references. For example, when querying a database that supports both schemas and catalogs, you can specify a three-level identifier in the form of CATALOG.SCHEMA.TABLE-NAME.

### **Organize Data with a Catalog**

A catalog is a named collection of logically related schemas. The catalog is the first-level (top) grouping mechanism in a data organization hierarchy that qualifies schemas. At least one schema is required for each catalog.

For the BASE data service, you must create catalogs and schemas in order to make data available. For all other data services, catalogs and schemas are defined in the data source, and catalog and schema names can be registered in SAS Federation Server to reflect those objects.

**Catalog Registration**

Catalog names for all data sources must be registered in SAS Federation Server and they must be unique within the system.

This is accomplished by using one of the following methods:

- Use the CATALOG keyword on the CREATE DATA SERVICE command. Do this when the data source does not support native catalogs.
- Use the REGISTER keyword on the CREATE DATA SERVICE command. Do this when the data source supports native catalogs.
- Use the CREATE CATALOG command. Do this to provide a mapped name for a native catalog that cannot be registered using the REGISTER keyword because it conflicts with an existing registered catalog.

The following is a sample of the CREATE CATALOG DDL statement:

```
CREATE CATALOG catalog UNDER data-service
[ NATIVE NAME native-name ]
[ create-catalog-options ]
```

A complete list of options is shown in the [CREATE CATALOG DDL statement on page 205](#).

**Organize Data with a Schema**

A schema is a data container object that groups logically related objects such as tables and views. The schema provides a unique namespace that is used along with a catalog to qualify names.

For SAS data sets, a schema identifies the physical location such as a UNIX® directory or a Windows® folder that contains a collection of tables. For SAS data, the relationship between a schema and its files is similar to that of an operating system file directory and the files that are contained within that directory. A schema is approximately equivalent to a SAS library.

**Schema Registration**

Unlike catalogs, schema registration is not required for all schemas in the data source. Schemas are registered only when the administrator wants to assign an owner to the schema. Schemas are also created and maintained internally as needed by the system, such as when assigning permissions to a user or group on a schema.

The following is an example of the CREATE SCHEMA DDL statement:

```
CREATE SCHEMA [ catalog.schema ]
[ AUTHORIZATION|OWNER owner ]
[ create-schema-options ]
```

A complete list of options is shown in [CREATE SCHEMA DDL on page 207](#).

**Schema Ownership**

All schemas have an owner. If an owner is not explicitly assigned to a schema, ownership defaults to the system user account. Definer's rights views require a non-system schema owner for proper operation. The schema owner is the owner of all objects contained in the schema, though the owner has particular relevance to definer's rights views. And as the owner, certain privileges are automatically granted to the schema owner.

### **Privilege Rules for Schema Ownership**

Here are additional rules that apply to schema ownership:

- The schema owner automatically has all SQL privileges on tables and views in the schema. They are reported with **GRANTOR=ADMINISTRATOR**. Also, the schema owner can alter the schema's configuration options.
- Administrators can **DENY** a privilege on the schema, and the owner will be denied the privilege. This feature can be used to downgrade schema ownership rights. Therefore, the schema owner has no explicit privileges on the schema, but has default **GRANT** for privileges on schema objects.
- Administrators can reverse denied privileges using **GRANT**.
- **GRANT** to schema owner is equivalent to a **REVOKE**. The command clears any explicit denied privileges on the schema, but does not add any explicit ones. That way, when a schema owner's privilege is cleared on the schema, it defaults back to implicit **GRANT**.
- When schema ownership changes, the previous owner receives default privileges from the schema's container, whether it is a catalog or a data service. However, explicitly denied privileges remain in tact for the schema.
- Schemas should not be owned by any user who is a system user.

---

## **Federated SQL Views**

### **Overview**

When there is a need to view information from multiple data sources or other source types, you can create a reusable federated SQL view (FedSQL view) to deliver data from multiple relational and non-relational sources. A federated SQL view contains the information required to access database sources and can be stored separately from the data. By creating a view definition, you are storing only the instructions for where to find the data and how it is formatted, not the actual data.

Views can reduce the impact of data design changes on users. For example, you can redirect data sources or change variables that are stored in a view without changing the characteristics of the view's result. The view remains consistent even if the data source changes. To create a federated SQL view, the FedSQL dialect must be selected in the DSN.

### **Federated SQL Views as Data Abstraction**

The concept of data abstraction is used in database systems to define user interfaces through the creation of database views. Based on the data abstraction layer concept, a federated SQL view hides the complexity of data by defining an organized data structure for presentation to an end user or calling application. The result is that a user or application can request data in the organized virtual format, without regard to the physical layout. Data is fetched from potentially many data sources, transformed into the virtual structure and returned to the user or calling application.

## Invoker and Definer's Rights Views

There are two types of federated SQL views for SAS Federation Server:

- **Invoker's rights view:** An invoker's rights view is run with the invoking user's credentials.
- **Definer's rights view:** A definer's rights view is run with the credentials of the schema owner.

The invoker's rights view is accessed using the current user's authorization, credentials, and login information while the definer's rights view is accessed using the schema owner's authorization, credentials and login information. A definer's rights view is always associated with a schema owner.

A definer's rights view allows security management from a single layer of data, which in turn provides for a more secured system. For example, there are 100 tables that provide data to a set of users. 10 views are created and their data is acquired from the 100 tables. The users are selecting from the 10 views to get their data. With invoker's rights views, each invoker must have access to the 100 tables. This includes setting privileges in SAS Federation Server, and ensuring that each invoker has a login to the data sources containing the 100 tables.

With definer's rights views, the data available through the view is accessed by a single user only: the schema owner. Therefore, only this user, the schema owner, needs server privileges and database logins to the data sources containing the 100 tables. View invokers do not need direct access to the underlying tables. The administrator can secure the definer's rights view using Federation Server SQL authorizations to control which users and groups have access to the view's result set. Unless explicitly specified as a definer's rights view, a view is created as an invoker's rights view by default.

Here is an example of creating an invoker's rights view:

```
CREATE VIEW
    view1 AS SELECT * FROM table1
```

Definer's rights views are required for data caching. Only a schema owner can create a definer's rights view for a schema that he owns. A non-schema owner cannot create a definer's rights view for a schema for which he has the SELECT privilege only. You can create a definer's rights view using the following example syntax:

```
CREATE VIEW
    view1 SECURITY DEFINER AS SELECT * FROM table1
```

The following example alters an existing view to be an invoker's rights view or a definer's rights view. If a definer's rights view has any associated cache, it is dropped when the view is changed to an invoker's rights view.

```
ALTER VIEW
    view1 SECURITY INVOKER
```

```
ALTER VIEW
    view1 SECURITY DEFINER
```

## Requirements for Definer's Rights Views

*Note:* This section uses the acronym, VSO, to indicate View Schema Owner, which is the owner of the schema where the view resides.

Here are the requirements for definer's rights views:

- Use of definer's rights views requires a trust relationship between SAS Federation Server and the designated Authentication Server. Trust is established in the Authentication Server when connected using a [Trusted User Account](#). This enables impersonation of schema owners during execution of SQL. Specifically, the capability is required to retrieve:
  1. group memberships required for SQL authorization enforcement of data accessed in execution of a definer's rights view, and
  2. outbound database credentials forwarded to make transient connections used during view execution.

You can define a Trusted User in the Authentication Server's configuration file. See *Authentication Server Administrator's Guide* for additional information.

After defining a Trusted User Account, specify the Trusted User in the Federation Server by setting the trusted user account and password using an **ALTER SERVER** command:

```
ALTER SERVER {OPTIONS ( TRUSTED_USER_UID uid, TRUSTED_USER_PWD pwd )}
```

- Register all database catalog names referenced from the definer's rights view.

In order for a definer's rights view to access data from a SQL Server data service (or any data service that supports native catalogs), the federation server administrator must have pre-registered the referenced catalog names.

To ease administrative burden, the federation server administrator can automatically register catalogs when creating SQL Server data services using the **REGISTER** keyword of the **CREATE DATA SERVICE** command:

```
CREATE DATA SERVICE data-service TYPE data-service-type REGISTER
```

- Any user that is allowed to create, alter, or drop a view within a schema should be granted **CREATE VIEW**, **ALTER VIEW**, or **DROP VIEW** privilege on the schema or an object in its inheritance hierarchy. The schema owner implicitly has these privileges.
- A schema owner must be assigned to the schema that the view resides in. This is the view schema owner, also referred to as the VSO. This user owns all definer's rights views within that schema, and the VSO user is used for executing the uncached view. Authorization enforcement for all SQL data accessed by the view query is performed using the identity of the owner rather than the invoker.
- The VSO must have **CONNECT** privilege for the data service where the view is located.
- The VSO must own a database login to the database in which he is the schema owner, assuming that the database requires a login.
- The VSO must own database logins necessary to connect to the database accessed by the view query.
- The VSO must have **CONNECT** privilege on all data services referenced in the view. Data service references are based on the catalog names that appear in the view query and any queries referenced indirectly from other views.
- The VSO must have **SELECT** privilege on all tables referenced in the view.

## Required Ownership for Federated SQL Views

Objects such as tables and views do not have owners in SAS Federation Server, so ownership is granted on the schema containing the view. For example,

```
ALTER SCHEMA LD_ORAI_SERVICE.TKTSTST1 OWNER TO USER1
```

A definer's rights view must be associated with a schema owner. If an invoker calls a definer view for which the schema has no owner, an error similar to the following is returned:

```
ERROR: Definer's security context for view "%.*s" cannot be established because
the schema container "%.*s"."%.*s" has no configured owner.
```

For additional information about schema ownership, see [“Catalogs and Schemas” on page 105](#).

## Creating Federated SQL Views

### Overview

Views are created using the CREATE VIEW statement or command. You can create views from a single data source or multiple data sources. To create a view the user must have the **CREATE VIEW** privilege on the view schema, or inherited from a parent object. The CREATE VIEW privilege is not necessary for users to create views on schemas where the user is the owner of the schema.

Here is the syntax to create a FedSQL view for a single data source using invoker's rights. To specify a definer's rights view, replace INVOKER with DEFINER:

```
CREATE VIEW MYVIEW SECURITY INVOKER AS SELECT * FROM
CAT1.S1.MYTABLE T1
```

### Create a Federated SQL View from Multiple Data Sources

You can create a FedSQL view across multiple data sources. Suppose that data resides in two separate data sources, one in Oracle and the other in DB2. Using CREATE VIEW, you can tie the two data sources together to create a single federated view of the data.

To create a FedSQL view across two data sources:

1. Create a data service for each of the data sources that you want to access. Make note of the catalog names associated with each of the data sources. In this example, CAT1 and CAT2 are the catalog names. Also make note of the schema within the catalog where your data resides.

```
Data source one: T1 CAT1.S1.MYTABLE
Data source two: T2 CAT2.S2.MYOTHERTABLE
```

2. Invoke the SQL Console window and connect to one or more data services on SAS Federation Server. The data service can be one of the two created above or any other data service associated with the SAS Federation Server that you are using. Normally, you would use a DSN to connect to the data service.
3. Create a statement using Submit:

```
CREATE VIEW MYVIEW SECURITY
INVOKER AS SELECT * FROM CAT1.S1.MYTABLE
T1, CAT2.S2.MYOTHERTABLE T2 WHERE T1.X=T2.X
```

This statement creates an invoker's rights view. For definer's rights view, replace INVOKER with DEFINER.

4. Grant **SELECT** privileges to the users or groups that will access the view.

All users with permissions can now read from the view.

## Dynamic Connections

### Overview

A dynamic connection is a connection made during the execution of a federated SQL view. Dynamic connections are created within the initial set of connections when a user connects to SAS Federation Server but does not include a connection to data referenced within a view. The dynamic connection feature allows an administrator to create views that reference data from any data service defined in SAS Federation Server, but does not require the user's DSN to reference all the data sources in the view.

For example, an administrator creates a view V1, which references data in catalog C1 and C2, where C1 and C2 were defined through separate data services. Without dynamic connections, the administrator would need to create a federated DSN that included a connection to:

- The data service where V1 was stored.
- The data service containing catalog C1.
- The data service containing catalog C2.

Assume the view definition in V1 was changed to reference additional data in catalogs C3 and C4, each coming from a different data service. Without dynamic connections, the administrator would need to modify the user DSN(s) to include references to the data services containing catalogs C3 and C4.

Dynamic connections ease this administration burden because the administrator can simply create a user DSN to include a connection to the data service containing the views. Any data required by the view is accessed through connections made dynamically during view execution. If the view definitions change, the user DSNs do not need to change. This feature is also very useful with data caching, as the data cache can be moved from one data service to another without requiring modifications to user DSNs.

Dynamic connections are transient, so they do not modify the capabilities of the user's original connection properties. Dynamic connections only occur within the context of federated SQL views for invoker's or definer's rights. They do not occur any other place in the system.

### Object Privileges and Required Logins

For dynamic connections to function properly, the invoker or definer, depending on the view type, must have CONNECT privileges on the data service that is referenced by the dynamic connection. Privileges on the DSN are insufficient for dynamic connections to succeed because the underlying connection is made to the data service, and not through the DSN.

Example: USER1 connects across 2 different DSNs: BASEDSN in data service BASE and ORADSN in data service ORA1. USER1 then creates an invoker's rights view that references tables in both DSNs, and stores the view in BASEDSN. USER2 then connects to BASEDSN and issues a 'select \* from VIEW'. The view will succeed only if USER2 has CONNECT privileges on the ORA1 data service.

Also, note that the invoker or definer of the view must have SELECT privileges on any data referenced in the view. Also, the invoker or definer must have the required logins to the data sources that require dynamic connections. If an invoker's rights view requires a dynamic connection to reference data from catalog ORACAT in an Oracle data service, the invoker must have a login to the Oracle database in order to make the connection. The login can either be a personal or shared login.

### Secured Connections

Connections made dynamically during view execution are secured. If a connection is made to an Oracle service known by the ORACAT catalog, and a FedSQL view is read (`select * from ORACAT.schema.view`), and the view query references another catalog (`select * from DB2CAT.schema.table ...`), then the dynamic connection to the DB2 service known by DB2CAT is secured. This is true even if the ORACAT connection is unsecured.

For example, if you connect via `DSN=ORACLE_DSN`, where the DSN is unsecured but configured to use the FEDSQL driver, the DSN might expand to a connection string like this:

```
DRIVER=FEDSQL;conopts=(driver=ORACLE;catalog=ORACAT;...)
```

The data that lives under ORACAT is accessed without additional SAS Federation Server authorization enforcement applied.

If an additional SELECT statement is made:

```
SELECT* from ORACAT.Schema.View.
```

No SELECT privilege check is made against the columns of `ORACAT.Schema.View`.

Now assume that the view content is just a simple indirection that expands to this:

```
SELECT* from DB2CAT.Schema.Table ...
```

If DB2CAT's data service is configured to use the ODBC driver, the server will attempt to dynamically connect to DB2CAT using a secured connection that is equivalent to:

```
DRIVER=FEDSQL;conopts=(driver=ODBC;catalog=DB2CAT;...;
odbc_dsn=DB2DSN)
```

## Data Caching

### Overview

Data caching can be used to manage the performance of frequently accessed data sources to minimize impact on the databases and operational servers. You can free up resources for the high-availability data sources by caching data that is used often and fairly constant. Overall, data caching can have a positive effect on user satisfaction and system performance.

When query optimization alone is not sufficient, caching provides an alternative with greater flexibility than traditional replication and consolidation techniques. Any FedSQLdefiner's rights view can be used to create a cache and caches can be refreshed periodically to remain synchronized with their parent views. Queries can be processed against caches just as if you were accessing the original data source.

A definer's rights view can be cached in any catalog defined in SAS Federation Server, with the exception of any catalog defined in an MDS service. When the cache is

refreshed, the view is executed and the results are stored in a table which is named and maintained by SAS Federation Server. When a user selects from the cached view, results are returned from the data stored in the table, and not from a dynamic execution of the view. This improves performance by eliminating the need to derive a new FedSQL execution plan, fetching data from slow or unavailable data sources, and performing SQL operations such as joins or function evaluations. Instead, the view execution merely fetches the data stored in the table created during the cache refresh.

Data caching can have a positive impact on server performance. You can use data caching to pre-calculate results. Using the following example, a **SELECT** from view V1 would fetch data and calculate the results, returning a single number. If you cache the result anyone selecting from the view will receive the result immediately. For example,

```
CREATE VIEW V1 AS SELECT AVG(B1) FROM A,B,C,D,E
where E .C1 < (SELECT AVG(C1) FROM E) AND B.C1 = A.C1 AND
C.C1 < B.C1 AND D.C1 * E.C1 > SUM(B.C1)
```

You can also cache tables or result sets so that they remain consistent during multiple queries. For example, an ORDERS table is updated continuously during the day as customers purchase products. Caching the data guarantees consistent results and reduces the load on the servers, freeing resources to process incoming orders.

## Views and Caching

SAS Federation Server through the use of FedSQL, allows users to cache data from a definer's rights view, creating a materialized view of the data, also known as the cache table. A cache table is a snapshot of the target view from a specific point in time.

Only definer's rights views can be used to cache data. When a definer's rights view is executed, it uses the credentials of the view's schema owner rather than the current user's credentials, to access catalogs that are referenced in the view. A definer's rights view returns the same result from the underlying database, no matter who is requesting the data. This allows a single copy of the review result set to be cached and consumed by all users. You can use SAS Federation Server authorization, including table, column and row-level security, to provide a granular and user-specific access to the view.

If a definer's rights view is altered to an invoker's rights view, the view is no longer cached.

Before you begin, ensure that the following prerequisites and configuration tasks have been addressed for both definer's rights views and cached views.

## Requirements for Cached Views

*Note:* This section uses the following acronyms:

- **VS** (View Schema) is the schema where the cached view resides.
- **CS** (Cache Schema) is the schema where the cache tables reside.
- **VSO** (View Schema Owner) is the owner of the schema where the cached view resides.
- **CSO** (Cache Schema Owner) is the owner of the schema where the cache tables reside.

Here are the requirements for definer's rights views that are cached:

- A user that is allowed to create or drop cached views within a schema should be granted **CREATE CACHE** privilege on the VS or an object in its inheritance hierarchy. This privilege also allows the user to refresh the cache.

- A user that is allowed to refresh cached views within the schema should be granted ALTER CACHE privilege on the VS, unless they already have CREATE CACHE privilege. ALTER CACHE also allows the user to enable or disable a cache.
- A user that is creating a cache also needs the CONNECT privilege on the ADMIN DSN.
- Only definer's rights views can be cached.
- Administrators implicitly have all privileges, including CREATE CACHE and ALTER CACHE.
- A schema owner must be assigned to the CS object. This is the cache table schema owner, also referred to as the CSO. The CSO effectively owns all cache tables in its schema, and this identity is used for executing the view and saving the cached data.
- The CSO must have CONNECT privilege on the data service the cache tables are in.
- The CSO must have a database login to the database of the cached tables if the cached location requires credentials. This can be a personal login or a shared login. This is needed for cache creation and refresh.
- The CSO must have a database login to the database of the cached view. This can be a personal login or a shared login. The server impersonates the CSO user during cache creation and refresh, and the CSO must be able to select from the original view.

*Note:* During data cache connections, the CSO connects to the databases that contain the CS and the VS using a Credential Search Order of "PERSONAL, SHARED". See [“Credentials Search Order \(CSO\) for DSN Connections” on page 104](#) for more information.

- The VSO must be granted CREATE TABLESPACE privilege on the CS for the cache tables. This privilege allows the server to cache results sets in the CS for views owned by the VSO. The server assumes the identity of the CSO user to create and drop cached tables in the CS, to insert and delete rows in the cache table, and to select data from the cache table during client access.

## Working with Cached Views

### Overview

You can configure cached views using one of these methods:

- Issue administration DDL statements such as CREATE CACHE, REFRESH CACHE, ALTER CACHE, and DROP CACHE.
- Use the Data Cache module in SAS Federation Server Manager.

Administration DDL statements are described in Appendix 1. Procedures for caching data in SAS Federation Server Manager are described in the application's online Help.

The following scenarios describe how data operations are performed using various DDL statements.

### Creating a Cache

Use the [CREATE CACHE DDL statement on page 215](#) to cache a definer's rights view or change an existing cache definition. The CREATE CACHE privilege is required on the view that is being cached.

A cache table is created when the `CREATE CACHE` statement is executed. A cache table is also created when the `ALTER CACHE` statement is executed with the `REFRESH` option. A cache table is a snapshot of the target view from a specific point in time.

The `CREATE CACHE` statement adds entries to the `OBJECTS` and `CONFIG_OBJECTS` information views. Several options are available with the `CREATE CACHE` statement.

### **Alter a Cache**

To alter an existing cache, use the [ALTER CACHE DDL statement on page 217](#) with available options `REFRESH`, `ENABLE` and `DISABLE`.

To refresh an existing cached view, use `ALTER CACHE` with the `REFRESH` option. The `REFRESH` option creates a new cache table which is a snapshot of the target view when the refresh is done. The required privileges are `CREATE CACHE` or `ALTER CACHE` on the view that is being cached or refreshed.

The following example refreshes an existing cached view using the most recent definition:

```
ALTER CACHE [view_catalog_name].[view_schema_name.]view_name REFRESH
```

You can also disable and enable a cached view using `ALTER CACHE`. For additional details, see [“Disabling and Enabling Cached Views” on page 116](#).

### **Purge Cache**

Purge Cache forces the removal of outdated cache tables that are no longer in use. Only system users or those with `ADMINISTER` privilege can execute this DDL statement.

There are two commands that you can use to purge cache tables:

- Use the explicit [PURGE CACHE DDL statement on page 218](#) to activate the cache table cleanup process for all cache tables created through SAS Federation Server. When `PURGE CACHE` is issued, messages are returned indicating the cache tables that were successfully removed and what problems were encountered. This command has no options.

```
PURGE CACHE
```

- Schedule cache table cleanup for certain intervals using the `ALTER SERVER` statement. The syntax to set the time-out is:

```
ALTER SERVER {OPTIONS(xset PURGE_CACHE XX)}
```

where **xx** is the time-out value in minutes.

- A negative value indicates that the cleanup thread will wake only when the `PURGE CACHE` command is issued. It never wakes up automatically.
- A value of `0` indicates that the cleanup thread wakes whenever a `CREATE CACHE`, `ALTER CACHE REFRESH`, `DROP CACHE`, or `PURGE CACHE` statement is issued, or when the view is dropped. Note that this might not clean up all old caches since some cache views might be in use at the time of cleanup.
- A positive value indicates how often (in minutes) the cleanup thread wakes up to remove orphaned cache tables.

*Note:* Cleanup is not run on deferred caches. A cache is deferred when `CREATE CACHE` includes an option value of `[DEFERRED]`.

**Drop Cache**

Use the [DROP CACHE DDL statement on page 218](#) or issue a DROP VIEW command to drop a cache. Invoking DROP VIEW also drops all of the view's associated cache tables. DROP CACHE requires CREATE CACHE privilege on the view.

```
DROP CACHE [view_catalog_name.[view_schema_name.]]view_name [FORCE]
DROP VIEW [view_catalog_name.[view_schema_name.]]view_name [FORCE]
```

**Disabling and Enabling Cached Views****Overview**

In the event that a cached view needs to be taken offline for any reason, it can be temporarily disabled. Disabling a cached view does not drop or delete a cached view. Instead, the cached view is temporarily suspended while the users are rerouted to the original definer's rights view that the cached view was built on. When the cached view is enabled, users are transparently directed back to the actual cached view.

**Disabling a Cached View**

Use the [ALTER CACHE DDL statement on page 217](#) with the DISABLE option to disable a cached view. ALTER CACHE requires the CREATE CACHE or ALTER CACHE privilege on the view.

The following example disables an existing cached view and redirects users to the original definer's rights view while the cached view is offline:

```
ALTER CACHE [view_catalog_name.[view_schema_name.]]view_name DISABLE
```

When the cached view is disabled, the original definer's rights view is used. During the time that the current cached view remains disabled but continues to be reported with a status of suspended, the disabled cache view displays the following behavior:

- An ALTER SERVER REFRESH refreshes the cached view but does not enable the cached view. It remains disabled with a status of suspended.
- A CREATE CACHE behaves normally and the cached view remains disabled in a suspended status.
- An ALTER CACHE ENABLE re-enables the cached view and drops the suspended status.
- An ALTER CACHE DISABLE on a cached view that is disabled in a suspended status, returns a success message.
- An ALTER CACHE ENABLE on a cached view that is not disabled also returns a success message.

**Enabling a Cached View**

Use the [ALTER CACHE DDL statement on page 217](#) with the ENABLE option to reinstate a disabled cached view. Once a cached view is enabled, users are redirected from the original definer's rights view to the actual cached data. ALTER CACHE requires the CREATE CACHE or ALTER CACHE privilege on the view.

The following example enables a cached view that is currently disabled and brings the cached data online:

```
ALTER CACHE [view_catalog_name.[view_schema_name.]]view_name ENABLE
```

---

# Memory Data Store

## Overview

Memory Data Store, or MDS, is a transactional in-memory data store that can be used with SAS Federation Server. The MDS driver supports most of the FedSQL functionality. MDS runs strictly in memory with no backup data store. Therefore, changes are lost when the database is dropped or the server is restarted.

The database is created in memory when the first user connects to the database. The database remains in memory until one of the following conditions is met:

- The server is shut down.
- The data service or the catalog associated with the data service is dropped.

*Note:* You cannot drop a MDS data service or catalog if users are connected to the data service.

You can rename the database and change the memory value while users are connected, but you cannot drop the database while users are connected. To drop or rename a schema, the table within the schema cannot be in use. Users can be connected to the database, but they cannot have a table open in the schema. Also, you cannot drop a table if it is referenced by a prepared statement or has a pending transaction with uncommitted changes.

MDS supports optimistic concurrency providing a serializable transaction, which means that a transaction sees a consistent version of the data when the transaction is started. When an MDS transaction starts, the state of the database is logically frozen at that point in time. The transaction sees the database consistently but changes made by other transactions are not visible until the transaction is committed or rolled back, and its state synchronized so that it sees the new state of the database.

## The MDS Data Service

You can configure an MDS data service and table using one of these methods:

- Use the [CREATE DATA SERVICE DDL statement on page 199](#).
- Use the **New Data Service** function in SAS Federation Server Manager.

You can create multiple data services if needed. See [“MDS Driver Reference” on page 136](#) for a list of connection options for MDS.

Each MDS data service catalog contains a pre-defined, read-only schema, named SYSTEMINFO. The SYSTEMINFO schema contains an auto-generated MEMORY table. You will need to define at least one additional schema to use MDS. To define a schema use the CREATE SCHEMA DDL statement or the **New Schema** function in SAS Federation Server Manager. Schemas cannot be modified, renamed or dropped while there are active connections to the database.

## The MDS Database

### Limiting the Database Memory Size

To limit the memory size for a database, use the connection string option `MAXDBMEM=` that specifies the maximum size of memory to be used to store all rows of data in the database. This includes committed rows and pending row versions (INSERT, UPDATE, and DELETE operations that have not yet been committed or rolled back). If an INSERT, UPDATE, or DELETE operation exceeds this limit, an out of memory error is returned. For additional information, see Connection Options in the “MDS Driver Reference” on page 136.

### The MEMORY Table

The MEMORY table contains information about memory usage and is always available. The first row contains statistics about the database. Additional rows provide information about each of the tables in the database.

The MEMORY table includes the following columns:

**Table 6.1** Columns in `SYSTEMINFO.MEMORY` Table

Column Name	Description
DB_NAME	The name of the current database. This will be the same as the catalog name.
SCHEMA_NAME	The name of the schema for the table (NULL for the database info row).
TABLE_NAME	The name of the table (NULL for the database info row).
ROW_COUNT	The number of rows in the table (NULL for the database info row).
ROW_SIZE	The size of a single row in the table (NULL for the database info row).
MEM_SIZE	The current memory used by the table and database for data.
MEM_PEAK	The peak memory used by the table and database since creation.
MEM_LIMIT	The maximum memory this database can use (NULL for table info rows). This value corresponds to the <code>MAXDBMEM=</code> option specified when the database was created.

To view the MEMORY table using SAS Federation Server Manager:

1. In the navigation tree, open the MDS data service.
2. Select the catalog associated with the MDS data service.
3. Select the SYSTEMINFO schema to open it. The MEMORY table appears under the SYSTEMINFO schema.

---

# Understanding Data Federation and Best Practices

## Overview

Successful data federation projects require careful preparation and attention in two areas:

- data model
- data security

Proper understanding of your organization's data access needs is key to a successful data federation deployment. The data model controls the type of underlying work required at run time to satisfy requests for data. A poorly constructed data model can result in inefficient performance and incorrect results if the data is not well understood. Data security design should consider how your users access the back-end data. Will users access the data under their individual authorizations? Will you establish one or more data owners who act as data access proxies for the end users? There are several options to consider in the design phase so that SAS Federation Server can be properly configured.

## Data Model

A good starting place is to identify your data sources, understand the relationships between different sets of data, and derive a data model that meets the needs of your business. This is the same type of background work that typically goes into a data warehousing project, where the end result is often a set of tables or views that are loaded in a data mart. However, unlike data warehousing, with SAS Federation Server, the data does not need to be copied from the source into a data mart. Instead, the data can be fetched from their source locations and processed in real time during the data requests. These operations should be as efficient as possible and require a well-planned data model.

Data caching should be considered when designing the data model. It can provide a vital role in optimizing query performance by pre-executing and storing intermediate results. This can be particularly useful for back-end data sources that have low availability, slow access speeds, expensive access fees, or contain data where real-time values are not required. A data model can consist of a set of user-visible FedSQL views that are dependent on other restricted FedSQL views, which can be dependent on back-end data, such as tables and database views. FedSQL views are very similar to relational database views, except that the data can come from a heterogeneous set of data sources.

Any combination of FedSQL views can be cached, although the views must be definer's rights views. The cached views can be refreshed on a periodic basis through the scheduler, or refreshed at any time through direct administration SQL commands. If certain data sets will be used frequently or involve complicated SQL operations to return the results, you should consider data caching for those results. Cached results can be joined with uncached data to provide quick responses to user queries.

Another consideration is to use the Memory Data Store (MDS) to store frequently used or temporary data. MDS allows tables to be stored in the memory of the server process. This allows for extremely fast data access performance. These tables must be managed manually and are automatically deleted when the server is shut down. FedSQL can join

an MDS table with tables from any other data source, whether they are cached or uncached, including other MDS tables.

## Data Security

### Overview

In a federated system, data can be acquired from a large variety of data sources, and users might not always have direct access to those systems. Even if they do, administering a large number of database credentials and setting up database privileges so that users can access data with personal credentials can be burdensome. The security capabilities of SAS Federation Server provide some alternatives that can help ease security administration.

### Invoker's Rights Views

If users already have direct access to the back-end systems and you want them to access the data under their individual or shared logins, then you can configure SAS Federation Server to access the data under the rights of the invoking user. This is done by creating invoker's rights FedSQL views. When these views execute, any connections to back-end data require the invoking user to have proper credentials to access the data. In addition, you should ensure that SAS Federation Server security settings allow invokers to access the required data. For example, if a FedSQL view is defined to select data from an Oracle data source and a DB2 data source, the administrator needs to ensure that proper privileges are assigned to each invoker of the view on the data source, as well as the FedSQL view itself. The privileges are granted to a group, or set of groups for which the invokers are members, instead of to each individual invoker.

Because users must be able to directly access all of their data, you should secure each object for each user. Also, invoker's rights views need to be sensitive to these security settings on underlying objects, particularly column level security. Consider a view that selects columns C1 and C2 from table T, as shown in the following example:

```
SELECT C1, C2 FROM T
```

If table T has been secured through the Federation Server such that User1 does not have SELECT privilege on the table, then when User1 selects from the view, User1 receives an error when attempting to access table T. In this case, you might consider denying SELECT privilege on the view to User1. Furthermore, if User2 has SELECT privilege on table T1 and column C1 but not on column C2, then User2 receives an error when selecting from the view. Assigning column-level security on the view that denies SELECT privilege on column C2 to User2 does not help, because the underlying view definition specifically requires access to column C2 in table T. When creating views that require data from other objects with column-level security, you might consider selecting all columns using an asterisk (\*). Here is an example:

```
SELECT * FROM T
```

When the Federation Server is configured with `SelectStarExpansion=VISIBLE`, the \* expands all columns for which the invoking user has SELECT privilege. This enables you to create a single view that can be used by all invokers, yet each invoker sees only the columns for which the invoker has SELECT privilege.

### Definer's Rights Views

If your users do not have direct access to the back-end systems, or you want to read objects under the authorizations of a single user, then you can configure SAS Federation Server to access the data under the rights of the defining user. This is accomplished by

creating definer's rights FedSQL views. When these views are executed by a user, connections to back-end data are made under the definer of the view. The definer is actually the owner of the schema that contains the view.

In this security model, you typically deny access to the back-end data to all users except the view definer. Any column-level, table-level, or row-level security will be set on the view itself. For example, if your view is defined as **SELECT C1, C2 FROM T**, you should ensure that the view definer has full access to columns C1 and C2 in table T.

All other users do not require access to back-end data. And in many cases, if your users are interfacing only with exposed top-level objects of the data model, then you can deny privileges to those users on back-end objects. Individual security settings can then be consolidated on the exposed objects of the data model, which are usually FedSQL views. Then users access the views only, and no other credentials are required. This greatly simplifies the security settings in the Federation Server and reduces administration on all the back-end data sources as well.

### **Mixed Models**

Note that you can use a combination of approaches with both invoker's and definer's rights views intermingled. FedSQL views can be nested, and view types are honored. For example, you might create view V1 that is owned by Owner1 and data is accessed through the credentials of Owner1. However, if view V1 selects from view V2, which is another definer's rights view but owned by Owner2, then all data within view V2 is accessed under the authorizations of Owner2. If view V2 selects from view V3, which is an invoker's rights view, then any data within view V3 continues to be accessed by Owner2, who is the invoker of the view.

### **Dropping Objects**

SAS Federation Server persists metadata for security settings in its set of system tables. The server attempts to synchronize this metadata with the actual objects that it represents. For example, if the administrator has granted SELECT privilege on table T1 to User1, and subsequently table T1 is dropped, the server will likewise drop its security metadata onto T1. If T1 is subsequently re-created, it will have no security set on the object itself. All security definitions for T1 come from its inheritance objects, including the schema, catalog, data service, and server.

If you prefer that the security metadata is retained on the server, there are a couple suggested possibilities. First, it might be that a job is dropping the table only to re-create it in a subsequent step. This is often the case when refreshing the contents of a table. Rather than dropping and re-creating the table, which might also affect indexes on the table, consider issuing a **DELETE FROM TABLE** command to delete all the rows from the table. In doing this, the underlying table remains intact, as well as the security metadata stored in the system tables.

A second approach is to use definer's rights views to assist with individual object security. Often, data in back-end data sources is being manipulated (created and dropped) by a few power users only. These users require CREATE, DROP TABLE, and VIEW privileges. The CREATE privilege set applies only to container objects (schema, catalog, data service, and server) and not individual tables or views. The DROP privilege can be applied at the object level. However, if designed correctly, the DROP privilege can be used from the container object.

Business users are typically the ones requiring table-level and column-level privileges to access the data. If you choose the security model using definer's rights views, then you can place your individual privileges on the views, which do not go away when a back-end table is dropped and re-created. Back-end data can be secured by denying privileges on the container object to business users, but granting SELECT to view owners. An

approach similar to this can be used to eliminate the need for table-level, column-level, and row-level security on back-end data source objects.

## Chapter 7

# Driver Reference for SAS Federation Server

---

<b>Database Functionality and Driver Performance</b> .....	<b>124</b>
<b>DB2 Driver Reference</b> .....	<b>124</b>
Understanding the SAS Federation Server Driver for DB2 .....	124
Data Service Connection Options for DB2 .....	125
DB2 Wire Protocol Driver Usage Notes .....	128
<b>FedSQL Driver Reference</b> .....	<b>129</b>
Overview .....	129
Connection Options .....	129
<b>Greenplum Driver Reference</b> .....	<b>132</b>
Understanding the SAS Federation Server Driver for Greenplum .....	132
Data Service Connection Options for Greenplum .....	132
Greenplum Wire Protocol Driver Usage Notes .....	136
<b>MDS Driver Reference</b> .....	<b>136</b>
About the MDS Driver .....	136
Data Service Connection Options for MDS .....	137
<b>Netezza Driver Reference</b> .....	<b>139</b>
Understanding the SAS Federation Server Driver for Netezza .....	139
Data Service Connection Options for Netezza .....	140
<b>ODBC Driver Reference</b> .....	<b>143</b>
About ODBC .....	143
Understanding the SAS Federation Server Driver for ODBC .....	143
Data Service Connection Options for ODBC .....	144
Configuring ODBC for Hadoop .....	149
Wire Protocol Driver Usage Notes .....	150
<b>Oracle Reference</b> .....	<b>151</b>
Understanding the SAS Federation Server Driver for Oracle .....	151
Data Service Connection Options for Oracle .....	151
Oracle Wire Protocol Driver Usage Notes .....	156
<b>PostgreSQL Driver Reference</b> .....	<b>156</b>
Understanding the SAS Federation Server Driver for PostgreSQL .....	156
Connection Options for PostgreSQL .....	157
<b>SAP Reference</b> .....	<b>161</b>
Understanding the SAS Federation Server Driver for SAP .....	161
Data Service Connection Options for SAP .....	161
Installing and Configuring the SAS Federation Server Driver for SAP .....	165
Installing SAP Components .....	168

<b>SAP HANA Driver Reference</b> .....	<b>175</b>
Understanding the SAS Federation Server Driver for SAP HANA .....	175
Data Service Connection Options for SAP HANA .....	175
Secure Sockets Layer (SSL) Connection Options .....	177
Advanced Connection String Options .....	178
<b>SAS Data Set Reference</b> .....	<b>182</b>
Overview .....	182
Understanding the Driver for Base SAS .....	182
Connection Options for SAS Data Sets .....	183
<b>SASHDAT Driver Reference</b> .....	<b>186</b>
About the SAS Federation Server Driver for SASHDAT .....	186
Understanding the SASHDAT Data Source Driver .....	186
Connection Options .....	186
Example Connection Strings .....	188
<b>Teradata Reference</b> .....	<b>189</b>
Understanding the SAS Federation Server Driver for Teradata .....	189
Data Service Connection Options for Teradata .....	189

---

## Database Functionality and Driver Performance

Because SAS Federation Server supports several data sources, a broad range of database functionality that is unique to each data source is provided. For example, a particular data source provides transaction support while another data source might not provide transaction support but supports indexes and integrity constraints.

You must understand database functionality and how its implementation affects processing, performance, and integrity of your data in order to determine which data sources are most appropriate for different types of applications. Because database functionality is unique to each data source, you cannot make assumptions about the data source to be accessed. For example, an application cannot request a locking level just because that locking level is more efficient. An application must respond to the attributes of a SAS Federation Server driver.

Database functionality is applied through the SAS Federation Server driver when the application submits requests. Requests can be in the form of FedSQL statements or the SQL statements that are the implementation of the data service. For additional information about database functionality including data types, see the *SAS FedSQL Reference Guide*.

---

## DB2 Driver Reference

### ***Understanding the SAS Federation Server Driver for DB2***

The SAS Federation Server Driver for DB2 (Driver for DB2) enables SAS Federation Server to read and update legacy DB2 tables. In addition, the driver creates DB2 tables that can be accessed by both SAS Federation Server and the DB2 database management system (DBMS).

The Driver for DB2 supports most of the FedSQL functionality. The driver also supports an application's ability to submit native DB2 SQL statements.

The SAS Federation Server Driver for DB2 is a remote driver, which means that it connects to a server process in order to access data. The process might be running on the same machine as SAS Federation Server, or it might be running on another machine in the network.

## Data Service Connection Options for DB2

### Overview

To access data that is hosted on SAS Federation Server, a client must submit a connection string, which defines how to connect to the data. The data service connection arguments for DB2 include connection options and advanced options.

*Note:* When performing connections through DSNs or connection strings, SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the *SAS FedSQL Reference Guide*.

### Connection Options

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

SAS Federation Server Driver for DB2 supports the following connection options for DB2 data sources.

Option	Description
CATALOG	<p><b>CATALOG=</b><i>catalog-identifier</i>;</p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups logically related schemas. Any identifier is valid (for example, <b>catalog=DB2</b>). You must specify a catalog. For the DB2 database, this is a logical catalog name to use as an SQL catalog identifier.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>
DATABASE DB	<p><b>DATABASE=</b><i>database-specification</i>;</p> <p>Specifies the name of the DB2 database, for example, <b>database=sample</b>, <b>DB=sample</b>.</p> <p><i>Note:</i> You must specify a database name.</p>
DRIVER	<p><b>DRIVER=DB2;</b></p> <p>Identifies the DB2 data source to which you want to connect.</p> <p><i>Note:</i> You must specify the driver.</p>

### Advanced Connection Options

SAS Federation Server Driver for DB2 supports the following advanced connection options for DB2 data sources.

Option	Description
CLIENT_ENCODING	<p data-bbox="443 369 906 394"><b>CLIENT_ENCODING=encoding-value</b></p> <p data-bbox="443 415 1318 470">Used to specify the encoding of the DB2CODEPAGE to the DB2 driver. When using this option, you must also set the DB2CODEPAGE environment variable on the client.</p> <p data-bbox="443 487 1347 684">When the encoding of the DB2 client layer (stored in DBCODEPAGE) is different from the encoding value of the DB2 operating system value, which is generally the SAS® session encoding value, the DB2 client layer attempts to convert incoming data to the DB2 encoding value that is stored in DB2CODEPAGE. To prevent the client layer from converting data incorrectly, you must first determine the correct value for DB2CODEPAGE and then set the CLIENT_ENCODING= option to match the corresponding encoding value in DB2CODEPAGE.</p> <p data-bbox="443 701 1378 844">For example, suppose you are storing Japanese characters in a DB2 database and the client machine where the DB2 driver is executing is a Windows® machine running CP1252 encoding. When the application tries to extract the data into SAS Federation Server, the DB2 client layer attempts to convert these Japanese characters into Latin1 representation, which does not contain Japanese characters. As a result, a garbage character appears to indicate a failure in transcoding.</p> <p data-bbox="443 861 1378 1087">To resolve this situation, you must first set the DB2CODEPAGE environment variable value to 1208 (the IBM® code page value that matches UTF-8 encoding) to specify that the DB2 client layer send the data to the application in UTF-8 instead of converting it into Latin1. In addition, you must specify the corresponding encoding value of DB2CODEPAGE because the SAS Federation Server Driver for DB2 cannot derive this information from a DB2 session. For this particular Windows case, set the CLIENT_ENCODING= option to the UTF-8 to match the DB2CODEPAGE value (1208) in order to specify the DB2CODEPAGE value to the DB2 driver.</p> <p data-bbox="443 1104 1350 1188">However, changing the value of DB2CODEPAGE affects all applications that run on that machine. You should reset the value to the usual DB2CODEPAGE value, which was derived when the database was created.</p> <p data-bbox="443 1205 1378 1289"><i>Note:</i> Setting the DB2CODEPAGE value or the CLIENT_ENCODING= value incorrectly can cause unpredictable results. You should set these values only when a situation such as the example above occurs.</p> <p data-bbox="443 1306 1270 1331"><i>Note:</i> You can specify any valid encoding value for CLIENT_ENCODING=option.</p>
CT_PRESERVE	<p data-bbox="443 1373 1212 1398"><b>CT_PRESERVE=STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p data-bbox="443 1419 1337 1503">Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul data-bbox="443 1520 1382 1902" style="list-style-type: none"> <li data-bbox="443 1520 1382 1575">• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li data-bbox="443 1591 1382 1675">• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="443 1692 1382 1810">• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="443 1827 1382 1902">• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>

Option	Description
DEFAULT_ATTR	<p><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p>Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul style="list-style-type: none"> <li>• <b>CURSORS=n</b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul style="list-style-type: none"> <li>0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li>1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li>2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p>Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> <ul style="list-style-type: none"> <li>• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li>• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>
DRIVER TRACE	<p><b>DRIVER_TRACE='API   SQL   ALL' ;</b></p> <p>Requests tracing information, which logs transaction records to an external file that can be used for debugging purposes. The SAS Federation Server driver writes a record of each command that is sent to the database to the trace log based on the specified tracing level, which determines the type of tracing information. The tracing levels are:</p> <ul style="list-style-type: none"> <li>• <b>ALL</b> Activates all trace levels.</li> <li>• <b>API</b> Specifies that API method calls be sent to the trace log. This option is most useful if you are having a problem and need to send a trace log to Technical Support for troubleshooting.</li> <li>• <b>DRIVER</b> Specifies that driver-specific information be sent to the trace log.</li> <li>• <b>SQL</b> Specifies that SQL statements that are sent to the database management system (DBMS) be sent to the trace log. Tracing information is DBMS specific, but most SAS Federation Server drivers log SQL statements such as SELECT and COMMIT.</li> </ul> <p><b>Default:</b> Tracing is not activated.</p> <p><i>Note:</i> If you activate tracing, you must also specify the location of the trace log with <b>DRIVER_TRACEFILE=</b>. Note that <b>DRIVER_TRACEFILE=</b> is resolved against the <b>TRACEFILEPATH</b> set in <b>ALTER SERVER</b>. <b>TRACEFILEPATH</b> is relative to the server's content root location.</p> <p>(Optional) You can control trace log formatting with <b>DRIVER_TRACEOPTIONS=</b>.</p> <p><b>Interaction:</b> You can specify one trace level, or you can concatenate more than one by including the   (OR) symbol. For example: <b>driver_trace='api   sql'</b> generates tracing information for API calls and SQL statements.</p>

Option	Description
DRIVER TRACE FILE	<p><b>DRIVER_TRACEFILE= 'filename';</b></p> <p>Used to specify the name of the text file for the trace log. Include the filename and extension in single or double quotation marks. For example:  <b>driver_tracefile= '\mytrace.log'</b></p> <p><b>Default:</b> The default TRACEFILE location applies to a relative filename, and it is placed relative to TRACEFILEPATH.</p> <p><b>Requirement:</b> DRIVER_TRACEFILE is required when activating tracing using DRIVER_TRACE.</p> <p><b>Interaction:</b> (Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p>
DRIVER TRACE OPTIONS	<p><b>DRIVER_TRACEOPTIONS=APPEND   THREADSTAMP   TIMESTAMP;</b></p> <p>Specifies options in order to control formatting and other properties for the trace log:</p> <ul style="list-style-type: none"> <li>• <b>APPEND</b> Adds trace information to the end of an existing trace log. The contents of the file are not overwritten.</li> <li>• <b>THREADSTAMP</b> Prepends each line of the trace log with a thread identification.</li> <li>• <b>TIMESTAMP</b> Prepends each line of the trace log with a time stamp.</li> </ul> <p><b>Default:</b> The trace log is overwritten with no thread identification or time stamp.</p>
PASSWORD	<p><b>PWD=password</b></p> <p>Specifies the password for DB2..</p>
USER ID	<p><b>UID=user-id;</b></p> <p>Specifies the DB2 login user ID.</p>

### DB2 Wire Protocol Driver Usage Notes

SAS Federation Server provides a number of wire protocol ODBC drivers that communicate directly with a database server, without having to communicate through a client library. When you configure the ODBC drivers on Windows or UNIX, you have the opportunity to set certain options. SAS products run best when these options are selected. Some, but not all, are selected by default.

Windows	The options are located on the <b>Advanced</b> or <b>Performance</b> tabs in the ODBC Administrator.
UNIX	The options are available when configuring data sources using the <b>dfdbconf</b> tool. Values can also be set by editing the <b>odbc.ini</b> file in which their data sources are defined.

*Note:* The behavior of a DSN using a wire protocol driver with the catalog option selected, returns only the schemas that have associated tables or views. To list all existing schemas, create a DSN without the catalog option selected.

When configuring an ODBC DSN using the DB2 Wire Protocol driver, set the following advanced options:

- **Application Using Threads**

---

## FedSQL Driver Reference

### Overview

The FedSQL language driver supports the FedSQL dialect, as documented in the *SAS FedSQL Language Reference Guide*. When loaded, the FedSQL driver parses SQL requests, and then sends the parsed query to the appropriate SAS Federation Server driver to determine whether the functionality can be handled by the data service. The FedSQL driver includes an SQL processor which supports the FedSQL dialect. The main emphasis of the FedSQL driver is to support federation of data sources. For example, if an SQL submission is requesting data from DB2 to be joined with data from Oracle, the SQL processor will request the data from the data sources and then perform the join in SAS Federation Server. The FedSQL driver supports the FedSQL dialect over any data source. For example, if the SQL request is from a single data source that does not support a particular SQL function, the FedSQL processor guarantees implementation of the request.

The FedSQL driver is also required for SAS Federation Server SQL Authorization Enforcement. If the DSN is configured to enable Federation Server SQL Authorization Enforcement, then the FedSQL driver is automatically loaded and used. The FedSQL dialect can also be requested when creating a DSN by choosing the FedSQL dialect for the DSN.

The FedSQL driver supports various connection options, as outlined below. You must use the `DEFAULT_ATTR` option in the `CREATE DSN` statement to specify these options. For example,

```
CREATE DSN MYDSN UNDER "Oracle Service" CONNECT
'DEFAULT_ATTR=(SQL_MAX_COL_SIZE=500);DRIVER=ORACLE'
```

### Connection Options

- `DEFAULT_CATALOG=`*catalog-name* - Used to specify the name of the catalog to set as the current catalog upon connecting. This option is useful for SQL Server connections and federated connections.
- `DEFAULT_ATTR=(`*attr=value*`;`...`)` - Used to specify connection handle or statement handle attributes supported for initial connect-time configuration. Where `attr=value` corresponds to any of the following options:

---

**SQL\_CURSORS=*n***

FedSQL connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2.

- A value of 0 causes the driver to use client side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.
- A value of 1 causes the driver to always use client side static cursor emulation if a scrollable cursor is requested. The database server's native cursor will never be used.
- A value of 2 (default) causes the driver to never use client side static cursor emulation if a scrollable cursor is requested. The database server's native cursor will be used if available, otherwise the cursor will be forward only.

DEFAULT\_ATTR= (SQL\_CURSORS=2)

---

**SQL\_AC\_BEHAVIOR=*n***

FedSQL connection handle option. Specifies whether FedSQL should use transactions when processing complex operations. For example, "**CREATE TABLE xxx AS SELECT yyy FROM zzz**" or a multi-row delete statement that requires multiple operations to delete the underlying rows. Possible values are 0 (default), 1 and 2.

- A value of **0** (default) means that no transactions are attempted under-the-covers and operations such as emulated UPDATE, DELETE or INSERT.
- A value of **1** means that FedSQL tries to use transaction to better support the correct behavior when AUTOCOMMIT is ON (where individual operations like UPDATE, DELETE and INSERT should be atomic).
- A value of **2** means that transactions are required. This option will fail if the underlying drivers do not support transactions.

DEFAULT\_ATTR= (SQL\_AC\_BEHAVIOR=0)

---

**SQL\_MAX\_COL\_SIZE=*n***

FedSQL statement handle option. Allows a user to specify the size of the **varchar** or **varbinary** that is used for the potentially truncated long data when direct bind is not possible.

- The default value is 32767.
- The limit for this size is 1 MG. If the value exceeds 1 MG, FedSQL resets the value and returns an **Option value changed** warning.

DEFAULT\_ATTR= (SQL\_MAX\_COL\_SIZE=1048576)

---

**SQL\_PUSHDOWN=*n***

FedSQL statement handle option. This option tells FedSQL if and when it should try to push down SQL to the underlying driver. The values are 8, 2 or 0 (default).

- A value of 8: (PLAN\_FORCE\_PUSHDOWN\_SQL) - Complete statement pushdown is required. If that is not possible, then the INSERT, UPDATE, DELETE, or CREATE TABLE AS statement should fail.
- A value of 2: (PLAN\_DISABLE\_PUSHDOWN\_SQL) - Specifies that the INSERT, UPDATE, DELETE, or CREATE TABLE AS statement not be pushed down to the underlying driver.
- A value of 0 (default): Specifies that the FedSQL processor decides whether the INSERT, UPDATE, DELETE, or CREATE TABLE AS statement should be pushed down to the underlying driver.

DEFAULT\_ATTR= (SQL\_PUSHDOWN=0)

---

---

**SQL\_STMT\_MEM\_LIMIT=*n***

FedSQL statement handle option. Used to control the amount of memory available to FedSQL to answer SQL requests.

- (*n*)umber is treated as an integer and is specified in bytes.
- The following example allows 200 MB of memory:

```
DEFAULT_ATTR= (SQL_STMT_MEM_LIMIT=209715200)
```

---

**SQL\_TXN\_EXCEPTIONS=*n***

FedSQL connection handle option. Supports dynamic connections regardless of the specified transaction isolation. Possible values are 0 or 2 (default).

- Specify a value of 0 to disable support for dynamic connections.
- Specify a value of 2 to enable support for dynamic connections.

```
DEFAULT_ATTR= (SQL_TXN_EXCEPTIONS=2)
```

---

**SQL\_USE\_EVP=*n***

FedSQL statement handle option. This option optimizes the driver for large result sets. The possible values are 0 or 1 (default) and are used as follows:

- Specify 0 to turn optimization OFF.
- Specify 1 to enable optimization (ON).

```
DEFAULT_ATTR= (SQL_USE_EVP=0)
```

---

**SQL\_VDC\_DISABLE=*n***

FedSQL statement handle option. This option is used to allow or disallow use of cached data for a statement. The possible values are 0 (default) or 1 and are used as follows::

- Specify a value of 0 to enable cached data.
- Specify a value of 1 to disable cached data.

```
DEFAULT_ATTR= (SQL_VDC_DISABLE=1)
```

---

**SQL\_XCODE\_WARN=*n***

FedSQL statement handle option. Used to warn if there is an error while transcoding data during row input or output operations. Possible values are 0 (default), 1 or 2 and are used as follows::

- Specify 0 to return an error if data cannot be transcoded.
- Specify 1 to return a warning if data cannot be transcoded.
- Specify 2 to ignore transcoding errors.

```
DEFAULT_ATTR= (SQL_XCODE_WARN=1)
```

---

## Greenplum Driver Reference

### Understanding the SAS Federation Server Driver for Greenplum

The SAS® Federation Server Driver for Greenplum (Driver for Greenplum) enables SAS Federation Server to read and update Greenplum tables. In addition, the driver creates Greenplum tables that can be accessed by both SAS Federation Server and Greenplum.

The Driver for Greenplum supports most of the FedSQL functionality. The driver also supports the application's ability to submit native Greenplum SQL statements.

The SAS® Federation Server Driver for Greenplum is a remote driver, which means that it connects to a server process in order to access data. The process might be running on the same machine as SAS Federation Server, or it might be running on another machine in the network.

### Data Service Connection Options for Greenplum

#### Connection Options

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

The Driver for Greenplum supports the following connection options.

Option	Description
CATALOG	<p><b>CATALOG=</b><i>catalog-identifier</i>;</p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups logically related schemas. Any identifier is valid (for example, <b>catalog=gps_test</b>). You must specify a catalog. For the Greenplum database, this is a logical catalog name to use as an SQL catalog identifier.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>
DATABASE	<p><b>DATABASE=</b><i>database-name</i>;</p> <p>Identifies the database to which you want to connect, which resides on the server previously specified through the SERVER option.</p>
DRIVER	<p><b>DRIVER=GREENPLUM;</b></p> <p>Specifies the data service for the Greenplum database to which you want to connect. Required. You must specify a driver.</p>
DSN	<p><b>DSN=</b><i>data_source_identifier</i>;</p> <p>Identifies the data source name to which you want to connect.</p>
SERVER	<p><b>SERVER=</b><i>server_name</i>;</p> <p>Identifies the name of the server where the Greenplum database resides.</p>

## Advanced Connection Options

The Driver for Greenplum supports the following advanced connection options.

Option	Description
ALLOW UNQUOTED NAMES	<p><b>ALLOW_UNQUOTED_NAMES=NO   YES;</b></p> <p>Specifies whether to enclose table and column names in quotation marks. Tables and columns are quoted when this option is set at NO (default). If set to YES, the driver will not automatically add quotation marks to table and column names if they are not specified. This allows Greenplum tables and columns to be created in the default lower case.</p>
CLIENT ENCODING	<p><b>CLIENT_ENCODING=cei;</b></p> <p>Specifies an encoding, different from the default, to use on the client.</p>
CT_PRESERVE	<p><b>CT_PRESERVE = STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p>Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li>• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>

Option	Description
DEFAULT_ATTR	<p data-bbox="456 243 799 268">DEFAULT_ATTR=(attr=value;...)</p> <p data-bbox="456 285 1382 342">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="456 359 1382 768" style="list-style-type: none"> <li data-bbox="456 359 1382 768">• <b>CURSORS=n</b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="746 449 1382 768" style="list-style-type: none"> <li data-bbox="746 449 1382 533">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="746 554 1382 638">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="746 659 1382 768">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="488 789 951 814">Example: <b>DEFAULT_ATTR = (CURSORS=2)</b></p> <ul data-bbox="456 835 1382 1050" style="list-style-type: none"> <li data-bbox="456 835 1382 919">• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR = (USE_EVP=0)</b></li> <li data-bbox="456 940 1382 1050">• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR = (XCODE_WARN=1)</b></li> </ul>
DRIVER TRACE	<p data-bbox="456 1094 927 1119"><b>DRIVER_TRACE='API   SQL   ALL';</b></p> <p data-bbox="456 1136 1382 1251">Requests tracing information, which logs transaction records to an external file that can be used for debugging purposes. The SAS Federation Server driver writes a record of each command that is sent to the database to the trace log based on the specified tracing level, which determines the type of tracing information. The tracing levels are:</p> <ul data-bbox="456 1268 1382 1535" style="list-style-type: none"> <li data-bbox="456 1268 1382 1293">• <b>ALL</b> Activates all trace levels.</li> <li data-bbox="456 1310 1382 1394">• <b>API</b> Specifies that API method calls be sent to the trace log. This option is most useful if you are having a problem and need to send a trace log to Technical Support for troubleshooting.</li> <li data-bbox="456 1411 1382 1436">• <b>DRIVER</b> Specifies that driver-specific information be sent to the trace log.</li> <li data-bbox="456 1453 1382 1535">• <b>SQL</b> Specifies that SQL statements that are sent to the database management system (DBMS) be sent to the trace log. Tracing information is DBMS specific, but most SAS Federation Server drivers log SQL statements such as SELECT and COMMIT.</li> </ul> <p data-bbox="456 1556 775 1581"><b>Default:</b> Tracing is not activated.</p> <p data-bbox="456 1598 1382 1713"><i>Note:</i> If you activate tracing, you must also specify the location of the trace log with DRIVER_TRACEFILE=. Note that DRIVER_TRACEFILE= is resolved against the TRACEFILEPATH set in ALTER SERVER. TRACEFILEPATH is relative to the server's content root location.</p> <p data-bbox="456 1730 1270 1755">(Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p> <p data-bbox="456 1772 1382 1850"><b>Interaction:</b> You can specify one trace level, or you can concatenate more than one by including the   (OR) symbol. For example: <b>driver_trace='api   sql'</b> generates tracing information for API calls and SQL statements.</p>

Option	Description
DRIVER TRACEFILE	<p><b>DRIVER_TRACEFILE='filename';</b></p> <p>Used to specify the name of the text file for the trace log. Include the filename and extension in single or double quotation marks. For example:  <b>driver_tracefile='\mytrace.log'</b></p> <p><b>Default:</b> The default TRACEFILE location applies to a relative filename, and it is placed relative to TRACEFILEPATH.</p> <p><b>Requirement:</b> DRIVER_TRACEFILE is required when activating tracing using DRIVER_TRACE.</p> <p><b>Interaction:</b> (Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p>
DRIVER TRACE OPTIONS	<p><b>DRIVER_TRACEOPTIONS=APPEND   THREADSTAMP   TIMESTAMP;</b></p> <p>Specifies options in order to control formatting and other properties for the trace log:</p> <ul style="list-style-type: none"> <li>• <b>APPEND</b> Adds trace information to the end of an existing trace log. The contents of the file are not overwritten.</li> <li>• <b>THREADSTAMP</b> Prepends each line of the trace log with a thread identification.</li> <li>• <b>TIMESTAMP</b> Prepends each line of the trace log with a time stamp.</li> </ul> <p><b>Default:</b> The trace log is overwritten with no thread identification or time stamp.</p>
MAX_BINARY_LEN	<p><b>MAX_BINARY_LEN=value;</b></p> <p>Specifies a value to limit the length of long binary fields (LONG VARBINARY). As opposed to other databases, Greenplum does not have a size limit for long binary fields.</p>
MAX_CHAR_LEN	<p><b>MAX_CHAR_LEN=value;</b></p> <p>Specifies a value to limit the length of character fields (CHAR and VARCHAR). As opposed to other databases, Greenplum does not have a size limit for character fields.</p>
MAX_TEXT_LEN	<p><b>MAX_TEXT_LEN=value;</b></p> <p>Specifies a value to limit the length of long character fields (LONG VARCHAR). As opposed to other databases, Greenplum does not have a size limit for long character fields.</p>
NUM BYTES PER CHAR	<p><b>NUMBYTESPERCHAR=value;</b></p> <p>Specifies the default number of bytes per character.</p>
PASSWORD	<p><b>PASSWORD=password;</b></p> <p>Specifies a password for the ID passed through the USER= option. The alias is PWD=.</p> <p><i>Note:</i> You must specify the PASSWORD= option.</p>
SCHEMA	<p><b>SCHEMA=value;</b></p> <p>Specifies the default schema for the connection. If not specified, the schema (or list of schemas) will be determined based on the value of the schema search path defined on the database server.</p>
STRIP BLANKS	<p><b>STRIP_BLANKS=value;</b></p> <p>Specifies whether to strip blanks from character fields.</p>

Option	Description
USER	<p><b>USER=user-id;</b></p> <p>Specifies a Greenplum user ID. If the ID contains blanks or national characters, enclose it in quotation marks. The alias is UID=.</p> <p><i>Note:</i> You must specify the USER= option.</p>

### Greenplum Wire Protocol Driver Usage Notes

SAS Federation Server provides a number of wire protocol ODBC drivers that communicate directly with a database server, without having to communicate through a client library. When you configure the ODBC drivers on Windows or UNIX, you have the opportunity to set certain options. SAS products run best when these options are selected. Some, but not all, are selected by default.

Windows	The options are located on the <b>Advanced</b> or <b>Performance</b> tabs in the ODBC Administrator.
UNIX	The options are available when configuring data sources using the <b>dfdbconf</b> tool. Values can also be set by editing the <b>odbc.ini</b> file in which their data sources are defined.

*Note:* The behavior of a DSN using a wire protocol driver with the catalog option selected, returns only the schemas that have associated tables or views. To list all existing schemas, create a DSN without the catalog option selected.

When configuring an ODBC DSN using the Greenplum Wire Protocol driver, select the following advanced options:

- **Application Using Threads**
- **Enable SQLDescribeParam**
- **Fetch TSFS as Time**
- **Fetch TSWTZ as Timestamp**

---

## MDS Driver Reference

### About the MDS Driver

The MDS driver provides access to an in-memory database for SAS Federation Server.

## Data Service Connection Options for MDS

MDS supports the following connection string options.

Option	Description
LOCALE	<b>LOCALE=<i>SAS locale identifier</i></b> Specifies the locale for message text and character conversion, both 'to' and 'from'.
DB   DATABASE	<b>DATABASE=<i>database name</i></b> Specifies the in-memory database instance. <b>DATABASE</b> must be specified if <b>CONOPTS=</b> is not specified. The database defaults to the catalog name if a database name is not specified.
CATALOG	<b>CATALOG=<i>catalog name</i>;</b> Specifies the catalog name. <b>CATALOG</b> must be specified if <b>CONOPTS=</b> is not specified.
CONOPTS	<b>CONOPTS=<i>connection string options</i></b> Specifies the connection string options for the driver to cache in memory. If a connection string is not specified, the default is in memory only.
COMMIT	<b>COMMIT=Y   N</b> Specifies if the in-memory changes are written to the <b>CONOPTS=</b> driver. <b>COMMIT</b> must be used with the <b>CONOPTS=</b> option.
BULKLOAD	<b>BULKLOAD=Y   N</b> Specifies if data is inserted immediately, which bypasses transactions. The <b>BULKLOAD</b> option is valid only when <b>CONOPTS=</b> is not specified.
NUMERICS	<b>NUMERICS=Y   N</b> Allows numeric data types or treats them as double precision. The default is Y (Yes).
RETAIN	<b>RETAIN=Y   N</b> Specifies if the in-memory database is dropped after the last client disconnects. The default is N (No).

Option	Description
CT_PRESERVE	<p data-bbox="432 243 1230 268"><b>CT_PRESERVE = STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p data-bbox="432 285 1326 369">Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul data-bbox="432 386 1374 772" style="list-style-type: none"> <li data-bbox="432 386 1374 443">• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li data-bbox="432 459 1374 543">• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="432 560 1374 674">• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="432 690 1374 772">• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>
IDCASE	<p data-bbox="432 816 895 842"><b>IDCASE=SENSITIVE   INSENSITIVE</b></p> <p data-bbox="432 858 1331 915">Specifies if schema, table, column, and alias identifiers are case-sensitive or insensitive. The default is case sensitive. <b>IDCASE</b> is valid only when <b>CONOPTS=</b> is not specified.</p>
DEFAULT_ATTR	<p data-bbox="432 959 874 984"><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p data-bbox="432 1001 1347 1058">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="432 1075 1374 1486" style="list-style-type: none"> <li data-bbox="432 1075 1374 1486">• <b>CURSORS=n</b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="722 1163 1374 1486" style="list-style-type: none"> <li data-bbox="722 1163 1374 1247">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="722 1264 1374 1348">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="722 1365 1374 1486">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="464 1503 927 1528">Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> <ul data-bbox="432 1545 1374 1766" style="list-style-type: none"> <li data-bbox="432 1545 1374 1629">• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li data-bbox="432 1646 1374 1766">• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>
DEFSHEMA	<p data-bbox="432 1810 756 1835"><b>DEFSHEMA= schema name</b></p> <p data-bbox="432 1852 1374 1908">Specifies the default schema for identifiers with no schema qualifier. The default is the first <b>SCHEMA=</b> in the connection string. This option is valid only when <b>CONOPTS=</b> is not specified.</p>

Option	Description
SCHEMAS	<p><b>SCHEMAS= ("schema1" ; "schema2" ; "schema3" )</b></p> <p>Specifies a list of schemas defined in the database. Identify schema names with double quotation marks and separate each name by a semicolon.</p>
SCHEMA	<p><b>SCHEMA= (NAME=schema-name1) ; SCHEMA= (NAME=schema-name2) ; . . .</b></p> <p>Defines one or more schemas in the database. The default is a single schema using the defined catalog name.</p>
REFTYPE	<p><b>REFTYPE=VARCHAR   NVARCHAR   VARBINARY</b></p> <p>Indicates that duplicate column data should be stored once and referenced by result sets rather than having separate instances in each row. This reduces memory usage with large numbers of duplicate data but might slow down performance.</p> <p><b>VARCHAR</b> Create a <b>REFCHAR</b> instead of a <b>VARCHAR</b> when specified. The default is create <b>VARCHAR</b>.</p> <p><b>NVARCHAR</b> Create an <b>NREFCHAR</b> instead of an <b>NVARCHAR</b> when specified. The default is create <b>NVARCHAR</b>.</p> <p><b>VARBINARY</b> Create a <b>REFBINARY</b> instead of a <b>VARBINARY</b> when specified. The default is create <b>VARBINARY</b>.</p> <p><i>Note:</i> A <b>REFCHAR (32)</b> uses less space than a <b>VARCHAR (32)</b> if there are many duplicate values in the table or if the data is less than 32 characters. However, a <b>REFCHAR (1)</b> generally uses more memory than a <b>VARCHAR (1)</b> because an extra pointer has to be stored instead of a single character.</p>
MAXDBMEM	<p><b>MAXDBMEM=number of bytes</b></p> <p>Specifies the maximum amount of memory the database can use to store all row data for all tables. The default is no limit.</p>

## Netezza Driver Reference

### Understanding the SAS Federation Server Driver for Netezza

The SAS Federation Server Driver for Netezza (Driver for Netezza) enables SAS Federation Server to read and update legacy Netezza tables. In addition, the driver creates Netezza tables that can be accessed by both SAS Federation Server and Netezza.

The Driver for Netezza supports most of the FedSQL functionality. The driver also supports the application's ability to submit native Netezza SQL statements.

The Driver for Netezza is a remote driver, which means that it connects to a server process in order to access data. The process might run on the same machine as SAS Federation Server, or it might run on another machine in the network.

## Data Service Connection Options for Netezza

### Overview

To access data that is hosted on SAS Federation Server, a client must submit a connection string, which defines how to connect to the data. The data service connection arguments for Netezza include connection options and advanced options.

### Connection Options

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

The Driver for Netezza supports the following connection options.

Option	Description
CATALOG	<p><b>CATALOG=</b><i>catalog-identifier</i>;</p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups logically related schemas. Any identifier is valid.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>
DATABASE	<p><b>DATABASE=</b><i>database-name</i>;</p> <p>Identifies the database to which you want to connect, which resides on the server previously specified through the SERVER option.</p>
DRIVER	<p><b>DRIVER=</b><i>NETEZZA</i>;</p> <p>Specifies the data service for the Netezza database to which you want to connect.</p> <p><i>Note:</i> You must specify the driver.</p>
CONNECTION OPTIONS	<p><b>CONOPTS=</b><i>(ODBC-compliant database connection string)</i>;</p> <p>Specifies an ODBC-compliant database connection string using ODBC-style syntax. These options, combined with the ODBC_DSN option, must specify a complete connection string to the data source. If you include a DSN= or FILEDSN= specification within the CONOPTS= option, do not use the ODBC_DSN= connection option. However, you can specify the ODBC database-specific connection options by using CONOPTS=. Then you can specify an ODBC DSN that contains other connection information by using the ODBC_DSN= connection option.</p>
DSN	<p><b>DSN=</b><i>data_source_identifier</i>;</p> <p>Identifies the data source name to which you want to connect.</p>
SERVER	<p><b>SERVER=</b><i>server_name</i>;</p> <p>Identifies the name of the server where the Netezza database resides.</p>
PORT	<p><b>PORT=</b><i>port_number</i></p> <p>Identifies the listen port of the server where the Netezza database resides.</p>

## Advanced Connection Options

The Driver for Netezza supports the following advanced connection options.

Option	Description
CLIENT ENCODING	<p><b>CLIENT_ENCODING=cei</b></p> <p>Used to specify encoding for the client.</p>
CT_PRESERVE	<p><b>CT_PRESERVE=STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p>Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li>• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>
DEFAULT_ATTR	<p><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p>Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul style="list-style-type: none"> <li>• <b>CURSORS=n</b>- Connection handle option. This option controls the driver’s use of client side result set cursors. The possible values are 0, 1 or 2. <ul style="list-style-type: none"> <li>0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li>1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server’s native cursor is not used.</li> <li>2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server’s native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> <p>Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> </li> <li>• <b>USE EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li>• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>

Option	Description
DRIVER TRACE	<p data-bbox="445 243 916 268"><b>DRIVER_TRACE='API   SQL   ALL';</b></p> <p data-bbox="445 289 1382 401">Requests tracing information, which logs transaction records to an external file that can be used for debugging purposes. The SAS Federation Server driver writes a record of each command that is sent to the database to the trace log based on the specified tracing level, which determines the type of tracing information. The tracing levels are:</p> <ul data-bbox="445 422 1342 684" style="list-style-type: none"> <li data-bbox="445 422 775 447">• <b>ALL</b> Activates all trace levels.</li> <li data-bbox="445 464 1342 541">• <b>API</b> Specifies that API method calls be sent to the trace log. This option is most useful if you are having a problem and need to send a trace log to Technical Support for troubleshooting.</li> <li data-bbox="445 558 1206 583">• <b>DRIVER</b> Specifies that driver-specific information be sent to the trace log.</li> <li data-bbox="445 600 1318 684">• <b>SQL</b> Specifies that SQL statements that are sent to the database management system (DBMS) be sent to the trace log. Tracing information is DBMS specific, but most SAS Federation Server drivers log SQL statements such as SELECT and COMMIT.</li> </ul> <p data-bbox="445 705 767 730"><b>Default:</b> Tracing is not activated.</p> <p data-bbox="445 747 1326 858"><i>Note:</i> If you activate tracing, you must also specify the location of the trace log with DRIVER_TRACEFILE=. Note that DRIVER_TRACEFILE= is resolved against the TRACEFILEPATH set in ALTER SERVER. TRACEFILEPATH is relative to the server's content root location.</p> <p data-bbox="445 875 1262 900">(Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p> <p data-bbox="445 917 1374 999"><b>Interaction:</b> You can specify one trace level, or you can concatenate more than one by including the   (OR) symbol. For example: <b>driver_trace='api   sql'</b> generates tracing information for API calls and SQL statements.</p>
DRIVER TRACE FILE	<p data-bbox="445 1045 863 1071"><b>DRIVER_TRACEFILE='filename';</b></p> <p data-bbox="445 1092 1382 1144">Used to specify the name of the text file for the trace log. Include the filename and extension in single or double quotation marks. For example:</p> <p data-bbox="445 1150 916 1176"><b>driver_tracefile='\\mytrace.log'</b></p> <p data-bbox="445 1192 1310 1245"><b>Default:</b> The default TRACEFILE location applies to a relative filename, and it is placed relative to TRACEFILEPATH.</p> <p data-bbox="445 1262 1230 1314"><b>Requirement:</b> DRIVER_TRACEFILE is required when activating tracing using DRIVER_TRACE.</p> <p data-bbox="445 1331 1086 1383"><b>Interaction:</b> (Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p>
DRIVER TRACE OPTIONS	<p data-bbox="445 1430 1254 1455"><b>DRIVER_TRACEOPTIONS=APPEND   THREADSTAMP   TIMESTAMP;</b></p> <p data-bbox="445 1476 1270 1501">Specifies options in order to control formatting and other properties for the trace log:</p> <ul data-bbox="445 1518 1374 1654" style="list-style-type: none"> <li data-bbox="445 1518 1374 1570">• <b>APPEND</b> Adds trace information to the end of an existing trace log. The contents of the file are not overwritten.</li> <li data-bbox="445 1587 1278 1612">• <b>THREADSTAMP</b> Prepends each line of the trace log with a thread identification.</li> <li data-bbox="445 1629 1150 1654">• <b>TIMESTAMP</b> Prepends each line of the trace log with a time stamp.</li> </ul> <p data-bbox="445 1671 1230 1696"><b>Default:</b> The trace log is overwritten with no thread identification or time stamp.</p>
USER ID	<p data-bbox="445 1738 671 1764"><b>USER="user-id";</b></p> <p data-bbox="445 1785 1310 1837">Specifies a Netezza user ID. If the ID contains blanks or national characters, enclose it in quotation marks. Alias: UID.</p> <p data-bbox="445 1854 855 1879"><i>Note:</i> You must specify the USER option.</p>

Option	Description
PASSWORD	<b>PASSWORD=<i>password</i>;</b> Specifies a password for the ID passed through the USER= option. Alias: PWD. <i>Note:</i> You must specify the PASSWORD option with USER.
STRIP_BLANKS	<b>STRIP_BLANKS=YES   NO;</b> Specifies whether to strip blanks from character fields.
READ ONLY	<b>READONLY=YES   NO;</b> Specifies whether to connect to the Netezza database in Read-Only mode. The default is NO. Alias: READ_ONLY
SHOW SYSTEM TABLES	<b>SHOWSYSTEMTABLES=YES   NO;</b> Specifies whether tables are included in the available table list. If set to YES or TRUE, system tables are included in the available table list. The default setting is NO. Alias: SST
NUMBER BYTES PER CHARACTER	<b>NUMBYTESPERCHAR=<i>value</i>;</b> Specifies the default number of bytes per character.

---

## ODBC Driver Reference

### About ODBC

This section provides functionality details and guidelines for the open database connectivity (ODBC) databases that are supported by the SAS Federation Server Driver for ODBC (Driver for ODBC).

ODBC standards provide a common interface to a variety of databases, including dBASE, Microsoft® Access®, Oracle®, Paradox, and Microsoft SQL Server databases. Specifically, ODBC standards define APIs that enable an application to access a database if both the application and the database conform to the specification. ODBC also provides a mechanism to enable dynamic selection of a database that an application is accessing, so that users have the flexibility of selecting databases other than those that are specified by the application developer.

### Understanding the SAS Federation Server Driver for ODBC

The Driver for ODBC enables SAS Federation Server to read and update legacy ODBC database tables. In addition, the driver creates tables that can be accessed by both SAS Federation Server and an ODBC database.

The Driver for ODBC supports most of the FedSQL functionality. The driver also supports an application's ability to submit native database-specific SQL statements.

The Driver for ODBC is a remote driver, which means that it connects to a server process in order to access data. The process might be running on the same machine as SAS Federation Server, or it might be running on another machine in the network.

## Data Service Connection Options for ODBC

### Overview

To access data that is hosted on SAS Federation Server, a client must submit a connection string, which defines how to connect to the data. The data service connection arguments for an ODBC-compliant database include connection options and advanced connection options.

### Connection Options

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

The Driver for ODBC supports the following connection options.

Option	Description
CATALOG	<p><b>CATALOG=<i>catalog-identifier</i>;</b></p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups logically related schemas. For databases that do not support native catalogs, any identifier is valid (for example, <code>catalog=myodbc</code>). For databases like SQL Server that do support native catalogs, CATALOG= is not required. The connection defaults to CATALOG=* unless you specify a logical name for the catalog and map it to the native catalog name in the database. For example, to map the logical catalog <code>mycat</code> to the native catalog named <code>newusers</code>, use the following command:  <b>catalog=(mycat=newusers) ;</b>. Catalog name maps can be used only with FedSQL. They are not valid with native SQL.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>
CONOPTS	<p><b>CONOPTS=(<i>ODBC-compliant database connection string</i>);</b></p> <p>Specifies an ODBC-compliant database connection string using ODBC-style syntax. These options, combined with the ODBC_DSN option, must specify a complete connection string to the data source. If you include a DSN= or FILEDSN= specification within the CONOPTS= option, do not use the ODBC_DSN= connection option. However, you can specify the ODBC database-specific connection options by using CONOPTS=. Then you can specify an ODBC DSN that contains other connection information by using the ODBC_DSN= connection option.</p> <p>Here is an example string using the CONOPTS option:  <b>DRIVER=ODBC;CONOPTS=(DRIVER={DataFlux 32-bit SQL Server Wire Protocol};SERVER=myserver;APP=Microsoft ODBC SDK;DATABASE=mydb)</b></p> <p>This example uses the ODBC_DSN option with the CONOPTS option:  <b>DRIVER=ODBC;ODBC_DSN=mydsn;CONOPTS=(DATABASE=mydb)</b></p>
DRIVER	<p><b>DRIVER=ODBC;</b></p> <p>Calls the SAS Federation Server Driver for ODBC. This specifies that the data service to which you want to connect must be an ODBC-compliant database.</p> <p><i>Note:</i> DRIVER is a required option. You must specify the driver.</p>

Option	Description
ODBC_DSN	<p><b>ODBC_DSN=odbc dsn name</b></p> <p>Specifies a valid ODBC-compliant database DSN that contains connection information for connecting to the ODBC-compliant database. You can use the CONOPTS= option in addition to ODBC_DSN= option to specify database-specific connection options not provided by SAS Federation Server. Do not specify the ODBC DSN in both CONOPTS= and ODBC_DSN=.</p>

### Advanced Connection Options

The Driver for ODBC supports the following advanced connection options for an ODBC-compliant database.

Option	Description
CT_PRESERVE	<p><b>CT_PRESERVE = STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p>Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li>• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>
ENABLE MULTIPLE ACTIVE RESULT SETS (MARS)	<p><b>ENABLE_MARS= NO   YES</b></p> <p>Enables or disables the use of multiple active result sets (MARS) on SQL Server. FedSQL cannot permit transactions on top of SQL Server because SQL Server only allows one cursor per transaction. Set this option to YES which gives FedSQL the ability to allow transactions under a given SQL Server connection.</p>

Option	Description
DEFAULT_ATTR	<p data-bbox="432 243 874 268"><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p data-bbox="432 285 1347 342">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="432 359 1362 768" style="list-style-type: none"> <li data-bbox="432 359 1362 415">• <b>CURSORS=n</b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="724 449 1342 768" style="list-style-type: none"> <li data-bbox="724 449 1342 533">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="724 554 1342 638">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="724 659 1342 768">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="464 789 927 814">Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> <ul data-bbox="432 835 1378 1050" style="list-style-type: none"> <li data-bbox="432 835 1378 919">• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li data-bbox="432 940 1378 1050">• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>

Option	Description
DEFAULT CURSOR TYPE	<p data-bbox="432 247 1350 296"><b>DEFAULT_CURSOR_TYPE=FORWARD_ONLY   KEYSSET_DRIVEN   DYNAMIC   STATIC;</b></p> <p data-bbox="432 317 1193 342">Specifies a valid default cursor type for new statements. The valid options are:</p> <p data-bbox="432 363 619 388"><b>FORWARD_ONLY</b></p> <p data-bbox="464 392 1382 506">Specifies a non-scrollable cursor that moves only forward through the result set. Forward-only cursors are dynamic in that all changes are detected as the current row is processed. If an application does not require scrolling, the forward-only cursor retrieves data quickly, with the least amount of overhead processing.</p> <p data-bbox="432 527 635 552"><b>KEYSET_DRIVEN</b></p> <p data-bbox="464 556 1382 695">Specifies a scrollable cursor that detects changes that are made to the values of rows in the result set but that does not always detect changes to deletion of rows and changes to the order of rows in the result set. A keyset-driven cursor is based on row keys, which are used to determine the order and set of rows that are included in the result set. As the cursor scrolls the result set, it uses the keys to retrieve the most recent values in the table.</p> <p data-bbox="464 716 1382 854">It is sometimes helpful to have a cursor that can detect changes in the rows of a result set. A keyset-driven cursor uses a row identifier rather than caching the entire row into memory. It therefore uses much less disk space than other row caching mechanisms. Deleted rows can be detected when a <b>SELECT</b> statement that references the bookmark, row ID, or key column values fails to return a row.</p> <p data-bbox="432 875 544 900"><b>DYNAMIC</b></p> <p data-bbox="464 905 1350 1018">Specifies a scrollable cursor that detects changes that are made to the rows in the result set. All <b>INSERT</b>, <b>UPDATE</b>, and <b>DELETE</b> statements that are made by all users are visible through the cursor. The dynamic cursor is good for an application that must detect all concurrent updates that are made by other users.</p> <p data-bbox="432 1039 528 1064"><b>STATIC</b></p> <p data-bbox="464 1068 1382 1148">Specifies a scrollable cursor that displays the result set as it existed when the cursor was first opened. The static cursor provides forward and backward scrolling. If the application does not need to detect changes but requires scrolling, the static cursor is a good choice.</p> <p data-bbox="432 1169 1382 1218"><i>Note:</i> The application can still override this value, but if the application does not explicitly set a cursor type, this value will be in effect</p>

Option	Description
DRIVER TRACE	<p><b>DRIVER_TRACE= 'API   SQL   ALL' ;</b></p> <p>Requests tracing information, which logs transaction records to an external file that can be used for debugging purposes. The SAS Federation Server driver writes a record of each command that is sent to the database to the trace log based on the specified tracing level, which determines the type of tracing information. The tracing levels are:</p> <ul style="list-style-type: none"> <li>• <b>ALL</b> Activates all trace levels.</li> <li>• <b>API</b> Specifies that API method calls be sent to the trace log. This option is most useful if you are having a problem and need to send a trace log to Technical Support for troubleshooting.</li> <li>• <b>DRIVER</b> Specifies that driver-specific information be sent to the trace log.</li> <li>• <b>SQL</b> Specifies that SQL statements that are sent to the database management system (DBMS) be sent to the trace log. Tracing information is DBMS specific, but most SAS Federation Server drivers log SQL statements such as SELECT and COMMIT.</li> </ul> <p><b>Default:</b> Tracing is not activated.</p> <p><i>Note:</i> If you activate tracing, you must also specify the location of the trace log with DRIVER_TRACEFILE=. Note that DRIVER_TRACEFILE= is resolved against the TRACEFILEPATH set in ALTER SERVER. TRACEFILEPATH is relative to the server's content root location.</p> <p>(Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p> <p><b>Interaction:</b> You can specify one trace level, or you can concatenate more than one by including the   (OR) symbol. For example: <b>driver_trace='api   sql'</b> generates tracing information for API calls and SQL statements.</p>
DRIVER TRACE FILE	<p><b>DRIVER_TRACEFILE= 'filename' ;</b></p> <p>Used to specify the name of the text file for the trace log. Include the filename and extension in single or double quotation marks. For example: <b>driver_tracefile='mytrace.log'</b></p> <p><b>Default:</b> The default TRACEFILE location applies to a relative filename, and it is placed relative to TRACEFILEPATH.</p> <p><b>Requirement:</b> DRIVER_TRACEFILE is required when activating tracing using DRIVER_TRACE.</p> <p><b>Interaction:</b> (Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p>
DRIVER TRACE OPTIONS	<p><b>DRIVER_TRACEOPTIONS=APPEND   THREADSTAMP   TIMESTAMP ;</b></p> <p>Specifies options in order to control formatting and other properties for the trace log:</p> <ul style="list-style-type: none"> <li>• <b>APPEND</b> Adds trace information to the end of an existing trace log. The contents of the file are not overwritten.</li> <li>• <b>THREADSTAMP</b> Prepends each line of the trace log with a thread identification.</li> <li>• <b>TIMESTAMP</b> Prepends each line of the trace log with a time stamp.</li> </ul> <p><b>Default:</b> The trace log is overwritten with no thread identification or time stamp.</p>
USER	<p><b>USER=user-ID ;</b></p> <p>Specifies the user ID for logging on to the ODBC-compliant database, such as Microsoft SQL Server, with a user ID that differs from the default ID.</p> <p><i>Note:</i> The alias is <b>UID=</b>.</p>

Option	Description
PASSWORD	<p><b>PASSWORD=</b><i>password</i>;</p> <p>Specifies the password that corresponds to the user ID in the database.</p> <p><i>Note:</i> The alias is <b>PWD=</b>.</p>

Here are example connection strings using the SAS Federation Server Driver for ODBC:

This connection string specifies an ODBC DSN:

```
driver=odbc; uid=scott; pw=myspw; odbc_dsn=myOracleDSN;
catalog=odbc_oracle;
```

This connection string specifies catalog name maps to access multiple catalogs on Microsoft SQL Server:

```
driver=odbc; uid=jfox; pw=myspw; odbc_dsn=mySQLdsn;
catalog=(cat1=mycat; cat2=testcat; cat3=users;
```

## Configuring ODBC for Hadoop

### Connection Options

In addition to the ODBC data service connection options above, the following option is available for Hive and Impala using ODBC:

Option	Description
SCHEMA	<p><b>SCHEMA=</b><i>schema-name</i>;</p> <p>This option is for use with Apache Hive or Cloudera Impala only. Specifies the name of the schema that is passed to FedSQL.</p> <p>This example connection string specifies a schema name that is passed to FedSQL for Hive or Impala connections:</p> <pre>DRIVER=ODBC;UID=hadoop;PWD=myspw;CONOPTS= (DSN=tktshive);SCHEMA=schema1;CATALOG=CATALOG_HIVE;</pre>

### Setting the Maximum Character String Size for Hive

Hive has a single data type for storing text, **STRING**, which is a variable-length character string with a maximum size of 2G. As a result, this can create very large character fields when processing data. Since Hive's string type is comparable to **VARCHAR** in other data sources, you can set the ODBC attribute, **Max Varchar Size** to specify the maximum character string size. Set the **Max Varchar Size** value using Advanced Options in Windows ODBC Administrator, or in Unix by editing `odbc.ini` in the specified path or \$HOME directory.

You can also specify this option in a connection string using the **CONOPTS** container. Here is an example:

```
DRIVER=ODBC;DB=hive;UID=dbitest;PWD=dbigrp1;SCHEMA=default;CON
OPTS=(MaxVarcharSize=300);CATALOG=FOO;
```

## Wire Protocol Driver Usage Notes

### Overview

SAS Federation Server provides a number of wire protocol ODBC drivers that communicate directly with a database server, without having to communicate through a client library. When you configure the ODBC drivers on Windows or UNIX, you have the opportunity to set certain options. SAS products run best when these options are selected. Some, but not all, are selected by default.

Windows	The options are located on the <b>Advanced</b> or <b>Performance</b> tabs in the ODBC Administrator.
UNIX	The options are available when configuring data sources using the <b>dfdbconf</b> tool. Values can also be set by editing the <b>odbc.ini</b> file in which their data sources are defined.

*Note:* The behavior of a DSN using a wire protocol driver with the catalog option selected, returns only the schemas that have associated tables or views. To list all existing schemas, create a DSN without the catalog option selected.

### MySQL

When configuring an ODBC DSN using the MySQL Wire Protocol driver, select the following advanced options:

- **Application Using Threads**
- **Enable SQLDescribeParam**

### SQL Server and SQL Server Legacy

Configure the following **Advanced** options for the SQL Server Wire Protocol driver and the SQL Server Legacy Wire Protocol driver:

- **Application Using Threads**
- **Enable Quoted Identifiers**
- **Fetch TWFS as Time**
- **Fetch TSWTZ as Timestamp**

*Note:*

1. Significant performance improvements have been realized when using the SQL Server Legacy Wire Protocol Driver, as compared to the SQL Server Wire Protocol Driver.
2. The SQL Server Legacy Wire Protocol Driver does not support transactions when used with FedSQL enabled because the driver only allows a single statement per connection while FedSQL requires multiple statements per connection when using transactions.

---

## Oracle Reference

### ***Understanding the SAS Federation Server Driver for Oracle***

The SAS Federation Server Driver for Oracle (Driver for Oracle) enables SAS Federation Server to read and update legacy Oracle tables. In addition, the driver creates Oracle tables that can be accessed by both SAS Federation Server and Oracle.

The Driver for Oracle supports most of the FedSQL functionality. The driver also supports the application's ability to submit native Oracle SQL statements.

The Driver for Oracle is a remote driver, which means that it connects to a server process in order to access data. The process might be running on the same machine as SAS Federation Server, or it might be running on another machine in the network.

### ***Data Service Connection Options for Oracle***

#### ***Overview***

To access data that is hosted on SAS Federation Server, a client must submit a connection string, which defines how to connect to the data. The data service connection arguments for Oracle Server include connection options and advanced options.

#### ***Connection Options***

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

The Driver for Oracle supports the following connection options.

Option	Description
CATALOG	<p><b>CATALOG=</b><i>catalog-identifier</i>;</p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups logically related schemas. Any identifier is valid such as <b>catalog=oracle_test</b>. You must specify a catalog. For the Oracle database, this is a logical catalog name to use as an SQL catalog identifier.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>
DRIVER	<p><b>DRIVER=ORACLE</b>;</p> <p>Identifies the data service to which you want to connect, which is an Oracle database.</p> <p><i>Note:</i> You must specify the driver.</p>
PATH	<p><b>PATH=</b><i>database-specification</i>;</p> <p>Specifies the Oracle connect identifier. A connect identifier can be a net service name, a database service name, or a net service alias.</p>

---

Option	Description
USERID (UID)	<p><b>UID=user-id;</b></p> <p>Specifies an optional Oracle user ID. If the user ID contains blanks or national characters, enclose it in quotation marks. If you omit an Oracle user ID and password, the default Oracle user ID <b>OPSS\$sysid</b> is used, if it is enabled.</p>
PASSWORD (PWD)	<p><b>PWD=password;</b></p> <p>Specifies an optional Oracle database password that is associated with the Oracle user ID. <b>PWD=</b> is always used with <b>UID=</b> and the associated password is case-sensitive. If you omit <b>PWD=</b>, the password for the default Oracle user ID <b>OPSS\$sysid</b> is used, if it is active.</p>

### Advanced Connection Options

The Driver for Oracle supports the following advanced connection options.

Option	Description
CT_PRESERVE	<p><b>CT_PRESERVE = STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p>Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li>• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>

Option	Description
DEFAULT_ATTR	<p data-bbox="432 243 874 268"><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p data-bbox="432 289 1347 344">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="432 365 1362 772" style="list-style-type: none"> <li data-bbox="432 365 1362 415">• <b>CURSORS=n</b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="724 449 1347 772" style="list-style-type: none"> <li data-bbox="724 449 1347 533">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="724 554 1347 638">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="724 659 1347 772">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="464 793 932 819">Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> <ul data-bbox="432 835 1378 1052" style="list-style-type: none"> <li data-bbox="432 835 1378 919">• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li data-bbox="432 940 1378 1052">• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>
DRIVER TRACE;	<p data-bbox="432 1094 890 1119"><b>DRIVER_TRACE='API   SQL   ALL'</b></p> <p data-bbox="432 1140 1378 1255">Requests tracing information, which logs transaction records to an external file that can be used for debugging purposes. The SAS Federation Server driver writes a record of each command that is sent to the database to the trace log based on the specified tracing level, which determines the type of tracing information. The tracing levels are:</p> <ul data-bbox="432 1276 1378 1507" style="list-style-type: none"> <li data-bbox="432 1276 772 1302">• <b>ALL</b> Activates all trace levels.</li> <li data-bbox="432 1318 1378 1373">• <b>API</b> Specifies that API method calls be sent to the trace log. This option is most useful if you are having a problem and need to send a trace log to Technical Support for troubleshooting.</li> <li data-bbox="432 1390 1187 1415">• <b>DRIVER</b> Specifies that driver-specific information be sent to the trace log.</li> <li data-bbox="432 1432 1378 1507">• <b>SQL</b> Specifies that SQL statements that are sent to the database management system (DBMS) be sent to the trace log. Tracing information is DBMS specific, but most SAS Federation Server drivers log SQL statements such as SELECT and COMMIT.</li> </ul> <p data-bbox="432 1528 756 1554"><b>Default:</b> Tracing is not activated.</p> <p data-bbox="432 1575 1315 1690"><i>Note:</i> If you activate tracing, you must also specify the location of the trace log with <b>DRIVER_TRACEFILE=</b>. Note that <b>DRIVER_TRACEFILE=</b> is resolved against the <b>TRACEFILEPATH</b> set in <b>ALTER SERVER</b>. <b>TRACEFILEPATH</b> is relative to the server's content root location.</p> <p data-bbox="432 1711 1251 1736">(Optional) You can control trace log formatting with <b>DRIVER_TRACEOPTIONS=</b>.</p> <p data-bbox="432 1757 1378 1843"><b>Interaction:</b> You can specify one trace level, or you can concatenate more than one by including the   (OR) symbol. For example: <b>driver_trace='api   sql'</b> generates tracing information for API calls and SQL statements.</p>

Option	Description
DRIVER TRACEFILE	<p><b>DRIVER_TRACEFILE= 'filename' ;</b></p> <p>Used to specify the name of the text file for the trace log. Include the filename and extension in single or double quotation marks. For example: <b>driver_tracefile= 'mytrace.log'</b></p> <p><b>Default:</b> The default TRACEFILE location applies to a relative filename, and it is placed relative to TRACEFILEPATH.</p> <p><b>Requirement:</b> DRIVER_TRACEFILE is required when activating tracing using DRIVER_TRACE.</p> <p><b>Interaction:</b> (Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p>
DRIVER TRACE OPTIONS	<p><b>DRIVER_TRACEOPTIONS=APPEND   THREADSTAMP   TIMESTAMP ;</b></p> <p>Specifies options in order to control formatting and other properties for the trace log:</p> <ul style="list-style-type: none"> <li>• <b>APPEND</b> Adds trace information to the end of an existing trace log. The contents of the file are not overwritten.</li> <li>• <b>THREADSTAMP</b> Prepends each line of the trace log with a thread identification.</li> <li>• <b>TIMESTAMP</b> Prepends each line of the trace log with a time stamp.</li> </ul> <p><b>Default:</b> The trace log is overwritten with no thread identification or time stamp.</p>
ORA_ENCODING	<p><b>ORA_ENCODING=UNICODE ;</b></p> <p>Specifies that the Oracle data be returned in Unicode to SAS Federation Server. <b>UNICODE</b> is the default setting and is independent of the <b>NLS_LANG</b> environment variable setting.</p>
ORNUMERIC	<p><b>ORANUMERIC=NO   YES</b></p> <p>Specifies how numbers read from or inserted into the Oracle NUMBER column will be treated. This option defaults to YES so that a NUMBER column with precision or scale is described as <b>TKTS_NUMERIC</b>. This option can be specified as both a connection option and a table option. When specified as both connection and table option, the table option value overrides the connection option.</p> <ul style="list-style-type: none"> <li>• <b>NO</b> Indicates that the numbers will be treated as <b>TKTS_DOUBLE</b> values. They might not have precision beyond 14 digits.</li> <li>• <b>YES</b> Indicates that non-integer values with explicit precision will be treated as <b>TKTS_NUMERIC</b> values. This is the default setting.</li> </ul>
USE_CACHED_CATALOG	<p><b>USE_CACHED_CATALOG=YES   NO ;</b></p> <p>Specifies whether to use the cached catalog rather than compiling a new catalog on every run. Setting this option to YES can improve the performance of the TKTSForeignKeys API. The default setting is YES.</p> <p><i>Note:</i> Before you can use this option, you must complete the following steps:</p> <ol style="list-style-type: none"> <li>1. Create a materialized view. See the example code in “<a href="#">Creating a Materialized View (USE_CACHED_CATALOG)</a>” on page 155.</li> <li>2. Use the ALTER DSN statement to add the USE_CACHED_CATALOG connection option. For more information about the ALTER DSN statement, see ALTER DSN Statement.</li> </ol>

## Creating a Materialized View (USE\_CACHED\_CATALOG)

The following example shows you how to create a materialized view. Use this script if `USE_CACHED_CATALOG` is set to YES above.

---

```

/*-----SAS_CACHED_CATALOG.SQL-----*/
/* This script is used to create the materialized and the synonym needed to
   get the ForeignKey metadata. Work with your DBA to set this up.
   Materialized views can be complex and so thorough understanding will help us
   use them effectively. Especially deciding how to do the refreshes.
   Here we provide the simplest possible steps to create the required materialized
   view and the command to refresh it manually. The materialized view below can
   be created in any schema with any name. Feel free to add whatever REFRESH
   options suits your purpose. Note that you might need additional steps based
   on the REFRESH option setting. Here we provide the simplest possible way to
   do this. The PUBLIC synonym pointing to this Materialized view must be
   named "SAS_CACHED_FK_CATALOG_PSYN". This synonym must be visible to
   PUBLIC (or the set of users who will be needing ForeignKey metadata) so that
   it is accessible from any schema.
*/

Create materialized view SAS_CACHED_FK_CATALOG_MATVIEW REFRESH ON DEMAND as SELECT
PKAC.OWNER as PKTABLE_SCHEM,
PKAC.TABLE_NAME as PKTABLE_NAME,
PKACC.COLUMN_NAME as PKCOLUMN_NAME,
FKAC.OWNER as FKTABLE_SCHEM,
FKAC.TABLE_NAME as FKTABLE_NAME,
FKACC.COLUMN_NAME as FKCOLUMN_NAME,
FKACC.POSITION as KEY_SEQ,
FKAC.CONSTRAINT_NAME as FK_NAME,
PKAC.CONSTRAINT_NAME as PK_NAME
from
sys.all_constraints PKAC, sys.all_constraints FKAC,
sys.all_cons_columns PKACC, sys.all_cons_columns FKACC

where

FKAC.r_constraint_name=PKAC.constraint_name and
FKAC.constraint_name=FKACC.constraint_name and
PKAC.constraint_name=PKACC.constraint_name and PKAC.constraint_type='P' and
FKAC.constraint_type='R' and FKAC.owner=FKACC.owner and PKAC.owner=PKACC.owner
and PKAC.table_name=PKACC.table_name and FKAC.table_name=FKACC.table_name and
FKACC.position = PKACC.position ;

/* The synonym name *must* be SAS_CACHED_FK_CATALOG_PUBLIC_SYNONYM */
create public synonym SAS_CACHED_FK_CATALOG_PSYN for SAS_CACHED_FK_CATALOG_MATVIEW;
grant all on SAS_CACHED_FK_CATALOG_PSYN to PUBLIC;

/*-----Manual REFRESH of the Materialized View-----*/
/* Note there are several ways to do this, consult with your DBA.
   Here are a couple of ways:
*/
execute DBMS_MVIEW.REFRESH('SAS_CACHED_FK_CATALOG_MATVIEW');
execute DBMS_SNAPSHOT.REFRESH('SAS_CACHED_FK_CATALOG_MATVIEW', '?');

```

---

## Oracle Wire Protocol Driver Usage Notes

SAS Federation Server provides a number of wire protocol ODBC drivers that communicate directly with a database server, without having to communicate through a client library. When you configure the ODBC drivers on Windows or UNIX, you have the opportunity to set certain options. SAS products run best when these options are selected. Some, but not all, are selected by default.

---

Windows	The options are located on the <b>Advanced</b> or <b>Performance</b> tabs in the ODBC Administrator.
UNIX	The options are available when configuring data sources using the <b>dfdbconf</b> tool. Values can also be set by editing the <b>odbc.ini</b> file in which their data sources are defined.

---

*Note:* When you use a wire protocol driver to create an ODBC connection, the following special considerations apply:

1. The behavior of a DSN using a wire protocol driver with the catalog option selected, returns only the schemas that have associated tables or views. To list all existing schemas, create a DSN without the catalog option selected.
2. Verify that the Enable Bulk Load option is turned on in the ODBC DSN for databases that support this option. The Bulk Load option is not enabled by default in the newer wire protocol drivers. As a result, insert performance suffers.

When configuring an ODBC DSN using the Oracle Wire Protocol driver, set the following advanced options:

- **Application Using Threads**
- **Enable SQLDescribeParam**
- **Describe at Prepare**
- **Enable N-CHAR Support**
- **Enable Scrollable Cursors**

---

## PostgreSQL Driver Reference

### *Understanding the SAS Federation Server Driver for PostgreSQL*

The SAS Federation Server Driver for PostgreSQL(Driver for PostgreSQL) enables SAS Federation Server to read and update legacy PostgreSQL tables. In addition, the driver creates PostgreSQL tables that can be accessed by both SAS Federation Server and the PostgreSQL data management system.

The Driver for PostgreSQL supports most of the FedSQL functionality. The driver also supports an application's ability to submit native SQL statements.

The Driver for PostgreSQL is a remote driver, which means that it connects to a server process in order to access data. The process might be running on the same machine as SAS Federation Server, or it might be running on another machine in the network.

## Connection Options for PostgreSQL

### Overview

To access data that is hosted on SAS Federation Server, a client must submit a connection string, which defines how to connect to the data. The data service connection arguments for PostgreSQL include connection options and advanced options.

### Connection Options

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

The following connection options are supported for Postgres data sources.

Option	Description
CATALOG	<p><b>CATALOG=</b><i>catalog-identifier</i>;</p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups schemas that are logically related. For example, <b>catalog=ptgtest</b>.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>
CONOPTS	<p><b>CONOPTS=(</b><i>ODBC-compliant database connection string</i><b>);</b></p> <p>Specifies an ODBC-compliant database connection string using ODBC-style syntax. These options, combined with the ODBC_DSN option, must specify a complete connection string to the data source. If you include a DSN= or FILEDSN= specification within the CONOPTS= option, do not use the ODBC_DSN= connection option. However, you can specify the ODBC database-specific connection options by using CONOPTS=. Then you can specify an ODBC DSN that contains other connection information by using the ODBC_DSN= connection option.</p> <p>Here is an example string using the CONOPTS option:  <b>DRIVER=ODBC;CONOPTS=(DRIVER={DataFlux 32-bit SQL Server Wire Protocol};SERVER=myserver;APP=Microsoft ODBC SDK;DATABASE=mydb)</b></p> <p>This example uses the ODBC_DSN option with the CONOPTS option:  <b>DRIVER=ODBC;ODBC_DSN=mydsn;CONOPTS=(DATABASE=mydb)</b></p>
DRIVER	<p><b>DRIVER=</b><i>postgres</i>;</p> <p>Specifies the data service for the PostgreSQL database to which you want to connect.</p> <p><i>Note:</i> DRIVER is a required option. You must specify a driver.</p>
DATABASE	<p><b>DATABASE=</b><i>database-name</i>;</p> <p>Specifies the name of the PostgreSQL database. Enclose the database name in single quotation marks if it contains spaces or non-alphanumeric characters. You can also specify DATABASE= with the DB= alias. <b>database=sample</b>, <b>DB=sample</b>.</p>
DSN	<p><b>DSN=</b><i>data-source-identifier</i>;</p> <p>Specifies the data source name to which you want to connect.</p>

Option	Description
PASSWORD	<b>PWD=<i>password</i>;</b> Specifies the password associated with the user ID. Enclose password in single quotation marks if it contains spaces or nonalphanumeric characters. You can also specify PASSWORD= with the PWD=, PASS=, and PW= aliases.
PORT	<b>PORT=<i>port_number</i></b> Specifies the port number that is used to connect to the specified PostgreSQL Server. If you do not specify a port, the default is 5432.
SERVER	<b>SERVER= '<i>server-name</i>'</b> Specifies the server name or IP address of the PostgreSQL server to which you want to connect. Enclose the server name in single quotation marks if the name contains spaces or non-alphanumeric characters: <b>SERVER=' <i>server name</i>'</b> .
USER	<b>USER=<i>user-name</i></b> Specifies the PostgreSQL user name (also called the user ID) that you use to connect to your database. If the user name contains spaces or non-alphanumeric characters, you must enclose it in quotation marks.

### **Advanced Options**

The following advanced options are supported for PostgreSQL data sources.

Option	Description
ALLOW UNQUOTED NAMES	<b>ALLOW_UNQUOTED_NAMES=NO   YES</b> Specifies whether to enclose table and column names in quotation marks. Tables and columns are quoted when this option is set at NO. If set to YES, the driver does not automatically add quotation marks to table and column names if they are not specified. This allows PostgreSQL tables and columns to be created in the default lower case. The default option is NO.
CLIENT ENCODING	<b>CLIENT_ENCODING=<i>cei</i></b> Used to specify encoding for the client.

Option	Description
CT_PRESERVE	<p data-bbox="445 243 1214 268"><b>CT_PRESERVE=STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p data-bbox="445 285 1337 369">Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul data-bbox="445 386 1382 772" style="list-style-type: none"> <li data-bbox="445 386 1382 445">• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li data-bbox="445 457 1382 541">• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="445 554 1382 676">• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="445 688 1382 772">• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>
DEFAULT_ATTR	<p data-bbox="445 814 884 840"><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p data-bbox="445 856 1358 915">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="445 932 1382 1344" style="list-style-type: none"> <li data-bbox="445 932 1382 1344">• <b>CURSORS=<i>n</i></b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="732 1024 1382 1344" style="list-style-type: none"> <li data-bbox="732 1024 1382 1108">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="732 1121 1382 1205">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="732 1218 1382 1344">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="475 1360 940 1386">Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> <ul data-bbox="445 1402 1382 1625" style="list-style-type: none"> <li data-bbox="445 1402 1382 1495">• <b>USE_EVP=<i>n</i></b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li data-bbox="445 1507 1382 1625">• <b>XCODE_WARN=<i>n</i></b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>

Option	Description
DRIVER TRACE	<p><b>DRIVER_TRACE='API   SQL   ALL';</b></p> <p>Requests tracing information, which logs transaction records to an external file that can be used for debugging purposes. The SAS Federation Server driver writes a record of each command that is sent to the database to the trace log based on the specified tracing level, which determines the type of tracing information. The tracing levels are:</p> <ul style="list-style-type: none"> <li>• <b>ALL</b> Activates all trace levels.</li> <li>• <b>API</b> Specifies that API method calls be sent to the trace log. This option is most useful if you are having a problem and need to send a trace log to Technical Support for troubleshooting.</li> <li>• <b>DRIVER</b> Specifies that driver-specific information be sent to the trace log.</li> <li>• <b>SQL</b> Specifies that SQL statements that are sent to the database management system (DBMS) be sent to the trace log. Tracing information is DBMS specific, but most SAS Federation Server drivers log SQL statements such as SELECT and COMMIT.</li> </ul> <p><b>Default:</b> Tracing is not activated.</p> <p><i>Note:</i> If you activate tracing, you must also specify the location of the trace log with DRIVER_TRACEFILE=. Note that DRIVER_TRACEFILE= is resolved against the TRACEFILEPATH set in ALTER SERVER. TRACEFILEPATH is relative to the server's content root location.</p> <p>(Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p> <p><b>Interaction:</b> You can specify one trace level, or you can concatenate more than one by including the   (OR) symbol. For example: <b>driver_trace='api   sql'</b> generates tracing information for API calls and SQL statements.</p>
DRIVER TRACE FILE	<p><b>DRIVER_TRACEFILE='filename';</b></p> <p>Used to specify the name of the text file for the trace log. Include the filename and extension in single or double quotation marks. For example: <b>driver_tracefile='\\mytrace.log'</b></p> <p><b>Default:</b> The default TRACEFILE location applies to a relative filename, and it is placed relative to TRACEFILEPATH.</p> <p><b>Requirement:</b> DRIVER_TRACEFILE is required when activating tracing using DRIVER_TRACE.</p> <p><b>Interaction:</b> (Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p>
DRIVER TRACE OPTIONS	<p><b>DRIVER_TRACEOPTIONS=APPEND   THREADSTAMP   TIMESTAMP;</b></p> <p>Specifies options in order to control formatting and other properties for the trace log:</p> <ul style="list-style-type: none"> <li>• <b>APPEND</b> Adds trace information to the end of an existing trace log. The contents of the file are not overwritten.</li> <li>• <b>THREADSTAMP</b> Prepends each line of the trace log with a thread identification.</li> <li>• <b>TIMESTAMP</b> Prepends each line of the trace log with a time stamp.</li> </ul> <p><b>Default:</b> The trace log is overwritten with no thread identification or time stamp.</p>
MAX_BINARY_LEN	<p><b>MAX_BINARY_LEN=value;</b></p> <p>Specifies a value, in bytes, that limits the length of long binary fields (LONG VARBINARY). Unlike other databases, PostgreSQL does not have a size limit for long binary fields. The default is 1048576.</p>

Option	Description
MAX_CHAR_LEN	<b>MAX_CHAR_LEN=value;</b> Specifies a value that limits the length of character fields (CHAR and VARCHAR). The default is 2000.
MAX_TEXT_LEN	<b>MAX_TEXT_LEN=value;</b> Specifies a value that limits the length of long character fields (LONG VARCHAR). The default is 409500.
SCHEMA	<b>SCHEMA=value;</b> Specifies the default schema for the connection. If not specified, the schema, or list of schemas, is determined based on the value of the schema search path defined on the database server.
STRIP_BLANKS	<b>STRIP_BLANKS=YES   NO;</b> Specifies whether to strip blanks from character fields.

## SAP Reference

### *Understanding the SAS Federation Server Driver for SAP*

The SAS Federation Server Driver for SAP enables SAS Federation Server to read tables from SAP systems. The SAS Federation Server Driver for SAP has read-only capabilities.

The SAS Federation Server Driver for SAP supports most of the FedSQL functionality. The driver does not support the application's ability to submit native SQL statements.

The SAS Federation Server Driver for SAP is a remote driver, which means that it connects to a server process in order to access data. The process might be running on the same machine as SAS Federation Server, or it might be running on another machine in the network.

### *Data Service Connection Options for SAP*

The table below describes the data service connection options for the SAS Federation Server Driver for SAP.

Option	Description
ABAPFM	<b>ABAPFM = abap_function_name</b> Specifies the name of the Advanced Business Application Programming (ABAP) function module that the driver uses internally. The default is <b>/SAS/Z_SAS_DIALOG</b> . Alias: ABAPFUNCTION, ABAPFUNC

Option	Description
ABAP_NAMESPACE	<p><b>ABAP_NAMESPACE =namespace</b></p> <p>Specifies the namespace for ABAP functions and programs that are used by the driver. If the ABAP programs are installed in the customer namespace rather than in the default namespace, this parameter identifies where the ABAP programs are installed. The default is <b>/SAS/</b></p> <p>Alias: ABAPNAMESPACE, ABAP_NAME_SPACE, ABAPNS, ABAP_NS</p>
ABAPPROG	<p><b>ABAPPROG = abap_program</b></p> <p>Specifies the name of the ABAP language that the driver uses internally. This value is set by the ABAP function module. The default is <b>/SAS/Z_SAS_READ</b>.</p> <p>Alias: ABAPREPORT, ABAPPROGRAM</p>
ASHOST	<p><b>ASHOST=application_server_host</b></p> <p>Specifies the host name of the server or IP address of a specific application server. There is no default.</p> <p>Alias: HST, RFCHOST, R3HOST</p>
BATCH	<p><b>BATCH = 0   1   Y   N</b></p> <p>Specifies whether the SAS Federation Server Driver for SAP should use SAP batch jobs for the data extracts. If set at <b>Y</b> (Yes), the SAS Federation Server Driver for SAP uses batch jobs to extract R/3 data. When set at <b>N</b> (No), the SAS Federation Server Driver for SAP uses dialog processes to extract R/3 data. The default is <b>N</b> (No).</p> <p>Alias: BATCH_MODE, BATCHMODE</p>
BUFFER_SIZE	<p><b>BUFFER_SIZE = buffersize</b></p> <p>Sets the minimum buffer size for data transfers in batch and dialog modes. The number of bytes should be greater than 10,000 and no more than eight digits. The default is 100,000 bytes.</p> <p>Alias: BUFFERSIZE, BUFFSIZE, BLOCK_SIZE, BLOCKSIZE</p>
CLIENT	<p><b>CLIENT=client_number</b></p> <p>Specifies the SAP logon parameter client. Examples for a client are 000 or 800. The default is the SAP system default.</p> <p><i>Note:</i> When you access the SAP system using the driver, specify valid logon information including <b>CLIENT</b>, <b>USER NAME</b>, <b>PASSWORD</b>, and <b>LANGUAGE</b>. The user ID and password can also be passed through single sign-on (SSO). The driver performs a logon check at OPEN time.</p> <p>Alias: CLI, RFCCLIENT, RFCCLI</p>

Option	Description
DEFAULT_ATTR	<p data-bbox="461 243 839 268"><b>DEFAULT_ATTR= (attr=value)</b></p> <p data-bbox="461 285 1374 342">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="461 359 1347 800" style="list-style-type: none"> <li data-bbox="461 359 1347 415">• <b>CURSORS=n</b> - Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="751 449 1326 800" style="list-style-type: none"> <li data-bbox="751 449 1302 537">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="751 554 1302 642">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="751 659 1326 800">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="493 821 956 846">Example: <b>DEFAULT_ATTR= (CURSORS=2)</b></p> <ul data-bbox="461 863 1362 1079" style="list-style-type: none"> <li data-bbox="461 863 1362 951">• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR= (USE_EVP=0)</b></li> <li data-bbox="461 968 1347 1079">• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR= (XCODE_WARN=1)</b></li> </ul>
DESTGROUP	<p data-bbox="461 1125 876 1150"><b>DESTGROUP=destination_group</b></p> <p data-bbox="461 1167 1302 1224">Specifies the name of the destination group for batch access to the SAP system. The destination groups are defined in the /SAS/DESTS table in SAP. The default is <b>SAS1</b>.</p>
DESTINATION	<p data-bbox="461 1266 815 1291"><b>DESTINATION=destination</b></p> <p data-bbox="461 1308 1366 1396">Specifies the destination in the <b>saprfc.ini</b> file, if working with a <b>saprfc.ini</b> file. If the RFC server is an R/3 system, this destination must also be defined in the <b>SIDEINFO</b> file for the SAP gateway.</p> <p data-bbox="461 1413 719 1438">Alias: DEST, DST, DSTN</p>
GROUP	<p data-bbox="461 1482 922 1507"><b>GROUP=application_server_group</b></p> <p data-bbox="461 1524 1171 1549">Specifies the name of the group of application servers for load balancing.</p>
GWHOST	<p data-bbox="461 1591 863 1617"><b>GWHOST = gateway_host_name</b></p> <p data-bbox="461 1633 1198 1659">Specifies the host name of the SAP gateway, if the server is R/2 or external.</p> <p data-bbox="461 1675 724 1701">Alias: GATEWAY_HOST</p>
GWSERV	<p data-bbox="461 1745 831 1770"><b>GWSERV = gateway_service</b></p> <p data-bbox="461 1787 1262 1812">Specifies the service of the SAP gateway, if the server is an R/2 server or external.</p> <p data-bbox="461 1829 762 1854">Alias: GATEWAY_SERVICE</p>

Option	Description
IEEE_REVERSE	<p><b>IEEE_REVERSE = Y   N</b></p> <p>Specifies whether floating point numbers are byte reversed. The possible values are Y (floating-point numbers are byte reversed) and N (floating-point numbers are not byte reversed). The default value for an R/3 application server on Windows NT is Y. On other platforms, the default value is N.</p>
INENCODING	<p><b>INENCODING = code_page</b></p> <p>Specifies the code page. Indicates the code page of the SAP server. The encoding is determined by the value returned by the SAP server. In some rare cases, it might be necessary to override this value by setting the <b>inencoding</b> = connection parameter.</p>
LANGUAGE	<p><b>LANGUAGE = language</b></p> <p>Specifies the SAP logon parameter language. The value for language is either the 2-byte ISO-language key or the 1-byte SAP-language. Examples for the language are EN, DE or E, D. If not specified, the language set on the SAP system is the default.</p> <p><i>Note:</i> When you access the SAP system using the driver, specify valid logon information including <b>CLIENT</b>, <b>USER NAME</b>, <b>PASSWORD</b>, and <b>LANGUAGE</b>. The user ID and password can also be passed through single sign-on (SSO). The driver performs a logon check at OPEN time.</p> <p>Alias: LANG, LNG, RFCLANG, RFCLNG</p>
MAX_TABLE_JOINS	<p><b>MAX_TABLE_JOINS = number</b></p> <p>Specifies the number of tables that can be used in a left outer join or an inner join in ABAP Open SQL. The default is 25.</p> <p>Alias: MAX_TABLES_JOIN, MAX_TABLES_JOINS, MAX_TABLE_JOIN</p>
MSHOST	<p><b>MSHOST = message_server_host</b></p> <p>Specifies the host name of the Message Server for load balancing.</p>
NAMESPACE	<p><b>NAMESPACE=namespace</b></p> <p>Specifies the name space to be viewed with directory services. Also used to apply table filters that limit the number of SAP namespaces returned by the SAS Federation Server Driver for SAP by specifying multiple items in a comma-delimited string: <b>NAMESPACE={T000, AAA, B1}</b>.</p>
PWD	<p><b>PWD=password</b></p> <p>Specifies the SAP logon parameter password.</p> <p><i>Note:</i> When you access the SAP system using the driver, specify valid logon information including <b>CLIENT</b>, <b>USER NAME</b>, <b>PASSWORD</b>, and <b>LANGUAGE</b>. The user ID and password can also be passed through single sign-on (SSO). The driver performs a logon check at OPEN time.</p> <p>Alias: PASSWORD, PASSWD, PW, PASS</p>
R3NAME	<p><b>R3NAME=system_name</b></p> <p>Specifies the name of the R/3 system, if load balancing is being used.</p>

Option	Description
RFC_STRING	<p><b>RFC_STRING = additional_rfc_options</b></p> <p>Specifies additional logon or connection parameters for the <b>RfcOpenEx ()</b> call. The SAS Federation Server Driver for SAP uses the <b>RfcOpenEx ()</b> call to log on to the SAP system. Using this option, parameters which are not connection attributes for the SAS Federation Server Driver for SAP, can be passed to the <b>RfcOpenEx</b> call. For example: <b>RFC_STRING = "ABAP_DEBUG=1"</b>. This option can be used to support future extensions of the <b>RfcOpenEx</b> call.</p> <p>Alias: RFCSTRING, RFC_OPTIONS_EXT, RFCOPENEX, ADDITIONAL_RFC_OPTIONS</p>
SAPLOGON_ID	<p><b>SAPLOGON_ID = saplogon_id</b></p> <p>Specifies the string defined for SAPLOGON on a Windows 32-bit system.</p>
SYSNR	<p><b>SYSNR=system_number</b></p> <p>Specifies the SAP system number, if load balancing is not being used. The number is the two-byte code that identifies the system on the host, for example, 00 and 01.</p> <p>Alias: SYS, SYSTEM, SYSNO</p>
TRACE	<p><b>TRACE = 0   1   Y   N</b></p> <p>Specifies if the SAS Federation Server Driver for SAP should trace requests. If the trace option is set to <b>1</b> or <b>Y</b>, the driver writes log information into a file. The RFC library logs messages in the <b>dev_rfc</b> file. The default setting is <b>0 (N)</b>, RFC trace is off.</p> <p><i>Note:</i> The RFC trace directory is set using the <b>RFC_TRACE_DIR</b> environment variable.</p>
UID	<p><b>UID = user</b></p> <p>Specifies the SAP logon parameter user.</p> <p><i>Note:</i> When you access the SAP system using the driver, specify valid logon information including <b>CLIENT</b>, <b>USER NAME</b>, <b>PASSWORD</b>, and <b>LANGUAGE</b>. The user ID and password can also be passed through single sign-on (SSO). The driver performs a logon check at OPEN time.</p> <p>Alias: USR, USER, RFCUSER, USERNAME, USERID</p>

## Installing and Configuring the SAS Federation Server Driver for SAP

### Overview

Installing the Driver for SAP involves several steps that you must complete in the appropriate sequence. Review the system requirements and authorization profiles, set up the SAS Federation Server Driver for SAP, and then install the SAP® components on the SAP system and SAS Federation Server.

After installing the SAS Federation Server Driver for SAP, you must complete additional steps to configure it to be used by SAS Federation Server.

### **SAS Federation Server Driver for SAP Requirements**

The following list identifies the system requirements for a SAS Federation Server Driver for SAP that are used by SAS Federation Server. After the driver is installed, verify that the following requirements have been met:

#### **SAP System**

The SAP Kernel Release 4.6C or higher is required.

The SAS Federation Server Driver for SAP for Windows® and UNIX® requires the 64-bit SAP Unicode RFC library, Release 7.10, which is provided by SAP AG. Refer to SAP Note 413708 for the current version, download and installation instructions.

#### **SAPGUI**

During the installation of the SAP components, a SAPGUI is required.

#### **User IDs**

An SAP user ID and password is required. The user ID must have appropriate authorizations to access data and use communication methods. For more information about customizing the authorization, see Authorization Profiles below.

To install and run SAS Federation Server Driver for SAP, the following SAP user IDs are required:

- **RFC user** This is an SAP user ID that is used for the communication link between the SAS Federation Server Driver for SAP and the SAP System Application Server. Typically, there are several RFC user IDs (one per person).
- **SAP System Administrator** An SAP System Administrator ID is required for the installation of ABAP programs and function modules, for the configuration of destinations and variant for batch operations, and for setting up authorizations for user IDs to use the SAS Federation Server Driver for SAP. This user ID is used only for the installation.

#### **Connectivity**

The SAS Federation Server Driver for SAP and the SAP Application Server usually use TCP/IP communication. Refer to the RFC documentation from SAP AG. The host of the SAP Application Server must be known by the host of the SAS Federation Server Driver for SAP. Alternatively, you can use the IP address to identify the SAP System Application Server. The TCP/IP services file must contain entries for the services, ports, and protocols used for the communication.

The following is an example for entries in the services file:

```
sapdp00 3200/tcp
sapdp01 3201/tcp
sapdp99 3299/tcp
sapgw00 3300/tcp
sapgw01 3301/tcp
...
sapgw99 3399/tcp
sapsp00 3400/tcp
sapsp01 3401/tcp
...
sapsp99 3499/tcp
```

*Note:* If the SAPGUI is installed on the system, the TCP/IP services file already contains these entries.

## Authorization Profiles

To install and use the SAS Federation Server Driver for SAP, a user ID with certain authorizations is required. An authorization has an authorization object. Several authorizations can be bundled together into an authorization profile.

If the batch functionality of the SAS Federation Server Driver for SAP is used, the RFC user ID needs to have authorization to submit batch jobs already released.

The RFC user IDs require authorizations for the following authorization objects:

Object	Minimum Requirement for Values	Example for Predefined Authorization
S_RFC (Authorization check for RFC access)	ACTVT: * RFC_NAME: * RFC_TYPE: *	S_RFC_ALL
S_TABU_DIS (Table maintenance via standard tools such as SM31)	ACTVT: 03 DICBERCLS: *	S_TABU_SHOW
S_BTCH_JOB (Background processing: Operations on Background Jobs) <sup>1</sup>	JOBACTION: RELE JOBGROUP: *	

<sup>1</sup> Required only if batch functionality of the RFC server is used.

The existing authorizations, for example S\_TABU\_SHOW, can be used. The S\_RFC and the S\_TABU\_DIS authorizations are in authorization profile A\_ANZEIGE.

## Setting Up the SAS Federation Server Driver for SAP

This section describes the set up for the SAS Federation Server Driver for SAP after the software has been installed.

Complete the following steps on SAS Federation Server:

1. Install the Unicode RFC libraries from SAP.

The SAS Federation Server Driver for SAP requires the 64-bit version of the SAP Unicode RFC libraries, Release 7.10. The libraries must be installed on SAS Federation Server. These shared libraries are provided by SAP AG. Download and install the 64-bit Unicode libraries for your platform, following the instructions in SAP Note 413708.

2. Set the environment variables.

The SAS Federation Server Driver for SAP executable uses the SAP shared libraries. You must add the location of the SAP RFC shared libraries to the shared library path environment variable specific to your operating system. For Windows, ensure that the shared libraries are installed in the system path, or add the directory of the installed SAP Unicode RFC libraries to the Path environment variable. For UNIX, replace rfclib\_directory in the table below with the directory where the RFC shared libraries are installed.

**Table 7.1** AIX

Bourne Shell	\$ LIBPATH=rfclib_directory:\$LIBPATH \$ export LIBPATH
C Shell	\$ setenv LIBPATH rfclib_directory:\$LIBPATH

**Table 7.2** HP-UX

Bourne Shell	\$ LD_LIBRARY_PATH=rfclib_directory:\$LD_LIBRARY_PATH \$ export LD_LIBRARY_PATH
C Shell	\$ setenv LD_LIBRARY_PATH=rfclib_directory:\$LD_LIBRARY_PATH

**Table 7.3** HP-UX for the Itanium Processor Family Architecture

Bourne Shell	\$ LD_LIBRARY_PATH=rfclib_directory:\$LD_LIBRARY_PATH \$ export LD_LIBRARY_PATH
C Shell	\$ setenv LD_LIBRARY_PATH=rfclib_directory:\$LD_LIBRARY_PATH

**Table 7.4** Linux for Intel Architecture, Linux for x64, Solaris, and Solaris for x64

Bourne Shell	\$ LD_LIBRARY_PATH=rfclib_directory:\$LD_LIBRARY_PATH \$ export LD_LIBRARY_PATH
C Shell	\$ setenv LD_LIBRARY_PATH=rfclib_directory:\$LD_LIBRARY_PATH

## Installing SAP Components

### Verify the Prerequisites

Make sure that you have fulfilled the following prerequisites:

#### SAPGUI Prerequisites

The installation of the SAS Federation Server Driver for SAP components require SAPGUI software to be installed on your PC or workstation.

*Note:* Although it is not absolutely necessary to have the SAPGUI installed on the same PC or workstation where the SAS Federation Server Driver for SAP is

going to be installed, you need access to the SAPGUI during the installation. Because the usage of the SAPGUI complements SAP functionality, it is recommended that the SAPGUI be installed on the same PC or workstation.

#### SAP Administrator ID Prerequisites

A valid SAP user ID and password is required. The user must have authorization to transport files and for RFC destination maintenance. It is strongly recommended to get assistance from your SAP System Administrator to perform these tasks.

### **Install ABAP Programs and Function Modules**

Delivery transport files are included in the SAS Federation Server Driver for SAP. These transport files include all of the components, ABAP programs and function modules needed to run the SAS Federation Server Driver for SAP.

The delivery transports have to be imported on each SAP application server that is going to be accessed by SAS Federation Server. If an SAP system is upgraded, the delivery transports have to be imported again.

Two sets of transports are included, one for releases prior to SAP Release 7.0 and one for SAP Release 7.0 and above. You must import the transport files that apply to your system.

Version	Transport	Purpose	To be applied to
SAP systems prior to SAP NetWeaver 7.0 (Kernel 6.40 or earlier)	SAPKA93120INSAS <i>Note:</i> This transport must be installed first.	Supports the SAS Federation Server Driver for SAP	All SAP systems to be accessed by the SAS Federation Server Driver for SAP
SAP NetWeaver 7.0 based systems and later	SAPKA93130INSAS <i>Note:</i> This transport must be installed first.	Supports the SAS Federation Server Driver for SAP	All SAP systems to be accessed by the SAS Federation Server Driver for SAP
	SAPKB93031INSAS	Supports new BI 7.0 authorization concept	Optional; SAP BI 7.0 systems and above; only apply if you are using the new authorization concept

To import the transport files to your SAP systems, follow the instructions below. The instructions are based on the usage of the tp program (a utility for transport between SAP systems) on the operating system level.

*Note:* Replace HOME in these instructions with the actual directory path where SAS Federation Server is installed.

1. Log on as SAP System Administrator to the SAP application server.
2. Move the transport files from SAS Federation Server into the appropriate directories on your SAP systems.

---

Windows	Copy the r3trans.exe file to your SAP application server and extract the files into the transport directory, for example, HOME:\share\SAP. The files for all transports will be put into the cofiles and data subdirectories.
---------	---

---

---

UNIX Copy the r3trans.tar file to your SAP application server and extract the files into the transport directory, for example, HOME:/share/SAP. Assuming that the TAR file is downloaded to the user's HOME directory, follow these procedures to extract the files into the cofiles and data subdirectory in /usr/sap/trans.

```
.$ cd /usr/sap/trans$ tar -xvf $HOME/r3trans.tar
```

---

3. Change to the transport program directory using the following command:

---

Windows <drive>:  
 cd \usr\sap\trans\bin

---

UNIX \$ cd /usr/sap/trans/bin

---

4. Load the transport into the transport buffer and import the transport into your SAP system with the following commands. Replace *SID* with the system ID for your SAP system.

```
tp addtobuffer SAPKA93120INSAS SID
tp import SAPKA93120INSAS SID U2
```

*Note:*

1. Make sure you are using the correct profile for the transport control program tp. In some cases it might be necessary to use the parameter pf= to specify the TPPARAM file.
2. Because the transport file uses a long name, the nbuffer=true TP option must be set. The option can either be maintained in the SAP system using transaction STMS, or it can be specified as a parameter to the tp command. Also, the TP option tp\_version= must be set to at least 264 to allow the long names.
3. The U2 option allows the originals to be overwritten if the user has previously installed these ABAP objects.
4. The transports contain only client-independent ABAP objects. The tp import can therefore use any existing client that is correctly set up for imports. Verify that the ABAP program RDDIMPDP is correctly scheduled in the client that you use for the import.
5. If the transport files are imported into a Unicode SAP system, use the transport profile parameter "setunicodeflag=true" to force setting the Unicode flags in the imported programs. Refer to SAP Note 330267 for more details. The "setunicodeflag=true" is not necessary if you are using the transports for SAP NetWeaver 7.0 based system and higher. Those transports have been created with the Unicode flag.

Considering these notes, the tp commands might require additional parameters. Replace *SID* with the system ID for the SAP system.

*Note:* The tp commands listed on several lines in the following examples should be entered on a single command line. Be sure to include a space before adding the text from each of the following lines.

---

```
Windows      tp addtobuffer
              SAPKA93120INSASSID
              pf=\usr\sap\trans\bin\TP_DOMAIN_SID.PFL
              -D"nbuffer=true" -D"tp_version=264"

              tp import
              SAPKA93120INSASSID
              pf=\usr\sap\trans\bin\TP_DOMAIN_SID.PFL
              -D"nbuffer=true" -D"tp_version=264"
```

---

```
UNIX         $ tp addtobuffer
              SAPKA93120INSAS SID
              pf=/usr/sap/trans/bin/TP_DOMAIN_sid.PFL
              -D"nbuffer=true" -D"tp_version=264"

              $ tp import
              SAPKA93120INSASSID
              pf=/usr/sap/trans/bin/TP_DOMAIN_sid.PFL
              -D"nbuffer=true" -D"tp_version=264"
```

---

#### SAP Release prior to SAP NetWeaver 7.0 (Kernel 6.40 or lower), Unicode SAP Server

---

```
Windows      tp addtobuffer
              SAPKA93120INSASSID
              pf=\usr\sap\trans\bin\TP_DOMAIN_SID.PFL
              -D"nbuffer=true" -D"tp_version=264" -D"setunicodedeflag=true"

              tp import
              SAPKA93120INSASSID
              pf=\usr\sap\trans\bin\TP_DOMAIN_SID.PFL
              -D"nbuffer=true" -D"tp_version=264" -D"setunicodedeflag=true"
```

---

```
UNIX         $ tp addtobuffer
              SAPKA93120INSASSID
              pf=/usr/sap/trans/bin/TP_DOMAIN_SID.PFL
              -D"nbuffer=true" -D"tp_version=264" -D"setunicodedeflag=true"

              $ tp import
              SAPKA93120INSASSID
              pf=/usr/sap/trans/bin/TP_DOMAIN_SID.PFL
              -D"nbuffer=true" -D"tp_version=264" -D"setunicodedeflag=true"
```

---

#### SAP NetWeaver 7.0 based systems and later, Unicode SAP Server

---

```
Windows      tp addtobuffer
              SAPKA93130INSASSID
              pf=\usr\sap\trans\bin\TP_DOMAIN_SID.PFL
              -D"nbuffer=true"

              tp import
              SAPKA93130INSASSID
              pf=\usr\sap\trans\bin\TP_DOMAIN_SID.PFL
```

---

```

UNIX
$ tp addtobuffer
SAPKA93130INSASSID
pf=/usr/sap/trans/bin/TP_DOMAIN_sid.PFL
-D"nbuffer=true"

$ tp import
SAPKA93130INSASSID
pf=/usr/sap/trans/bin/TP_DOMAIN_SID.PFL
-D"nbuffer=true"

```

### **Maintaining RFC Destinations**

*Note:* If the SAS Federation Server Driver for SAP will execute requests using the SAP batch processing facility (recommended), you must complete this section.

The SAS Federation Server Driver for SAP uses multiple RFC destinations (TCP/IP connection type) for accessing an SAP System in batch. The number of destinations setup for the SAS Federation Server Driver for SAP limits the number of concurrent requests to the SAP application server.

For example, create six destinations with connection type T and activation type Registered Server Program that can be used by the SAS® server. The program ID for the registered server program must be unique on the SAP gateway.

<b>RFC Destination Name</b>	<b>Program ID</b>
SAS1	RFC.SAS1
SAS2	RFC.SAS2
SAS3	RFC.SAS3
SAS4	RFC.SAS4
SAS5	RFC.SAS5
SAS6	RFC.SAS6

Complete the following steps:

1. Call transaction SM59 in SAP. Specify transaction code /nsm59 in the command field.
2. Click **Create**.
3. Enter SAS1 as the RFC destination.
4. Enter T as the Connection type.
5. Enter a description for the destination.
6. Click **Enter**.
7. Choose **Registration for the Activation Type** or **Registered Server Program** in the **Technical Settings** tab.
8. Enter the RFC.SAS1 as the program ID.

9. If required, enter the gateway host and gateway service in the Gateway Options panel. The gateway host is the host name of the local gateway and gateway service is usually `sapgwsysnr`, where `sysnr` is replaced by the system number of the SAP system.
10. (Unicode SAP systems only) Select the Unicode on the **MDMP & Unicode** tab. Ignore the message about performing the Unicode test. The Unicode test cannot be performed with the destinations created for the SAS Federation Server Driver for SAP.
11. Save the destination.
12. Repeat step 1 through 11 for each of the new RFC destinations.

### **Maintaining the /SAS/DESTS Table**

The RFC destinations defined in the previous step must be grouped into destination groups. The groups are defined in table /SAS/DESTS which is used for controlling the access to the destinations from the SAS Federation Server that accesses the SAP system. The destination group is a parameter of the SAS Federation Server Driver for SAP. The default is "SAS1".

Complete the following steps:

1. Call transaction SM30 in SAP. On the command line, type transaction code `/nsm30`.
2. In the **Table** field, enter the table name /SAS/DESTS.
3. In the **Restrict Data Range** field, select **No Restrictions**.
4. Click **Maintain**.
5. An information message appears. Click **OK**.
6. Click **New Entries**.
7. For each of the RFC destinations that you defined in step 2, enter the destination group ID as the SAS ID and the RFC destination name. The following examples define the destinations for destination group SAS1:

SAS ID	RFC Destination	Used
SAS1	SAS1	
SAS1	SAS2	
SAS1	SAS3	
SAS1	SAS4	
SAS1	SAS5	
SAS1	SAS6	

8. Save the table.

### **Activating BAdI Implementation**

The SAS Federation Server Driver for SAP has three basic implementations for table access authorization checks. The default implementation uses the SAP authorization

object S\_TABU\_DIS to check the authorization. If you want to use any of the other two implementations you have to activate the appropriate BAdI implementation.

**Table 7.5** BAdI Implementation

Default	Authorization object S_TABU_DIS
Classic BAdI /SAS/AUTHBW01	For BW and BI: User authorization checks at the InfoCube, InfoObject and ODS level using the reporting authorization (SAP standard authorization concept).
New BAdI enhancement / SAS/IM_AUTHBI01	For BI 7.0+ only: User authorization checks using the analysis authorization. This not only provides an authorization check for the infoProvider (infoCube, infoObject, DSO) but also column level restrictions on master data attributes and key figures, and row-level restrictions on attributes.

In releases prior to SAP NW BI 7.0, SAP uses the reporting authorization concept that uses the SAP standard authorization concept. If you want to activate the SAS implementation for those authorization checks:

1. Call transaction SE19 in SAP. In the command field, type transaction code `/nse19`.
2. Enter `/SAS/AUTHBW01` as the implementation.
3. Click **Activate**.

In BI 7.0, SAP introduces a new authorization concept for analysis authorization. If you want to use the SAS implementation for those authorization checks, import the appropriate transport (SAPKB92331INSAS). The implementation is activated by default. If you want to deactivate the implementation:

1. Call transaction SE19 in SAP. In the command field, type transaction code `/nse19`.
2. In the **Edit Implementation** field, select the **New BAdI** check box.
3. Enter `/SAS/IM_AUTHBI01` as the enhancement implementation.
4. Click **Change**.
5. Double-click the **BAdI Implementation** to deactivate (such as `/SAS/BADI_CHECK_FILTER`) and clear the **Implementation is active** check box in the **Runtime Behavior** field. Repeat for each of the implementations listed in the left hand side of the **Enh.Implementation Elements** tab.
6. Save and activate the changes.

---

## SAP HANA Driver Reference

### Understanding the SAS Federation Server Driver for SAP HANA

The SAS Federation Server Driver for SAP HANA (Driver for SAP HANA) enables Read and Write access to SAP HANA data sources. The driver supports both native SQL and FedSQL dialects.

Installation of the SAP HANA ODBC driver is required for the Driver for SAP HANA. See the installation documentation for the recommended version.

### Data Service Connection Options for SAP HANA

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

The following table describes the data service connection options for SAP HANA. When configuring a data service, you must include one of the following configurations to establish connection to an SAP HANA system:

- SERVER and INSTANCE
- SAPHANA\_DSN or a DSN in CONOPTS
- SERVER and PORT
- SERVER with a full server name and port

Option	Description
CATALOG	<p><b>CATALOG=<i>mysaphanacatalog</i></b></p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups logically related schemas. CATALOG is a required option.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>
DRIVER	<p><b>DRIVER=<i>SAPHANA</i></b></p> <p>Identifies the type of data service to which you want to connect. The data service SAPHANA represents the SAP HANA database.</p>

---

Option	Description
SAPHANA_DSN, DB, DATABASE	<p>Specifies the configured SAP HANA ODBC datasource to which you want to connect. Use this option if you have existing SAP HANA ODBC datasources that are configured on your client. This method requires additional setup, either through the ODBC Administrator control panel on Windows platforms or through the <code>odbc.ini</code> file on UNIX platforms.</p> <p>Here is an example of an <code>odbc.ini</code> entry in UNIX:</p> <pre>[SAPHANADSN] SERVERNODE=107.20.242.225:30015</pre> <p>Connection options specified in <code>CONOPTS=</code> are appended to the connection string. Use <code>CONOPTS</code> or <code>SAPHANA_DSN</code> to specify the DSN. Do not use both of these options to specify the DSN.</p>
SERVER	<p>Specifies the server name or IP address of the SAP HANA server. The port can be included in the specified value. The port number is <code>3[instance-number]15</code> (for example, <code>30015</code> for instance number <code>00</code>).</p> <p>You can specify a list of hostnames separated by a semicolon to support failover. If a host is not available, the next host from the list is used.</p> <p>alias: <code>SERVERNODE</code>, <code>SERVER</code>, <code>HOST</code></p> <p>Here are some examples using the <code>SERVER=</code> option:</p> <pre>SERVER=&lt;'&gt;server-name&lt;'&gt; SERVER=&lt;'&gt;server-name:port&lt;'&gt; SERVER='server-name:port;failover-server-name1:port;failover-server-name2:port'</pre>
PORT	<p><b>PORT=30015</b></p> <p>Specifies the port number that is used to connect to the specified SAP HANA server. If you do not specify the port, the instance number, or include the port number in the server specification, the default <code>30015</code> is used.</p> <p><i>Note:</i> <code>3[instance]15</code> is the port for the standard SQL communication for client access. This is the only port required for client access.</p>
INSTANCE	<p>Specifies the instance number of the SAP HANA database engine. The port number is <code>3[instance-number]15</code>. For example, <code>30015</code> is the port number for instance number <code>00</code>. If the port number is explicitly specified in either the <code>PORT=</code> or the <code>SERVER=</code> option, the <code>INSTANCE=</code> option is ignored. A warning is written to the server log.</p>
CONOPTS	<p><b>CONOPTS=(ODBC-compliant connection string)</b></p> <p>Specifies an ODBC-compliant database connection string using ODBC-style syntax. These options, combined with the <code>ODBC_DSN</code> option, must specify a complete connection string to the data source. If you include a <code>DSN=</code> or <code>FILEDSN=</code> specification within the <code>CONOPTS=</code> option, do not use the <code>ODBC_DSN=</code> connection option. However, you can specify the ODBC database-specific connection options by using <code>CONOPTS=</code>. Then you can specify an ODBC DSN that contains other connection information by using the <code>ODBC_DSN=</code> connection option.</p>
UID	<p><b>UID='user-ID'</b></p> <p>Specifies the SAP HANA user name, or user ID that you use to connect to a database. If the user name or ID contains spaces or nonalphanumeric characters, enclose it in quotation marks.</p>

Option	Description
PWD	<p><b>PWD=' user password'</b></p> <p>Specifies the password that is associated with your SAP HANA user name. If the password contains spaces or non-alphanumeric characters, you must enclose it in quotation marks. You can also specify PASSWORD= with the PWD=, PASS=, and PW= aliases.</p>

### Secure Sockets Layer (SSL) Connection Options

The Driver for SAP HANA supports SSL. Here are the connection options for SSL.

Option	Description
ENCRYPT	<p><b>ENCRYPT=0   1</b></p> <p>Used to enable or disable SSL encryption. The default is 0 (NO).</p>
SSLCRYPTOPROVIDER	<p><b>SSLCRYPTOPROVIDER=SAPCRYPTO   OPENSLL   MSCRYPTO</b></p> <p>Specifies the cryptographic library provider for SSL connectivity.</p> <p>Alias: SSLPROVIDER</p>
SSLKEYSTORE	<p><b>SSLKEYSTORE=' file_path'</b></p> <p>Specifies the path to the keystore file that contains the server's private key. If a value is not specified, the ODBC driver uses the default <b>\$HOME/.ssl/key.pem</b>.</p>
SSLTRUSTSTORE	<p><b>SSLTRUSTSTORE=' file_path'</b></p> <p>Specifies the path to the trust store file that contains the server's certificate. If a value is not specified, the ODBC driver uses the default <b>\$HOME/.ssl/trust.pem</b>.</p> <p><i>Note:</i> Leave this option empty if you are using the cryptographic library, mscrypto.</p>
SSLVALIDATECERTIFICATE	<p><b>SSLVALIDATECERTIFICATE=NO   YES   0   1</b></p> <p>Set this option to validate the server's certificate. Setting to YES or 1 activates validation. The default is NO or 0. If this option is not specified, the ODBC driver uses the default and does not validate certificates.</p>
SSLHOSTNAMEINCERTIFICATE	<p><b>SSLHOSTNAMEINCERTIFICATE=' string'</b></p> <p>Specifies the host name to use for validation. Use this host name when validating a server's certificate using SSLVALIDATECERTIFICATE.</p> <p>Alias: SSLHOSTNAMEINCERT</p>

Option	Description
SSLCREATESELSIGNEDCERTIFICATE	<p><b>SSLCREATESELSIGNEDCERTIFICATE=NO   YES   0   1</b></p> <p>Specifies if a self-signed certificate is created if the keystore cannot be found. If set to YES, a self-signed certificate is created in the event that the keystore is not found. If this option is not specified, the driver uses the default, which is NO.</p> <p>Alias: SSLCREATECERT</p>

### Advanced Connection String Options

The Driver for SAP HANA supports the following advanced connection options.

Option	Description
CT_PRESERVE	<p><b>CT_PRESERVE = STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p>Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li>• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>

Option	Description
DEFAULT_ATTR	<p data-bbox="445 243 884 268"><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p data-bbox="445 285 1358 342">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="445 359 1374 768" style="list-style-type: none"> <li data-bbox="445 359 1374 415">• <b>CURSORS=n</b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="735 449 1337 768" style="list-style-type: none"> <li data-bbox="735 449 1286 533">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="735 554 1326 638">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="735 659 1337 768">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="477 789 940 814">Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> <ul data-bbox="445 835 1366 1050" style="list-style-type: none"> <li data-bbox="445 835 1366 919">• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li data-bbox="445 940 1366 1050">• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>

Option	Description
DEFAULT CURSOR TYPE	<p data-bbox="445 243 1362 296"><b>DEFAULT_CURSOR_TYPE=FORWARD_ONLY   KEYSET_DRIVEN   DYNAMIC   STATIC;</b></p> <p data-bbox="445 317 1310 342">Specifies a valid default cursor type for new statements. The valid options are as follows:</p> <ul data-bbox="445 359 1382 1077" style="list-style-type: none"> <li data-bbox="445 359 1382 474">• <b>FORWARD_ONLY</b> Specifies a non-scrollable cursor that moves only forward through the result set. Forward-only cursors are dynamic in that all changes are detected as the current row is processed. If an application does not require scrolling, the forward-only cursor retrieves data quickly, with the least amount of overhead processing.</li> <li data-bbox="445 491 1382 657">• <b>KEYSET_DRIVEN</b> Specifies a scrollable cursor that detects changes that are made to the values of rows in the result set but that does not always detect changes to deletion of rows and changes to the order of rows in the result set. A keyset-driven cursor is based on row keys, which are used to determine the order and set of rows that are included in the result set. As the cursor scrolls the result set, it uses the keys to retrieve the most recent values in the table.  It is sometimes helpful to have a cursor that can detect changes in the rows of a result set. A keyset-driven cursor uses a row identifier rather than caching the entire row into memory. Therefore, it uses much less disk space than other row caching mechanisms. Deleted rows can be detected when a SELECT statement that references the bookmark, row ID, or key column values fails to return a row.</li> <li data-bbox="445 835 1382 951">• <b>DYNAMIC</b> Specifies a scrollable cursor that detects changes that are made to the rows in the result set. All INSERT, UPDATE, and DELETE statements that are made by all users are visible through the cursor. The dynamic cursor is good for an application that must detect all concurrent updates that are made by other users.</li> <li data-bbox="445 968 1382 1077">• <b>STATIC</b> Specifies a scrollable cursor that displays the result set as it existed when the cursor was first opened. The static cursor provides forward and backward scrolling. If the application does not need to detect changes but requires scrolling, the static cursor is a good choice.</li> </ul> <p data-bbox="445 1094 1337 1150"><i>Note:</i> The application can override this value, but if the application does not explicitly set a cursor type, the value specified in <b>DEFAULT_CURSOR_TYPE</b> is in effect.</p>

Option	Description
DRIVER TRACE	<p data-bbox="445 243 916 270"><b>DRIVER_TRACE='API   SQL   ALL';</b></p> <p data-bbox="445 289 1382 401">Requests tracing information, which logs transaction records to an external file that can be used for debugging purposes. The SAS Federation Server driver writes a record of each command that is sent to the database to the trace log based on the specified tracing level, which determines the type of tracing information. The tracing levels are:</p> <ul data-bbox="445 420 1342 688" style="list-style-type: none"> <li data-bbox="445 420 775 447">• <b>ALL</b> Activates all trace levels.</li> <li data-bbox="445 464 1342 543">• <b>API</b> Specifies that API method calls be sent to the trace log. This option is most useful if you are having a problem and need to send a trace log to Technical Support for troubleshooting.</li> <li data-bbox="445 560 1206 588">• <b>DRIVER</b> Specifies that driver-specific information be sent to the trace log.</li> <li data-bbox="445 604 1318 688">• <b>SQL</b> Specifies that SQL statements that are sent to the database management system (DBMS) be sent to the trace log. Tracing information is DBMS specific, but most SAS Federation Server drivers log SQL statements such as SELECT and COMMIT.</li> </ul> <p data-bbox="445 705 767 732"><b>Default:</b> Tracing is not activated.</p> <p data-bbox="445 749 1326 858"><i>Note:</i> If you activate tracing, you must also specify the location of the trace log with DRIVER_TRACEFILE=. Note that DRIVER_TRACEFILE= is resolved against the TRACEFILEPATH set in ALTER SERVER. TRACEFILEPATH is relative to the server's content root location.</p> <p data-bbox="445 875 1262 903">(Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p> <p data-bbox="445 919 1374 1003"><b>Interaction:</b> You can specify one trace level, or you can concatenate more than one by including the   (OR) symbol. For example: <b>driver_trace='api   sql'</b> generates tracing information for API calls and SQL statements.</p>
DRIVER TRACE FILE	<p data-bbox="445 1045 871 1073"><b>DRIVER_TRACEFILE='filename';</b></p> <p data-bbox="445 1092 1382 1146">Used to specify the name of the text file for the trace log. Include the filename and extension in single or double quotation marks. For example:</p> <p data-bbox="445 1150 916 1178"><b>driver_tracefile='mytrace.log'</b></p> <p data-bbox="445 1194 1310 1249"><b>Default:</b> The default TRACEFILE location applies to a relative filename, and it is placed relative to TRACEFILEPATH.</p> <p data-bbox="445 1266 1230 1320"><b>Requirement:</b> DRIVER_TRACEFILE is required when activating tracing using DRIVER_TRACE.</p> <p data-bbox="445 1337 1086 1392"><b>Interaction:</b> (Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p>
DRIVER TRACE OPTIONS	<p data-bbox="445 1430 1254 1457"><b>DRIVER_TRACEOPTIONS=APPEND   THREADSTAMP   TIMESTAMP;</b></p> <p data-bbox="445 1476 1270 1503">Specifies options in order to control formatting and other properties for the trace log:</p> <ul data-bbox="445 1520 1374 1659" style="list-style-type: none"> <li data-bbox="445 1520 1374 1575">• <b>APPEND</b> Adds trace information to the end of an existing trace log. The contents of the file are not overwritten.</li> <li data-bbox="445 1591 1278 1619">• <b>THREADSTAMP</b> Prepends each line of the trace log with a thread identification.</li> <li data-bbox="445 1635 1150 1663">• <b>TIMESTAMP</b> Prepends each line of the trace log with a time stamp.</li> </ul> <p data-bbox="445 1680 1230 1707"><b>Default:</b> The trace log is overwritten with no thread identification or time stamp.</p>

Option	Description
TABLE TYPE	<p><b>TABLE_TYPE=ROW   COLUMN   LOCAL   LOCAL TEMPORARY   GLOBAL   GLOBAL TEMPORARY</b></p> <p>Specifies the default table type when creating tables using FedSQL (CREATE TABLE). This option can be overridden by the TABLE_TYPE table option. If the table store type is not specified in connection options nor in the table options, then the default SAP HANA store type is used.</p> <p><b>ROW</b> Creates a table using ROW-based storage in SAP HANA.</p> <p><b>COLUMN</b> Creates a table using COLUMN-based storage in SAP HANA.</p> <p><b>LOCAL   LOCAL TEMPORARY</b> Creates a local temporary table in SAP HANA. The table definition and data are visible only in the current session.</p> <p><b>GLOBAL   GLOBAL TEMPORARY</b> Creates a global temporary table in SAP HANA. The global temporary tables are globally available, and the data is visible only in the current session.</p>

---

## SAS Data Set Reference

### Overview

The SAS data set is a SASProprietary file format, which contains data values organized as a table of rows (SAS observations) and columns (SAS variables). The supported file format is the same as a SAS data set that is created by the BASE engine in SAS for Version 7 and later. A supported SAS data set uses the extension **.sas7bdat**.

### Understanding the Driver for Base SAS

The SAS Federation Server Driver for Base SAS is a SASProprietary driver that provides Read and Update access to legacy SAS data sets. With the SAS Federation Server Driver for Base SAS you can create SAS data sets that can be accessed by both the legacy and SAS Federation Server data access services.

The driver supports much of the Base SAS functionality, such as SAS indexing and general integrity constraints, as well as much of the Federated Query Language (FedSQL) functionality.

The SAS Federation Server Driver for Base SAS is an in-process driver, which means that it accesses data in the same process that executes the data access services. All server connections made with the SAS Federation Server Driver for Base SAS use **LOCKTABLE=SHARED** and **PATH\_BIND=ACCESS** connection options.

## Connection Options for SAS Data Sets

### Connection Options

To access data that is hosted on SAS Federation Server, a client must submit a connection string, which defines how to connect to the data. The data service connection arguments for a SAS data set include connection options and advanced options.

The following connection options are supported for SAS data sets:

Option	Description
CATALOG	<p><b>CATALOG=</b><i>catalog-identifier</i>;</p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups logically related schemas. A catalog name can be up to 32 characters long. You must specify a catalog.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>
DRIVER	<p><b>DRIVER=BASE</b>;</p> <p>Identifies the data service to which you want to connect, which is a SAS data set.</p> <p><i>Note:</i> You must specify DRIVER=BASE to access a SAS data set.</p>
(SCHEMA) NAME	<p><b>NAME=</b><i>schema-identifier</i>;</p> <p>Specifies an arbitrary identifier for an SQL schema. Any identifier is valid (for example, name=myfiles). The schema identifier is an alias for the physical location of the SAS library, which is much like the Base SAS libref. A schema name must be a valid SAS name and can be up to 32 characters long. You must specify a schema identifier.</p>
PRIMARY PATH	<p><b>PRIMARYPATH=</b><i>physical-location</i>;</p> <p>Specifies the physical location for the SAS library, which is a collection of one or more SAS files. For example, in directory-based operating environments, a SAS library is a group of SAS files that are stored in the same directory.</p> <p><i>Note:</i> You must specify a primary path.</p>
SCHEMA ATTRIBUTES	<p><b>SCHEMA=</b><i>(attributes)</i>;</p> <p>Specifies schema attributes that are specific to a SAS data set. A schema is a data container object that groups tables. The schema contains a name, which is unique within the catalog that qualifies table names. For a SAS data set, a schema is similar to a SAS library, which is a collection of tables with assigned attributes.</p>

### Advanced Options

Advanced driver options are additional options that are not required in order to connect to the data source. They are used to establish connections to catalogs, data source names (DSNs), and schemas. Although advanced options can also be used when connecting to a data service, doing so will cause the specified options to apply to all data service connections.

The following advanced options are supported for SAS data sets:

Option	Description
ACCESS	<p data-bbox="418 296 783 323"><b>ACCESS=READONLY   TEMP ;</b></p> <ul data-bbox="418 338 1390 472" style="list-style-type: none"> <li data-bbox="418 338 1390 401">• <b>READONLY</b> Assigns a read-only attribute to the schema. You cannot open a SAS data set to update or write new information.</li> <li data-bbox="418 411 1390 472">• <b>TEMP</b> specifies that the SAS data sets be treated as scratch files. That is, the system will not consume CPU cycles to ensure that the files do not become corrupted.</li> </ul> <p data-bbox="724 485 1390 625"><b>TIP</b> Use ACCESS=TEMP to save resources only when the data is recoverable. If TEMP is specified, data in memory might not be written to disk on a regular basis. This saves I/O, but could cause a loss of data if there is a crash.</p>
CT_PRESERVE	<p data-bbox="418 642 1230 669"><b>CT_PRESERVE = STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p data-bbox="418 684 1390 772">Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul data-bbox="418 787 1390 1192" style="list-style-type: none"> <li data-bbox="418 787 1390 850">• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li data-bbox="418 861 1390 949">• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="418 959 1390 1073">• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="418 1083 1390 1192">• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>
COMPRESS	<p data-bbox="418 1209 951 1236"><b>COMPRESS=NO   YES   CHAR   BINARY ;</b></p> <p data-bbox="418 1251 1007 1278">Controls the compression of rows in created SAS data sets.</p> <ul data-bbox="418 1293 1390 1633" style="list-style-type: none"> <li data-bbox="418 1293 1390 1356">• <b>NO</b> Specifies that the rows in a newly created SAS data set are uncompressed (fixed-length records). This setting is the default.</li> <li data-bbox="418 1367 1390 1480">• <b>YES   CHAR</b> Specifies that the rows in a newly created SAS data set are compressed (variable-length records) by using RLE (Run Length Encoding). RLE compresses rows by reducing repeated consecutive characters (including blanks) to two- or three-byte representations.</li> </ul> <p data-bbox="724 1495 1390 1535"><b>TIP</b> Use this compression algorithm for character data.</p> <ul data-bbox="418 1545 1390 1633" style="list-style-type: none"> <li data-bbox="418 1545 1390 1633">• <b>BINARY</b> Specifies that the rows in a newly created SAS data set are compressed (variable-length records) by using RDC (Ross Data Compression). RDC combines run-length encoding and sliding-window compression to compress the file.</li> </ul> <p data-bbox="724 1648 1390 1810"><b>TIP</b> This method is highly effective for compressing medium to large (several hundred bytes or larger) blocks of binary data (numeric columns). Because the compression function operates on a single record at a time, the record length must be several hundred bytes or larger for effective compression.</p>

Option	Description
DEFAULT_ATTR	<p data-bbox="432 243 874 268"><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p data-bbox="432 289 1347 344">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="432 365 1362 772" style="list-style-type: none"> <li data-bbox="432 365 1362 772">• <b>CURSORS=n</b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="724 453 1347 772" style="list-style-type: none"> <li data-bbox="724 453 1347 537">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="724 558 1347 642">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="724 663 1347 772">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="464 793 927 819">Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> <ul data-bbox="432 840 1378 1052" style="list-style-type: none"> <li data-bbox="432 840 1378 924">• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li data-bbox="432 945 1378 1052">• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>
ENCODING	<p data-bbox="432 1098 799 1123"><b>ENCODING=encoding-value;</b></p> <p data-bbox="432 1144 1273 1169">Overrides and transcodes the encoding for input or output processing of SAS data sets.</p> <p data-bbox="432 1190 1046 1215"><i>Note:</i> The default value is the current operating system setting.</p>
LOCKTABLE-	<p data-bbox="432 1255 831 1281"><b>LOCKTABLE=SHARED   EXCLUSIVE</b></p> <p data-bbox="432 1302 1362 1377">Places exclusive or shared locks on SAS data sets. You can lock tables only if you are the owner or have been granted the necessary privilege. The default value for SAS Federation Server is SHARED.</p> <ul data-bbox="432 1398 1362 1520" style="list-style-type: none"> <li data-bbox="432 1398 1362 1453">• <b>SHARED</b> Locks tables in shared mode, allowing other users or processes to read data from the tables, but preventing other users from updating.</li> <li data-bbox="432 1474 1362 1520">• <b>EXCLUSIVE</b> Locks tables exclusively, preventing other users from accessing any table that you open.</li> </ul>
PATH_BIND	<p data-bbox="432 1566 799 1591"><b>PATH_BIND=CONNECT   ACCESS</b></p> <p data-bbox="432 1612 1362 1749">Specifies when and how schemas are validated during connection. CONNECT validates the entire connection string at the time of connection and returns an error if one or more schemas is invalid. ACCESS validates schemas when they are accessed so that processing continues regardless of errors in the schema portion of the connection string. ACCESS is the default for SAS Federation Server.</p>

---

## SASHDAT Driver Reference

### About the SAS Federation Server Driver for SASHDAT

The SAS Federation Server Driver for SASHDAT (Driver for SASHDAT) is a write-only driver designed for use with Hadoop on the SAS LASR Analytic Server. SAS LASR Analytic Server integrates with Hadoop by storing SAS data in the Hadoop Distributed File system (HDFS). Using the Driver for SASHDAT, you can access the SAS LASR Analytic Server and transfer data to HDFS. Because the data volumes in HDFS are usually very large, the driver is not designed to read from HDFS and transfer data back to the client.

### Understanding the SASHDAT Data Source Driver

The Driver for SASHDAT enables a user to create a table and insert into it multiple times. Data is written when the connection is closed or when the user issues a commit or rollback. When a table is created, the Driver for SASHDAT stores the table definition in the connection. Catalog functions return information about the table, but the information is not written to HDFS until the first row is inserted. The table remains open and available for data until one of the following conditions are met:

- The SQL command COMMIT WORK is executed. This command closes and finalizes all tables with inserted data that are open for connection.
- A COMMIT or ROLLBACK is received from the client.
- The client issues a disconnect event at which time all tables with inserted data are closed.
- When the COMMIT=STATEMENT option is specified in the connection string, the table is closed when the INSERT statement changes to unprepared. If a user creates a table, prepares an INSERT, and executes it multiple times, the table remains open and available for more data. When the statement changes to unprepared, the table is closed.

When a table is closed, no further data can be written to it. To change the data in an HDFS table, it must be dropped and re-created.

### Connection Options

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

The Driver for SASHDAT supports the following connection options.

Option	Description
CATALOG	<p><b>CATALOG=</b><i>catalog name</i>;</p> <p>Specifies the catalog name for the connection.</p> <p><i>Note:</i> SAS Federation Server automatically quotes SQL identifiers that do not meet the regular naming convention as defined in the <i>SAS FedSQL Reference Guide</i>.</p>

---

Option	Description
COMMIT	<p><b>COMMIT=S   STATEMENT   C   CONNECTION;</b></p> <p>Specifies when the SASHDAT file will be closed. Use S to close when the statement is unprepared, or C when the connection is disconnected. The default is to close when the connection (C) is disconnected.</p>
COPIES	<p><b>COPIES=number-of-copies;</b></p> <p>Specifies how many copies are made when file blocks are written to HDFS. Note that specifying COPIES=0 is valid and signals the engine that you do not want any replicate copies of the data in HDFS. Defaults for this option depend on the setting for NODIST. The default is 1 when NODIST=NO is specified and 2 when NODIST=YES is specified.</p>
CT_PRESERVE	<p><b>CT_PRESERVE = STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p>Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul style="list-style-type: none"> <li>• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li>• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li>• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>
DEFAULT SCHEMA	<p><b>DEFSHEMA= schema-name;</b></p> <p>Specifies the default schema for the connection. When using the DEFSHEMA option, the default schema and path must also be specified with SCHEMA=(NAME= schema-name PATH  PRIMARYPATH= path) in the connection string.</p>
ENCODING	<p><b>ENCODING=SAS-NLS-encoding-identifier;</b></p> <p>Specifies the encoding for SASHDAT data and character conversions, both to and from. If not specified, the default encoding is inherited from SAS Federation Server.</p> <p><i>Note:</i> Tables created with the SASHDAT Data Source driver use the encoding specified in the connection string. If the encoding option is not specified, encoding defaults to the character set associated with the operating system for SAS Federation Server.</p>
HASH	<p><b>HASH=Y   YES   N   NO;</b></p> <p>Specifies the algorithm that determines the distribution of partitions to nodes of the LASR server proxy. The default is <b>HASH=NO</b>, which specifies that the distribution scheme depends on a binary tree. <b>HASH=YES</b> indicates that the distribution scheme depends on a hash function. As a result, the distribution properties of the partitions are not as balanced, but result in less memory usage. <b>HASH=YES</b> is recommended when working with high-cardinality partition keys (in the order of millions of partitions).</p>

Option	Description
HOST	<b>HOST</b>   <b>SERVER=grid-server-name;</b> Specifies the name of the SAS SAS LASR Analytic Server server or host. This option is required in the connection string. There is no default.
INSTALL	<b>INSTALL=path;</b> Specifies the path to the SAS SAS LASR Analytic Server server. This option is required in the connection string. There is no default.
LOCALE	<b>LOCALE=SAS-locale-identifier;</b> Specifies the locale for message text and character conversions, both to and from. The default locale is acquired from the server operating system.
NODIST	<b>NODIST</b>   <b>INNAMEONLY=Y</b>   <b>YES</b>   <b>N</b>   <b>NO</b> ; Specifies whether to place small tables into HDFS. Use NODIST=Y as the mode for placing small tables into HDFS. The default is N (No).
PASSWORD	<b>PWD=alternate-password;</b> Specifies the password for the alternate user when connecting to the SAS SAS LASR Analytic Server server.
SCHEMA	<b>SCHEMA= (NAME= schema-name PATH   PRIMARYPATH=path) ;</b> This option maps a logical schema name to a specific path for the SAS SAS LASR Analytic Server server. This option can be specified multiple times in a single connection string to define multiple schemas. At least one schema is required. There is no default if a schema is not specified. However, once specified, the first schema listed in the connection string is designated as the default schema if <b>DEFSHEMA=</b> is not used.
TIMEOUT	<b>TIMEOUT=timeout-in-seconds;</b> Specifies the amount of time to wait while trying to establish a connection before terminating the attempt and generating an error. The default is 20 seconds.
USER	<b>UID=alternate-userid;</b> Specifies the ID of an alternate user when connecting to the SAS SAS LASR Analytic Server server.

### Example Connection Strings

The following connection string connects a default user and closes files on disconnect, commit, or rollback:

```
CATALOG=HDAT;DRIVER=TSHDAT;HOST=rdcgrd001;INSTALL="/opt/v940m1/laxnd/TKGrid";
SCHEMA=(name=SCHEMA1;PATH="/user/dbitest");
```

This connection string connects the user **dbitest**, defines two schemas, stores data in UTF8, and closes files on statement unprepare, disconnect, commit, or rollback:

```
CATALOG=HDAT;DRIVER=TSHDAT;COMMIT=S;UID=dbitest;
ENCODING=UTF8;HOST=rdcgrd001;INSTALL="/opt/v940m1/laxnd/TKGrid";
SCHEMA=(name=CUSTOMERS;PATH="/user/custs");SCHEMA=(name=Accounts;PATH="/user/accts");
```

---

## Teradata Reference

### ***Understanding the SAS Federation Server Driver for Teradata***

The SAS Federation Server Driver for Teradata provides read and update access to Teradata database tables and creates tables that can be accessed by both SAS Federation Server and Teradata.

The SAS Federation Server Driver for Teradata supports most of the FedSQL functionality. The driver also supports an application's ability to submit native Teradata SQL statements.

The SAS Federation Server Driver for Teradata is a remote driver, which means that it connects to a server process to access data. The process might be running on the same machine as SAS Federation Server, or it might be running on another machine in the network.

### ***Data Service Connection Options for Teradata***

#### ***Overview***

To access data that is hosted on SAS Federation Server, a client must submit a connection string, which defines how to connect to the data. The data service connection arguments for a Teradata database include connection options and advanced options.

#### ***Connection Options***

Connection options are used to establish a connection to a data source. Specify one or more connection options when defining a data service using the [CREATE DATA SERVICE DDL statement on page 199](#).

The following connection options are supported for a Teradata database.

Option	Description
CATALOG	<p><b>CATALOG=<i>catalog-identifier</i>;</b></p> <p>Specifies an arbitrary identifier for an SQL catalog, which groups logically related schemas. Any identifier is valid (for example, catalog=tera).</p> <p><i>Note:</i> You must specify a catalog.</p>
DATABASE	<p><b>DATABASE=<i>database-name</i>;</b></p> <p>Specifies the Teradata database. If you do not specify DATABASE=, you connect to the default Teradata database, which is often named the same as your user ID. If the database value that you specify contains spaces or non-alphanumeric characters, you must enclose it in quotation marks.</p>
DRIVER	<p><b>DRIVER=TERA;</b></p> <p>Identifies the data service to which you want to connect, which is a Teradata database.</p> <p><i>Note:</i> You must specify the driver.</p>

---

Option	Description
SERVER	<b>SERVER=server-name;</b> Specifies the Teradata server identifier.

### **Advanced Connection Options**

The following advanced options are supported for Teradata database.

Option	Description
ACCOUNT	<b>ACCOUNT=account-ID;</b> Specifies an optional account number that you want to charge for the Teradata session.
CLIENT ENCODING	<b>CLIENT_ENCODING=encoding-value</b> Used to specify the character set for the session. UTF8 is the default if encoding is not specified. The character sets supported are:  ASCII EBCDIC EBCDIC037_0E KATAKANAEBCDIC KANJIUC_0U LATIN9_0A THAI874_4A0 LATIN1250_1A0 CYRILLIC1251_2A0 LATIN1254_7A0 HEBREW1255_5A0 ARABIC1256_6A0 LATIN1258_8A0 TCHBIG5_1R0 SCHINESE936_6R0 KANJI932_1S0 HANGUL949_7R0 TCHINESE950_8R0 LATIN1252_3A0 SCHEBCDIC935_2IJ TCHEBCDIC937_3IB HANGULEBCDIC933_1II EBCDIC273_0E EBCDIC277_0E KANJIEBBCDIC5035_0I KANJIEBBCDIC5026_0I UTF8 UTF16

Option	Description
CT_PRESERVE	<p data-bbox="432 243 1230 268"><b>CT_PRESERVE = STRICT   SAFE   FORCE   FORCE_COL_SIZE</b></p> <p data-bbox="432 289 1326 369">Allows users to control how data types are mapped. Note that data type mapping is disabled when CT_PRESERVE is set to STRICT. If the requested type does not exist on the target database, an error is returned. The options are as follows:</p> <ul data-bbox="432 390 1374 772" style="list-style-type: none"> <li data-bbox="432 390 1374 449">• <b>STRICT</b> The requested type must exist in the target database. No type promotion occurs. If the type does not exist, an error is returned.</li> <li data-bbox="432 464 1374 543">• <b>SAFE</b> Target data types are upscaled only if they do not result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="432 558 1374 676">• <b>FORCE</b> This is the default for all drivers. The best corresponding target data type is chosen, even if it could potentially result in a loss of precision or scale. When character encodings are changed, the new column size is recalculated to ensure all characters can be stored in the new encoding.</li> <li data-bbox="432 690 1374 772">• <b>FORCE_COL_SIZE</b> This option is the same as FORCE, except that the column size for the new encoding is the same as the original encoding. This option can be used to avoid column size creep. However, the resulting column might be too large or too small for the target data.</li> </ul>
DEFAULT_ATTR	<p data-bbox="432 816 874 842"><b>DEFAULT_ATTR=(attr=value;...)</b></p> <p data-bbox="432 863 1342 921">Used to specify connection handle or statement handle attributes supported for initial connect-time configuration, where <b>attr=value</b> corresponds to any of the following options:</p> <ul data-bbox="432 942 1374 1346" style="list-style-type: none"> <li data-bbox="432 942 1374 1001">• <b>CURSORS=n</b>- Connection handle option. This option controls the driver's use of client side result set cursors. The possible values are 0, 1 or 2. <ul data-bbox="719 1022 1374 1346" style="list-style-type: none"> <li data-bbox="719 1022 1374 1102">0 Causes the driver to use client-side static cursor emulation if a scrollable cursor is requested but the database server cannot provide one.</li> <li data-bbox="719 1123 1374 1203">1 Causes the driver to always use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is not used.</li> <li data-bbox="719 1224 1374 1346">2 (Default) Causes the driver to never use client-side static cursor emulation if a scrollable cursor is requested. The database server's native cursor is used if available. Otherwise, the cursor will be forward-only.</li> </ul> </li> </ul> <p data-bbox="464 1367 927 1392">Example: <b>DEFAULT_ATTR=(CURSORS=2)</b></p> <ul data-bbox="432 1413 1374 1625" style="list-style-type: none"> <li data-bbox="432 1413 1374 1493">• <b>USE_EVP=n</b> - Statement handle option. This option optimizes the driver for large result sets. The possible values are 0 (OFF) or 1 (ON), which is the default. Example: <b>DEFAULT_ATTR=(USE_EVP=0)</b></li> <li data-bbox="432 1514 1374 1625">• <b>XCODE_WARN=n</b> - Statement handle option. Used to warn on character transcoding errors that occur during row input or output operations. Possible values are 0 (returns an error), 1 (returns a warning), or 2 (ignore transaction errors). 0 is the default. Example: <b>DEFAULT_ATTR=(XCODE_WARN=1)</b></li> </ul>

Option	Description
DRIVER TRACE	<p><b>DRIVER_TRACE= 'API   SQL   ALL' ;</b></p> <p>Requests tracing information, which logs transaction records to an external file that can be used for debugging purposes. The SAS Federation Server driver writes a record of each command that is sent to the database to the trace log based on the specified tracing level, which determines the type of tracing information. The tracing levels are:</p> <ul style="list-style-type: none"> <li>• <b>ALL</b> Activates all trace levels.</li> <li>• <b>API</b> Specifies that API method calls be sent to the trace log. This option is most useful if you are having a problem and need to send a trace log to Technical Support for troubleshooting.</li> <li>• <b>DRIVER</b> Specifies that driver-specific information be sent to the trace log.</li> <li>• <b>SQL</b> Specifies that SQL statements that are sent to the database management system (DBMS) be sent to the trace log. Tracing information is DBMS specific, but most SAS Federation Server drivers log SQL statements such as SELECT and COMMIT.</li> </ul> <p><b>Default:</b> Tracing is not activated.</p> <p><i>Note:</i> If you activate tracing, you must also specify the location of the trace log with DRIVER_TRACEFILE=. Note that DRIVER_TRACEFILE= is resolved against the TRACEFILEPATH set in ALTER SERVER. TRACEFILEPATH is relative to the server's content root location.</p> <p>(Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p> <p><b>Interaction:</b> You can specify one trace level, or you can concatenate more than one by including the   (OR) symbol. For example: <b>driver_trace='api   sql'</b> generates tracing information for API calls and SQL statements.</p>
DRIVER TRACE FILE	<p><b>DRIVER_TRACEFILE= 'filename' ;</b></p> <p>Used to specify the name of the text file for the trace log. Include the filename and extension in single or double quotation marks. For example: <b>driver_tracefile='mytrace.log'</b></p> <p><b>Default:</b> The default TRACEFILE location applies to a relative filename, and it is placed relative to TRACEFILEPATH.</p> <p><b>Requirement:</b> DRIVER_TRACEFILE is required when activating tracing using DRIVER_TRACE.</p> <p><b>Interaction:</b> (Optional) You can control trace log formatting with DRIVER_TRACEOPTIONS=.</p>
DRIVER TRACE OPTIONS	<p><b>DRIVER_TRACEOPTIONS=APPEND   THREADSTAMP   TIMESTAMP ;</b></p> <p>Specifies options in order to control formatting and other properties for the trace log:</p> <ul style="list-style-type: none"> <li>• <b>APPEND</b> Adds trace information to the end of an existing trace log. The contents of the file are not overwritten.</li> <li>• <b>THREADSTAMP</b> Prepends each line of the trace log with a thread identification.</li> <li>• <b>TIMESTAMP</b> Prepends each line of the trace log with a time stamp.</li> </ul> <p><b>Default:</b> The trace log is overwritten with no thread identification or time stamp.</p>
PASSWORD	<p><b>PASSWORD=password ;</b></p> <p>Specifies a Teradata password. The password must correlate to your USER= value. The alias is PWD=.</p> <p><i>Note:</i> You must specify the PASSWORD= option.</p>
ROLE	<p><b>ROLE=security-role ;</b></p> <p>Specifies a security role for the session.</p>

Option	Description
USER	<b>USER=user-id;</b> Specifies a Teradata user ID. If the ID contains blanks or national characters, enclose it in quotation marks. The alias is UID= <i>Note:</i> You must specify the USER= option.

---



*Appendix 1*

# Administration DDL Statements Reference

---

<b>ALTER SERVER Statement</b> . . . . .	<b>195</b>
<b>CREATE DATA SERVICE Statement</b> . . . . .	<b>199</b>
<b>DROP DATA SERVICE Statement</b> . . . . .	<b>202</b>
<b>ALTER DATA SERVICE Statement</b> . . . . .	<b>203</b>
<b>CREATE CATALOG Statement</b> . . . . .	<b>205</b>
<b>DROP CATALOG Statement</b> . . . . .	<b>206</b>
<b>ALTER CATALOG Statement</b> . . . . .	<b>207</b>
<b>CREATE SCHEMA Statement</b> . . . . .	<b>207</b>
<b>DROP SCHEMA Statement</b> . . . . .	<b>209</b>
<b>ALTER SCHEMA Statement</b> . . . . .	<b>209</b>
<b>CREATE DSN Statement</b> . . . . .	<b>211</b>
<b>DROP DSN Statement</b> . . . . .	<b>214</b>
<b>ALTER DSN Statement</b> . . . . .	<b>214</b>
<b>CREATE CACHE Statement</b> . . . . .	<b>215</b>
<b>REFRESH CACHE Statement</b> . . . . .	<b>217</b>
<b>ALTER CACHE Statement</b> . . . . .	<b>217</b>
<b>DROP CACHE Statement</b> . . . . .	<b>218</b>
<b>PURGE CACHE Statement</b> . . . . .	<b>218</b>
<b>DROP AUTHID Statement</b> . . . . .	<b>218</b>
<b>GENERIC OPTIONS Syntax</b> . . . . .	<b>219</b>
<b>ALTER GENERIC OPTIONS Syntax</b> . . . . .	<b>219</b>
<b>GRANT and DENY Statements</b> . . . . .	<b>220</b>
<b>REVOKE Statement</b> . . . . .	<b>221</b>

---

## ALTER SERVER Statement

Enables you to change the server configuration by specifying server options. Here is the syntax:

```
ALTER SERVER
  [ alter-server-options ]
```

Here are the options for the ALTER SERVER statement:

#### **alter-server-options**

Specifies the list of server options to alter.

```
alter-server-options ::=
  "{" OPTIONS [{" alter-server-option
                [{" "," alter-server-option } ... ] [{"}"]}"
```

#### *alter-server-option*

Specifies the server option to alter.

```
alter-server-option ::=
  [alter-operation ]server-option | cache-option
```

#### *alter-operation*

Indicates the required action for the specified options. **ADD** is the default operation if a value is not set. The possible values are:

#### **ADD**

Use **ADD** to add the specified option.

#### **SET**

Use **SET** to change an option that already exists.

#### **XSET**

Use **XSET** to set the option if it has previously been added. Otherwise, **XSET** adds the option if it does not exist.

#### **DROP**

Use **DROP** to drop the specified option.

#### *server-option*

Specifies the server configuration as one of the following options:

#### **FILEDSPATH | FILEDSN\_ROOT**

```
FILEDSPATH | FILEDSN_ROOT directory-path
```

File DSN content location root for **FILEDSN** and **SAVEFILE** keywords. This path can be absolute or relative to the server's root content location specified in the server configuration file.

#### **PURGE CACHE**

```
PURGE_CACHE time-out-value
```

Specifies in minutes, how often old data cache tables are removed from the server. A positive time value indicates the interval at which the cleanup thread will check for items to purge. A negative time-out-value means that cleanup happens only in response to an explicit **PURGE CACHE** command. A value of **0** indicates that old cache tables are removed after a **CREATE CACHE**, **REFRESH CACHE**, or **PURGE CACHE** command is issued.

#### **TRACEFILEPATH | TRACEFILE\_ROOT**

```
TRACEFILEPATH | TRACEFILE_ROOT directory-path
```

Specifies the trace file content location root for **TRACEFILE** keywords and the **TRACEFILE** environment handle attribute. The trace file path is used to store all trace files that are created, either from SAS Federation Server start up or when enabled on a connection. This path can be absolute or relative to the server's ContentRoot that is specified in the server configuration file. The default trace file path is **drive:\Program Files\SASHome**

`\FedServer\server instance\var`. The `DRIVER_TRACEFILE=` path that is set in the connection option, `DRIVER TRACE`, is resolved against the path that is set here.

#### SHARED LOGIN MANAGER

`SHAREDLOGINMANAGER` *manager*

The Authentication Server manager account that will be used to retrieve the shared login map credentials. The shared login manager is paired with a password.

#### SHARED LOGIN PASSWORD

`SHAREDLOGINPASSWORD` *password*

The Authentication Server password for the manager account used to retrieve the shared login map credentials.

*Note:* Shared login manager and password options are paired.

#### SHARED LOGIN KEY

`SHAREDLOGINKEY` *key*

The Authentication Server grouping key used to search for shared login map credentials.

#### TRUSTED USER ID

`TRUSTED_USER_UID` *userid*

The Authentication Server trusted user account that is used to act on behalf of other users for purposes of retrieving user-owned logins and group memberships (trusting that those users have already been authenticated). The Federation Server makes a trusted user connection to Authentication Server to process queries on definer's rights views since these views retrieve SQL data under the identity of the view's schema owner (the definer) rather than the invoking user.

#### TRUSTED USER PWD

`TRUSTED_USER_PWD` *password*

Specifies the password for the `TRUSTED_USER_UID`.

*Note:* The Shared Login and Trusted User passwords are encrypted. SAS Proprietary Encryption is the default and the method of encryption can vary depending on the security platform where SAS Federation Server is installed. If DataFlux Secure for SAS Federation Server has been installed, AES encryption can be used.

#### CONNECTION POOLING

`CONNECTION_POOLING` [N[O] | F[ALSE] | OFF | Y[ES] | T[RUE] | O[N]]

This option controls connection pooling for the server. If connection pooling is enabled, database connections are not immediately disconnected upon client request. Instead, the connections are put into a connection pool so that they can be reused by subsequent connection requests for the same database with the same attributes and credentials.

#### CONNECTION POOLING TIMEOUT

`CONNECTION_POOL_TIMEOUT` *seconds*

This option identifies the time in seconds an unused connection stays in the connection pool. The default is 60 seconds. If the option is unset or set to 0, the connection stays in the pool for 60 seconds. If the time is exceeded, the connection is removed from the pool and the connection is closed.

## CONNECTION POOLING MAXSIZE

```
CONNECTION_POOL_MAXSIZE maxsize
```

This option identifies the maximum number of unused connections in the connection pool. The default is 50. If the option is not set or set to 0, a maximum number of 50 connections are kept in the pool. If the maximum number of connections is reached and a new connection is added to the connection pool, the oldest connection is removed from the pool and the connection is closed.

**cache-option :=**

Specifies a list of cache options and can be one of the following:

```
CACHE ( NAME cache-name, cache-property cache-property value )
```

The **NAME** option specifies the name of the cache. Cache properties for that particular cache are altered or created within the sublist. Normal generic SQL options syntax and rules apply to the CACHE option and its suboptions.

*cache name* :=

Specifies the name of the Authentication Server (AS) cache.

**AS**: All AS cached resources.

**AS.Name**: All AS.Name mappings.

**AS.Name.Subjects**: User name to AS identifier cache.

**AS.Name.Groups**: Group name to AS identifier cache.

**AS.Subject**: All AS.Subject cache resources.

**AS.Subject.Groups**: User group memberships cache.

**AS.Subject.Principals**: User principal listings cache.

**AS.List**: Listings.

**AS.List.Subjects**: User listings cache.

**AS.List.Groups**: Group listings cache.

**Authorization**: Privileges.

**ResultSet**: Result set caches.

**ResultSet.View**: Materialized view cache (VDC).

*cache-property* :=

Used to specify a property of the cache.

## TIMEOUT

```
TIMEOUT=n
```

**TIMEOUT** is the number of seconds before the cache data becomes stale after a refresh. After timing out, the cache is emptied and refreshed on demand or emptied automatically, depending on the cache implementation. **TIMEOUT** can be set for multiple related caches by specifying a non-terminal cache namespace for the name suboption such as **AS.List** instead of **AS.List.Groups**. A value of **-1** corresponds to infinite and a value of **0** corresponds to immediate. The default timeout value is **0** for all caches if not set through a parent namespace.

**CONOPTS**

Use the **CONOPTS** connection string option to call the FedSQL driver and set SQL statement limits.

**FedSQL Driver**

```
CONOPTS(driver FEDSQL)
```

FedSQL is the required driver for SQL requests on SAS Federation Server.

**SQL Statement Limit**

```
DEFAULT_ATTR(SQL_STMT_MEM_LIMIT = n)
```

Use the **SQL\_STATEMENT\_LIMIT** connection string option to control the amount of memory available to answer SQL requests, and enforce this limit for all connections. If the option is specifically set on a particular DSN, then the DSN value should override the system setting. (*n*)umber is treated as an integer and is specified in bytes.

**Examples:**

```
ALTER SERVER {OPTIONS add TRACEFILEPATH "C:\tracefiles"}
ALTER SERVER {OPTIONS add TRACEFILEPATH "logs\tracefiles"}
ALTER SERVER {OPTIONS xset SHAREDLOGINKEY 'DefaultKey'}
ALTER SERVER {OPTIONS( CACHE(NAME AS.Subject, TIMEOUT 300),
  CACHE(NAME AS.List.Subjects, TIMEOUT 60) )}
ALTER SERVER {OPTIONS(xset PURGE_CACHE 30)}
ALTER SERVER {options CONOPTS(driver FEDSQL,
xset DEFAULT_ATTR(SQL_STMT_MEM_LIMIT 8000000))}
```

---

## CREATE DATA SERVICE Statement

The **CREATE [DATA] SERVICE** statement enables you to create a data service using the following syntax:

```
CREATE [DATA] SERVICE data-service-name
  TYPE data-service-type
  [CATALOG [NAME] catalog-name]
  [DOMAIN [NAME] domain-name]
  [REGISTER [( catalog-name1 [,catalog-name2 ...]) | ALL] [register-options]]
  [data-service-options]
```

**CREATE [DATA] SERVICE *data-service-name***

Specifies the name of the data service. The following rules apply:

- The specified name cannot match the reserved name of the internal data service, BASE.
- The specified name cannot match the name of an existing data service.
- The specified name cannot match the name of an existing catalog for defined data services that do not support catalog names, unless the CATALOG option is used to specify a different name.

**TYPE *data-service-type***

Specifies the data service type. Here is a list of valid data service types:

```
DB2UNXPC
GENERIC
```

GENERIC\_FED  
 GREENPLUM  
 MDS  
 NETEZZA  
 ODBC  
 ODBC\_FED  
 ORACLE  
 POSTGRES  
 SAP  
 SAPHANA  
 SASHDAT  
 SQLSVR  
 TERADATA

**CATALOG** [NAME] *catalog-name*

Specifies the logical catalog name associated with the data service. The catalog name must be unique. If an identical catalog name is encountered, a warning message is issued. The default logical catalog name matches the data service name if omitted for data sources that do not support catalogs.

**DOMAIN** [NAME] *domain-name*

Specifies the Authentication Server domain name associated with the specified data service. If omitted, the default domain name matches the data service name. The domain is checked against the authentication server and if not valid, will return an error only if **VALIDATE** is specified for the data service. If **VALIDATE** is not specified, a warning is presented but the data service is still created.

**REGISTER** [(*catalog-name1* [, *catalog-name2*...]) | *ALL*]

**REGISTER** is an optional catalog registration specification. Specify a list of catalogs to register or use the **ALL** keyword to register all catalogs visible to the connection.

```

register-options  :=
  [ UID "userid" ]
  [ PWD 'password' ]
  [ VALIDATE [N[0] | F[ALSE] | OFF|0 | Y[ES] | T[RUE] | ON|1] ]
  
```

*Note:* When using register options, the **"userid"** requires double quotation marks and the **'password'** requires single quotation marks:

```

  UID "userid"
  PWD 'password'
  
```

The data service user id and password are used for making the connection to the data source so that the connection and catalog names may be validated.

**VALIDATE** specifies whether the connection, domain and catalog names are to be validated:

```

  [ VALIDATE [N[0] | F[ALSE] | OFF|0 | Y[ES] | T[RUE] | ON|1] ]
  
```

If **UID** or **PWD** is specified, but **VALIDATE** is not, **VALIDATE** runs with a default of **TRUE**. If a **UID** or **PWD** is specified and **VALIDATE** is specified

without a value, it defaults to TRUE. If VALIDATE is specified as TRUE but UID or PWD are not specified, then personal credentials are extracted on behalf of the caller, from the domain associated with the data service. The statement returns **TKTS\_ERROR** if credentials do not exist in the domain.

If an explicit catalog list is specified and VALIDATE is TRUE, then each catalog must exist or the statement returns **TKTS\_ERROR**. If VALIDATE is TRUE but no catalogs are specified (including REGISTER ALL), then only the connection itself is validated since the catalog list is either empty or supplied by the connection. If an attempt is made to register a catalog that has already been registered or associated with a data service, then the statement returns **TKTS\_SUCCESS\_WITH\_INFO**.

### data-service-options

Specifies the list of data service options.

```
data-service-options ::=
    "{" OPTIONS [{" data-service-option
                  [{" data-service-option } ... ] [{" }]" }
```

### data-service-option

Specifies the data service option.

```
data-service-option ::=
    data-service-dependent-option | data-service-independent-option
```

### data-service-dependent-option

Data service specific options. For details about the specific options for your data source, see [“Database Functionality and Driver Performance” on page 124](#).

### GENERIC

The following options apply for GENERIC data services:

```
LOCAL N[O] | F[ALSE] | OFF | 0 | Y[ES] | T[RUE] | ON | 1
```

Specifies if the data service refers to a local data source or an external database. Local sources do not require secondary authentication; and as such, specification of the **DOMAIN** clause results in an error if the **LOCAL** option is specified as one of the true values. The default is **NO**. This option is not persisted.

### GENERIC\_FED

Any options that apply to multi-catalog data services will apply to **GENERIC\_FED** data services.

```
GENERIC_FED-option ::= GENERIC-option
```

### data-service-independent-option

```
data-service-independent-option ::=
    conopts-configuration-list | case-sensitivity-option
conopts-configuration-list ::=
    CONOPTS "(" [DRIVER driver-name] [{" driver-connection-string-option ... }]" ...
```

### driver-name

If the **driver-name** option is omitted, the default driver for the data service is assumed. Associated options within the **CONOPTS** list are used for connections using the appropriate driver. Some data services such as ORACLE accept connections from the ODBC driver as well. For these data services, two **CONOPTS** lists can be configured, one per driver to accept connections for the two drivers. The ODBC driver accepts a **CONOPTS** driver connection string option. To configure this option and suboptions within it,

the configuration format is **CONOPTS**( **driver** **ODBC**, **CONOPTS**( ... ) ). The inner **CONOPTS** option, within the parenthesis, is a list-valued driver connection string, while the outer **CONOPTS** option groups arbitrary driver connection string options configured for the service.

*driver-connection-string-option*

Specifies the connection options that correspond to the driver which is specified in **DRIVER driver-name**. The **DATA\_SERVICE** and **CATALOG** connection string options should not be specified here since they are implied by the data service and its configured catalogs.

```
case-sensitivity-option ::=
CASE_SENSITIVITY "("
OBJECT N[O] | F[ALSE] | OFF|0 | Y[ES] | T[RUE] | ON|1 " , "
COLUMN N[O] | F[ALSE] | OFF|0 | Y[ES] | T[RUE] | ON|1 " ) "
```

This option specifies the case sensitivity to use when comparing identifiers for security purposes. False indicates case insensitive compares while True indicates case sensitive compares are used. If not specified, the value for case sensitivity is set to the default for the data service. The defaults are **TRUE** (sensitive) for DB2, Oracle, and GreenPlum, and **FALSE** (insensitive) for all other data services. Both **OBJECT** and **COLUMN** are required when specifying **CASE\_SENSITIVITY**.

*Note:* SCHEMA, CATALOG, DATA SERVICE, DSN, USER, and GROUP identifiers are always compared in a case insensitive manner.

**Examples:**

```
CREATE SERVICE ORASERV TYPE ORACLE domain ORA1 {OPTIONS ( conopts
( Driver odbc, conopts(DSN tktsora)), conopt (Driver oracle, PATH tktsora) ) }

CREATE DATA SERVICE SQLServer1 TYPE SQLSVR domain SQLSERVER {OPTIONS
( conopts ( conopts(DSN tktssql) ) ) }

CREATE SERVICE DB2_SERVICE TYPE DB2UNXPC domain DB2 {OPTIONS
( conopts (DB DEV1) ) }

CREATE SERVICE TERA_SERVICE TYPE TERADATA domain TERA {OPTIONS
( conopts (Server kaching.unx.df.com) ) }

CREATE SERVICE SAPSERV TYPE SAP DOMAIN SAPDOMAIN {OPTIONS
conopts(ashost apsrv.sup.com, sysnr 03, batch 1)}
```

---

## DROP DATA SERVICE Statement

Enables you to drop a data service. Here is the syntax:

```
DROP [DATA] SERVICE data-service-name [drop-disposition]
```

**data-service-name**

Specifies the name of the data service to drop.

*drop-disposition*

Specifies the drop disposition as one of the following values:

```
drop-disposition ::=
{RESTRICT | CASCADE} [FORCE]
```

RESTRICT	Specifies that the drop target is empty. This is the default value.
CASCADE	Specifies that contained objects are dropped.
FORCE	Specifies the optional FORCE keyword that will suppress error messages when the data service does not exist. This additional option does not affect the performance of the RESTRICT or CASCADE options.

### Examples

```
drop DATA SERVICE ORACLE3

drop service "MYSQL_SERVICE" cascade force

drop data service ORACLE1 cascade
```

---

## ALTER DATA SERVICE Statement

Enables you to change the name of a data service. You can also change common data service attributes such as version, catalog, and domain, and other data service specific options.

```
ALTER [DATA] SERVICE data-service RENAME TO newname

ALTER [DATA] SERVICE data-service

  [ CATALOG [NAME] catalog-name ]

  [ DOMAIN [NAME] domain-name ]

  [ REGISTER [( catalog-name1 [,catalog-name2 ...]) | ALL]

  [alter-data-service-options ]
```

### *data-service*

Specifies the data service name.

### *newname*

Specifies the new data service name.

### *catalog-name*

Specifies the domain name.

### *domain-name*

Specifies the domain name.

### REGISTER [(*catalog-name1* [, *catalog-name2*...]) | ALL]

Catalog registration specification. (Optional) Specify a list of catalogs to register or ALL keyword to register all catalogs visible to the connection.

```
register-options ::=
  [ UID principal ]
  [ PWD password ]
  [ VALIDATE [N|O] | F[ALSE] | OFF|0|Y[ES] | T[RUE] | ON|1]]
```

### UID *principal*

Data Service user principal name.

### PWD *password*

Data service user password.

VALIDATE [N[O] | F[ALSE] | OFF | 0 | Y[ES] | T[RUE] | ON | 1]

Specifies whether the connection and the catalog names are to be validated.

VALIDATE defaults to TRUE if not specified or if no Boolean value keyword is specified and either UID or PWD is specified. If VALIDATE is true and neither UID, nor PWD are specified, then a data service user principal name and password (personal credentials) are extracted on behalf of the caller from the domain associated with the data service. The statement returns TKTS\_ERROR if no such credentials exist.

If an explicit catalog list is specified and VALIDATE is TRUE, then each catalog must exist or the statement returns TKTS\_ERROR. If VALIDATE is TRUE but no catalogs are specified (including REGISTER ALL), then only the connection itself is validated since the catalog list is either empty or supplied by the connection.

### **alter-data-service-options**

Specifies the list of data service options to alter.

```
alter-data-service-options ::=
    "{ " OPTIONS ["(" alter-data-service-option
        [{ ", " alter-data-service-option } ...]
        [")"] " }
```

### *alter-data-service-option*

Specifies the data service option to alter.

```
alter-data-service-option ::=
    [ alter-operation ] data-service-option
```

### *data-service-option*

Specifies the data service option.

```
data-service-option ::=
    conopts-configuration-list
```

### *conopts-configuration-list*

If **DRIVER** *driver-name* is omitted, the default driver for the data service is assumed. Associated options within the **CONOPTS** list are used for connections using the appropriate driver. Some data services such as ORACLE accept connections from the SAS Federation Server Driver for ODBC as well. For these data services, two **CONOPTS** lists can be configured, one per driver to accept connections for the two drivers. The SAS Federation Server Driver for ODBC accepts a **CONOPTS** driver connection string option. To configure this option and suboptions within it, the configuration format is **CONOPTS( driver ODBC, CONOPTS( ... ) )**. The inner **CONOPTS** option, within the parenthesis, is a list-valued driver connection string, while the outer **CONOPTS** option groups arbitrary driver connection string options configured for the service.

```
conopts-configuration-list ::=
    CONOPTS "(" [DRIVER driver-name] [", "
    driver-connection-string-option ...] ")"...
```

### *driver-name*

Specifies the driver name.

*driver-connection-string-option*

Specifies the connection options that correspond to the driver which is specified in DRIVER driver-name. The DATA\_SERVICE and CATALOG connection string options should not be specified here since they are implied by the data service and its configured catalogs.

Both connection options and advanced options are valid to use as driver connection string options. For more information about which connection options and advanced options are supported for each data service, see the [SAS Federation Server Driver Reference on page 124](#) for your driver.

## Autoindex ON|OFF

This option is transient and valid for the SQL\_LOG data service only. Specifies whether to create indexes (ON) or to drop indexes (OFF) for the EVENTS table used for SQL Logging. The default is ON.

**Examples**

```
ALTER DATA SERVICE ORACLE3 {OPTIONS conopts(Driver odbc,
  ODBC_DSN tktsora) }

ALTER DATA SERVICE ORACLE3 RENAME TO ORACLE3_RENAME

ALTER DATA SERVICE ORACLE3_RENAME {OPTIONS DROP conopts(driver odbc) }

alter service service1 catalog newcatalog

alter service ORACLE4 DOMAIN ORA8

ALTER DATA SERVICE ORACLE3_RENAME {OPTIONS SET conopts(driver odbc,
  ODBC_DSN tktsora3) }

ALTER SERVICE SAPSERV {OPTIONS conopts(xset batch 1, xset destgroup SAS)};
```

---

## CREATE CATALOG Statement

The CREATE CATALOG statement enables you to create a catalog using the syntax below.

```
CREATE CATALOG "catalog" UNDER data-service
  [ NATIVE NAME native-name ]
  [ create-catalog-options ]
```

**“catalog”**

Specifies the catalog name. The catalog name is surrounded in double quotation marks.

**data-service**

Specifies the data service name under which the catalog is to be created.

**native-name**

Specifies the native catalog name. Specified when the native catalog name is not unique within the server. The native name should be used to resolve catalog name collisions between multiple data services that support catalogs. Client SQL always references the catalog via the logical catalog name (catalog) regardless of whether a native name is specified. Specifying a native name that matches the logical name does nothing.

**create-catalog-options**

Specifies the options to create a catalog. This option only applies to the BASE data service.

```
create-catalog-options ::=
create-catalog-options ::=
conopts-configuration-list
```

*conopts-configuration-list*

If **DRIVER***driver-name* is omitted, the default driver for the data service is assumed. Associated options within the **CONOPTS** list are used for connections using the appropriate driver. The multiple driver syntax is not supported.

```
conopts-configuration-list ::=
CONOPTS "(" [DRIVERdriver-name] ["," driver-connection-string-option ...] ")" ...
```

*driver-name*

Specifies the driver name.

*driver-connection-string-option*

Specifies the connection options that correspond to the driver which is specified in **DRIVER** *driver-name*.

Advanced options are valid to use as driver connection string options; connection options are invalid. For more information about which advanced options are supported for each data service, see the data service reference chapter for your driver. For information specific to the data service that you are configuring, refer to the section that corresponds to your data source.

**Examples:**

```
CREATE CATALOG "catalog1_BASE" UNDER BASE
CREATE CATALOG "TKTEST" UNDER SQLServer1
CREATE CATALOG "Catalog1" UNDER SQLServer1 NATIVE NAME "TKTEST"
CREATE CATALOG "c1" UNDER BASE {OPTIONS conopts (COMPRESS YES)}
```

## DROP CATALOG Statement

The DROP CATALOG statement enables you to drop a catalog by using this syntax:

```
DROP CATALOG "catalog" [ drop-disposition ]
```

**“*catalog*”**

Specifies the catalog name. Use double quotation marks when specifying the catalog name.

***drop-disposition***

```
drop-disposition ::=
{RESTRICT | CASCADE} [FORCE]
```

Specifies the drop disposition as one of the following values:

- RESTRICT Specifies that the drop target is empty. This is the default value.
- CASCADE Specifies that contained objects are dropped.
- FORCE Specifies the optional FORCE keyword that will suppress error messages when the data service does not exist. This additional

option does not affect the performance of the RESTRICT or CASCADE options.

**Examples:**

```
drop CATALOG "Catalog3"

drop catalog "catalog1_BASE" cascade
```

## ALTER CATALOG Statement

ALTER CATALOG enables you to change the name of a catalog. You can also change the native catalog name and the advanced options that are driver-specific. For information about which advanced options are supported for each data service, see the driver reference chapter for your data source.

```
ALTER CATALOG "catalog" RENAME TO "newcatalogname"

ALTER CATALOG "catalog"

    [ NATIVE NAME "native-name" ]

    [ alter-catalog-options ]
```

**“catalog”**

Specifies the name of the catalog.

**“newcatalogname”**

Specifies the new catalog name.

**“native-name”**

Specifies the name of the native catalog.

**alter-catalog-options**

Specifies the options to alter the catalog. This option only applies to the BASE data service. The syntax for alter-catalog-options is the same as the syntax for alter-generic-options. All create-catalog-options are also supported.

**Examples:**

```
ALTER CATALOG "catalog3_BASE" RENAME TO "catalog3_BASE_RENAME"

ALTER CATALOG "Catalog3" NATIVE NAME "TKTEST3_RENAME"

ALTER CATALOG "catalog1_BASE" {OPTIONS add CONOPTS(DRIVER BASE, ACCESS READONLY)}

ALTER CATALOG "catalog1_BASE" {OPTIONS set (CONOPTS(DRIVER BASE, ACCESS READONLY))}

ALTER CATALOG "catalog1_BASE" {OPTIONS xset CONOPTS(DRIVER BASE, COMPRESS YES)}

ALTER CATALOG "catalog1_BASE" {OPTIONS drop CONOPTS(DRIVER BASE)}
```

## CREATE SCHEMA Statement

Use the CREATE SCHEMA statement to create a schema and designate an owner. Here is the syntax:

```
CREATE SCHEMA [ "catalog"."schema"

    [ AUTHORIZATION|OWNER owner ]
```

```
[ create-schema-options ]
```

**“catalog”**

Specifies the optional catalog name under which to create the schema. This is useful for data services that are defined for data sources that support catalog names. For those that do not, the catalog name must be the logical catalog name which defaults to the name of the data service.

**“schema”**

Specifies the name of the schema.

**owner**

Authorization identifier of the schema owner. If the AUTHORIZATION clause is not specified, schema ownership defaults to ‘system’. However, schema ownership as ‘system’ could present problems with FedSQL views and data cache.

**create-schema-options**

Specifies an option list for the schema.

```
create-schema-options ::=
"{" OPTIONS ["(" schema-option
[ { ", " schema-option } ... ] ["")" ]}"
```

*schema-option*

Specifies the syntax for schema options. This option only applies to the BASE data service.

*conopts-configuration-list*

If **DRIVER driver-name** is omitted, the default driver for the data service is assumed. Associated options within the **CONOPTS** list are used for connections using the appropriate driver.

```
conopts-configuration-list ::=
CONOPTS "(" [DRIVER driver-name] [", "driver-connection-string-option ...] ")"...
```

*driver-name*

Specifies the driver name.

*driver-connection-string-option*

Specifies the connection options that correspond to the driver which is specified in **DRIVER driver-name**.

Advanced options are valid to use as driver connection string options; connection options are invalid. For more information about which advanced options are supported for each data service, see the data service reference chapter for your driver. For information specific to the data service that you are configuring, go to the section that corresponds to your data service.

**PRIMARYPATH path**

Specifies the physical location for the SAS library, which is a collection of one or more SAS files. For example, in directory-based operating environments, a SAS library is a group of SAS files that are stored in the same directory. This option is required for BASE schemas.

```
PRIMARYPATH path ::=
quoted-identifier
```

*quoted-identifier*

Specifies a single quoted or double quoted name.

**Examples:**

```
CREATE SCHEMA "catalog1_BASE"."schema1_BASE" {OPTIONS (primarypath
'C:\schema1_BASE')}

CREATE SCHEMA "ORACLE1"."TKTSTST1"

CREATE SCHEMA "catalog1"."schema1" {OPTIONS primarypath 'C:\my_schema',
conopts (LOCKTABLE EXCLUSIVE)}
```

---

## DROP SCHEMA Statement

Enables you to drop a schema.

```
DROP SCHEMA [ "catalog"."schema" ] [ drop-disposition ]
```

**“catalog”**

Specifies the catalog name.

**“schema”**

Specifies the schema name.

**drop-disposition**

Specifies the drop disposition and is one of the following values:

```
drop-disposition ::=
{RESTRICT | CASCADE} [FORCE]
```

**RESTRICT** Specifies that the drop target is empty. This is the default value.

**CASCADE** Specifies that contained objects are dropped.

**FORCE** Specifies the optional FORCE keyword that will suppress error messages when the data service does not exist. This additional option does not affect the performance of the RESTRICT or CASCADE options.

**Examples:**

```
DROP SCHEMA "catalog1_BASE"."schema1_BASE"

DROP SCHEMA "catalog1_BASE"."schema1_BASE" force
```

---

## ALTER SCHEMA Statement

Enables you to change the name of a schema. You can also alter advanced options that are driver-specific.

```
ALTER SCHEMA [ "catalog"."schema" ] RENAME TO "newschema"

ALTER SCHEMA [ "catalog"."schema" ] AUTHORIZATION|OWNER TO owner

[ create-if option ]

ALTER SCHEMA [ "catalog"."schema" ]

[ alter-schema-options ]
```

**“catalog”**

Specifies the catalog name.

**“schema”**

Specifies the schema name.

**“newschema”**

Specifies the new schema name.

**alter-schema-options**

Specifies what options to alter in the schema.

```
alter-schema-options ::=
  "{" OPTIONS ["("] alter-schema-option
  [{"," alter-schema-option} ... ] [")"]}"
```

*alter-schema-option*

Specifies the schema option to alter. This option only applies to the BASE data service.

```
alter-schema-option ::=
  [DROP schema-option-name ]
  [{ADD | SET} schema-option ]
  [create-if-option
```

*schema-option*

Specifies the syntax for schema options.

*conopts-configuration-list*

If **DRIVER***driver-name* is omitted, the default driver for the data service is assumed. Associated options within the **CONOPTS** list are used for connections using the appropriate driver.

```
conopts-configuration-list ::= CONOPTS "(" [DRIVER driver-name]
  [{"," driver-connection-string-option ...] ")" ...
```

*driver-name*

Specifies the driver name.

*driver-connection-string-option*

Specifies the connection options that correspond to the driver which is specified in **DRIVER** *driver-name*.

Advanced options are valid to use as driver connection string options; connection options are invalid. For more information about which advanced options are supported for each data service, see the data service reference chapter for your driver. For information specific to the data service that you are configuring, go to the section that corresponds to your data service.

**PRIMARYPATH** *path*

Specifies the physical location for the SAS library, which is a collection of one or more SAS files. For example, in directory-based operating environments, a SAS library is a group of SAS files that are stored in the same directory. This option applies to BASE schemas only and is required.

```
path ::= quoted-identifier
```

*quoted-identifier*

Specifies a single quoted or double quoted name.

*create-if-option*

Creates the schema if it does not already exist using the remaining options.

```
create-if-option
CREATE_IF N[O] | F[ALSE] | OFF|0 | Y[ES] | T[RUE] | ON|1
```

**Examples:**

```
ALTER SCHEMA "catalog1_BASE"."schema3_BASE" RENAME TO "schema3_BASE_RENAME"
ALTER SCHEMA "catalog1_BASE"."schema3_BASE" {OPTIONS set primarypath 'C:\mydir'}
ALTER SCHEMA "catalog1_BASE"."schema3_BASE" {OPTIONS add conopts (LOCKTABLE SHARE)}
```

---

## CREATE DSN Statement

The CREATE DSN statement enables you to create a standard, single-service DSN or a federated DSN. A federated DSN is created to group one or more standard DSNs.

Here is the syntax to create a standard DSN:

```
CREATE DSN dsn-name
UNDER data-service
[DESCRIPTION 'description text']
[CONNECT 'driver-connection-string-options']
[create-dsn-options]
[AS ADMINISTRATOR]
```

Here is the syntax to create a federated DSN:

```
CREATE DSN dsn-name
[DESCRIPTION 'description text']
[create-dsn-options]
ADD "(" dsn-name
["," ...] ")"
[AS ADMINISTRATOR]
```

***dsn-name***

Specifies the DSN name (required). Quotation marks surrounding the DSN name are optional.

```
CREATE DSN dsn-name
```

***data-service***

Specifies the data service name. This option only applies to a standard DSN and is required.

```
UNDER data-service-name
```

The following naming rules apply:

- The specified name must not match the reserved names of the internal BASE data service.
- The specified name must not match the name of any existing data service.
- The specified name must not match the name of an existing catalog for defined data services that do not support catalog names, unless the CATALOG option is used to specify a different name.

**DESC[RIPTION] '*description-text*'**

Description of the DSN surrounded in single quotation marks. Use for a standard or federated DSN (optional).

```
[DESC 'description text']
```

### **CONNECT *driver-connection-string-options***

Specifies the connection string options.

```
[CONNECT driver-connection-string-options']
```

The Federation Server driver connection string options are an extension of the ODBC syntax that specifies options as semicolon-delimited **key=value** pairs. For more information about which connection options and advanced options are supported for each data service, see the [driver reference topic on page 124](#) for your data source.

### **create-dsn-options**

Specifies what options are included with the DSN.

#### *dsn-config-options*

Specifies the options to configure with the DSN.

```
dsn-config-options ::=
"{" OPTIONS [{" dsn-config-option
  [{" dsn-config-option } ... ] [{" } ] }"
```

#### *dsn\_config\_option*

Specifies the DSN configuration option as one of the following:

FEDSQL Y[ES] | T[RUE] | 1 | N[O] | F[ALSE] | 0

Specifies whether to use the FedSQL dialect. BASE DSN connections always use FedSQL. Therefore, FedSQL cannot be turned off for BASE DSNs. The FEDSQL option applies to both standard and federated DSNs.

SECURITY Y[ES] | T[RUE] | 1 | N[O] | F[ALSE] | 0

YES is the default value. Specifies whether to secure SQL statements before processing them. For example, if a DSN is defined to use SECURITY NO, Federation Server security is bypassed. Therefore, when you connect with the DSN, you are connecting with the privileges granted at the data source level. If a DSN is defined to use SECURITY YES, privileges granted through the Federation Server will be enforced in addition to those of the underlying data source. Used in conjunction with CSO SHARED, this feature facilitates management of a more granular security policy in the Federation Server over a less granular one in the back-end database.

If SECURITY is set to YES, FEDSQL is automatically set to YES.

The SECURITY option applies to both standard and federated DSNs. It corresponds to the Federation Server SQL Authorization Enforcement setting that is displayed for the DSN through the Federation Server Manager.

On a standard DSN, SECURITY NO will ignore any SQL privileges configured in the Federation Server. SECURITY YES enforces SQL privileges. See Security Permissions in the Federation Server Authorization section for a list of privileges that are affected by the SECURITY setting.

On a federated DSN, SECURITY NO indicates that the security setting of the child DSN is used. A setting of NO allows each child DSN to operate under the security settings that are configured for it. Setting SECURITY to YES activates SQL privilege enforcement for all of the child DSNs affiliated with the federated DSN. Effectively, the SECURITY setting on a federated DSN can demand privilege enforcement on child DSNs, but cannot be used to remove it.

```
CREDENTIALS_SEARCH_ORDER | CSO "(" cso-value [ {" cso-value} ... ]
)"
cso-value ::= PERSONAL | SHARED
```

Specifies whether to use back-end credentials owned by the current user (**PERSONAL**) or shared among many users (**SHARED**). The DSN can be configured to search for either in the order specified. If a search is not specified, the default is **CSO (PERSONAL, SHARED)**. Credentials Search Order applies to a standard DSN only.

For example, if a user owns a database login and a DSN is configured with a CSO value of **PERSONAL**, the DSN will use that user's database login to connect to the database. However, when the DSN is configured using **CSO SHARED** and the user is configured as an authorized consumer of a shared login, the DSN connects using the shared login credentials. Users, logins, and shared logins are created and managed using the Authentication Server. For more information, see the *DataFlux Authentication Server User's Guide*.

### AS ADMINISTRATOR

```
[ AS ADMINISTRATOR ]
```

Creates the DSN using the **ADMINISTRATOR** role as the owner. With the **ADMINISTRATOR** role, the DSN is owned by the individual user. If the user is **SYSTEM**, the DSN is owned by **SYSTEM**. 'AS ADMINISTRATOR' is optional and can be used in a standard or federated DSN.

### Examples:

```
CREATE DSN "DSN1" UNDER BASE DESCRIPTION 'creating DSN1' NOPROMPT
'DRIVER=BASE; CATALOG="catalog1_BASE"; SCHEMA="schema1" {OPTIONS (FEDSQL
NO, SECURITY NO)}

CREATE DSN BASEDSN under BASE NOPROMPT 'DATA_SERVICE=BASE;
LOCKTABLE=EXCLUSIVE'

CREATE DSN BASEDSN under BASE CONNECT 'CATALOG="catalog1_BASE";
LOCKTABLE=EXCLUSIVE'

CREATE DSN BASEDSN under BASE NOPROMPT '(CATALOG="catalog1_BASE";
LOCKTABLE=SHARE); (CATALOG="catalog2_BASE"; LOCKTABLE=EXCLUSIVE) '

CREATE DSN BASEDSN under BASE NOPROMPT 'CATALOG="catalog1_BASE";
LOCKTABLE=SHARE; SCHEMA=(NAME="schema1_BASE"; LOCKTABLE=EXCLUSIVE) '

CREATE DSN BASEDSN under BASE CONNECT 'CATALOG="catalog1_BASE";
LOCKTABLE=SHARE; SCHEMA=(NAME="schema1_BASE"; LOCKTABLE=EXCLUSIVE);
SCHEMA=(NAME="schema2_BASE"; ACCESS=TEMP) '

CREATE DSN MYDSN under MYSERV {OPTIONS CREDENTIALS_SEARCH_ORDER
(PERSONAL)}

CREATE DSN MYDSN under MYSERV {OPTIONS CREDENTIALS_SEARCH_ORDER
(PERSONAL, SHARED)}

CREATE DSN ORADSN UNDER ORASERVICE CONNECT 'ORA ENCODING=UNICODE;
ORANUMERIC=YES'

CREATE DSN MYFEDERATED_DSN ADD mydsn1, mydsn2, mydsn3)
AS ADMINISTRATOR
```

---

## DROP DSN Statement

The DROP DSN statement enables you to drop a server-based DSN.

```
DROP DSN dsn-name [FORCE]
```

### *dsn-name*

Specifies the name of the DSN.

**FORCE** Specifies the optional FORCE keyword that will suppress error messages when the DSN does not exist.

### Examples:

```
DROP DSN "DSN1"
```

```
DROP DSN "DSN1" FORCE
```

---

## ALTER DSN Statement

Enables you to change the server-based DSN. You can change the name and alter advanced options. For information about which advanced options are supported for each data source, see the Driver Reference chapter for your driver.

### Syntax

```
ALTER DSN dsn-name alter-dsn-options
```

```
ALTER DSN dsn-name RENAME TO new-dsn-name
```

```
ALTER DSN dsn-name ADD "(" dsn-name ["," ...] ")"
```

```
ALTER DSN dsn-name DROP "(" dsn-name ["," ...] ")"
```

### *dsn-name*

Specifies the DSN name.

### *alter-dsn-options*

Specifies the options to alter.

```
alter-dsn-options
```

```
 ::= create-dsdsn-options
```

### *new-dsn-name*

Specifies the new DSN name.

### Examples:

```
ALTER DSN "DSN1" DESC 'altering DSN1 description' NOPROMPT
'DRIVER=BASE;CATALOG="catalog1_BASE";SCHEMA=(name="schema1_BASE")'
```

```
ALTER DSN "DSN5" RENAME to DSN7
```

```
ALTER DSN "DSN7" {OPTIONS set (FEDSQL
YES, SECURITY YES)}
```

```
ALTER DSN "DSN7" {OPTIONS xset CREDENTIALS_SEARCH_ORDER(SHARED),
xset FEDSQL NO, xset SECURITY NO}
```

```
ALTER DSN "DSN7" {OPTIONS DROP FEDSQL, DROP SECURITY}
```

---

## CREATE CACHE Statement

This statement creates a cache definition for a specified view in the cache catalog and schema and specifies options for the cache.

```
CREATE CACHE "catalog"."schema"."view"
IN "cache-catalog"."cache-schema"
[ {OPTIONS SAS-table-option=value [ ... SAS-table-option=value ] } ]
[ USING ( fedsql_syntax_sql ) ]
[ WITH COMMENT '...' ]
[ DEFERRED ]
[ FORCE ]
[ EXEC | EXECUTE ]
[ BEFORE ( fedsql_syntax_sql ) [ FORCE ]
      [, ( fedsql_syntax_sql ) [ FORCE ] [...] ] ]
[ AFTER ( fedsql_syntax_sql ) [ FORCE ]
      [, ( fedsql_syntax_sql ) [ FORCE ] [...] ] ]
[ CLEANUP ( fedsql_syntax_sql ) [ FORCE ]
      [, ( fedsql_syntax_sql ) [ FORCE ] [...] ] ]
[ SET ] option_name = 'string_literal' | numeric_literal
      | ON | OFF | YES | NO | TRUE | FALSE ]
```

### ***“catalog”***

Specifies the catalog name of the view to be cached.

### ***“schema”***

Specifies the schema name of the view to be cached.

### ***“view”***

Specifies the name of the cache view.

### ***“cache-catalog”***

Specifies the catalog name under which to create the cache.

### ***“cache-schema”***

Specifies the schema name under which to create the cache.

```
{OPTIONS SAS-table-option=value... }
```

Specifies the table options to use when the data cache table is created and populated.

```
[USING ( fedsql_syntax_sql )]
```

The optional USING clause provides a way for users to control how the cache table is created. The user must ensure that it is compatible with the view data cache table.

```
WITH COMMENT '...'
```

Text comments stored with the cache definition.

**DEFERRED**

The cache definition is stored but a cache table is not created or populated with data. A separate REFRESH CACHE command is needed to create and populate the cache table.

**FORCE**

The cache table and definition is retained even if an error occurs during REFRESH.

**EXEC | EXECUTE**

Execute command used with BEFORE, AFTER, and CLEANUP options.

**BEFORE**

```
BEFORE (fedsql_syntax) [FORCE]
```

The statements in `fedsql_syntax` will be executed before the cache table is created and populated. The FORCE option suppresses any errors.

**AFTER**

```
AFTER (fedsql_syntax) [FORCE]
```

The statements in `fedsql_syntax` will be executed after the cache table is created and populated. The FORCE option suppresses any errors.

**CLEANUP**

```
CLEANUP (fedsql_syntax) [FORCE]
```

The statements in `fedsql_syntax` will be executed only in the event of an error during creation or population of the cache table. The FORCE option will suppress any errors.

Use the following keywords with the USING, BEFORE, AFTER, and CLEANUP clauses. The keywords must be in UPPER CASE and contain no blank spaces with the brackets.

```
{CACHE}
```

Expands to a fully qualified cache table name using “double quotation marks”.

```
{CACHE_CATALOG}
```

Expands to a cache catalog name. Does not use quotation marks.

```
{CACHE_SCHEMA}
```

Expands to a cache schema name. Does not use quotation marks.

```
{CACHE_TABLE}
```

Expands to a cache table name. Does not use quotation marks.

**SET**

```
SET option_name=value
```

Specifies additional options and values to use during cache creation and population. The options are listed below.

**ERRLIMIT**

Sets a limit on the number of errors to allow before a statement stops inserting data.

**DBCMMIT**

Sets a limit on the number of modified rows to commit at one time, which affects transaction logging limits on the back-end database. This option overrides the ERRLIMIT option.

**INSERTBUFF**

Sets a limit for the number of rows inserted at a time which places a limit on a driver's row array size when inserting data.

**CT\_PRESERVE**

Sets the CT\_PRESERVE connection string option which controls how data types are mapped between the source view and the cache table. Valid options are **FORCE** | **FORCE\_COL** | **FORCE\_COL\_SIZE** | **STRICT** | **SAFE**.

**Examples:**

```
CREATE CACHE ORACLE_SERVICE.TKTSTST1.view_red in "SAMPLE"."tktstst1" set CT_PRESERVE='SA
```

## REFRESH CACHE Statement

The REFRESH CACHE statement refreshes the data cache on a specified view.

**Syntax**

```
ALTER CACHE "catalog"."schema"."view" REFRESH
```

**“catalog”.”schema”.”view”**

Specifies the catalog and schema and name of the data cache view.

**REFRESH** Refreshes cache of the specified cache view.

## ALTER CACHE Statement

With the ALTER CACHE statement, you can disable, enable or refresh cache tables. This requires the CREATE CACHE or ALTER CACHE privilege.

```
ALTER CACHE "catalog"."schema"."view" OPTION
```

**“catalog”.’schema”.”view”**

Specifies the catalog, schema and name of the view.

**OPTION**

Specifies an option for the statement as one of the following:

**DISABLE**

Disables use of the specified cache. References to the view use the view rather than the cached data.

**ENABLE**

Enables use of the specified cache after having been disabled.

**REFRESH**

Refreshes the cache table for the specified view.

**Examples**

```
ALTER CACHE "catalog1"."schema1"."view1" DISABLE
```

```
ALTER CACHE [ "catalog2"."schema2". ] view2 ENABLE
```

```
ALTER CACHE [ "catalog3"."schema3". ] view3 REFRESH
```

---

## DROP CACHE Statement

Enables you to drop a cached view.

```
DROP CACHE [ "catalog"."schema". ] view [FORCE ]
```

### **“catalog”.”schema”**

Specifies the catalog and schema of the data cache view.

### **view**

Specifies the name of the data cache view.

**FORCE** Suppresses error messages if the cache to be dropped does not exist.

---

## PURGE CACHE Statement

PURGE CACHE forces the removal of outdated cache tables that are no longer needed. PURGE CACHE is also used with options in the ALTER SERVER statement. Only system users or those with ADMINISTER privilege on the SAS® Federation Server can execute PURGE CACHE through the ALTER SERVER statement.

```
PURGE CACHE
```

---

## DROP AUTHID Statement

Enables you to drop authorization identifiers.

```
DROP { AUTHID | AUTHORIZATION [IDENTIFIER] } "ID" [TRANSFER TO  
name] [CASCADE|RESTRICT] [FORCE]
```

### **AUTHORIZATION IDENTIFIER “ID”**

Specifies the authorization identity to drop. The ID must be surrounded in double quotation marks. This is the user ID or the group ID that are found by using the AUTHORIZATION\_IDENTIFIERS information view.

### **TRANSFER TO name**

Specifies the authorization identity that will receive object ownership from the dropped identity. This user name is created using the Authentication Server.

### **Drop disposition**

Used to specify additional options.

#### **CASCADE**

Specifies the entities are dropped unconditionally. All records that reference the entity are removed. This option is invalid if the TRANSFER TO option is specified.

#### **RESTRICT**

Specifies that the drop fails if the entity is the grantor of any privilege. The drop also fails if the entity is the owner of a DSN or schema. This option is ignored if the TRANSFER TO option is specified.

**FORCE**

Specifies the optional FORCE keyword that will suppress error messages when the user does not exist.

**Examples:**

```
drop authid "5E563F78B0D70854086FB3D8441EF9AA" transfer to user2
drop authid "F135005B80DED494E996F70DCC53790D" cascade
drop authid "B8A105927F25B1A47AE8198D1E3C4B86" transfer to user2 force
```

## GENERIC OPTIONS Syntax

Generic options is a set of comma separated name or name-value pairs, within an OPTIONS list.

```
generic-options ::=
    "{" OPTIONS ["("] generic-option-list [")"] "}"
generic-option-list ::=
    generic-option [ {"," generic-option} ... ]
generic-option ::=
    option-name [ {option-value | option-list} ]

option-list ::=
    "("
option-value [ {"," option-value} ... ] ")"
option-value ::=
    quoted-identifier
```

## ALTER GENERIC OPTIONS Syntax

A set of comma separated name or name-value pairs with optional operation keywords, within an OPTIONS list.

```
alter-generic-options ::=
    "{" OPTIONS ["("] alter-generic-option-list [")"] "}"
alter-generic-option-list ::=
    alter-generic-option [ {"," alter-generic-option} ... ]
alter-generic-option ::=
    [alter-operation] option-name
    [ {option-value | option-list} ]
alter-operation ::= ADD | SET | XSET | DROP
```

**alter-operation**

Indicates the required action for the specified options. If a value is omitted for alter-operation, the default operation is ADD. The possible values are:

**ADD**

Adds the specified option.

**SET**

Changes an option that already exists.

**XSET**

Sets the option if it has already been added. Otherwise, adds the option if it does not already exist.

**DROP**

Drops the specified option.

---

## GRANT and DENY Statements

Enables you to give privileges to a specific user or all users to perform actions on objects. When submitting a grant, revoke, or deny request, surround all identifiers in double quotation marks, including table and column names.

*Note:* You cannot grant(deny) CREATE DSN and ADMINISTER privileges for the PUBLIC and USERS groups.

```
GRANT | DENY { { "objectpriv" | "containerpriv" |
"serverpriv" [,...] } |
                ALL [ PRIVILEGES ] }
[ ON { SCHEMA "schemaname" | CATALOG "catalogname" |
[DATA] SERVICE "servicename" | DSN "dsname" | SERVER } ]
TO { "authid" | PUBLIC | USERS } [, ...]
[ AS ADMINISTRATOR ]
```

**“objectpriv”**

Specifies the name of an object-level privilege as one of the following values:

- SELECT
- INSERT
- UPDATE
- DELETE
- REFERENCES

**“containerpriv”**

Specifies the name of a container-level privilege as one of the following values:

- CREATE VIEW
- ALTER VIEW
- DROP VIEW
- CREATE TABLE
- ALTER TABLE
- DROP TABLE

**“serverpriv”**

Specifies the name of the server-level privilege to grant or deny, as one of the following values:

- ADMINISTER
- TRACE
- CREATE DSN
- CONNECT

**SCHEMA “schemaname”**

Specifies the name of the schema.

**CATALOG “catalogname”**

Specifies the name of the catalog.

**[DATA] SERVICE “servicename”**

Specifies the name of the data service.

**DSN “dsnname”**

Specifies the name of the DSN.

**“authid”**

Specifies the user or group name for which the privileges are granted or denied.

**[ AS ADMINISTRATOR ]**

Grants privileges using the ADMINISTRATOR role. If this option is not used, the privilege is granted as the individual user. If the user is a system user, privileges are assigned with SYSTEM as the grantor.

**Examples:**

```
GRANT INSERT ON SCHEMA "BASE_CATALOG1"."schemal_BASE" TO "user1"
GRANT CONNECT ON SERVER TO "user1"
DENY CONNECT ON DSN "SQLSRVDSN1" TO "user1"
GRANT CREATE DSN ON DATA SERVICE "SQLSRV1" TO "user1"
GRANT ADMINISTER ON SERVER TO "user1"
DENY ALL ON SCHEMA "BASE_CATALOG1"."schemal_BASE" TO "user1"
```

---

## REVOKE Statement

Enables you to remove explicitly granted or denied privileges from the specified object.

```
REVOKE { { "objectpriv" | "containerpriv" |
"serverpriv" [,...] } |
ALL [ PRIVILEGES ] }
[ ON { SCHEMA "schemaname" | CATALOG "catalogname" |
[DATA] SERVICE "servicename" | DSN "dsnname" | SERVER } ]
FROM { "authid" | PUBLIC | USERS } [, ...]
```

**“objectpriv”**

Specifies the name of an object-level privilege as one of the following values:

- SELECT

- INSERT
- UPDATE
- DELETE
- REFERENCES

**“containerpriv”**

Specifies the name of a container-level privilege as one of the following values:

- CREATE VIEW
- ALTER VIEW
- DROP VIEW
- CREATE TABLE
- ALTER TABLE
- DROP TABLE

**“serverpriv”**

Specifies the name of the server-level privilege to grant or deny, as one of the following values:

- ADMINISTER
- TRACE
- CREATE DSN
- CONNECT

**ALL [PRIVILEGES]**

Specifies that (all) object, container, and server privileges be removed from the specified object.

**SCHEMA “schemaname”**

Specifies the name of the schema.

**CATALOG “catalogname”**

Specifies the name of the catalog.

**[DATA] SERVICE “servicename”**

Specifies the name of the data service.

**DSN “dsnname”**

Specifies the name of the DSN.

**“authid”**

Specifies the user or group name for which the privileges are granted or denied.

**[ AS ADMINISTRATOR ]**

Grants privileges using the ADMINISTRATOR role. If this option is not used, the privilege is granted as the individual user. If the user is a system user, privileges are assigned with SYSTEM as the grantor.

**Examples:**

```
REVOKE ALL ON SCHEMA "BASE_CATALOG1"."schema1_BASE" FROM "user1"
REVOKE INSERT ON DATA SERVICE BASE FROM "USER1"
REVOKE all on server from "user1"
```

## Appendix 2

# Information Views

---

<b>Visibility Rules for Information Views</b> .....	<b>224</b>
<b>AUTHORIZATION_IDENTIFIERS View</b> .....	<b>226</b>
<b>CACHES View</b> .....	<b>227</b>
<b>CATALOGS View</b> .....	<b>228</b>
<b>CATALOG_PRIVILEGES View</b> .....	<b>228</b>
<b>COLUMNS View</b> .....	<b>229</b>
<b>COLUMN_PRIVILEGES View</b> .....	<b>230</b>
<b>CONFIG_CATALOGS View</b> .....	<b>230</b>
<b>CONFIG_DATA_SERVICES View</b> .....	<b>231</b>
<b>CONFIG_DSNS View</b> .....	<b>231</b>
<b>CONFIG_OBJECTS View</b> .....	<b>232</b>
<b>CONFIG_SCHEMAS View</b> .....	<b>233</b>
<b>DATA_SERVICES View</b> .....	<b>233</b>
<b>DATA_SOURCE_NAMES View</b> .....	<b>234</b>
<b>DS_PRIVILEGES View</b> .....	<b>235</b>
<b>DSN_CONTENT View</b> .....	<b>236</b>
<b>DSN_LINEAGE View</b> .....	<b>237</b>
<b>DSN_PRIVILEGES and EFFECTIVE_DSN_PRIVILEGES Views</b> .....	<b>237</b>
<b>IDENTITY View</b> .....	<b>239</b>
<b>MESSAGES View</b> .....	<b>239</b>
<b>OBJECTS View</b> .....	<b>240</b>
<b>OBJECT_PRIVILEGES View</b> .....	<b>240</b>
<b>PRIVILEGES and EFFECTIVE_PRIVILEGES Views</b> .....	<b>241</b>
<b>SCHEMAS View</b> .....	<b>243</b>
<b>SCHEMA_PRIVILEGES View</b> .....	<b>243</b>
<b>X_COLUMN_PRIVILEGES/ X_EFFECTIVE_COLUMN_PRIVILEGES View</b> .....	<b>244</b>
<b>X_OBJECT_PRIVILEGES/X_EFFECTIVE_OBJECT_PRIVILEGES View</b> ..	<b>246</b>

---

## Visibility Rules for Information Views

Data visible in the information views is based on the user object and the privileges associated with the object. Therefore, user privileges determine what records are visible to the user. All users can query views but an empty result set is returned if the user does not have privileges to a specific view. Use the ADMIN DSN to connect to the information views.

A majority of the information views return system-level data that is relevant only to administrators or to technical support staff working with customers. The exceptions are certain information views that return privilege information, since users should be able to see what privileges they are granted on objects for which they have at least a single privilege.

The following information summarizes user visibility and the data returned from SAS Federation Server.

### Administrators and System Users

System users and server administrators can view all data in all information views. The following related views are restricted to system users and administrators only:

**Table A2.1** Administrators and System Users

Views	Visibility
AUTHORIZATION_IDENTIFIERS OBJECTS COLUMNS	System user and SAS® Federation Server administrators only

### Data Services

The following table lists the visibility rules that are associated with information views that are related to data services:

**Table A2.2** Data Services

Views	Visibility
DATA_SERVICES CONFIG_DATA_SERVICES	<p>A data service is visible to a user if:</p> <ul style="list-style-type: none"> <li>the user has CONNECT, ADMINISTER, or CREATE DSN privileges on the data service, or</li> <li>the user has CONNECT privilege on any data service DSN.</li> </ul>

### DSN

The following table lists the visibility rules that are associated with information views for data sources names:

**Table A2.3** DSN

Views	Visibility
DATA_SOURCE_NAMES	A data source name (DSN) is visible to a user if:
DSN_CONTENT	<ul style="list-style-type: none"> <li>the user is the owner of the DSN, or</li> </ul>
DSN_LINEAGE	<ul style="list-style-type: none"> <li>has CONNECT privilege on the DSN.</li> </ul>
CONFIG_DSNS	

### Catalogs and Schemas

SAS Federation Server needs to display catalogs and schemas for the BASE service without connecting to the data service first. This is different from other data services because SAS Federation Server Manager can connect to a data service and query it for an associated list of catalogs and schemas. Non-administrator users must be able to see BASE objects. One example is if the user has CREATE CACHE privilege and needs to be able to cache views from the user interface. Creating views from SAS Federation Server Manager is another example. Results from the catalogs and schemas information views will be filtered depending on the user's privileges.

**Table A2.4** Catalogs and Schemas

Views	Visibility
CATALOGS	A catalog is visible to a user if:
CONFIG_CATALOGS	<ul style="list-style-type: none"> <li>the data service is visible.</li> </ul>
SCHEMAS	A schema is visible to a user if:
CONFIG_SCHEMAS	<ul style="list-style-type: none"> <li>the data service is visible.</li> </ul>

### Object Privileges

The following table lists the visibility rules that are associated with information views for object privileges:

**Table A2.5** Object Privileges

Views	Visibility
DSN_PRIVILEGES	Privilege rows are visible to a user if:
DS_PRIVILEGES	<ul style="list-style-type: none"> <li>the user is the grantor of the privilege, or</li> </ul>
CATALOG_PRIVILEGES	<ul style="list-style-type: none"> <li>the user is the grantee of the privilege, or</li> </ul>
SCHEMA_PRIVILEGES	<ul style="list-style-type: none"> <li>one of the user's groups is the grantee of the privilege (including the USERS or PUBLIC group)</li> </ul>
OBJECT_PRIVILEGES	AND
COLUMN_PRIVILEGES	<ul style="list-style-type: none"> <li>the user has at least one privilege on the object in the view (DSN/data service/catalog/schema/object/column)</li> </ul>

### Data Cache

Data cache metadata is distributed between the CACHES, MESSAGES and CONFIG\_OBJECTS information views. Users with CREATE CACHE or ALTER CACHE privilege will need to see data from these information views.

**Table A2.6** Data Cache

Views	Visibility
CACHES	Data items are visible to a user if:
MESSAGES	<ul style="list-style-type: none"> <li>the item is a data cache item, and</li> </ul>
CONFIG_OBJECTS	<ul style="list-style-type: none"> <li>the user has CREATE CACHE or ALTER CACHE privilege on the item</li> </ul>

### Container and Object Privileges

Privileges in the container and object categories pertain to server, data services, catalogs, schemas, objects, and columns.

**Table A2.7** Container and Object Privileges

Views	Visibility
DSN_PRIVILEGES and EFFECTIVE_DSN_PRIVILEGES	Privileges for these items are visible to a user if:
PRIVILEGES and EFFECTIVE_PRIVILEGES	<ul style="list-style-type: none"> <li>the user is the grantee of the privilege.</li> </ul>
X_COLUMN_PRIVILEGES/ X_EFFECTIVE_COLUMN_PRIVILEGES	<ul style="list-style-type: none"> <li>one of the user's groups is the grantee of the privilege, including the USERS group.</li> <li>the privilege is granted in the PUBLIC group.</li> </ul>

## AUTHORIZATION\_IDENTIFIERS View

The AUTHORIZATION\_IDENTIFIERS view displays Authentication Server objects (users and groups) that have been resolved on SAS Federation Server. This view also shows inactive records for SAS Federation Server which are records created when an object is removed from the Authentication Server. Removing an object from the Authentication Server does not remove it from the AUTHORIZATION\_IDENTIFIERS view. Use the [DROP AUTHID | AUTHORIZATION IDENTIFIER on page 218](#) DDL statement to remove these objects from the AUTHORIZATION\_IDENTIFIERS view.

The following table lists the columns that will be displayed:

Name	Type	Description
NAME	VARCHAR(256)	Specifies the name of a group, role, or user that was created using the Authentication Server. NULL is displayed if the group, role, or user is orphaned.
ID	VARCHAR(256)	Specifies the authorization identifier from the Authentication Server.

Name	Type	Description
TYPE	CHAR(1)	Specifies the type of entity as G (Group) or U (User).

Only System or Federation Server administrators can view data in the AUTHORIZATION\_IDENTIFIERS view. All users can query the AUTHORIZATION\_IDENTIFIERS view, but the query will return an empty result set if the user is not a system user or SAS Federation Server administrator.

## CACHES View

The CACHES view displays information for all defined caches. The table below lists the columns that are associated with the CACHES view:

Name	Type	Description
CATALOG_NAME	WVARCHAR(256) Not Null	Specifies the catalog containing the object being cached.
SCHEMA_NAME	WVARCHAR(256) Not Null	Specifies the schema containing the object being cached.
OBJECT_NAME	WVARCHAR(256) Not Null	Specifies the name of the object being cached.
CACHE_STATUS	WCHAR(1) Not Null	Specifies the current status of the cache:  E – An ERROR row indicates that the last operation failed - there should only ever be one ERROR row (or 0 if the last command succeeded).  D - A DEFERRED row indicates that the last operation was specified as DEFERRED and has not yet been successfully refreshed - there should only ever be one DEFERRED row (or 0 if the view was successfully refreshed).  A – An ACTIVE row indicates that valid or complete cache data for the object exists, and is available for use. New users should always "see" the ACTIVE row with the latest timestamp. Old ACTIVE rows will go away as their user's pending transactions end. At start-up, only the newest ACTIVE row will remain.  S - SUSPENDED indicates that a cache has been disabled.
START_TS	TIMESTAMP Nullable	Specifies when the refresh cache operation began.
END_TS	TIMESTAMP Nullable	Specifies when the refresh cache operation completed.
MESSAGE_ID	INTEGER Nullable	Shows any message generated when the cache was defined or refreshed. (See the <a href="#">“MESSAGES View”</a> on page 239.
NUM_ROWS	BIGINT Nullable	Specifies the number of rows inserted into the data cache table. Note that this should be considered a rough estimate.

Name	Type	Description
NUM_BYTES	BIGINT Nullable	Specifies the number of bytes inserted into the data cache table. Note that this should be considered a rough estimate.
USER_NAME	WVARCHAR(256)	Shows the user who performed the CREATE CACHE or REFRESH cache operation that caused this row.

*Note:* Only FedSQL views with definer's rights can be cached.

Status Rules for rows in the CACHES Table

The most recent row for a cache is returned, no matter the status. If the most recent row was Active, that should be the only row returned. If the most recent row was Deferred or Error, the next most recent Active row is returned if one exists.

---

## CATALOGS View

The CATALOGS view contains the name and system identifier for each catalog.

Name	Type	Description
CATALOG_NAME	VARCHAR(256)	Specifies the name of the catalog.
DATA_SERVICE_NAME	VARCHAR(256)	Specifies the name of the data service associated with catalog.
NATIVE_CATALOG_NAME	VARCHAR(256)	Specifies the name of the native catalog. If not applicable, the value is NULL.

---

## CATALOG\_PRIVILEGES View

The CATALOG\_PRIVILEGES view displays the catalog-level privileges for each catalog. The table below lists the columns that are associated with the CATALOG\_PRIVILEGES view:

Name	Type	Description
CATALOG_NAME	VARCHAR(256)	The catalog name.
GRANTOR	VARCHAR(256)	The user name of the grantor.
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as R (Role) or U (User).
GRANTEE	VARCHAR(256)	Specifies the user name of the grantee.
GRANTEE_TYPE	CHAR(1)	Specifies the grantee type as G (Group) or U (User).

Name	Type	Description
PRIVILEGE_NAME	VARCHAR(20)	Privilege name specified as one of the following: SELECT UPDATE INSERT DELETE EXECUTE REFERENCES CREATE TABLE ALTER TABLE DROP TABLE CREATE VIEW ALTER VIEW DROP VIEW CREATE CACHE ALTER CACHE CREATE TABLESPACE
PRIVILEGE_TYPE	VARCHAR(5)	Specifies the privilege type as GRANT or DENY.
GRANTABLE	CHAR(1)	Specifies if the privilege can be granted to others. The only valid value is N (No).

## COLUMNS View

The COLUMNS view contains the name and system identifier for each column. The table below lists the columns that are associated with the COLUMNS view:

Name	Type	Description
CATALOG_NAME	VARCHAR(256)	Specifies the catalog name.
SCHEMA_NAME	VARCHAR(256)	Specifies the schema name.
OBJECT_NAME	VARCHAR(256)	Specifies the object name.
COLUMN_NAME	VARCHAR(256)	Specifies the column name.

---

## COLUMN\_PRIVILEGES View

The COLUMN\_PRIVILEGES view displays the privilege descriptors for all column-level privileges. The table below lists the columns that are associated with the COLUMN\_PRIVILEGES view:

Name	Type	Description
CATALOG_NAME	VARCHAR(256)	Specifies the name of the catalog.
SCHEMA_NAME	VARCHAR(256)	Specifies the name of the schema.
OBJECT_NAME	VARCHAR(256)	Specifies the name of the object.
COLUMN_NAME	VARCHAR(256)	Specifies the column name.
GRANTOR	VARCHAR(256)	Specifies the user name of the grantor.
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as U (User) or R (Role).
GRANTEE	VARCHAR(256)	Specifies the name of the user who is granted privileges.
GRANTEE_TYPE	CHAR(1)	Specifies the grantee type as G (Group) or U (User).
PRIVILEGE_NAME	VARCHAR(20)	Specifies the privilege name as one of the following values: SELECT UPDATE INSERT REFERENCES
PRIVILEGE_TYPE	VARCHAR(5)	Specifies the privilege type as GRANT or DENY.
GRANTABLE	CHAR(1)	Specifies if the privilege is grantable. The only valid value is N (No).

---

## CONFIG\_CATALOGS View

The CONFIG\_CATALOGS view contains generic configuration variables for each defined catalog. All configuration settings for a single catalog can be obtained by concatenating rows matching the correct CATALOG\_NAME and ordering the results by sequence.

The table below lists the columns that are associated with the CONFIG\_CATALOGS view:

Name	Type	Description
DATA	VARCHAR(128)	Specifies the configuration data as not NULL.
CATALOG_NAME	VARCHAR(256)	Specifies the name of the catalog.
SEQUENCE	SMALLINT	Specifies the configuration chunk sequence.

## CONFIG\_DATA\_SERVICES View

The CONFIG\_DATA\_SERVICES view contains generic configuration variables for each defined data service. All configuration settings for a single service can be obtained by concatenating rows matching the correct DATA\_SERVICE\_NAME and ordering the results by sequence.

The table below lists the columns that are associated with the CONFIG\_DATA\_SERVICES view:

Name	Type	Description
DATA	VARCHAR(128) Not Null	Specifies the configuration data.
DATA_SERVICE_NAME	VARCHAR(256)	Specifies the unique name of the data service.
SEQUENCE	SMALLINT	Specifies the configuration chunk sequence.

## CONFIG\_DSNS View

The CONFIG\_DSN view displays generic configuration variables for each DSN defined in the definition schema. All configuration settings for a single DSN can be obtained by concatenating rows that match the correct DSN\_NAME and ordering the results by sequence.

The table below lists the columns that are associated with the CONFIG\_DSNS view:

Name	Type	Description
DATA	VARCHAR(128) Not Null	Specifies the configuration data.
DSN_NAME	VARCHAR(256)	Specifies the unique name of the DSN.
SEQUENCE	SMALLINT	Specifies the configuration chunk sequence.

## CONFIG\_OBJECTS View

The CONFIG\_OBJECTS view contains the configuration statement for the specified object. Stored information varies by object. The configuration statement text is broken up into pieces by type based on how FedSQL parses the statement. Sub-text for each type is broken up to fit into the DATA column ordered by CFG\_TYPE\_SEQUENCE.

Concatenating the DATA entries for a given OBJECT\_NAME in SEQUENCE order will produce the original configuration statement syntax with the following exceptions:

- It will be normalized and broken up into pieces based on what options are specified.
- It can be reordered, although positional clauses like “EXEC” will always remain in sequence
- It will contain any comments that were in the original create cache statement that was submitted
- It might be modified to contain information specified by later statements (like ALTER).

*Note:* A view cache will return a single entry containing the most recent CREATE CACHE statement, an OBJECT\_TYPE of 2 (SAS View) or 3 (CACHE), and the catalog/schema/name of the view being cached.

The following table lists the available columns for this view. Data is visible only if the user has CREATE CACHE or ALTER CACHE privilege on the referenced view.

Name	Type	Description
CATALOG_NAME	WVARCHAR(256) Not Null	Specifies the catalog containing the object.
SCHEMA_NAME	WVARCHAR(256) Not Null	Specifies the schema containing the object.
OBJECT_NAME	WVARCHAR(256) Not Null	Specifies the object name.
OBJECT_TYPE	INTEGER Not Null	Specifies the object type: 1 – Table, 2 – SAS View, 3 – Cache, 4 – SAS Package
DATA	WVARCHAR(128) Not Null	A piece of text from the configuration statement.
SEQUENCE	INTEGER Not Null	Orders the entries for the entire configuration statement for a single object.
CFG_TYPE	INTEGER Not Null	Indicates the type for this piece of the configuration statement as parsed by FedSQL.
CFG_TYPE_SEQUENCE	INTEGER	Orders the entries within each type.

---

## CONFIG\_SCHEMAS View

The CONFIG\_SCHEMAS view contains generic configuration variables for schemas defined in the definition schema. All configuration settings for a single schema can be obtained by concatenating rows matching the correct CATALOG\_NAME and SCHEMA\_NAME and ordering the results by sequence.

The table below lists the columns that are associated with the CONFIG\_SCHEMAS view:

Name	Type	Description
DATA	VARCHAR(128) Not Null	Specifies the configuration data.
CATALOG_NAME	VARCHAR(256)	Specifies the name of the catalog. If not applicable, the value is NULL.
SCHEMA_NAME	VARCHAR(256)	Specifies the name of the schema.
SEQUENCE	SMALLINT	Specifies the configuration chunk sequence.

---

## DATA\_SERVICES View

The Data Services view displays information about each data service. It also shows a single entry for SAS Federation Server with a value in the data\_service\_name column of \_SERVER\_ and a value in the type column of SERVER.

The DATA\_SERVICES table contains one entry per configured data service, both internal and external. The table below lists the columns that are associated with the DATA\_SERVICES view:

Name	Type	Description
DATA_SERVICE_NAME	VARCHAR(256)	The unique name of the data service.
VERSION	CHAR(32)	The version of the data service.

Name	Type	Description
TYPE	CHAR(32)	Keyword for the data type. Valid values include: BASE DB2 GREENPLUM MDS ODBC ODBC_FED ORACLE TRAN TERADATA SQLSVR Other values are possible.
DOMAIN	VARCHAR(256)	The Authentication Server domain that is associated with the data service.

## DATA\_SOURCE\_NAMES View

The DATA\_SOURCE\_NAMES view contains one entry per configured DSN and includes the following:

- This view contains a default BASE DSN.
- This view also contains an entry for the ADMIN DSN which is a DSN generated by the system to be used with server administration DDL and system catalog queries. The value of the DATA\_SERVICE\_NAME column is \_SERVER\_.
- If SQL Logging is enabled, this view also shows the SQL\_LOG DSN.

The table below lists the columns that are associated with the DATA\_SOURCE\_NAMES view:

Name	Type	Description
DSN_NAME	VARCHAR(256)	Specifies the unique name of the DSN.
DATA_SERVICE_NAME	VARCHAR(256)	Specifies the unique name of the data service.
DESC	VARCHAR(256)	Specifies the descriptive text.
FORMAT	CHAR(32)	Specifies the format of the content as STANDARD, which is the standard driver connection string.
OWNER_NAME	VARCHAR(256)	Specifies the name of the user that owns the DSN.

Name	Type	Description
OWNER_ID	VARCHAR(256)	Specifies the ID of the user that owns the DSN.
OWNER_TYPE	CHAR(1)	Specifies the owner type as U (User) or R (Role).
DSN_TYPE	VARCHAR(256)	Specifies the DSN type as one of the following: FEDERATED NOPROMPT CONNECT FILE

## DS\_PRIVILEGES View

The DS\_PRIVILEGES view displays the privilege descriptors for all data source-level privileges.

The table below lists the columns that are associated with the DS\_PRIVILEGES view:

Name	Type	Description
DATA_SERVICE_NAME	VARCHAR(256)	Specifies the unique name of the data service.
GRANTOR	VARCHAR(256)	Specifies the user name of the grantor.
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as U (User) or R (Role).
GRANTEE	VARCHAR(256)	Specifies the user name of the grantee.
GRANTEE_TYPE	CHAR(1)	Specifies the grantee type as U (User) or G (Group).

Name	Type	Description
PRIVILEGE_NAME	VARCHAR(20)	Indicates the name of privilege as reflected in the following list: SELECT UPDATE INSERT DELETE EXECUTE REFERENCES CREATE TABLE ALTER TABLE DROP TABLE CREATE VIEW ALTER VIEW DROP VIEW CREATE CACHE ALTER CACHE CREATE TABLESPACE  If the value of DATA_SERVICE_NAME is _SERVER_, which corresponds to the SAS® Federation Server, the ADMINISTER and TRACE privileges can also be displayed. The ADMINISTER and TRACE privileges can be set on the SAS® Federation Server only.
PRIVILEGE_TYPE	VARCHAR(5)	Specifies the privilege type can be GRANT or DENY.
GRANTABLE	CHAR(1)	Specifies if the privilege is grantable. The only valid value is N (No).

---

## DSN\_CONTENT View

The DSN\_CONTENT view contains one or more rows per configured DSN. Each row contains a portion of the DSN content, ordered by sequence. If DATA\_SOURCE\_NAMES.format is STANDARD, then the content column contains driver connection string syntax.

The table below lists the columns that are associated with the DSN\_CONTENT view:

Name	Type	Description
DSN_NAME	VARCHAR(256)	Specifies the unique name of the DSN.
SEQUENCE	INTEGER	Specifies the configuration chunk sequence.
CONTENT	VARCHAR(1024)	Specifies the DSN content.

---

## DSN\_LINEAGE View

The DSN\_LINEAGE view contains information for federated DSNs. There is one entry per referenced DSN, so a federated DSN containing a reference to two DSNs would have two entries in this view.

The table below lists the columns that are associated with the DSN\_LINEAGE view:

Name	Type	Description
DSN_NAME	VARCHAR(256)	Specifies the unique name of the DSN.
CHILD_DSN_NAME	VARCHAR(256)	Specifies the unique name of the child DSN.

---

## DSN\_PRIVILEGES and EFFECTIVE\_DSN\_PRIVILEGES Views

Displays privileges for users and groups on each data source name (DSN) and indicates inheritance. Both views show all direct (explicit) and inherited privileges based on the privileges of the user and group, or its group membership.

The DSN\_PRIVILEGES result set contains rows for users and groups that have the CONNECT privilege explicitly set on either the server, service, or DSN. If a user or group does not have any direct privilege, it will not be shown in this view. It is a condensed view of the EFFECTIVE\_DSN\_PRIVILEGES view.

The EFFECTIVE\_DSN\_PRIVILEGES result set contains rows for all users and groups that have any privilege directly set or a privilege can be derived from its group membership. For example, if a user does not have any privileges set on any of the SAS® Federation Server objects, the user will still be in the result set if the user is a member of a group that has a direct privilege set.

*Note:* Both of these views can return very large result sets depending on the configuration of SAS® Federation Server. Subsetting on DATA\_SERVICE, CATALOG\_NAME, and/or SCHEMA\_NAME can reduce the size of the result set.

Name	Type	Description
DSN_NAME	VARCHAR(256)	Specifies the unique name of the DSN.
DATA_SERVICE	VARCHAR(256)	Specifies the name of the data service.
GRANTOR_ID	VARCHAR(256)	Specifies the AuthID of the user that granted or denied the privilege.
GRANTOR	VARCHAR(256)	Specifies the name of the user who is granted or denied the privilege.

Name	Type	Description
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as U (User) or R (Role).
GRANTEE_ID	VARCHAR(256)	Specifies the AuthID of the user that is granted or denied the privilege.
GRANTEE	VARCHAR(256)	Specifies the name of the user who is granted or denied the privilege.
GRANTEE_TYPE	CHAR(1)	Specifies the grantee type as U (User) or G (Group).
PRIVILEGE_NAME	VARCHAR(20)	Specifies the privilege name: SELECT UPDATE INSERT DELETE EXECUTE REFERENCES CREATE TABLE ALTER TABLE DROP TABLE CREATE VIEW ALTER VIEW DROP VIEW
PRIVILEGE_TYPE	VARCHAR(5)	Specifies the privilege type as GRANT or DENY.
GRANTABLE	CHAR(1)	Specifies if the user can grant this privilege to other users. The only valid value is N (No).
INHERITED	CHAR(1)	Indicates if the privilege is inherited as either Y or N.
SOURCE_OBJECT_LEVEL	INTEGER	Specifies the object level where the privilege is inherited, as one of the following values: 0 — Server 1 — Data service 2 — DSN
SOURCE_GRANTEE_ID	VARCHAR(256)	Specifies the AuthID of a group or user from which the privilege is derived.
SOURCE_GRANTEE	VARCHAR(256)	Specifies the name of the group or user from which the privilege is derived.
SOURCE_GRANTEE_TYPE	CHAR(1)	Specifies the source_grantee type U (User) or G (Group).

---

## IDENTITY View

The IDENTITY view returns identity information for a user connected to SAS Federation Server. The table below lists the columns that are associated with this view:

Name	Type	Description
USER_NAME	VARCHAR(256)	Specifies the name of the user that was created with the Authentication Server.
AUTH_DOMAIN	VARCHAR(256)	Specifies the domain name of the authenticated user. For example, if you connect to SAS Federation Server with the <b>local\myuser</b> account, the auth_domain is local.
AUTH_ID	VARCHAR(256)	Specifies the user name for the authenticated user. For example, if you connect to SAS Federation Server with the <b>local\myuser</b> account, the auth_id is myuser.

---

## MESSAGES View

The MESSAGES view contains one or more messages associated with an operation. Each MESSAGE\_ID represents one or more messages. Each message in a MESSAGE\_ID has a unique number MESSAGE\_NUM, and is separated into pieces that will fit within DATA. The SEQUENCE column can be used to order the entries within a MESSAGE\_ID, and MESSAGE\_NUM identifies individual messages.

*Note:* Rows are visible only if the rows referenced by the message are visible to the user.

Name	Type	Description
DATA	WVARCHAR(128)	Specifies a piece of the message text. Not NULL.
MESSAGE_ID	INTEGER	The unique ID of the messages stored in the table. Not NULL.
SEQUENCE	INTEGER	Orders the message entries for the entire MESSAGE_ID. Not NULL.
MESSAGE_NUM	INTEGER	Used to indicate separate messages within a MESSAGE_ID. Not NULL.

---

## OBJECTS View

The OBJECTS view displays the name and system identifier for each schema object. This view is visible to SAS Federation Server System Users and Administrators. The table below lists the columns that are associated with this view:

Name	Type	Description
CATALOG_NAME	VARCHAR(256)	Specifies the catalog name.
SCHEMA_NAME	VARCHAR(256)	Specifies the schema name.
OBJECT_NAME	VARCHAR(256)	Specifies the object name.
OBJECT_TYPE	INTEGER	Object type: 1 – Table 2 – SAS View 3 – Cache 4 – SAS Package 5 – Relation
FLAGS	INTEGER	Flags  1 - Definer's rights. For FedSQL views, the definer's rights bit indicates the view executes under the privileges of the view's schema owner rather than the invoker.  2 - Object was implicitly added as a result of a GRANT (as opposed to explicitly added through a user DDL). The object can be automatically removed if the grant is revoked.

---

## OBJECT\_PRIVILEGES View

The OBJECT\_PRIVILEGES view displays the privilege descriptors for all object-level privileges. All privileges for a single table can be obtained by concatenating rows matching the correct priv\_id and ordering the results by sequence. The table below lists the columns that are associated with the OBJECT\_PRIVILEGES view:

Name	Type	Description
CATALOG_NAME	VARCHAR(256)	Specifies the name of the catalog.
SCHEMA_NAME	VARCHAR(256)	Specifies the name of the schema.
OBJECT_NAME	VARCHAR(256)	Specifies the name of the object.

---

Name	Type	Description
GRANTOR	VARCHAR(256)	Specifies the user name of the grantor.
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as U (User) or R (Role).
GRANTEE	VARCHAR(256)	Specifies the name of the user who is granted privileges.
GRANTEE_TYPE	CHAR(1)	Specifies the grantee type as U (User) or G (Group).
PRIVILEGE_NAME	VARCHAR(20)	Specifies the privilege name as one of the following values: SELECT UPDATE INSERT DELETE EXECUTE REFERENCES CREATE TABLE ALTER TABLE DROP TABLE CREATE VIEW ALTER VIEW DROP VIEW CREATE CACHE ALTER CACHE
PRIVILEGE_TYPE	VARCHAR(5)	Specifies the privilege type as GRANT or DENY.
GRANTABLE	CHAR(1)	Specifies if the privilege is grantable. The only valid value is N (No).
PREDICATE	VARCHAR(128)	Portion of the row-level security (RLS) predicate. The predicate might spawn multiple rows; is nullable.
PREDICATE_SEQUENCE	SMALLINT	specifies the sequence number of the portion of the RLS predicate; is nullable.
PRIV_ID	BIGINT	Privilege identifier.

## PRIVILEGES and EFFECTIVE\_PRIVILEGES Views

Displays the privileges, including inheritance, for users and groups on schemas, catalogs and data services. Both views show all direct (explicit) and inherited privileges based on the privileges of the user and group, or its group membership.

The PRIVILEGES result set contains rows for users and groups that have any privilege directly set. If a user or group does not have any direct privilege, it will not be shown in this view. It is a condensed view of the EFFECTIVE\_PRIVILEGES view.

The EFFECTIVE\_PRIVILEGES result set contains rows for all users and groups that have any privilege directly set or a privilege can be derived from its group membership. For example, if a user does not have any privileges set on any of the Federation Server objects, the user will still be in the result set if the user is a member of a group that has a direct privilege set.

By default, if a privilege is not explicitly listed in the result sets, it is denied.

*Note:* Both of these views can return very large result sets depending on the configuration of SAS Federation Server. Subsetting on DATA\_SERVICE, CATALOG\_NAME, and/or SCHEMA\_NAME can reduce the size of the result set.

The table below lists the columns that are associated with the PRIVILEGES and EFFECTIVE\_PRIVILEGES view:

Name	Type	Description
DATA_SERVICE	VARCHAR(256)	Specifies the name of the data service.
CATALOG_NAME	VARCHAR(256)	Specifies the name of the catalog.
SCHEMA_NAME	VARCHAR(256)	Specifies the name of the schema.
GRANTOR_ID	VARCHAR(256)	Specifies the AuthID of the user that granted or denied the privilege.
GRANTOR	VARCHAR(256)	Specifies the name of the user who is granted or denied the privilege.
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as U (User) or R (Role).
GRANTEE_ID	VARCHAR(256)	Specifies the AuthID of the user that is granted or denied the privilege.
GRANTEE	VARCHAR(256)	Specifies the name of the user who is granted or denied the privilege.
PRIVILEGE_NAME	VARCHAR(20)	Specifies the privilege name as one of the following values: SELECT UPDATE INSERT REFERENCES
PRIVILEGE_TYPE	VARCHAR(5)	Specifies the privilege type as GRANT or DENY.
GRANTABLE	CHAR(1)	Specifies if the privilege can be granted. The only valid value is or N (No).
INHERITED	CHAR(1)	Indicates if the privilege is inherited as either Y or N.

Name	Type	Description
SOURCE_OBJECT_LEVEL	INTEGER	Specifies the object level where the privilege is inherited, as one of the following values: 0 — server 1 — data service 2 — catalog 3 — schema
SOURCE_GRANTEE_ID	VARCHAR(256)	Specifies the AuthID of a group or user from which the privilege is derived.
SOURCE_GRANTEE	VARCHAR(256)	Specifies the name of the group or user from which the privilege is derived.
SOURCE_GRANTEE_TYPE	CHAR(1)	Specifies the source_grantee type as U (User) or G (Group).

## SCHEMAS View

The SCHEMAS view contains the name and system identifier for each schema. The table below lists the columns that are associated with the SCHEMAS view:

Name	Type	Description
SCHEMA_NAME	VARCHAR(256)	Specifies the schema name.
CATALOG_NAME	VARCHAR(256)	Specifies the name of the catalog. If not applicable, the value is NULL.
USER_NAME	VARCHAR(256)	Specifies the user name of the schema owner.

## SCHEMA\_PRIVILEGES View

The SCHEMA\_PRIVILEGES view displays the privilege descriptors for all schema-level privileges. The table below lists the columns that are associated with the SCHEMA\_PRIVILEGES view:

Name	Type	Description
CATALOG_NAME	VARCHAR(256)	Specifies the name of the catalog. If not applicable, the value is NULL.
SCHEMA_NAME	VARCHAR(256)	Specifies the name of the schema.

Name	Type	Description
GRANTOR	VARCHAR(256)	Specifies the user name of the grantor.
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as U (User) or R (Role).
GRANTEE	VARCHAR(256)	Specifies the name of the user who is granted privileges.
GRANTEE_TYPE	CHAR(1)	Specifies the grantee type as U (User) or G (Group).
PRIVILEGE_NAME	VARCHAR(20)	Specifies the privilege name as one of the following values: SELECT UPDATE INSERT DELETE EXECUTE REFERENCES CREATE TABLESPACE CREATE TABLE ALTER TABLE DROP TABLE CREATE VIEW ALTER VIEW DROP VIEW CREATE CACHE ALTER CACHE CREATE TABLESPACE
PRIVILEGE_TYPE	VARCHAR(5)	Specifies the privilege type as GRANT or DENY.
GRANTABLE	CHAR(1)	Specifies if the privilege is grantable. The only valid value is N (No).

---

## X\_COLUMN\_PRIVILEGES/ X\_EFFECTIVE\_COLUMN\_PRIVILEGES View

The X\_COLUMN\_PRIVILEGES and the X\_EFFECTIVE\_COLUMN\_PRIVILEGES views contain both the privileges for users and groups on all objects<sup>2</sup>, and indicate inheritance. They show all direct (explicit) and inherited privileges based on the privileges of the user or group or its group membership. Unlike most other views, the views do not strictly derive the information from system tables. It will merge metadata from the physical data sources with metadata in system tables to produce a complete result set for all objects.

The X\_COLUMN\_PRIVILEGES result set contains rows for users and groups that have any privilege directly set. If a user or group does not have any direct privilege, it will not

be shown in this view. It is a condensed view of the X\_EFFECTIVE\_COLUMN\_PRIVILEGES view.

The X\_EFFECTIVE\_COLUMN\_PRIVILEGES result set contains rows for all users and groups that have any privilege directly set or a privilege can be derived from its group membership. For example, even if a user does not have any privileges directly set, records for this user will be in the result set if any of the groups in its group hierarchy has a privilege directly set.

*Note:* Both of these views can return very large result sets depending on the configuration of Federation Server. Subsetting on DATA\_SERVICE, CATALOG\_NAME, and/or SCHEMA\_NAME can reduce the size of the result set.

Name	Type	Description
DATA_SERVICE	VARCHAR(256)	Specifies the data service name.
CATALOG_NAME	VARCHAR(256)	Specifies the catalog name.
COLUMN_NAME	VARCHAR(256)	Specifies the column name.
SCHEMA_NAME	VARCHAR(256)	Specifies the schema name.
GRANTOR_ID	VARCHAR(256)	Specifies the AuthID of the grantor.
GRANTOR	VARCHAR(256)	Specifies the name of the grantor. This field could be NULL if the user no longer exists.
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as R (Role) or U (User).
GRANTEE_ID	VARCHAR(256)	Specifies the AuthID of the grantee.
GRANTEE	VARCHAR(256)	Specifies the name of the grantee.
GRANTEE_TYPE	CHAR(1)	Specifies the grantee type as U (User) or G (Group).
PRIVILEGE_NAME	VARCHAR(256)	Name of privilege as reflected in the following list: SELECT UPDATE INSERT DELETE EXECUTE REFERENCES CREATE TABLE ALTER TABLE DROP TABLE CREATE VIEW ALTER VIEW DROP VIEW
PRIVILEGE_TYPE	VARCHAR(256)	Specifies the privilege type as GRANT or DENY.

Name	Type	Description
GRANTABLE	CHAR(1)	Specifies if the privilege is grantable. The only valid value is N (No).
INHERITED	CHAR(1)	Specifies if the privilege is inherited with Y (Yes) or N (No).
SOURCE_OBJECT_LEVEL	INTEGER	Specifies the object level where the privilege is inherited from: 0 – server 1 – data service 2 – catalog 3 – schema 4 – object 5 – column
SOURCE_GRANTEE_ID	VARCHAR(256)	AuthID of group or user the privilege is derived from.
SOURCE_GRANTEE	VARCHAR(256)	Specifies the group or user name the privilege is derived from.
SOURCE_GRANTEE_TYPE	CHAR(1)	Specifies the grantee as U (User) or G (Group).
OBJECT_NAME	VARCHAR(256)	Specifies the name of the object.

<sup>2</sup>Current list of objects includes:

- table server
- data services
- catalogs
- schemas

---

## X\_OBJECT\_PRIVILEGES/ X\_EFFECTIVE\_OBJECT\_PRIVILEGES View

The X\_OBJECT\_PRIVILEGES and the X\_EFFECTIVE\_OBJECT\_PRIVILEGES views contain both the privileges for users and groups on all objects<sup>2</sup> and indicates inheritance. They show all direct (explicit) and inherited privileges based on the privileges of the user and group or its group membership. Unlike most other views, the views do not strictly derive the information from system tables. It will merge metadata from the physical data sources with metadata in system tables to produce a complete result set for all objects.

The X\_OBJECT\_PRIVILEGES result set contains rows for users and groups that have any privilege directly set. If a user or group does not have any direct privilege, it will not be shown in this view. It is a condensed view of the X\_EFFECTIVE\_OBJECT\_PRIVILEGES view.

The X\_EFFECTIVE\_OBJECT\_PRIVILEGES result set contains rows for all users and groups that have any privilege directly set or a privilege can be derived from its group membership. For instance, even if a user does not have any privileges directly set, records for this user will be in the result set if any of the groups in its group hierarchy has a privilege directly set.

If a privilege is not explicitly listed in the result sets, it is DENIED by default.

Name	Type	Description
DATA_SERVICE	VARCHAR(256)	Specifies the data service name.
CATALOG_NAME	VARCHAR(256)	Specifies the catalog name.
SCHEMA_NAME	VARCHAR(256)	Specifies the schema name.
GRANTOR_ID	VARCHAR(256)	Specifies the AuthID of the grantor.
GRANTOR	VARCHAR(256)	Specifies the name of the grantor. This field could be NULL if the user no longer exists.
GRANTOR_TYPE	CHAR(1)	Specifies the grantor type as U (User) or R (Role).
GRANTEE_ID	VARCHAR(256)	Specifies the AuthID of the grantee.
GRANTEE	VARCHAR(256)	Specifies the grantee name.
GRANTEE_TYPE	CHAR(1)	Specifies the grantee type as U (User) or G (Group).
PRIVILEGE_NAME	VARCHAR(256)	Name of privilege as reflected in the following list: SELECT UPDATE INSERT DELETE EXECUTE REFERENCES CREATE TABLE ALTER TABLE DROP TABLE CREATE VIEW ALTER VIEW DROP VIEW CREATE CACHE ALTER CACHE
PRIVILEGE_TYPE	VARCHAR(256)	Specifies the privilege type as GRANT or DENY.
GRANTABLE	CHAR(1)	Specifies if the privilege is grantable. The only valid value is N (No).

Name	Type	Description
INHERITED	CHAR(1)	Specifies if the privilege is inherited with Y or N.
SOURCE_OBJECT_LEVEL	INTEGER	Specifies the object level where the privilege is inherited from: 0 – server 1 – data service 2 – catalog 3 – schema 4 – object 5 – column
SOURCE_GRANTEE_ID	VARCHAR(256)	AuthID of group or user the privilege is derived from.
SOURCE_GRANTEE	VARCHAR(256)	Specifies the group or user name the privilege is derived from.
SOURCE_GRANTEE_TYPE	CHAR(1)	Specifies the grantee as U - User or G - Group.
OBJECT_NAME	VARCHAR(256)	Specifies the name of the object.

<sup>2</sup>Current list of objects includes:

- server
- data services
- catalogs
- schemas

# Glossary

---

**ACID**

See atomicity, consistency, isolation, durability

**American National Standards Institute**

the organization that coordinates the development of voluntary consensus standards for products, services, processes, systems, and personnel in the United States. ANSI works with the International Organization for Standardization to establish global standards. Short form: ANSI.

**ANSI**

See American National Standards Institute

**API**

See application programming interface

**application programming interface**

a set of software functions that facilitate communication between applications and other types of programs or services. Short form: API.

**Application Response Measurement**

the name of an application programming interface that was developed by an industry partnership and which is used to monitor the availability and performance of software applications. ARM monitors the application tasks that are important to a particular business. Short form: ARM.

**ARM**

See Application Response Measurement

**atomicity, consistency, isolation, durability**

the characteristics of a transaction, such as a group of SQL statements, in an RDBMS that support commit and rollback operations. The characteristics include atomicity (the execution of a transaction, either committed or rolled back); consistency (the successful application of the consistency rules of an RDBMS that commits only valid data to the database); isolation (the separation of a transaction from all other concurrent processes in an RDBMS); and durability (the persistence or repeatability of a transaction under all conditions, including system failure, after which a transaction can be re-created from a log that contains committed transactions). Short form: ACID

**authentication**

See client authentication

**authorization**

the process of determining the permissions that particular users have for particular resources. Authorization either permits or denies a specific action on a specific resource, based on the user's identity and on group memberships.

**client authentication**

the process of verifying the identity of a person or process for security purposes.

**connection string**

information that defines how to connect an application to the data. In SAS Federation Server, a connection string identifies the query language syntax that the application submits, as well as the information that is required to connect to a data source or data sources.

**data source name**

a persistent identifier that is associated with a data source definition. The data source definition specifies how to locate and access a data source, including any authentication (such as a user name and password) that a user must provide. Short form: DSN.

**data type**

an attribute of every column in a table or database, indicating the type of data in the column and how much physical storage it occupies.

**definer's rights view**

a view that is created by a schema owner. Definer's rights views are required for data caching in SAS Federation Server.

**driver**

a special-purpose software program that enables two disparate software programs, such as an application and an API, to interact.

**DSN**

See data source name

**encryption**

the act or process of converting data to a form that is unintelligible except to the intended recipients.

**federated DSN**

a data source name that references multiple data sources. The data sources can be on the same DBMS, or on a different one.

**grouping data source name**

See federated DSN

**grouping DSN**

See federated DSN

**Integrated Object Model**

the set of distributed object interfaces that make SAS software features available to client applications when SAS is executed as an object server. Short form: IOM.

**invoker's rights view**

a federated view or cache that is accessed using the current user's authorization instead of the schema owner's authorization. See also "definer's rights view."

**IOM**

See Integrated Object Model

**Java Virtual Machine**

a software application that can execute Java bytecode, on either a client or a server, enabling Java programs to be run on many different hardware and software platforms. Short form: JVM.

**join**

an operation that combines data from two or more tables. A join is typically created by means of SQL (Structured Query Language) code or a user interface.

**JVM**

See Java Virtual Machine

**MDS**

See Memory Data Store

**Memory Data Store**

a transactional data cache that runs strictly in-memory. Because there is no back up data storage, changes are lost when the in-memory database is closed.

**result set**

the set of rows or records that a server or other application returns in response to a query.

**RLS**

See row-level security

**RLS predicate**

See row-level security predicate

**row-level security**

a security feature that controls access to rows and columns in a table in order to prevent users from accessing restricted data.

**row-level security predicate**

a query that restricts the rows that are available to grantees for specified operations. Only rows that match the predicate can be accessed by the grantees.

**scrollable cursor**

a device that enables an application to set a position on any row in a result set. For example, a scrollable cursor can back up and revisit a row, start at the end of the file and work backward, skip some rows, or go directly to a specific row.

**serializability**

a capability commonly required in database processing that ensures the highest level of isolation between transactions for the purposes of concurrency control.

**SQL**

See Structured Query Language

**Structured Query Language**

a standardized, high-level query language that is used in relational database management systems to create and manipulate objects in a database management system. SAS implements SQL through the SQL procedure. Short form: SQL.

**thread**

the smallest unit of processing that can be scheduled by an operating system.

**threaded processing**

processing that is performed in multiple threads in order to improve the speed of CPU-bound applications.

**time-out**

an error condition that is produced when a required response from a device or program is not received after a specified length of time.

**transactional data store**

a storage mechanism for transactional data that is characterized by ACID features (atomicity, consistency, isolation and durability).

**type**

See data type

**Unicode**

a 16-bit encoding that is the industry standard for supporting the interchange, processing, and display of characters and symbols from most of the world's writing systems.