



Spotfire® Data Virtualization

Oracle Adapter Guide

Version 8.9.0 | April 2026

Contents

Contents	2
SDV Oracle Adapter	4
Introduction	4
Obtain and Install the Drivers for Oracle Applications	4
Obtain and Install the Driver for Oracle OCI Client	5
Datasource Configuration	11
Basic Tab	11
Advanced Tab	17
Oracle Data Source Limitations	36
Oracle Caching Limitations	36
Oracle to SDV Data Types	37
Oracle NUMBER Data Types and SDV Data Types	37
Oracle to Data Types Common to All Versions	38
Oracle 11g to SDV Data Types	40
Oracle 19c to SDV Data Types	41
Oracle Cache Mapping	42
Oracle Function Support	44
Oracle Aggregate Function Support	45
Oracle Analytic Function Support	46
Oracle Analytic Aggregate Function Support	48
Oracle Binary Function Support	49
Oracle Character Function Support	49
Oracle Conditional Function Support	51
Oracle Conversion Function Support	51
Oracle Date Function Support	52
Oracle Geospatial Function Support	53
Oracle Encryption Function Support	57

Oracle Numeric Function Support	57
Oracle Specific Connection Properties	59
Oracle Time Function Support	62
Oracle XML Function Support	62
References	63
Spotfire Documentation and Support Services	65
Legal and Third-Party Notices	67

SDV Oracle Adapter

Introduction

This section explains the various connection and configuration options of the datasource Oracle as well as the capabilities:

[Obtain and Install the Drivers for Oracle Applications](#)

[Datasource Configuration](#)

[Obtain and Install the Driver for Oracle OCI Client](#)

[Oracle Data Source Limitations](#)

[Oracle to SDV Data Types](#)

[Oracle Function Support](#)

[References](#)

Obtain and Install the Drivers for Oracle Applications

SDV Supports the Oracle versions 11g, 12c and 19c. Refer to the Installation guide for more details about the supported versions.

Obtain and install the JDBC driver for Oracle 11g and 12c type 4 (11.2.0.2 or higher)

1. Navigate to the Oracle Database JDBC Drivers web page. For example:

<https://www.oracle.com/database/technologies/jdbc-ucp-122-downloads.html>

2. Locate and download the ojdbc8.jar and xdb6.jar. Oracle 11.2.0.2 or higher version is required.
3. Copy them to the SDV installation directory:

```
<TDV_install_dir>/apps/dlm/cis_ds_oracle/lib
```

4. You will also need the file `xmlparserv2.jar` for handling XML datatypes with Oracle databases. You can get the file for your version of Oracle in the following location:

<https://mvnrepository.com/artifact/com.oracle.database.xml/xmlparserv2>

5. Copy it to the SDV installation directory:

```
<TDV_install_dir>/apps/dlm/cis_ds_oracle/lib
```

6. Stop and restart SDV Server.

Obtain and install the JDBC driver for Oracle 19c type 4

7. Navigate to the Oracle Database 19c Version 19.3.0.0 JDBC Drivers web page.

<https://www.oracle.com/database/technologies/appdev/jdbc-ucp-19c-downloads.html>

8. Locate the `ojdbc10.jar` and `xdb.jar` files.
9. Copy it to the SDV installation directory:

```
<TDV_install_dir>/conf/adapters/system/oracle_19c_thin_driver.
```

10. Stop and restart SDV Server.

Obtain and Install the Driver for Oracle OCI Client

Use the instructions in this section to set up the connections between an Oracle data source and SDV, and between SDV and a client interface.

For Oracle versions 11g and 12c

These instructions are for the 11g and 12c versions of Oracle. SDV uses the Oracle 11g R2 Type 2 client to connect to Oracle Type 2 data sources. If possible, use the Oracle 11g driver to connect to all Oracle versions.

Note: If you expect to use XML types in your SDV environment, download and install the Oracle OCI Type 2 driver. The Oracle Thin driver does not support schema-based XML documents.

Instructions for installing and configuring Oracle drivers are provided for both Windows and UNIX.

To configure an Oracle OCI type 2 JDBC adapter for SDV (on Windows)

1. Install an Oracle client instance on the machine where SDV server is running.

You can download the appropriate Oracle client for your platform from this URL:

<http://www.oracle.com/technetwork/database/features/instant-client/index-097480.html>

The Oracle client instance must have Oracle net services components for the OCI adapter. Oracle administrator and runtime clients contain net services components by default. Use a 32-bit or 64-bit Oracle client, as appropriate for the SDV installation. See Oracle documentation about how to set up an Oracle client.

The Oracle client instance contains a Java driver (ojdbc8.jar) and a native library (ocijdbc11.dll or ocijdbc12.dll).

2. Set up TNS names.

TNS names are created in Oracle Net Configuration Assistant. These TNS names contain information about how to connect to a physical Oracle database. See Oracle documentation for details on how to set up TNS names.

3. Copy the Java driver ojdbc8.jar file to:

```
<TDV_install_dir>\apps\d1m\cis_ds_oracle_type2\lib
```

4. Copy the native library ocijdbc11.dll file to:

```
<TDV_install_dir>\jdk\bin
```

5. Set up environment variables on the machine that hosts SDV Server.

- Set the ORACLE_HOME environment variable to something like: <ORACLE_INSTALLATION>\product\11.2.0\client_1.
- Add ORACLE_HOME\bin to system PATH.
- Set the CLASSPATH environment variable to use Oracle JDBC adapter from ORACLE_HOME\jdbc\lib.

6. If you expect to use XML types in your SDV environment:

Obtain xdb6.jar from the Oracle website. For example:

<https://www.oracle.com/database/technologies/jdbc-ucp-122-downloads.html>

Copy xdb6.jar to <TDV_install_dir>\conf\adapters\system\oracle_<version>_oci_driver

Obtain xmlparserv2.jar from the following location:

<https://mvnrepository.com/artifact/com.oracle.database.xml/xmlparserv2>

Download the Oracle Database 11g Release 2 Client (11.2.0.1.0) for Microsoft Windows (x64), choose "Runtime" during installation. The final installation has the jar file under <INSTALL_DIR>/product/11.2.0/client_1/lib.

Copy xmlparserv2.jar to <TDV_install_dir>\conf\adapters\system\oracle_<version>_oci_driver

7. Stop and restart SDV Server.

You can now add Oracle data sources that use the OCI adapter.

To configure an Oracle OCI Type 2 JDBC adapter for SDV (on UNIX)

8. Export the path for OCI libraries.

9. Install an Oracle client instance on the machine where SDV Server is running. You can download the appropriate Oracle client for your platform from this URL:

<http://www.oracle.com/technetwork/database/features/instant-client/index-097480.html>

The Oracle client instance must have Oracle net services components for the OCI adapter. Oracle administrator and runtime clients contain net services components by default. Use a 32-bit or 64-bit Oracle client, as appropriate for the SDV installation. See Oracle documentation about how to set up an Oracle client.

The Oracle client instance contains a Java driver (ojdbc8.jar) and the native library files for the specific platforms:

10. Set up TNS names.

TNS names are created in Oracle Net Configuration Assistant. These TNS names contain information about how to connect to a physical Oracle database. See Oracle documentation for details on how to set up TNS names.

11. Copy the Java driver ojdbc8.jar file to:

```
<TDV_install_dir>/apps/dlm/cis_ds_oracle_type2/lib
```

12. Set up environment variables on the machine that hosts SDV Server.

- Set the ORACLE_HOME environment variable to something like <ORACLE_INSTALLATION>/product/11.2.0/client_1. It is recommended that you set your environment variables in a system init script rather than use the EXPORT command.
- Add ORACLE_HOME/lib to the library path environment variable for your platform:
- Linux: LD_LIBRARY_PATH
- AIX: LIBPATH
- Set the CLASSPATH environment variable to use Oracle JDBC adapter from ORACLE_HOME/jdbc/lib.

13.If you expect to use XML types in your SDV environment:

Obtain xdb6.jar from the Oracle website. For example:

<https://www.oracle.com/database/technologies/jdbc-ucp-122-downloads.html>

Copy xdb6.jar to <TDV_install_dir>/conf/adapters/system/oracle_<version>_oci_driver

Obtain xmlparserv2.jar from the Oracle website. For example:

<https://mvnrepository.com/artifact/com.oracle.database.xml/xmlparserv2>

Download the Oracle Database 11g Release 2 Client (11.2.0.1.0) for Microsoft Windows (x64), choose “Runtime” during installation. The final installation has the JAR file under <INSTALL_DIR>/product/11.2.0/client_1/lib.

Copy xmlparserv2.jar to <TDV_install_dir>/conf/adapters/system/oracle_<version>_oci_driver

14.Stop and restart the SDV Server.

You can now add Oracle data sources that use the OCI adapter.

For Oracle version 19c

These instructions are for the 19c versions of Oracle.

Instructions for installing and configuring Oracle drivers are provided for both Windows and UNIX.

To configure an Oracle OCI type 2 JDBC adapter for SDV (on Windows)

15.Download the following files for configuring the Oracle 19c OCI environments:

- https://download.oracle.com/otn_software/nt/instantclient/19500/instantclient-basic-windows.x64-19.5.0.0.0dbru.zip
- https://download.oracle.com/otn_software/nt/instantclient/19500/instantclient-sqlplus-windows.x64-19.5.0.0.0dbru.zip
- https://download.oracle.com/otn_software/nt/instantclient/19500/instantclient-tools-windows.x64-19.5.0.0.0dbru.zip
- https://download.oracle.com/otn_software/nt/instantclient/19500/instantclient-sdk-windows.x64-19.5.0.0.0dbru.zip

The Oracle client instance must have Oracle net services components for the OCI adapter. Oracle administrator and runtime clients contain net services components by default. Use a 32-bit or 64-bit Oracle client, as appropriate for the SDV installation. See Oracle documentation about how to set up an Oracle client.

The Oracle client instance contains a Java driver (ojdbc10.jar and xdb.jar) and the following native libraries:

- oci.dll
- ocijdbc19.dll
- ociw32.dll
- oramysql19.dll
- orannzsbb19.dll
- oraocci19.dll
- oraocci19d.dll
- oraociei19.dll
- oraons.dll
- orasql19.dll
- orasqlplusic19.dll

Copy the Java driver files (ojdbc10.jar and xdb.jar) to:

```
<TDV_install_dir>\apps\d1m\cis_ds_oracle_type2\lib
```

Copy the native library files to:

```
<TDV_install_dir>\jdk\bin
```

16. Set up TNS names.

TNS names are created in Oracle Net Configuration Assistant. These TNS names contain information about how to connect to a physical Oracle database. See Oracle documentation for details on how to set up TNS names.

17. Set up environment variables on the machine that hosts SDV Server.

Set the ORACLE_HOME environment variable to:

```
<ORACLE_INSTALLATION>\instantclient_19_5
```

Add ORACLE_HOME to system PATH.

Set the CLASSPATH environment variable to use Oracle JDBC adapter from ORACLE_HOME

18. Stop and restart SDV Server.

You can now add Oracle data sources that use the OCI adapter.

To configure an Oracle OCI type 2 JDBC adapter for SDV (on Unix)

19. Download the following files for configuring the Oracle 19c OCI environments:

- <https://www.oracle.com/database/technologies/oracle19c-linux-downloads.html>

The Oracle client instance must have Oracle net services components for the OCI adapter. Oracle administrator and runtime clients contain net services components by default. Use a 32-bit or 64-bit Oracle client, as appropriate for the SDV installation. See Oracle documentation about how to set up an Oracle client.

The Oracle client instance contains Java drivers (ojdbc10.jar and xdb.jar) and the following native libraries:

- libipc1.so
- libmq11.so
- libnnz19.so
- libocci.so
- libociei.so
- libocijdbc19.so
- liboramysql19.so
- libsqlplusic.so
- libsqlplus.so

Copy the Java driver files (ojdbc10.jar and xdb.jar) to:

```
<TDV_install_dir>/apps/dlm/cis_ds_oracle_type2/lib
```

Copy the native library files to:

```
<TDV_install_dir>/apps/server/lib/svn/lin64
```

20. Set up TNS names.

TNS names are created in Oracle Net Configuration Assistant. These TNS names contain information about how to connect to a physical Oracle database. See Oracle documentation for details on how to set up TNS names.

21. Set the following environment variables in the .bash_profile and make it available:

```
export ora_home=<ORACLE_INSTALLATION>
```

```
export PATH=$PATH:$ora_home/instantclient_19_5
```

```
export ORACLE_BASE=$ora_home
```

```
export ORACLE_HOME=$ORACLE_BASE/instantclient_19_5
```

```
export LD_LIBRARY_PATH=$ORACLE_HOME
```

```
export NLS_LANG=<LANGUAGE>
```

```
export TNS_ADMIN=$ORACLE_HOME/network/admin
```

```
export CLASSPATH=$ORACLE_HOME
```

22. Stop and restart SDV Server.

You can now add Oracle data sources that use the OCI adapter.

Datasource Configuration

This section explains the connection properties that are defined while setting up a datasource.

Basic Tab

The following table and the sections below lists and explains the connection properties that are in the Basic Tab of the New Data Source Window.

Datasource Name	The name of the Datasource.
Host	Name of the host machine or the host machine's IP address.

Port	The Port number.
Service Name	This property is available for an Oracle data source with Thin driver. Name of the database service.
Net Service Name	This property is available for an Oracle data source with OCI driver. The TNS name that is set up through Oracle Net Configuration Assistant.
Login/User, Password	User name and password required to access the data source.
Database Name	Name or alias of the underlying data source. SDV Server uses this name to find and connect to the data source.
Pass-through Login	Flag to indicate whether pass-through login is enabled or not.
Transaction Isolation	The degree to which transactions are isolated from data modifications made by other transactions.
Authentication Type	The type of Authentication used by the datasource.
Ticket Cache	Specify the location of where your ticket cache is stored in this field when using Kerberos authentication. Leave the Login and Password fields blank when using ticket cache Note: Credential cache must be generated with a forwardable option set. To do this, KRB5.INI or KRB5.CONF in the client machine needs to have the parameter libdefaults: forwardable set to true.

Datasource Name

The name of the data source.

Data Type

string

Default Value

""

Host

Name of the host machine or the host machine's IP address.

Data Type

string

Default Value

""

Port

The Port number

Data Type

string

Default Value

""

Service Name

This property is available for an Oracle data source with Thin driver.

Name of the database service.

Data Type

String

Default Value

""

Remarks

This field is available in Oracle(Thin) driver.

Net Service Name

This property is available for an Oracle data source with OCI driver.

The TNS name that is set up through Oracle Net Configuration Assistant.

Data Type

String

Default Value

""

Remarks

This field is available in Oracle(OCI) driver.

Database Name

Name or alias of the underlying data source. SDV Server uses this name to find and connect to the data source.

Data Type

string

Default Value

""

Login/User, Password

User name and password required to access the data source.

Data Type

string

Default Value

""

Remarks

When the data source is used as a target for cache tables or data ship, the user must also have permission to create tables, execute DDL, and perform other required tasks. Refer to the individual data source descriptions for details.

Transaction Isolation

The degree to which transactions are isolated from data modifications made by other transactions.

Data Type

string

Default Value

NONE

Remarks

Valid values are:

- Read Committed (default) –Nonrepeatable reads and phantom reads can occur.
- Serializable–Dirty reads, nonrepeatable reads, and phantom reads are prevented.
- None

Pass-through Login

Flag to indicate whether pass-through login is enabled or not.

Data Type

string

Default Value

""

Remarks

Disabled (default)—This allows automated provisioning of a connection pool. Open connection threads can be used by authorized users after the validation query verifies connection status. If pass-through login is disabled, the Save Password check box is not available.

Enabled—A new connection to the data source uses the credentials supplied by the client when data is requested from that data source for the first time. Subsequent requests by the same user reuse the existing connection. When another user attempts to connect to a data source, a new connection is created.

See “Managing Security for SDV Resources” in the SDV Administration Guide for details.

Authentication Type

Indicates the type of authentication used by the data source.

Data Type

String

Default Value

BASIC

Remarks

Select BASIC or Kerberos authentication method, where offered.

See the *SDV Administration Guide* for more information about Kerberos authentication.

Ticket Cache

Specify the location of where your ticket cache is stored in this field when using Kerberos authentication. Leave the Login and Password fields blank when using ticket cache

Note: Credential cache must be generated with a forwardable option set. To do this, KRB5.INI or KRB5.CONF in the client machine needs to have the parameter libdefaults: forwardable set to true.

Data Type

String

Default Value

""

Remarks

None

Advanced Tab

The following table and the sections below lists and explains the connection properties that are in the Advanced Tab of the New Data Source Window.

Connection URL Pattern	A template for generating a URL to connect to the physical data source.
Connection URL String	The URL string generated from the connection URL pattern with the connection information you provide.

JDBC Connection Properties	Lets you specify property-value pairs to pass to the JDBC data source
Connection Pool Maximum Size	Maximum number of connections (both active and idle) allowed for the data source. When the maximum is reached, new requests must wait until a connection is available.
Connection Pool Minimum Size	Minimum number of connections in the pool even when the pool is inactive.
Connection Pool Idle Timeout	Number of seconds that a connection can remain idle without being dropped from the pool when there are more than the minimum number of connections.
Maximum Connection Lifetime	The number of minutes that a connection that was returned to the pool persists if there are more open connections than the minimum pool size.
Connection Validation Query	A data-source-specific query that the SDV query engine sends to see if the data source connection is valid. This query is executed every time a connection is checked out from the pool. Enter a query that returns quickly.
Execution Timeout	The number of seconds an execution query on the data source can run before being canceled.
Execute SELECTs Independently	Lets a SELECT statement be executed using a new connection from the connection pool, and committed immediately after completion. INSERT, UPDATE, and DELETE statements are executed using the same connection as part of the transaction.
Dynamic Configuration Procedure	Specify the path for the Dynamic Configuration Procedure callback that is invoked before the connection is established. It allows for the runtime modification of the configuration of a data source.

Connection Checkout Procedure	A procedure that returns a valid SQL statement that can be used to initialize the connection.
Supports Star Schema	Check only if this data source supports very large predicates and very large cardinalities for star schema semijoins.
Max Source Side Cardinality for Semi Join	See the documentation for semijoins and the SDV Administration Guide for more information.
Min Target to Source Ratio for Semi Join	Sets the minimum target-to-source ratio of cardinality for semijoins. Refer to the SDV Administration Guide for more information.
Max Source Side of Semi Join To Use OR Syntax	See the documentation for semijoins and the SDV Administration Guide for more information.
Enable Native Data Loading	Lets the data source use its proprietary functionality to optimize performance.
Collation Sensitive	SDV does not use the SORT MERGE join algorithm if any data source involved in the join is marked Collation Sensitive.
Concurrent Request Limit	Works with the Massively Parallel Processing engine configuration parameters to control the amount of parallelization for the queries for a particular data source.
Enable Pass-Through Prepared Statements	For pass-through to work, the prepared statement must call data from only one Oracle database instance. Prepared statements can use data from multiple tables within a single Oracle database instance.
Introspect Using DBA_* Views	Oracle maintains multiple metadata views. By default, SDV introspection uses ALL_* views, which list resources for which the user has access to both data and metadata. DBA_* views show all resources in the database regardless of data access permissions. Refer

	to Oracle documentation for differences and privileges.
Introspect Procedures	Checked by default. Ignoring procedures speeds up the initial introspection when only tables are wanted.
Introspect Comments	During the introspection process, SDV can retrieve table and column level comments and add them to the annotations field for each resource. See the User Guide Chapter <i>Introspecting Data Source Table and Column Comment Metadata</i> .
Include Invalid Introspection Objects	Check to return all objects during introspection, including invalid objects.
Use pass-through user's certificate for encryption	When pass-through login is configured for use with an Oracle data source, check this box to include the user's certificate in the SSL negotiation (handshake).
Is dataship source	Indicates whether the physical data source might be used as a source of shipped tables to another data ship enabled data source
Is dataship target	Indicates whether the physical data source might be used to receive shipped tables from another data ship enabled data source.
Lower bound/Upper bound for dataships	SDV uses Explain Plan to arrive at a numeric estimate of the cost of shipping data from a node to the Data Virtualizer. When the cost of shipping a federated query node falls between the limits of the Lowerbound and Upperbound, it is considered eligible for shipment so that it can be processed locally.
Schema Path for Temp Tables	A relative path to set the location of the temp tables on the data source. It is the name of a schema in the data source.
Temp Table Prefix	A character string addition to temporary table names so that they are recognized if they are needed.

Enable Oracle Database Links

Check to improve performance if you plan to use this data source for data caching or data ship optimization. Also add one or more Oracle database links. See the User Guide chapters *Configuration Native Caching for Oracle* and *Data Ship Performance Optimization*.

Database Link List

Add your database links separated with semicolons, using the following syntax:

```
[DBLINK NAME]@[DBLINK OWNER DS PATH]
```

For example:

```
oral2_
dblink@/users/composite/test/sources/dship/DEV-DSJ-
ORA11G-2
```

Connection URL Pattern

A template for generating a URL to connect to the physical data source.

Data Type

string

Default Value

jdbc:<DATA SOURCE>://<HOST>:<PORT>/<DATABASE_NAME>

Remarks

SDV does not validate modifications at the time of configuration. The data source adapter might not validate changes.

Connection URL String

The URL string generated from the connection URL pattern with the connection information you provide.

Data Type

string

Default Value

""

Remarks

This string is used by the JDBC adapter to connect to the physical data source. This field cannot be edited. For details, see the section “Connecting through JDBC Adapters” in the *SDV Administration Guide*.

JDBC Connection Properties

Lets you specify property-value pairs to pass to the JDBC data source.

Data Type

string

Default Value

""

Remarks

Click to add custom connection properties for any JDBC data source. Commonly used properties are populated with default values. Use the Add Argument button to specify other properties and values.

SDV does not validate property names. Some data source adapters ignore invalid property names or values; others return an error.

The driver properties specify connection timeout settings required by specific drivers. To avoid leaving connections open indefinitely, specify properties explicitly for your data source.

Connection Pool Maximum Size

Maximum number of connections (both active and idle) allowed for the data source. When the maximum is reached, new requests must wait until a connection is available.

Data Type

Numeric

Default Value

100

Remarks

If the maximum number of connections is in use when a request comes in (even with pass-through authentication), the new request is blocked and queued until a connection is available or the Connection Pool Idle Timeout is reached.

If no connection was made available within the specified timeout, a check is made for an available connection by the same user. If none is available, the least recently used connection for another user is dropped and a new connection is opened.

Studio reuses pooled connections if they continue to be valid after changes (such as connection name), but JDBC requests are forced to use new connections if any part of the data source connection configuration has changed.

Connection Pool Minimum Size

Minimum number of connections in the pool even when the pool is inactive.

Data Type

Numeric

Default Value

0

Remarks

When a connection has been idle, a validation query is used to verify whether an open connection is still valid just prior to submission of a request. If the connection is invalid, the connection is discarded and another is used.

Connection Pool Idle Timeout

Number of seconds that a connection can remain idle without being dropped from the pool when there are more than the minimum number of connections.

Data Type

Numeric

Default Value

30

Maximum Connection Lifetime

The number of minutes that a connection that was returned to the pool persists if there are more open connections than the minimum pool size.

Data Type

Numeric

Default Value

30

Remarks

The duration is calculated from connection creation. Default value is 60 minutes. Set a smaller value if the pool is likely to run out of connections. Be sure to add a validation query. Set a larger value if you want the connections to be held for a longer period. Set a value of 0 to keep connections alive indefinitely.

Connection Validation Query

A data-source-specific query that the SDV query engine sends to see if the data source connection is valid. This query is executed every time a connection is checked out from the pool. Enter a query that returns quickly.

Data Type

string

Default Value

""

Remarks

If this query returns a non-error result, the data source connection is considered valid. If this query fails, the connection is discarded and a new connection is checked out from the available pool.

No one SELECT statement works with all data sources. To verify that SDV is running and that it can connect to the data source, devise a query against a published table from that data source.

Enable Native Data Loading

Lets the data source use its proprietary functionality to optimize performance.

Data Type

Bool

Default Value

True

Remarks

See the User Guide, Chapter About Data Source Native Load Performance Options” for more details,.

Collation Sensitive

SDV does not use the SORT MERGE join algorithm if any data source involved in the join is marked Collation Sensitive.

Data Type

Bool

Default Value

False

Remarks

None

Concurrent Request Limit

Works with the Massively Parallel Processing engine configuration parameters to control the amount of parallelization for the queries for a particular data source.

Data Type

Numeric

Default Value

0

Remarks

None

Enable Pass-Through Prepared Statements

For pass-through to work, the prepared statement must call data from only one Oracle database instance. Prepared statements can use data from multiple tables within a single Oracle database instance.

Data Type

Bool

Default Value

False

Remarks

None

Introspect Using DBA_* Views

Oracle maintains multiple metadata views. By default, SDV introspection uses ALL_* views, which list resources for which the user has access to both data and metadata. DBA_* views show all resources in the database regardless of data access permissions. Refer to Oracle documentation for differences and privileges.

Data Type

Bool

Default Value

False

Remarks

None

Introspect Procedures

Checked by default. Ignoring procedures speeds up the initial introspection when only tables are wanted.

Data Type

Bool

Default Value

True

Remarks

None

Introspect Comments

During the introspection process, SDV can retrieve table and column level comments and add them to the annotations field for each resource. See the User Guide Chapter *Introspecting Data Source Table and Column Comment Metadata*.

Data Type

Bool

Default Value

False

Remarks

None

Include Invalid Introspection Objects

Check to return all objects during introspection, including invalid objects.

Data Type

Bool

Default Value

False

Remarks

None

Use pass-through user's certificate for encryption

When pass-through login is configured for use with an Oracle data source, check this box to include the user's certificate in the SSL negotiation (handshake).

Data Type

Bool

Default Value

False

Remarks

None

Is dataship source

Indicates whether the physical data source might be used as a source of shipped tables to another data ship enabled data source

Data Type

Bool

Default Value

False

Remarks

None

Is dataship target

Indicates whether the physical data source might be used to receive shipped tables from another data ship enabled data source.

Data Type

Bool

Default Value

False

Remarks

None

Lower bound/Upper bound for dataships

SDV uses Explain Plan to arrive at a numeric estimate of the cost of shipping data from a node to the Data Virtualizer. When the cost of shipping a federated query node falls between the limits of the Lowerbound and Upperbound, it is considered eligible for shipment so that it can be processed locally.

Data Type

Numeric

Default Value

Lower bound - 50000

Upper bound - 5000000

Remarks

To make changes to this field Is dataship source must be enabled.

Execution Timeout

The number of seconds an execution query on the data source can run before being canceled.

Data Type

Numeric

Default Value

0

Remarks

None

Execute SELECTs Independently

Lets a SELECT statement be executed using a new connection from the connection pool, and committed immediately after completion. INSERT and UPDATE statements are executed using the same connection as part of the transaction.

Data Type

Bool

Default Value

True

Remarks

None

Dynamic Configuration Procedure

Specify the path for the Dynamic Configuration Procedure callback that is invoked before the connection is established.

It allows for the runtime modification of the configuration of a data source.

Data Type

string

Default Value

""

Remarks

- The procedure, which must be owned by a member of the Admin group, will accept no parameters and return no results.
- Refer to the **API Guide -> GetEnvironment** chapter to see the list of system variables that can be used to dynamically change the configuration of the relational data source.

Connection Checkout Procedure

A procedure that returns a valid SQL statement that can be used to initialize the connection.

Data Type

string

Default Value

""

Remarks

The signature of the initialization procedure should be:

```
(IN ds_name VARCHAR, OUT sqlText VARCHAR)
```

Give the full path to the procedure in the Connection Check-out Procedure box.

Max Source Side Cardinality for Semi Join

See the documentation for semijoins and the SDV Administration Guide for more information.

Data Type

Numeric

Default Value

""

Remarks

None

Max Source Side of Semi Join To Use OR Syntax

See the documentation for semijoins and the SDV Administration Guide for more information.

Data Type

Numeric

Default Value

2147483647

Remarks

None

Min Target to Source Ratio for Semi Join

Sets a minimum ratio to trigger use of semi join optimization.

Data Type

Numeric

Default Value

""

Remarks

None

Supports Star Schema

Check only if this data source supports very large predicates and very large cardinalities for star schema semijoins.

Data Type

Bool

Default Value

False

Remarks

Refer to the section Star Schema Semijoin in the User Guide, for more information.

Schema Path for Temp Tables

A relative path to set the location of the temp tables on the data source. It is the name of a schema in the data source.

Data Type

String

Default Value

""

Remarks

This field is available only if you enable Is dataship target.

Temp Table Prefix

A character string addition to temporary table names so that they are recognized if they are needed.

Data Type

String

Default Value

T

Remarks

This field is available only if you enable Is dataship target.

Enable Oracle Database Links

Check to improve performance if you plan to use this data source for data caching or data ship optimization. Also add one or more Oracle database links. See the User Guide chapters *Configuration Native Caching for Oracle* and *Data Ship Performance Optimization*.

Data Type

Bool

Default Value

False

Remarks

None

Oracle Data Source Limitations

This section describes limitations of Oracle data sources that use an OCI or thin driver.

- In a MERGE statement, you cannot modify columns that are also in the ON clause. For example, the column `vendor_name` cannot be updated:

```
MERGE INTO output_table USING vendor_list ON vendor_name = Widgets
```

```
UPDDATE SET vendor_name = V9187
```

- SDV never pushes a MERGE statement with a DELETE branch, because Oracle's behavior is not ANSI-compliant. (Oracle can only delete a row if the row is first updated and then meets certain conditions.)
- To see what JDBC drivers support various versions of Oracle databases, search the Oracle web site.
- Inconsistent capitalization of projection names in query text can result in an Oracle Invalid Identifier error.

Note: Oracle supports all ANSI-supported joins including LEFT OUTER JOIN and RIGHT OUTER JOIN.

Oracle Caching Limitations

If the same Oracle database instance is acting as the source of and target for your cached data:

- LONG and LONG RAW data types. SQL*Plus is unable to SELECT a LONG RAW column. To work around the issue, you can convert the LONG RAW data types to BLOB.
- INTERVAL DAY TO SECOND and INTERVAL YEAR TO MONTH lose precision when they are cached using INSERT and SELECT statements.

Oracle to SDV Data Types

This section shows the mapping from Oracle data types to SDV data types.

- [Oracle NUMBER Data Types and SDV Data Types](#)
- [Oracle to Data Types Common to All Versions](#)
- [Oracle 11g to SDV Data Types](#)

Oracle NUMBER Data Types and SDV Data Types

Static Mapping

The following details apply to the mapping of Oracle NUMBER data types to SDV data types:

- The NUMBER data type is mapped to arbitrary Numeric in SDV.
- You may change this behavior by setting the Oracle-specific property "Arbitrary Numeric" to False. You can tune this via Administration -> Configuration -> Data sources -> Oracle Sources -> Arbitrary NUMERIC. By default, this is TRUE. When False, then:
 - If the scale of the NUMBER column is not specified, it is mapped as DOUBLE.
 - If either the precision or the scale is NULL, the data type is mapped to DOUBLE.
- When casting a value as DECIMAL(p, s), (for example, CAST (Oracle_column AS DECIMAL(40))):
 - Where the precision (p) is greater than 38, it is processed in SDV.
 - The maximum precision supported in SDV is Integer.MAX_VALUE, which is 2147483647.
 - The maximum scale that SDV supports is 255.
 - Any scale larger than 255 is automatically reduced to 255.

Results Mapping

Oracle has no SQL-standard equivalent of INTEGER. An INTEGER in Oracle is NUMBER(38), but division results in an arbitrary NUMBER value (for example, 32/5 will 6.4).

Oracle to Data Types Common to All Versions

This table lists the base Oracle data type to SDV data type mappings. Specific Oracle database versions might have additional data types. The additional data type mappings are listed in the sections for specific Oracle database versions.

Note: While Oracle honors trailing spaces in general, it ignores them when comparing CHARs. When SDV is set to honor trailing spaces, a filter on a CHAR column might return different results when executed in Oracle vs. SDV.

Oracle Base Data Type	SDV Data Type	Notes
BFILE	BLOB	
BLOB	BLOB	
CHAR	CHAR	Char and Varchar type will include special characters. The datatype length need not be explicitly increased to include the special characters.
CLOB	CLOB	
DATE	TIMESTAMP	
FLOAT	FLOAT	
LONG	CLOB	
LONG RAW	BLOB	

Oracle Base Data Type	SDV Data Type	Notes
LONG VARCHAR	CLOB	
NCHAR	CHAR	
NCLOB	CLOB	
NUMBER	DECIMAL	See Oracle NUMBER Data Types and SDV Data Types .
NUMBER(2,0)	TINYINT	In these examples, a hyphen indicates that the value is not specified in Oracle.
NUMBER(4,0)	SMALLINT	
NUMBER(8,0)	INTEGER	
NUMBER(15,0)	BIGINT	
NUMBER(22,0)	NUMERIC(22,0)	
NUMBER(10,3)	DECIMAL(10,3)	
NUMBER(-,0)	NUMERIC(38,0)	
NUMBER(-,2)	DECIMAL(38,2)	
NUMBER(12,-)	DOUBLE	
NUMBER(-,-)	DOUBLE	
NVARCHAR	VARCHAR	
NVARCHAR2	VARCHAR	
RAW	VARBINARY	
ROWID	VARCHAR	

Oracle Base Data Type	SDV Data Type	Notes
UROWID	VARCHAR	
VARCHAR	VARCHAR	Char and Varchar type will include special characters. The datatype length need not be explicitly increased to include the special characters.
VARCHAR2	VARCHAR	

Oracle 11g to SDV Data Types

The maximum VARBINARY length in Oracle 11g is 2000.

The following table shows the mapping from Oracle 11g data types to SDV data types.

Oracle 11g Data Type	SDV Data Type
ANYDATA	OTHER
ANYDATASET	OTHER
ANYTYPE	OTHER
BINARY DOUBLE	DOUBLE
BINARY FLOAT	FLOAT
INTERVAL DAY(0) TO SECOND(0) - INTERVAL DAY(9) TO SECOND(9)	INTERVAL DAY
INTERVAL YEAR(0) TO MONTH - INTERVAL YEAR(9) TO MONTH	INTERVAL YEAR

Oracle 11g Data Type	SDV Data Type
SDO_GEORASTER	OTHER
SI_STILLIMAGE	VARBINARY
TIMESTAMP	TIMESTAMP
TIMESTAMP(0) - TIMESTAMP(9)	TIMESTAMP
TIMESTAMP(0) WITH LOCAL TIME ZONE - TIMESTAMP(9) WITH LOCAL TIME ZONE	TIMESTAMP
TIMESTAMP(0) WITH TIME ZONE - TIMESTAMP(9) WITH TIME ZONE	TIMESTAMP
URITYPE	OTHER
UROWID	VARCHAR
XMLTYPE	XML

Oracle 19c to SDV Data Types

The maximum VARBINARY length in Oracle 19c is 2000.

The following table shows the mapping from Oracle 19c data types to SDV data types.

Oracle 19c Data Type	SDV Data Type
ANYDATA	OTHER
ANYDATASET	OTHER
ANYTYPE	OTHER

Oracle 19c Data Type	SDV Data Type
BINARY DOUBLE	DOUBLE
BINARY FLOAT	FLOAT
INTERVAL DAY(0) TO SECOND(0) - INTERVAL DAY(9) TO SECOND(9)	INTERVAL DAY
INTERVAL YEAR(0) TO MONTH - INTERVAL YEAR(9) TO MONTH	INTERVAL YEAR
SDO_GEORASTER	OTHER
SI_STILLIMAGE	VARBINARY
TIMESTAMP	TIMESTAMP
TIMESTAMP(0) - TIMESTAMP(9)	TIMESTAMP
TIMESTAMP(0) WITH LOCAL TIME ZONE - TIMESTAMP(9) WITH LOCAL TIME ZONE	TIMESTAMP
TIMESTAMP(0) WITH TIME ZONE - TIMESTAMP(9) WITH TIME ZONE	TIMESTAMP
URITYPE	OTHER
UROWID	VARCHAR
XMLTYPE	XML

Oracle Cache Mapping

This section discusses the data type mappings and restrictions for caches stored on Oracle.

Oracle changes any empty string stored in a VARCHAR2 or NVARCHAR2 column to NULL. This can alter empty string data stored in such columns.

FLOAT columns have a maximum of 126 digits, equivalent to a floating-point number with exponent E125. SDV FLOAT values have a maximum of E38 and DOUBLE values have a maximum of E308. This is why VARCHAR is used to store SDV DOUBLE values by default. However, you can use the FLOAT type if your values fit within that range.

Data Type	Preferred Data Type	Other Allowed Native Types
BIGINT	NUMBER(19, 0)	NUMBER(19+, 0), VARCHAR (20+), NVARCHAR(20+)
BINARY(n)	RAW(n); BLOB [if n > 255]	RAW(n+), BLOB
BIT	NUMBER(1, 0)	NUMBER(1+, 0)
BLOB	BLOB	
BOOLEAN	NUMBER(1,0)	NUMBER(1+,0)
CHAR(n)	CHAR(n); CLOB [if n > 2000]	CHAR(n+), VARCHAR2(n+), NVARCHAR2(n+), CLOB
CLOB	CLOB	
DATE	VARCHAR2(10)	VARCHAR2(10+), NVARCHAR2(10+)
DECIMAL(p,s)	NUMBER(p,s); CLOB [if p > 38]	NUMBER(p+,s+), VARCHAR2 (p+ 3+), NVARCHAR2(p+3+), CLOB
DOUBLE	VARCHAR(24)	VARCHAR(24+), FLOAT, BINARY DOUBLE
FLOAT	FLOAT	VARCHAR(24+), FLOAT, BINARY FLOAT, BINARY DOUBLE

Data Type	Preferred Data Type	Other Allowed Native Types
INTEGER	NUMBER(10, 0)	NUMBER(10+, 0), VARCHAR(20+), NVARCHAR(20+)
NUMERIC(p,s)	NUMBER(p,s); CLOB [if p > 38]	NUMBER(p+,s+), VARCHAR2(p+ 3+), NVARCHAR2(p+3+), CLOB
OTHER	[cannot be cached]	
SMALLINT	NUMBER(5, 0)	NUMBER(5+, 0), VARCHAR(20+), NVARCHAR(20+)
TIME	VARCHAR2(15)	VARCHAR2(15+), NVARCHAR2(15+)
TIMESTAMP	TIMESTAMP(9)	
TINYINT	NUMBER(3, 0)	NUMBER(3+, 0), VARCHAR(20+), NVARCHAR(20+)
VARBINARY(n)	RAW(n); BLOB [if n > 255]	RAW(n+), BLOB
VARCHAR(n)	VARCHAR(n); CLOB [if n > 4000]	VARCHAR2(n+), NVARCHAR2(n+), CLOB
XML	CLOB	VARCHAR(*), NVARCHAR2(*) [Truncates data if column is too small]

Oracle Function Support

SDV has made every effort to support all of the aggregate and analytic functions that Oracle supports. The following table lists the Oracle functions and notes describing how SDV interprets the functions.

Aggregate functions return a single result row based on groups of rows, rather than based on single rows. Aggregate functions can appear in SELECT lists and in ORDER BY and HAVING clauses. They are commonly used with the GROUP BY clause in a SELECT statement, where Oracle Database divides the rows of a queried table or view into groups.

SDV supports the following types of functions for Oracle:

- [Oracle Aggregate Function Support](#)
- [Oracle Analytic Function Support](#)
- [Oracle Analytic Aggregate Function Support](#)
- [Oracle Binary Function Support](#)
- [Oracle Character Function Support](#)
- [Oracle Conditional Function Support](#)
- [Oracle Conversion Function Support](#)
- [Oracle Date Function Support](#)
- [Oracle Geospatial Function Support](#)
- [Oracle Encryption Function Support](#)
- [Oracle Numeric Function Support](#)
- [Oracle Time Function Support](#)
- [Oracle XML Function Support](#)

Oracle Aggregate Function Support

SDV supports the aggregate functions listed in the table below for Oracle.

Oracle Aggregate Function	Notes
AVG	Does not support whole numbers, because Oracle returns floating-point numbers instead of integers in the result.
CORR	

Oracle Aggregate Function	Notes
COUNT	BLOB and CLOB not supported. In version 11g, also not supported: BFILE, LONG, LONG RAW, LONGVARCHAR, NCLOB, or VARBINARY.
COVAR_POP	
COVAR_SAMP	
factorial sign (!)	Argument can be any whole number.
LISTAGG	
MAX	
MIN	
PERCENTILE_CONT	
PERCENTILE_DISC	
SUM	

Oracle Analytic Function Support

SDV supports the analytic functions listed in the table below for Oracle.

Analytic functions are commonly used in data warehousing environments. They are all push-only functions.

Oracle Analytic Function	Notes
CORR	
COUNT	

Oracle Analytic Function	Notes
COVAR_POP	
COVAR_SAMP	
CUME_DIST	
DENSE_RANK	
FIRST_VALUE	
LAG	
LAST_VALUE	
LEAD	
LISTAGG	Version 11g R2. Pushed. (Aggregate LISTAGG function is supported in SDV.
MAX	
MEDIAN	MEDIAN (DISTINCT) not supported.
MIN	
NTILE	
PERCENT_RANK	
PERCENTILE_CONT	
PERCENTILE_DISC	
RANK	
RATIO_TO_REPORT	

Oracle Analytic Function	Notes
ROW_NUMBER	
SUM	

Oracle Analytic Aggregate Function Support

SDV supports the analytic aggregate functions listed in the table below for Oracle.

Analytic functions are commonly used in data warehousing environments. They are all push-only functions.

Oracle Analytic Aggregate Function	Notes
LAST_VALUE	
PERCENTILE_CONT	
PERCENTILE_DISC	
STDDEV	
STDDEV_POP	
STDDEV_SAMP	
VAR_POP	Same as VARIANCE_POP.
VAR_SAMP	Same as VARIANCE_SAMP.
VARIANCE	
VARIANCE_POP	
VARIANCE_SAMP	

Oracle Binary Function Support

SDV supports the binary functions listed in the table below .

Oracle Binary Function	Notes
INT1AND	
INT2AND	
INT4AND	
INT8AND	

Oracle Character Function Support

SDV supports the character functions listed in the table below for Oracle.

Oracle Character Function	Notes
ASCII	
BTRIM	
CHR	
CONCAT	
GREATEST	Version 11g.
INITCAP	
INSTR	Case-sensitive by default.
LEAST	Version 11g.
LENGTH	

Oracle Character Function	Notes
LOWER	
LPAD	
LTRIM	
POSITION	Follows SQL-92 (STRICT). Not supported: mixing string, number, or date with NCHAR, NVARCHAR, or NVARCHAR2.
REPLACE	
RPAD	
RTRIM	
SOUNDEX	Returns a phonetic representation of a string.
SPACE	BIT not supported. Oracle returns SPACE(0) as NULL, but SQL-92 calls for ' '.
STRPOS	
SUBSTR	
SUBSTRING	Oracle does not follow SQL-92 standard. STRICT forces use of SDV, which follows the standard.
TRANSLATE	
TRIM	
UNICHR	
UPPER	

Oracle Conditional Function Support

SDV supports the conditional functions listed in the table below for Oracle.

Oracle Conditional Function	Notes
COALESCE	
DECODE	
NULLIF	NULL cannot be the first argument. Does not support BFILE, BLOB, CLOB, LONG, or LONGVARCHAR.
NVL	
NVL2	

Oracle Conversion Function Support

SDV supports the conversion functions listed in the table below for Oracle.

Oracle Conversion Function	Notes
CAST	
FORMAT_DATE	For timestamps, Oracle omits fractional parts unless the format string “ff” is used. If “ff” does not specify a precision, Oracle returns the precision of the data type used to store the returned value.
PARSE_DATE	Does not push.
PARSE_TIMESTAMP	Format “ff” (fractional part) is valid only for TO_TIMESTAMP.

Oracle Conversion Function	Notes
TO_CHAR	For timestamps, Oracle omits fractional parts unless the format string “ff” is used. If “ff” does not specify a precision, Oracle returns the precision of the data type used to store the returned value.
TO_DATE	Format “ff” (fractional part) does not work for this function.
TO_NUMBER	
TO_TIMESTAMP	For timestamps, Oracle omits fractional parts unless the format string “ff” is used. If “ff” does not specify a precision, Oracle returns the precision of the data type used to store the returned value.

Oracle Date Function Support

SDV supports the date functions listed in the table below for Oracle.

Oracle Date Function	Notes
ADD_MONTHS	
CURRENT_DATE	
CURRENT_TIME	Not supported.
CURRENT_TIMESTAMP	
DATE_TRUNC	
DAY	

Oracle Date Function	Notes
DAYS_BETWEEN	
LAST_DAY	
MONTH	
MONTHS_BETWEEN	
NEXT_DAY	
TRUNC	
YEAR	
SYSDATE	

Oracle Geospatial Function Support

SDV supports the geospatial functions listed in the table below for Oracle.

Oracle Geospatial Function	Notes
SDO_GEOMETRY	Refer https://docs.oracle.com/database/121/SPATL/st_geometry-and-sdo_geometry-interopability.htm#SPATL563
ST_AREA	
ST_ASBINARY	
ST_ASTEXT	
ST_BOUNDARY	

Oracle Geospatial Function	Notes
ST_BUFFER	
ST_CENTROID	
ST_CIRCULARSTRING	
ST_COMPOUNDCURVE	
ST_CONTAINS	
ST_CONVEXHULL	
ST_CROSSES	
ST_CURVE	
ST_CURVEPOLYGON	
ST_DIFFERENCE	
ST_DIMENSION	
ST_DISJOINT	
ST_DISTANCE	
ST_ENDPOINT	
ST_ENVELOPE	
ST_EQUALS	
ST_EXTERIORRING	
ST_GEOMCOLLECTION	Will be introspected as SDV geometry
ST_GEOMETRY	Will be introspected as SDV geometry

Oracle Geospatial Function	Notes
ST_GEOMETRYTYPE	
ST_GEOMFROMTEXT	
ST_GEOMFROMWKB	
ST_INTERIORRINGN	
ST_INTERSECTS	
ST_ISCLOSED	
ST_ISEMPTY	
ST_ISSIMPLE	
ST_ISVALID	
ST_LENGTH	
ST_LINEFROMTEXT	
ST_LINEFROMWKB	
ST_LINestring	Will be introspected as SDV geometry
ST_MLINEFROMTEXT	
ST_MLINEFROMWKB	
ST_MPOINTFROMTEXT	
ST_MPOINTFROMWKB	
ST_MPOLYFROMTEXT	
ST_MPOLYFROMWKB	

Oracle Geospatial Function	Notes
ST_MULTICURVE	
ST_MULTILINESTRING	Will be introspected as SDV geometry
ST_MULTIPPOINT	Will be introspected as SDV geometry
ST_MULTIPOLYGON	Will be introspected as SDV geometry
ST_MULTISURFACE	
ST_NUMPOINTS	
ST_OVERLAPS	
ST_POINT	Will be introspected as SDV geometry
ST_POINTFROMTEXT	
ST_POINTFROMWKB	
ST_POINTN	
ST_POLYFROMTEXT	
ST_POLYFROMWKB	
ST_POLYGON	Will be introspected as SDV geometry
ST_SRID	
ST_SRSID	
ST_STARTPOINT	
ST_SYMDIFFERENCE	
ST_TOUCHES	

Oracle Geospatial Function	Notes
ST_X	
ST_Y	
ST_Z	
ST_UNION	
ST_WITHIN	
ST_GEOGRAPHYFROMTEXT	

Oracle Encryption Function Support

SDV supports the encryption functions listed in the table below for Oracle version 10g with DBMS_CRYPTO package.

Oracle Encryption Function	Notes
HASHMD5	
HASHSHA1	

Oracle Numeric Function Support

SDV supports the numeric functions listed in the table below for Oracle.

Oracle Numeric Function	Notes
ABS	
ACOS	

Oracle Numeric Function	Notes
ASIN	
ATAN	
ATAN2	
CEIL	
CEILING	
COS	
COT	
DEGREES	Not available.
DENSE_RANK	
EXP	
FLOOR	
LN	BIT, NCHAR, and NVARCHAR not supported.
LOG	
MEDIAN	MEDIAN (DISTINCT) not supported.
MOD	
PI	Not available.
POW	Not supported: mixing string, number, or NULL with BIT, NCHAR, or NVARCHAR.
POWER	

Oracle Numeric Function	Notes
RADIANS	Not available.
RANDOM	
ROUND	
SIGN	
SIN	
SQRT	
TAN	
TRUNC	
VARIANCE	

Oracle Specific Connection Properties

This section describes the connection properties that are specific to the Oracle data source.

Port	Port number for the data source to connect with the host. The Port number for Oracle is 1521.
Net Service Name	This property is available for an Oracle data source with OCI driver. The TNS name that is set up through Oracle Net Configuration Assistant.
Pass-through Login	To use Kerberos tokens, Oracle data sources must

	Kerberos security module must be running on SDV so that Kerberos-authenticated users with session tokens can use it when submitting queries.
Service Name	Name of the database service.
Ticket Cache	Specify the location of where your ticket cache is stored in this field when using Kerberos authentication. Leave the Login and Password fields blank when using ticket cache Note: Credential cache must be generated with a forwardable option set. To do this, KRB5.INI or KRB5.CONF in the client machine needs to have the parameter libdefaults: forwardable set to true.
Enable Oracle Database Link	Check to improve performance if you plan to use this data source for data caching or data ship optimization. Also add one or more Oracle database links. See the User Guide chapters <i>Configuration Native Caching for Oracle</i> and <i>Data Ship Performance Optimization</i> .
Enable Pass-Through Prepared Statements	For pass-through to work, the prepared statement must call data from only one Oracle database instance. Prepared statements can use data from multiple tables within a single Oracle database instance.
Include Invalid Introspection Objects	Check to return all objects during introspection, including invalid objects.
Introspect Procedures	Checked by default. Ignoring procedures speeds up the initial introspection when only tables are wanted.
Introspect comments	During the introspection process, SDV can retrieve table and column level comments and add them to the annotations field for each resource. See the User Guide Chapter <i>Introspecting Data Source Table and Column Comment Metadata</i> .

Introspect Using DBA_* Views	Oracle maintains multiple metadata views. By default, SDV introspection uses ALL_* views, which list resources for which the user has access to both data and metadata. DBA_* views show all resources in the database regardless of data access permissions. Refer to Oracle documentation for differences and privileges.
Use pass-through user's certificate for encryption	When pass-through login is configured for use with an Oracle data source, check this box to include the user's certificate in the SSL negotiation (handshake).
Is dataship source	Indicates whether the physical data source might be used as a source of shipped tables to another data ship enabled data source
Lower bound/Upper bound	SDV uses Explain Plan to arrive at a numeric estimate of the cost of shipping data from a node to the Data Virtualizer. When the cost of shipping a federated query node falls between the limits of the Lowerbound and Upperbound, it is considered eligible for shipment so that it can be processed locally.
Is dataship target	Indicates whether the physical data source might be used to receive shipped tables from another data ship enabled data source.
Lower bound/Upper bound	SDV uses Explain Plan to arrive at a numeric estimate of the cost of shipping data from a node to the Data Virtualizer. When the cost of shipping a federated query node falls between the limits of the Lowerbound and Upperbound, it is considered eligible for shipment so that it can be processed locally.
Schema path for Temp Tables	A relative path to set the location of the temp tables on the data source. It is the name of a schema in the data source.
Temp Table Prefix	A character string addition to temporary table names so that they are recognized if they are needed.

Database Link List

Add your database links separated with semicolons, using the following syntax:

```
[DBLINK NAME]@[DBLINK OWNER DS PATH]
```

For example:

```
oral2_  
dblink@/users/composite/test/sources/dship/DEV-DSJ-  
ORA11G-2
```

Oracle Time Function Support

SDV supports the time functions listed in the table below for Oracle.

Oracle Time Function	Notes
EXTRACT	
NOW	
TIMEOFDAY	

Oracle XML Function Support

SDV supports the XML functions listed in the table below for Oracle.

Oracle XML Function	Notes
XMLAGG	
XMLATTRIBUTES	
XMLCOMMENT	Not supported, because it returns the following error:

Oracle XML Function	Notes
	ORA-00932: inconsistent data types: expected - got CHAR.
XMLCONCAT	
XMLDOCUMENT	Oracle does not support this function.
XMLELEMENT	Version 11g.
XMLFOREST	
XMLTEXT	Oracle does not support this function.

References

Refer to the following Guides for further details about the capabilities of the data source:

Capabilities	Section
Query Engine	User Guide, Chapter <i>SDV Query Engine Optimizations</i>
Data ship	User Guide, Chapter <i>Data Ship Performance Optimization</i>
Caching	User Guide, Chapter <i>SDV Caching</i>
Performance Optimization	User Guide, Chapter <i>Performance Tuning</i>
SDV Massively Parallel Processing Engine	User Guide, Chapter <i>Configuring the SDV MPP Engine</i>
Kerberos	Administration Guide Chapter <i>Configuring Kerberos</i>

Spotfire Documentation and Support Services

For information about this product, you can read the documentation, contact Spotfire Support, and join Spotfire Community.

How to Access Spotfire Documentation

Documentation for Spotfire products is available on the [Product Documentation website](#), mainly in HTML and PDF formats.

The [Product Documentation website](#) is updated frequently and is more current than any other documentation included with the product.

Product-Specific Documentation

The following documentation for this product is available on the [Spotfire® Data Virtualization](#) page.

Users

- SDV Getting Started Guide
- SDV User Guide
- SDV Web UI User Guide
- SDV Client Interfaces Guide
- SDV Tutorial Guide
- SDV Northbay Example

Administration

- SDV Installation and Upgrade Guide
- SDV Administration Guide
- SDV Active Cluster Guide
- SDV Security Features Guide

Data Sources

- SDV Adapter Guides

SDV Data Source Toolkit Guide (Formerly Extensibility Guide)

References

SDV Reference Guide

SDV Application Programming Interface Guide

Other

SDV Business Directory Guide

SDV Discovery Guide

SDV and Business Directory Release Notes - Read the release notes for a list of new and changed features. This document also contains lists of known issues and closed issues for this release.

Release Version Support

Spotfire 8.9 is designated as Long Term Support (LTS) versions. Some release versions of Spotfire® Data Virtualization products are selected to be long-term support (LTS) versions. Defect corrections will typically be delivered in a new release version and as hotfixes or service packs to one or more LTS versions. See also [Long Term Support](#).

How to Contact Support for Spotfire Products

You can contact the Support team in the following ways:

- To access the Support Knowledge Base and getting personalized content about products you are interested in, visit our [product Support website](#).
- To create a Support case, you must have a valid maintenance or support contract with a Cloud Software Group entity. You also need a username and password to log in to the [product Support website](#). If you do not have a username, you can request one by clicking **Register** on the website.

How to Join Spotfire Community

Spotfire Community is the official channel for Spotfire customers, partners, and employee subject matter experts to share and access their collective experience. Spotfire Community offers access to Q&A forums, product wikis, and best practices. It also offers access to extensions, adapters, solution accelerators, and tools that extend and enable customers to gain full value from Spotfire products. In addition, users can submit and vote on feature requests from within the [Spotfire Ideas Portal](#). For a free registration, go to [Spotfire Community](#).

Legal and Third-Party Notices

SOME CLOUD SOFTWARE GROUP, INC. ("CLOUD SG") SOFTWARE AND CLOUD SERVICES EMBED, BUNDLE, OR OTHERWISE INCLUDE OTHER SOFTWARE, INCLUDING OTHER CLOUD SG SOFTWARE (COLLECTIVELY, "INCLUDED SOFTWARE"). USE OF INCLUDED SOFTWARE IS SOLELY TO ENABLE THE FUNCTIONALITY (OR PROVIDE LIMITED ADD-ON FUNCTIONALITY) OF THE LICENSED CLOUD SG SOFTWARE AND/OR CLOUD SERVICES. THE INCLUDED SOFTWARE IS NOT LICENSED TO BE USED OR ACCESSED BY ANY OTHER CLOUD SG SOFTWARE AND/OR CLOUD SERVICES OR FOR ANY OTHER PURPOSE.

USE OF CLOUD SG SOFTWARE AND CLOUD SERVICES IS SUBJECT TO THE TERMS AND CONDITIONS OF AN AGREEMENT FOUND IN EITHER A SEPARATELY EXECUTED AGREEMENT, OR, IF THERE IS NO SUCH SEPARATE AGREEMENT, THE CLICKWRAP END USER AGREEMENT WHICH IS DISPLAYED WHEN ACCESSING, DOWNLOADING, OR INSTALLING THE SOFTWARE OR CLOUD SERVICES (AND WHICH IS DUPLICATED IN THE LICENSE FILE) OR IF THERE IS NO SUCH LICENSE AGREEMENT OR CLICKWRAP END USER AGREEMENT, THE LICENSE(S) LOCATED IN THE "LICENSE" FILE(S) OF THE SOFTWARE. USE OF THIS DOCUMENT IS SUBJECT TO THOSE SAME TERMS AND CONDITIONS, AND YOUR USE HEREOF SHALL CONSTITUTE ACCEPTANCE OF AND AN AGREEMENT TO BE BOUND BY THE SAME.

This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of Cloud Software Group, Inc.

Spotfire, Spotfire logo, Spotfire O logo, ActiveSpaces, Enterprise Messaging Service, Spotfire, TERR, S-PLUS, and S+ are either registered trademarks or trademarks of Spotfire Software Inc. in the United States and/or other countries.

Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

All other product and company names and marks mentioned in this document are the property of their respective owners and are mentioned for identification purposes only. You acknowledge that all rights to these third party marks are the exclusive property of their respective owners. Please refer to Cloud SG's Third Party Trademark Notices (<https://www.cloud.com/legal>) for more information.

This document includes fonts that are licensed under the SIL Open Font License, Version 1.1, which is available at: <https://scripts.sil.org/OFL>

Copyright (c) Paul D. Hunt, with Reserved Font Name Source Sans Pro and Source Code Pro.

Cloud SG software may be available on multiple operating systems. However, not all operating system platforms for a specific software version are released at the same time. See the “readme” file for the availability of a specific version of Cloud SG software on a specific operating system platform.

THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS DOCUMENT COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THIS DOCUMENT. CLOUD SG MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S), THE PROGRAM(S), AND/OR THE SERVICES DESCRIBED IN THIS DOCUMENT AT ANY TIME WITHOUT NOTICE.

THE CONTENTS OF THIS DOCUMENT MAY BE MODIFIED AND/OR QUALIFIED, DIRECTLY OR INDIRECTLY, BY OTHER DOCUMENTATION WHICH ACCOMPANIES THIS SOFTWARE, INCLUDING BUT NOT LIMITED TO ANY RELEASE NOTES AND "README" FILES.

This and other products of Cloud SG may be covered by registered patents. For details, please refer to the Virtual Patent Marking document located at <https://www.Spotfire.com/patents>.

Copyright © 2002-2026. Cloud Software Group, Inc. All Rights Reserved.