



Progress DataDirect for JDBC for Greenplum User's Guide

Release 6.0.0

Copyright

Visit the following page online to see Progress Software Corporation's current Product Documentation Copyright Notice/Trademark Legend: <https://www.progress.com/legal/documentation-copyright>.

Updated: 2025/05/14

Table of Contents

Welcome to the Progress DataDirect for JDBC for Greenplum Driver.....9

What's new in this release?.....	10
Requirements.....	11
Installing and setting up the driver.....	11
Driver and DataSource classes.....	13
Connection URL examples.....	13
User ID/password authentication.....	13
Kerberos authentication.....	14
Proxy server.....	15
Data types.....	16
getTypeInfo().....	18
SQL escape sequences.....	32
Supported scalar functions.....	32
DataDirect tools.....	34
Troubleshooting.....	34
Additional information	34
Contacting Technical Support.....	34

Tutorials37

Tableau	37
DbVisualizer	38
Adding a driver	38
Connecting and executing SQL statements	39

Configuring and connecting41

Setting the classpath	42
Connecting using the JDBC Driver Manager.....	42
Passing the connection URL.....	42
Testing the connection.....	43
Connecting using data sources.....	47
How data sources are implemented.....	47
Creating data sources.....	47
Calling a data source in an application.....	48
Testing a data source connection.....	49
Authentication.....	51
Configuring the driver for user ID/password authentication.....	51
Configuring the driver for Kerberos authentication.....	52
Data Encryption.....	56

Configuring TLS/SSL encryption.....	57
Configuring TLS/SSL Server Authentication.....	58
Configuring TLS/SSL Client Authentication.....	59
FIPS (Federal Information Processing Standard).....	60
Performance considerations.....	60

Additional features and functionality61

Statement pooling.....	62
Connection pooling.....	62
Using client information.....	62
Bulk load support.....	63
Failover.....	63
Configuring failover.....	63
Refcursors.....	66
Isolation levels.....	66
Using scrollable cursors.....	66
JTA support.....	66
Large object (LOB) support.....	67
Using COPY command.....	67
Parameter metadata support.....	68
User-defined functions' results.....	68
ResultSet metadata support.....	70
Rowset support.....	71

Connection property descriptions.....73

AccountingInfo.....	81
AlternateServers.....	82
ApplicationName.....	83
AuthenticationMethod.....	83
BulkLoadBatchSize.....	84
CatalogOptions.....	85
ClientHostName.....	85
ClientUser.....	86
ConnectionRetryCount.....	87
ConnectionRetryDelay.....	88
ConvertNull.....	89
CryptoProtocolVersion.....	89
Database.....	91
DatabaseName.....	91
EnableCancelTimeout.....	92
EncryptionMethod.....	92
ExtendedColumnMetadata.....	93
HostNameInCertificate.....	94

ImportStatementPool.....	95
InitializationString.....	96
InsensitiveResultSetBufferSize.....	96
JavaDoubleToString.....	97
KeyPassword.....	98
KeyStore.....	99
KeyStorePassword.....	99
LoadBalancing.....	100
LoginConfigName.....	101
LoginTimeout.....	102
MaxLongVarcharSize.....	102
MaxNumericPrecision.....	103
MaxNumericScale.....	103
MaxPooledStatements.....	104
MaxStatements.....	105
MaxVarcharSize.....	105
Password.....	106
PortNumber.....	106
PrepareThreshold.....	107
ProgramID.....	108
ProxyHost.....	109
ProxyPassword.....	109
ProxyPort.....	110
ProxyUser.....	110
QueryTimeout.....	111
RegisterStatementPoolMonitorMBean.....	112
ResultSetMetaDataOptions.....	112
ServerName.....	113
ServicePrincipalName.....	114
SpyAttributes.....	115
SupportsCatalogs.....	116
TransactionErrorBehavior.....	117
TrustStore.....	118
TrustStorePassword.....	118
User.....	119
ValidateServerCertificate.....	120
VarcharClobThreshold.....	120

Welcome to the Progress DataDirect for JDBC for Greenplum Driver

The Progress® DataDirect® for JDBC™ for Greenplum™ driver (Greenplum driver) supports the JDBC API for SQL read-write access to Greenplum.

The documentation for the driver also includes the *Progress DataDirect for JDBC Drivers Reference*. The reference provides general reference information for all DataDirect drivers for JDBC, including content on troubleshooting, supported SQL escapes, and DataDirect tools.

For the complete documentation set, visit the Progress DataDirect Connectors Documentation Hub: <https://docs.progress.com/category/datadirect-pivotal-greenplum>.

For details, see the following topics:

- [What's new in this release?](#)
- [Requirements](#)
- [Installing and setting up the driver](#)
- [Driver and DataSource classes](#)
- [Connection URL examples](#)
- [Data types](#)
- [SQL escape sequences](#)
- [DataDirect tools](#)
- [Troubleshooting](#)

- [Additional information](#)
- [Contacting Technical Support](#)

What's new in this release?

Support and certification

Visit the following web pages for the latest support and certification information.

- Release Notes: <https://www.progress.com/datadirect-connectors/whats-new#jdbc>
- DataDirect Product Compatibility Guide:
<https://docs.progress.com/bundle/datadirect-product-compatibility/resource/datadirect-product-compatibility.pdf>

Changes Since 6.0.0 GA

- **Driver enhancements**
 - The driver has been enhanced to comply with FIPS standards for data encryption. As part of this enhancement, the driver was tested with FIPS 140-3 enabled using a Red Hat OpenJDK 21 on a Red Hat Universal Base Image 9 instance. See [FIPS \(Federal Information Processing Standard\)](#) on page 60 for details.
 - The driver has been enhanced to support the TLSv1.3 cryptographic protocol. See the [CryptoProtocolVersion](#) on page 89 connection property for details. The TLSv1.3 protocol requires Java SE 11 or higher.

Changes for 6.0.0 GA

- **Driver enhancements**
 - The driver has been enhanced to support the following types of data types:
 - Geometric types
 - Network address types
 - Monetary types
 - Date/time types
 - Range types

Note: For more information, see [Data types](#) on page 16.

- The driver has been enhanced to support the TXID_SNAPSHOT data type.
- The driver has been enhanced to support a customized version of the Greenplum COPY command. It provides an additional keyword, LOCALFILE, to allow you to copy data from or to standard file-system files that are stored anywhere on your network, not just on the database server. For more information, see [Using COPY command](#) on page 67.

- The driver has been enhanced to leverage server-side prepared statements via the `PrepareThreshold` connection property. `PrepareThreshold` determines the number of `PreparedStatement` executions that must occur before the driver begins using server-side prepared statements. The default value is 0. When set to a value greater than 1, the driver starts using server-side prepared statements upon reaching the configured threshold value.
- The driver has been enhanced to support the following data types: Citext, JSON, and UUID.
- The driver has been enhanced to support the Kerberos authentication protocol with the following connection properties:
 - `AuthenticationMethod`
 - `ServicePrincipalName`
- **Changed behavior**
 - The driver does not support access to Pivotal HAWQ anymore.
 - The default value for the `SupportsCatalogs` connection property has been changed to `true`.

Requirements

The driver is compatible with JDBC 2.0, 3.0, and 4.0.

The driver requires a Java Virtual Machine (JVM) that is Java SE 8 or higher, including Oracle JDK, OpenJDK, and IBM SDK (Java) distributions.

Installing and setting up the driver

This section provides you with an overview of the steps required to install and set-up the driver. After completing this procedure, you will be able to begin accessing data with your application.

To begin accessing data with the driver:

1. Install the driver:
 - a) After downloading the product, unzip the installer files to a temporary directory.
 - b) From the installer directory, run the appropriate installer file to start the installer.
 - **Windows:** `PROGRESS_DATADIRECT_JDBC_INSTALL.exe`
 - **Non-Windows:** `PROGRESS_DATADIRECT_JDBC_INSTALL.jar`

- c) Follow the prompts to complete installation.

The installer program supports multiple installation methods, including command-line and silent installations. For detailed instructions, refer to the *Progress DataDirect for JDBC Drivers Installation Guide*.

2. Set your system `CLASSPATH` to include the driver `.jar` file. The `CLASSPATH` is the search string your Java Virtual Machine (JVM) uses to locate JDBC drivers on your computer. The following examples demonstrate setting the `CLASSPATH` from a command line using the default installation directory.

- **Windows Example**

```
CLASSPATH=.;C:\Program Files\Progress\DataDirect\JDBC\lib\60\greenplum.jar
```

- **UNIX/LINUX Example**

```
CLASSPATH=./opt/Progress/DataDirect/JDBC/lib/60/greenplum.jar
```

3. Configure your driver using one of the following methods:

- **Connection URL:** You can begin using the driver immediately by passing a connection URL with your application or tool. The following examples show how to connect using either user ID/password or Kerberos authentication.

User ID/password authentication

```
jdbc:datadirect:greenplum://server1:5432;  
AuthenticationMethod=userIdPassword;User=test;  
Password=secret;
```

Note: See [Configuring the driver for user ID/password authentication](#) on page 51 for details.

Kerberos authentication

```
jdbc:datadirect:greenplum://server1:5432;  
DatabaseName=greenplumDB;AuthenticationMethod=kerberos;  
ServicePrincipalName=postgres/myserver.example.com@EXAMPLE.COM;
```

Note: See [Configuring the driver for Kerberos authentication](#) on page 52 for details.

- **Data sources:** The driver also supports connecting using JDBC data sources. A JDBC data source is a Java object, specifically a DataSource object, that defines connection information required for a JDBC driver to connect to the database. See [Connecting using data sources](#) for more information.

Note: For most connections, specifying the minimum required connection properties is sufficient to begin accessing data; however, you can provide values for optional properties to use additional supported features and improve performance.

4. Set the values for any optional properties that you want to configure. For additional information on optional features and functionality, see the following resources:

- [Connection URL Examples](#) provides connection string examples that can be used to configure common functionality and features. You can modify and combine these examples to create a string that best suits your environment.
- [Connection property descriptions](#) provides a complete list of supported properties by functionality.
- [Performance considerations](#) describes connection properties that affect performance, along with recommended settings.

5. Connect to your service and begin accessing data with your applications, BI tools, database tools, and more. To help you get started, the following resources guide you through accessing data with some common tools:

- [Tableau](#): Tableau is a business intelligence software program that allows you to easily create reports and visualized representations of your data.

- [DbVisualizer](#): DB Visualizer is a database tool that allows you to connect and execute SQL statements against your data.
- [DataDirect Test](#): DataDirect Test allows you to test connect, execute SQL statements, and practice using the JDBC API right out of the box.

This completes the deployment of the driver.

Driver and DataSource classes

The following are the `Driver` and `DataSource` classes used by the driver:

Driver class:

`com.ddtek.jdbc.greenplum.GreenplumDriver`

DataSource class:

`com.ddtek.jdbcx.greenplum.GreenplumDataSource`

Connection URL examples

After setting the CLASSPATH, the connection information needs to be passed in the form of a connection URL. This section provides examples of connection strings configured to use common features and functionality. You can modify and/or combine these examples to create a connection string for your environment.

Note:

- Connection property names are case-insensitive. For example, `Password` is the same as `password`.
 - For connection properties that support string values, use the following escape sequence to specify values containing leading or trailing spaces and curly brackets: `{value}`. For example: `User={hello }` or `Password={{hello}}`.
-

User ID/password authentication

This string includes the properties used to connect with the user ID/password authentication.

```
jdbc:datadirect:greenplum://servername:port;  
AuthenticationMethod=userIdPassword;User=userID;  
Password=password;[property=value[;...]];
```

where:

servername

is the IP address or name of the server to which you are connecting.

port

is the number of the TCP/IP port.

userID

specifies the user ID that is used to connect to the Greenplum database.

password

specifies a password that is used to connect to your Greenplum database.

property=value

specifies connection property settings. Multiple properties are separated by a semi-colon.

The following example connection string includes the properties required for connecting with the user ID/password authentication.

```
Connection conn = DriverManager.getConnection
( "jdbc:datadirect:greenplum://server1:5432;
  AuthenticationMethod=userIdPassword;User=test;
  Password=secret;" );
```

See also

[Connection property descriptions](#) on page 73

[Configuring the driver for user ID/password authentication](#) on page 51

Kerberos authentication

This string includes the properties used to connect with the Kerberos authentication.

```
jdbc:datadirect:greenplum://servername:port;
DatabaseName=databasename;AuthenticationMethod=kerberos;
ServicePrincipalName=serviceprincipalname[property=value[;...]];
```

where:

servername

is the IP address or name of the server to which you are connecting.

port

is the number of the TCP/IP port.

databasename

specifies the name of the database to which you want to connect.

serviceprincipalname

specifies the three-part service principal name registered with the key distribution center (KDC).

property=value

specifies connection property settings. Multiple properties are separated by a semi-colon.

The following example connection string includes the properties required for connecting with the user ID/password authentication.

```
Connection conn = DriverManager.getConnection
( "jdbc:datadirect:greenplum://server1:5432;
  DatabaseName=greenplumDB;AuthenticationMethod=kerberos;
  ServicePrincipalName=postgres/myserver.example.com@EXAMPLE.COM;" );
```

See also

[Connection property descriptions](#) on page 73

[Configuring the driver for Kerberos authentication](#) on page 52

Proxy server

This string includes the properties you may need to connect through a proxy server with user ID/password authentication.

```
jdbc:datadirect:greenplum://servername:port;ProxyHost=proxy_host;
ProxyPassword=proxy_password;ProxyPort=proxy_port;ProxyUser=proxy_user;
AuthenticationMethod=userIdPassword;User=userID;
Password=password;[property=value[;...]];
```

where:

servername

is the IP address or name of the server to which you are connecting.

port

is the number of the TCP/IP port.

proxy_host

specifies the proxy server to use for the first connection.

proxy_password

specifies the password needed to connect to a proxy server for the first connection.

proxy_port

specifies the port number where the proxy server is listening for requests for the first connection. The default is 0.

proxy_user

specifies the user name needed to connect to a proxy server for the first connection.

userID

specifies the user ID that is used to connect to the Greenplum database.

password

specifies a password that is used to connect to your Greenplum database.

property=value

specifies connection property settings. Multiple properties are separated by a semi-colon.

The following example connection string includes the properties required for using a proxy server with user ID/password authentication.

```
Connection conn = DriverManager.getConnection
("jdbc:datadirect:greenplum://server1:5432;ProxyHost=pserver;
ProxyPassword=secret;ProxyPort=808;ProxyUser=jsmith;
AuthenticationMethod=userIdPassword;User=test;Password=secret;");
```

See also

[Connection property descriptions](#) on page 73

Data types

The following table lists the data types supported by the Greenplum driver and describes how they are mapped to JDBC data types.

Table 1: Greenplum Data Types

Greenplum Data Type	JDBC Data Type
Bigint	BIGINT
Bigserial	BIGINT
Bit	BIT or BINARY ¹
Bit varying	VARBINARY
Boolean	BOOLEAN
Box	VARCHAR
Bytea	LONGVARBINARY
Character	CHAR
Character varying	VARCHAR or LONGVARCHAR ²
CIDR	VARCHAR
Circle	VARCHAR
Citext ^{3, 4}	LONGVARCHAR

¹ Bit maps to -7 (BIT) when the length for the bit is 1. If the length is greater than 1, the driver maps the column to BINARY.

² You may determine how these columns are described by setting the [Data types](#) on page 16 connection property.

³ The Citext data type behaves the same as the Text data type, except that it is case-insensitive. The select operations performed on Citext columns return case-insensitive results.

⁴ Supported for Greenplum versions 5.3 and higher.

Greenplum Data Type	JDBC Data Type
Date	DATE
DATERANGE ⁵	VARCHAR
Double precision	DOUBLE
INET	VARCHAR
INT4RANGE ⁵	VARCHAR
INT8RANGE ⁵	VARCHAR
Interval	VARCHAR
Integer	INTEGER
Json ⁶	VARCHAR
Line ⁵	VARCHAR
LSEG	VARCHAR
MACADDR	VARCHAR
Money	VARCHAR
Numeric	NUMERIC
NUMRANGE ⁵	VARCHAR
Path	VARCHAR
Point	VARCHAR
Polygon	VARCHAR
Real	REAL
Serial	INTEGER
Smallint	SMALLINT
Text	LONGVARCHAR
Time	TIMESTAMP
Time with time zone	TIMESTAMP

⁵ Supported for Greenplum versions 6.0 and higher.

⁶ Supported for Greenplum versions 5.0 and higher.

Greenplum Data Type	JDBC Data Type
Timestamp	TIMESTAMP
Timestamp with time zone	TIMESTAMP
TSRANGE ⁵	VARCHAR
TSTZRANGE ⁵	VARCHAR
TXID_SNAPSHOT ⁶	VARCHAR
UUID ⁶	VARCHAR

getTypeInfo()

The following table provides getTypeInfo() results for Greenplum databases supported by the driver.

Table 2: getTypeInfo() for Greenplum

<p>TYPE_NAME = bigint</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = -5 (BIGINT) FIXED_PREC_SCALE = false LITERAL_PREFIX = NULL LITERAL_SUFFIX = NULL LOCAL_TYPE_NAME = Bigint MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 19 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = bigserial</p> <p>AUTO_INCREMENT = true CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = -5 (BIGINT) FIXED_PREC_SCALE = false LITERAL_PREFIX = NULL LITERAL_SUFFIX = NULL LOCAL_TYPE_NAME = Bigserial MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 0 NUM_PREC_RADIX = 10 PRECISION = 19 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

<p>TYPE_NAME = bit⁷</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = -2 (BINARY) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Bit MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 83886080 SEARCHABLE = 2 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = bit varying</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>max length</i> DATA_TYPE = -3 (VARBINARY) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Bit varying MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 83886080 SEARCHABLE = 2 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = boolean</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 16 (BOOLEAN) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Boolean MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 1 SEARCHABLE = 2 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

⁷ Bit maps to -7 (BIT) when the length for the bit is 1. If the length is greater than 1, the driver maps the column to BINARY.

<p>TYPE_NAME = box</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Box MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10485760 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>
<p>TYPE_NAME = bytea</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = -4 (LONGVARBINARY) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Bytea MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 2147483647 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = character</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>length</i> DATA_TYPE = 1 (CHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Character MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 10485760 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

<p>TYPE_NAME = character varying⁸</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>max length</i> DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Character varying MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 10485760 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = CIDR</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = INET MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 128 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = circle</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Circle MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10485760 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>

⁸ Columns of this type will be described as VARCHAR when precision is 4000 or less. If precision is greater than 4000, columns will be described as LONGVARCHAR.

<p>TYPE_NAME = citext^{9, 10}</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>length</i> DATA_TYPE = -1 (LONGVARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Citext MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 1073741823 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = date</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 91 (DATE) FIXED_PREC_SCALE = false LITERAL_PREFIX = {d' LITERAL_SUFFIX = }' LOCAL_TYPE_NAME = DATE MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 10 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = daterange</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = DATERANGE MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 25 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>

⁹ The Citext data type behaves the same as the Text data type, except that it is case-insensitive. The select operations performed on Citext columns return case-insensitive results.

¹⁰ Supported for Greenplum versions 5.3 and higher.

<p>TYPE_NAME = double precision</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 8 (DOUBLE) FIXED_PREC_SCALE = false LITERAL_PREFIX = NULL LITERAL_SUFFIX = NULL LOCAL_TYPE_NAME = Double precision MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = 2 PRECISION = 53 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = INET</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = INET MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 128 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = integer</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 4 (INTEGER) FIXED_PREC_SCALE = false LITERAL_PREFIX = NULL LITERAL_SUFFIX = NULL LOCAL_TYPE_NAME = Integer MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

<p>TYPE_NAME = interval</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Interval MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10485760 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>
<p>TYPE_NAME = INT4RANGE</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = INT4RANGE MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 38 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>
<p>TYPE_NAME = INT8RANGE</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = INT8RANGE MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 56 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>

<p>TYPE_NAME = json¹¹</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = json MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 2147483647 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = line</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Line MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10485760 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>
<p>TYPE_NAME = LSEG</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = LSEG MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10485760 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>

¹¹ Supported for Greenplum versions 5.0 and higher.

<p>TYPE_NAME = MACADDR</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = MACADDR MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 17 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = money¹²</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = true LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Money MAXIMUM_SCALE = 2</p>	<p>MINIMUM_SCALE = 2 NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 27 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = numeric</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>precision, scale</i> DATA_TYPE = 2 (NUMERIC) FIXED_PREC_SCALE = false LITERAL_PREFIX = NULL LITERAL_SUFFIX = NULL LOCAL_TYPE_NAME = Numeric MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 999 NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 1000 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

¹² For Greenplum version 5.0 and higher, the precision is 27. For Greenplum versions before 5.0, the precision is 15.

<p>TYPE_NAME = NUMRANGE</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = NUMRANGE MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 1350 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>
<p>TYPE_NAME = path</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Path MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10485760 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>
<p>TYPE_NAME = point</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Point MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10485760 SEARCHABLE = 0 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>

<p>TYPE_NAME = polygon</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Polygon MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 10485760 SEARCHABLE = 0 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>
<p>TYPE_NAME = real</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = <i>precision</i> DATA_TYPE = 7 (REAL) FIXED_PREC_SCALE = false LITERAL_PREFIX = NULL LITERAL_SUFFIX = NULL LOCAL_TYPE_NAME = Real MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = 2 PRECISION = 24 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = serial</p> <p>AUTO_INCREMENT = true CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 4 (INTEGER) FIXED_PREC_SCALE = false LITERAL_PREFIX = NULL LITERAL_SUFFIX = NULL LOCAL_TYPE_NAME = Serial MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 0 NUM_PREC_RADIX = 10 PRECISION = 10 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

<p>TYPE_NAME = smallint</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 5 (SMALLINT) FIXED_PREC_SCALE = false LITERAL_PREFIX = NULL LITERAL_SUFFIX = NULL LOCAL_TYPE_NAME = Smallint MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 5 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = text</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>length</i> DATA_TYPE = -1 (LONGVARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = Text MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 1073741823 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = time</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>fractional_seconds_precision</i> DATA_TYPE = 93 (TIME) FIXED_PREC_SCALE = false LITERAL_PREFIX = '{' LITERAL_SUFFIX = '}' LOCAL_TYPE_NAME = Time MAXIMUM_SCALE = 6</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 15 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

<p>TYPE_NAME = time with time zone</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>fractional_seconds_precision</i> DATA_TYPE = 93 (TIMESTAMP) FIXED_PREC_SCALE = false LITERAL_PREFIX = {t' LITERAL_SUFFIX = }' LOCAL_TYPE_NAME = Time with time zone MAXIMUM_SCALE = 6</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 22 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = timestamp</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>fractional_seconds_precision</i> DATA_TYPE = 93 (TIMESTAMP) FIXED_PREC_SCALE = false LITERAL_PREFIX = {ts' LITERAL_SUFFIX = }' LOCAL_TYPE_NAME = Timestamp MAXIMUM_SCALE = 6</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 26 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = timestamp with time zone</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = <i>fractional_seconds_precision</i> DATA_TYPE = 93 (TIMESTAMP) FIXED_PREC_SCALE = false LITERAL_PREFIX = {ts' LITERAL_SUFFIX = }' LOCAL_TYPE_NAME = Timestamp with time zone MAXIMUM_SCALE = 6</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 33 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

<p>TYPE_NAME = TSRANGE</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = TSRANGE MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 57 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>
<p>TYPE_NAME = TSTZRANGE</p> <p>AUTO_INCREMENT = false CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = TSTZRANGE MAXIMUM_SCALE = 0</p>	<p>MINIMUM_SCALE = 0 NULLABLE = 1 NUM_PREC_RADIX = 10 PRECISION = 65 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = true</p>

<p>TYPE_NAME = TXID_SNAPSHOT</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = true CREATE_PARAMS = NULL DATA_TYPE = 12 (VARCHAR) FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = TXID_SNAPSHOT MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 10485760 SEARCHABLE = 0 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>
<p>TYPE_NAME = UUID¹¹</p> <p>AUTO_INCREMENT = NULL CASE_SENSITIVE = false CREATE_PARAMS = NULL DATA_TYPE = 12 FIXED_PREC_SCALE = false LITERAL_PREFIX = ' LITERAL_SUFFIX = ' LOCAL_TYPE_NAME = uuid MAXIMUM_SCALE = NULL</p>	<p>MINIMUM_SCALE = NULL NULLABLE = 1 NUM_PREC_RADIX = NULL PRECISION = 36 SEARCHABLE = 3 SQL_DATA_TYPE = NULL SQL_DATETIME_SUB = NULL UNSIGNED_ATTRIBUTE = NULL</p>

SQL escape sequences

The driver supports the following SQL escape sequences.

- Date, Time, and Timestamp Escape Sequences
- Scalar Functions
- Outer Join Escape Sequences
- LIKE Escape Character Sequence for Wildcards

Refer to "SQL escape sequences" in the *Progress DataDirect for JDBC Drivers Reference* for information about SQL escape sequences.

Supported scalar functions

You can use scalar functions in SQL statements with the following syntax:

{fn *scalar-function*}

where:

scalar-function

is a scalar function supported by the driver, as listed in the following table.

Example:

```
SELECT id, name FROM emp WHERE name LIKE {fn UCASE('Smith')}
```

Table 3: Supported Scalar Functions

String Functions	Numeric Functions	Timedate Functions	System Functions
ASCII	ABS	CURDATE	USERNAME
BIT_LENGTH	ACOS	CURRENT_DATE	DBNAME
CHAR	ASIN	CURRENT_TIME	IFNULL
CHAR_LENGTH	ATAN	CURRENT_TIMESTAMP	
CHARACTER_LENGTH	ATAN2	CURTIME	
CONCAT	CEILING	EXTRACT	
LCASE	COS	NOW	
LENGTH	COT		
LOCATE	DEGREES		
LTRIM	EXP		
OCTET_LENGTH	FLOOR		
POSITION	LOG		
REPEAT	LOG10		
REPLACE	MOD		
RIGHT	PI		
RTRIM	POWER		
SUBSTRING	RADIANS		
UCASE	RAND		
	ROUND		
	SIGN		
	SIN		
	SQRT		
	TAN		
	TRUNCATE		

DataDirect tools

Progress DataDirect for JDBC drivers install the set of tools described in this section. For detailed instructions on using these tools, refer to the corresponding topics in the *Progress DataDirect for JDBC Drivers Reference*.

- DataDirect Test allows you to test your JDBC driver and learn the JDBC API.
- DataDirect Connection Pool Manager allows you to pool connections when accessing databases. When your applications use connection pooling, connections are reused rather than created each time a connection is requested. Because establishing a connection is among the most costly operations an application may perform, using Connection Pool Manager to implement connection pooling can significantly improve performance.
- Statement Pool Monitor loads statements into and remove statements from the statement pool as well as generate information to help you troubleshoot statement pooling performance.
- DataDirect Spy logs detailed information about calls your driver makes that can be used for troubleshooting.

Troubleshooting

The *Progress DataDirect for JDBC Drivers Reference* provides information on troubleshooting problems should they occur. Refer to the "Troubleshooting" section in the *Reference* for details.

Additional information

In addition to the content provided in this guide, the documentation set also contains detailed conceptual and reference information that applies to all the drivers. For more information in these topics, refer the *Progress DataDirect for JDBC Drivers Reference* or use the links below to view some common topics:

- "JDBC support" describes support for JDBC interfaces and methods for the Progress DataDirect for JDBC drivers.
- "JDBC extensions" describes the JDBC extensions provided by the `com.ddtek.jdbc.extensions` package.
- "SQL escape sequences for JDBC" provides an overview of SQL escape sequences for JDBC. In addition, it documents the scalar functions that you use in SQL statements.
- "Security best practices for JDBC applications" describes the security best practices you should employ when developing and deploying your application with the driver.

Contacting Technical Support

Progress DataDirect offers a variety of options to meet your support needs. Please visit our Web site for more details and for contact information:

<https://www.progress.com/support>

The Progress DataDirect Web site provides the latest support information through our global service network. The SupportLink program provides access to support contact details, tools, patches, and valuable information, including a list of FAQs for each product. In addition, you can search our Knowledgebase for technical bulletins and other information.

When you contact us for assistance, please provide the following information:

- Your number or the serial number that corresponds to the product for which you are seeking support, or a case number if you have been provided one for your issue. If you do not have a SupportLink contract, the SupportLink representative assisting you will connect you with our Sales team.
- Your name, phone number, email address, and organization. For a first-time call, you may be asked for full information, including location.
- The Progress DataDirect product and the version that you are using.
- The type and version of the operating system where you have installed your product.
- Any database, database version, third-party software, or other environment information required to understand the problem.
- A brief description of the problem, including, but not limited to, any error messages you have received, what steps you followed prior to the initial occurrence of the problem, any trace logs capturing the issue, and so on. Depending on the complexity of the problem, you may be asked to submit an example or reproducible application so that the issue can be re-created.
- A description of what you have attempted to resolve the issue. If you have researched your issue on Web search engines, our Knowledgebase, or have tested additional configurations, applications, or other vendor products, you will want to carefully note everything you have already attempted.
- A simple assessment of how the severity of the issue is impacting your organization.

Tutorials

The following sections guide you through using the driver to access your data with some common third-party applications. For information on installing your driver and setting the CLASSPATH, see "Installing and setting-up the driver."

For details, see the following topics:

- [Tableau](#)
- [DbVisualizer](#)

Tableau

After you have installed your driver and defined it on the CLASSPATH, you can use the driver to access your data with Tableau. Tableau is a business intelligence software program that allows you to easily create reports and visualized representations of your data. By using the driver with Tableau, you can improve performance when retrieving data while leveraging the driver's relational mapping tools.

To use the driver to access data with Tableau:

1. Navigate to the `\lib\xx` subdirectory of the Progress DataDirect installation directory; then, locate the `jar` file for your driver:

```
greenplum.jar
```

2. Copy the `.jar` file for your driver into the following directory:

```
Windows: C:\Program Files\Tableau\Drivers
```

Linux: `/opt/tableau/tableau_driver/jdbc`

3. Open Tableau. From the **Connect** menu, select **Other Databases (JDBC)**.
4. In the **Other Databases (JDBC)** dialog, provide values for the following fields; then, click **Sign In**.
 - **URL:** Copy and paste your connection URL into this field. The following examples show how to connect using either user ID/password or Kerberos authentication.

User ID/password authentication

```
jdbc:datadirect:greenplum://server1:5432;  
AuthenticationMethod=userIdPassword;User=test;  
Password=secret;
```

Note: See [Configuring the driver for user ID/password authentication](#) on page 51 for details.

Kerberos authentication

```
jdbc:datadirect:greenplum://server1:5432;  
DatabaseName=greenplumDB;AuthenticationMethod=kerberos;  
ServicePrincipalName=postgres/myserver.example.com@EXAMPLE.COM;
```

Note: See [Configuring the driver for Kerberos authentication](#) on page 52 for details.

- **Dialect:** Select **SQL92** (the default) from the drop-down box.
 - **Username:** If required by the authentication method being used, enter the user name. Alternatively, this value can be specified with the `User` property in the connection string.
 - **Password:** If required by the authentication method being used, enter the password. Alternatively, this value can be specified with the `Password` property in the connection string.
5. The **Data Source** window appears. In the **Schema** field, select the schema for the service you want to use.
 6. In the **Table** field, the tables stored in the selected schema are now exposed and available for selection.

You have successfully accessed your data and are now ready to create reports with Tableau. For detailed information, refer to the Tableau product documentation at: <https://www.tableau.com/support/help>.


DbVisualizer

After you have installed your driver and defined it on the CLASSPATH, you can use the driver to access your data with the third-party DbVisualizer tool. The following topics guide you through using DbVisualizer to add your driver, connect, and execute SQL statements.

Adding a driver

To add a driver with DbVisualizer:

1. Open DbVisualizer.
2. From the menu, select **Tools>Driver Manager**. The Driver Manager window opens.

- From the Driver Manager menu, select **Driver>Create Driver**.
- Click the  button to navigate to the location of the driver jar file; then, click **OK**. The following are the default locations for the driver:

Windows

```
C:\Program Files\Progress\DataDirect\JDBC\lib\60\greenplum.jar
```

Linux

```
/opt/Progress/DataDirect/JDBC/lib/60/greenplum.jar
```

- Provide values for the following fields; then, close the Driver Manager window.

- Name:** Type an alias for your driver. For example:

```
Greenplum
```

- URL Format:** Optionally, specify the format of the connection string for your driver. For example:

```
jdbc:datadirect:greenplum:
```

- Driver Class:** From the drop down menu, select the driver class for your driver. For example:

```
com.ddtek.jdbcx.greenplum.GreenplumDataSource
```

You can now use your driver with DbVisualizer. Proceed to "Connecting and executing SQL statements" for information on connecting and executing SQL statements.

Connecting and executing SQL statements

To use the driver to access data with DbVisualizer:

- Open DbVisualizer.
- From the menu, select **Database>New Connection**. When prompted to use the Connection Wizard, click **OK**.
- Provide the following information when prompted; then, click **Next** to proceed:
 - Connection alias:** Type the name to be used when referring to this connection.
 - Driver:** Select the alias that you provided for your driver from the drop-down menu.
- Provide values for the following fields; then, click **Finish**.
 - Database URL:** Copy and paste your connection URL into this field. The following examples show how to connect using either user ID/password or Kerberos authentication.

User ID/password authentication

```
jdbc:datadirect:greenplum://server1:5432;
AuthenticationMethod=userIdPassword;User=test;
Password=secret;
```

Note: See [Configuring the driver for user ID/password authentication](#) on page 51 for details.

Kerberos authentication

```
jdbc:datadirect:greenplum://server1:5432;  
DatabaseName=greenplumDB;AuthenticationMethod=kerberos;  
ServicePrincipalName=postgres/myserver.example.com@EXAMPLE.COM;
```

Note: See [Configuring the driver for Kerberos authentication](#) on page 52 for details.

5. To execute SQL statements, select **SQL Commander>New SQL Commander**. A SQL Commander tab opens.
6. Select values for the following fields:
 - **Database Connection:** Select connection alias you provided for the connection from the drop-down menu.
 - **Schema:** Select the schema you want to execute queries against from the drop-down menu.
7. In the SQL Commander tab, enter SQL commands you want to execute; then select **SQL Commander>Execute**. For example:

To select all of the rows from the BLOGS table:

```
SELECT * FROM BLOGS
```

To select the details of a specified company:

```
SELECT * FROM COMPANY WHERE COMPANYID = <company_id>
```

Note: If you are fetching large sets of data, you may want to limit the results using the Max Rows and Max Chars fields.

You have successfully accessed your data with DbVisualizer.

Configuring and connecting

This section provides information on how to connect to your data store using either the JDBC Driver Manager or DataDirect JDBC data sources, as well as information on how to implement and use functionality supported by the driver.

After the driver has been installed and defined on your classpath, you can connect from your application to your data in either of the following ways.

- Using the JDBC `DriverManager` by specifying the connection URL in the `DriverManager.getConnection()` method.
- Creating a JDBC data source that can be accessed through the Java Naming Directory Interface (JNDI).

For details, see the following topics:

- [Setting the classpath](#)
- [Connecting using the JDBC Driver Manager](#)
- [Connecting using data sources](#)
- [Authentication](#)
- [Data Encryption](#)
- [Performance considerations](#)

Setting the classpath

The driver must be defined on your CLASSPATH before you can connect. The CLASSPATH is the search string your Java Virtual Machine (JVM) uses to locate JDBC drivers on your computer. If the driver is not defined on your CLASSPATH, you will receive a `class not found` exception when trying to load the driver. Set your system CLASSPATH to include the driver jar file as shown, where `install_dir` is the path to your product installation directory.

```
install_dir/lib/60/greenplum.jar
```

Windows Example

```
CLASSPATH=.;C:\Program Files\Progress\DataDirect\JDBC\lib\60\greenplum.jar
```

UNIX Example

```
CLASSPATH=./opt/Progress/DataDirect/JDBC/lib/60/greenplum.jar
```

Connecting using the JDBC Driver Manager

One way to connect to a service is through the JDBC DriverManager using the `DriverManager.getConnection()` method. As the following examples show, this method specifies a string containing a connection URL.

User ID/password authentication

```
Connection conn = DriverManager.getConnection  
( "jdbc:datadirect:greenplum://server1:5432;  
  AuthenticationMethod=userIdPassword;User=test;  
  Password=secret;" );
```

Kerberos authentication

```
Connection conn = DriverManager.getConnection  
( "jdbc:datadirect:greenplum://server1:5432;  
  DatabaseName=greenplumDB;AuthenticationMethod=kerberos;  
  ServicePrincipalName=postgres/myserver.example.com@EXAMPLE.COM;" );
```

Passing the connection URL

After setting the CLASSPATH, the required connection information needs to be passed in the form of a connection URL. The following example includes the properties required for connecting with user ID/password authentication.

Connection URL Syntax

The connection URL takes the following form:

```
jdbc:datadirect:greenplum://servername:port;
AuthenticationMethod=userIdPassword;User=userID;
Password=password;[property=value[;...]];
```

where:

servername

is the IP address or name of the server to which you are connecting.

port

is the number of the TCP/IP port.

userID

specifies the user ID that is used to connect to the Greenplum database.

password

specifies a password that is used to connect to your Greenplum database.

property=value

specifies connection property settings. Multiple properties are separated by a semi-colon.

The following example connection string includes the properties required for connecting with the user ID/password authentication.

```
Connection conn = DriverManager.getConnection
("jdbc:datadirect:greenplum://server1:5432;
 AuthenticationMethod=userIdPassword;User=test;
 Password=secret;");
```

See also

[Connection property descriptions](#) on page 73

[Connection URL examples](#) on page 13

[Configuring the driver for user ID/password authentication](#) on page 51

Testing the connection

You can use DataDirect Test™ to verify your connection. The screen shots in this section were taken on a Windows system.

To test the connection from the driver to your data source, follow these steps:

1. Navigate to the installation directory. The default location is:


- Windows systems: Program Files\Progress\DataDirect\JDBC\testforjdbc
- UNIX and Linux systems: /opt/Progress/DataDirect/JDBC/testforjdbc

Note: For UNIX/Linux, if you do not have access to /opt, your home directory will be used in its place.

2. From the `testforjdbc` folder, run the platform-specific tool:


- `testforjdbc.bat` (on Windows systems)
- `testforjdbc.sh` (on UNIX and Linux systems)

The **Test for JDBC Tool** window appears:



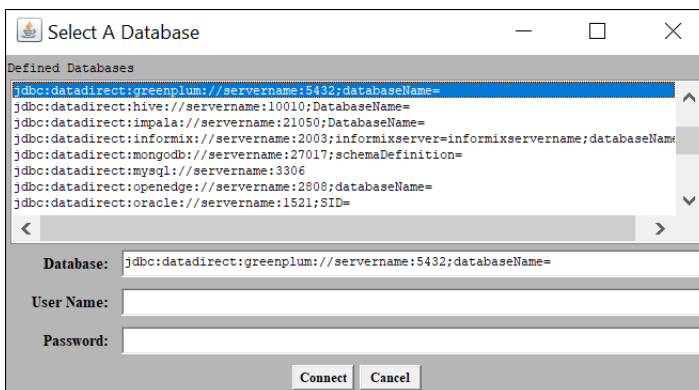
3. Click **Press Here to Continue**.

The main dialog appears:



- From the menu bar, select **Connection > Connect to DB**.

The **Select A Database** dialog appears:



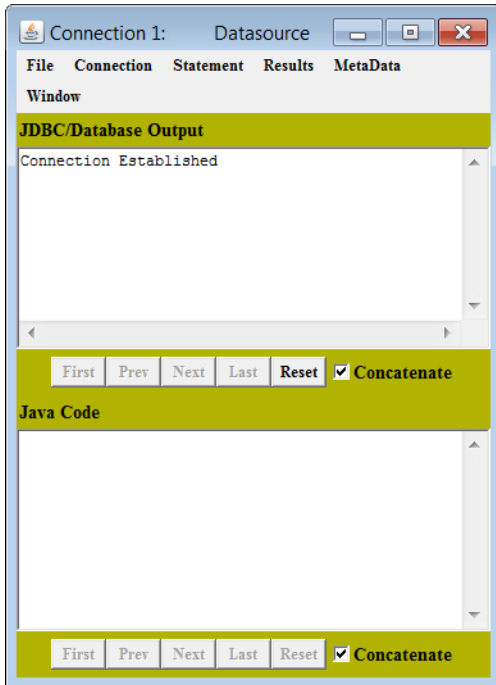
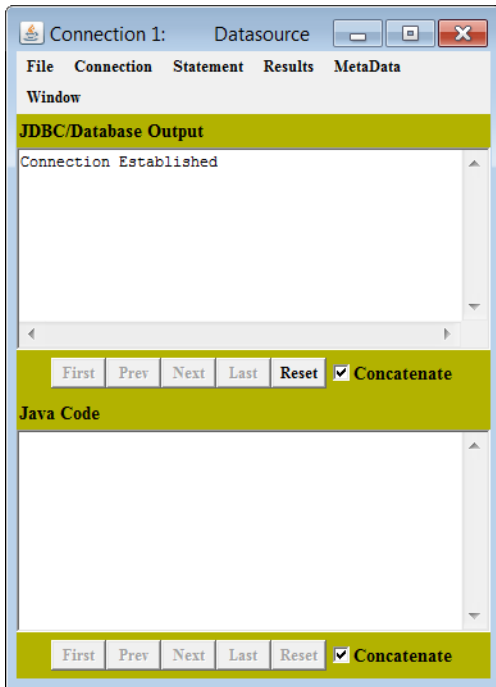
- Select the appropriate database template from the **Defined Databases** field.
- In the **Database** field, specify all required connection properties.

For example:

```
jdbc:datadirect:greenplum://server1:5432;
AuthenticationMethod=userIdPassword;User=test;
Password=secret;
```

- Click **Connect**.

If the connection information is entered correctly, the **JDBC/Database Output** window reports that a connection has been established. (If a connection is not established, the window reports an error.)



Refer to "DataDirect Test" in the *Progress DataDirect for JDBC Drivers Reference* for more information about using DataDirect Test.

Connecting using data sources

A *JDBC data source* is a Java object, specifically a `DataSource` object, that defines connection information required for a JDBC driver to connect to the database. Each JDBC driver vendor provides their own data source implementation for this purpose. A Progress DataDirect data source is Progress DataDirect's implementation of a `DataSource` object that provides the connection information needed for the driver to connect to a database.

Because data sources work with the Java Naming Directory Interface (JNDI) naming service, data sources can be created and managed separately from the applications that use them. Because the connection information is defined outside of the application, the effort to reconfigure your infrastructure when a change is made is minimized. For example, if the database is moved to another database server, the administrator need only change the relevant properties of the `DataSource` object. The applications using the database do not need to change because they only refer to the name of the data source.

How data sources are implemented

Data sources are implemented through a data source class. A data source class implements the following interfaces.

- `javax.sql.DataSource`
- `javax.sql.ConnectionPoolDataSource` (allows applications to use connection pooling)

Refer to "Connection Pool Manager" in the *Progress DataDirect for JDBC Drivers Reference* for more information.

See also

[Driver and DataSource classes](#) on page 13

Creating data sources

The following example files provide details on creating and using Progress DataDirect data sources with the Java Naming Directory Interface (JNDI), where `install_dir` is the product installation directory.

- `install_dir/Examples/JNDI/JNDI_LDAP_Example.java` can be used to create a JDBC data source and save it in your LDAP directory using the JNDI Provider for LDAP.
- `install_dir/Examples/JNDI/JNDI_FILESYSTEM_Example.java` can be used to create a JDBC data source and save it in your local file system using the File System JNDI Provider.

See "Example data source" for an example data source definition for the example files.

To connect using a JNDI data source, the driver needs to access a JNDI data store to persist the data source information. For a JNDI file system implementation, you must download the File System Service Provider from the [Oracle Technology Network Java SE Support downloads page](#), unzip the files to an appropriate location, and add the `fscontext.jar` and `providerutil.jar` files to your CLASSPATH. These steps are not required for LDAP implementations because the LDAP Service Provider is included with supported versions of Java SE.

Example data source

To configure a data source using the example files, you will need to create a data source definition. The content required to create a data source definition is divided into three sections.

First, you will need to import the data source class. For example:

```
import com.ddtek.jdbcx.greenplum.GreenplumDataSource;
```

Next, you will need to set the values and define the data source. For example, the following definition contains the minimum properties required to establish connection:

Note:

- Setting the password using a data source is generally not recommended. The data source persists all properties, including the Password property, in clear text.
 - In a JDBC data source, string values must be enclosed in double quotation marks, for example, `setUser("abc@defcorp.com")`.
-

```
GreenplumDataSource mds = new GreenplumDataSource();
mds.setDescription("My Greenplum Data Source");
mds.AuthenticationMethod("useridpassword");
mds.setServerName("server1");
mds.setPortNumber(5432);
mds.setUser("User123");
mds.setPassword("secret");
```

Finally, you will need to configure the example application to print out the data source attributes. Note that this code is specific to the driver and should only be used in the example application. For example, you would add the following section for the minimum properties required to establish a connection:

```
if (ds instanceof GreenplumDataSource)
{
    GreenplumDataSource jmds = (GreenplumDataSource) ds;
    System.out.println("description=" + jmds.getDescription());
    System.out.println("authenticationmethod=" + jmds.getAuthenticationMethod());
    System.out.println("serverName=" + jmds.getServerName());
    System.out.println("portNumber=" + jmds.getPortNumber());
    System.out.println("user=" + jmds.getUser());
    System.out.println("password=" + jmds.getPassword());
    System.out.println();
}
```

Calling a data source in an application

Applications can call a Progress DataDirect data source using a logical name to retrieve the `javax.sql.DataSource` object. This object loads the specified driver and can be used to establish a connection to the database.

Once the data source has been registered with JNDI, it can be used by your JDBC application as shown in the following code example.

```
Context ctx = new InitialContext();
DataSource ds = (DataSource)ctx.lookup("EmployeeDB");
Connection con = ds.getConnection("domino", "spark");
```

In this example, the JNDI environment is first initialized. Next, the initial naming context is used to find the logical name of the data source (`EmployeeDB`). The `Context.lookup()` method returns a reference to a Java object, which is narrowed to a `javax.sql.DataSource` object. Then, the `DataSource.getConnection()` method is called to establish a connection.

Testing a data source connection

You can use DataDirect Test™ to establish and test a data source connection. The screen shots in this section were taken on a Windows system.

Take the following steps to establish a connection.

1. Navigate to the installation directory. The default location is:

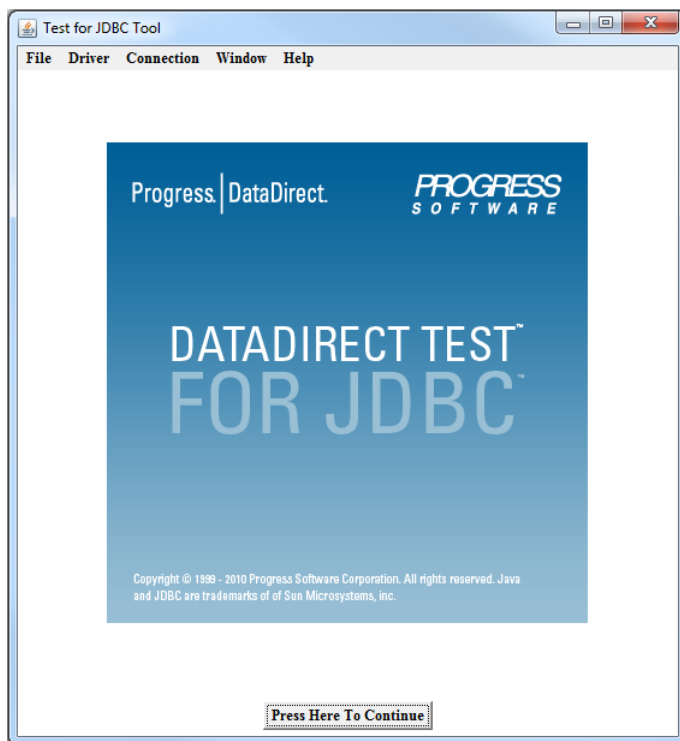
- Windows systems: `Program Files\Progress\DataDirect\JDBC\testforjdbc`
- UNIX and Linux systems: `/opt/Progress/DataDirect/JDBC/testforjdbc`

Note: For UNIX/Linux, if you do not have access to `/opt`, your home directory will be used in its place.

2. From the `testforjdbc` folder, run the platform-specific tool:


- `testforjdbc.bat` (on Windows systems)
- `testforjdbc.sh` (on UNIX and Linux systems)

The **Test for JDBC Tool** window appears:

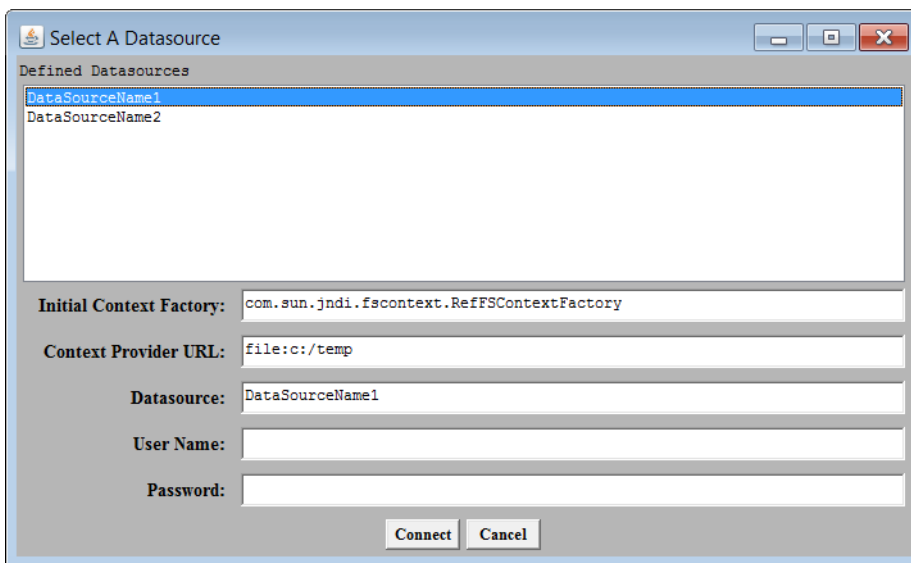


3. Click **Press Here to Continue**.

The main dialog appears:




4. From the menu bar, select **Connection > Connect to DB via Data Source**.
The **Select A Database** dialog appears:



5. Select a datasource template from the **Defined Datasources** field.
6. Provide the following information:
 - a) In the **Initial Context Factory**, specify the location of the initial context provider for your application.
 - b) In the **Context Provider URL**, specify the location of the context provider for your application.
 - c) In the **Datasource** field, specify the name of your datasource.
7. If you are using user ID/password authentication, enter your user ID and password in the corresponding fields.
8. Click **Connect**.

If the connection information is entered correctly, the **JDBC/Database Output** window reports that a connection has been established. If a connection is not established, the window reports an error.



Authentication

The driver supports the following authentication methods:

- *User ID/password authentication* authenticates using a database user name and password provided by the application.
- *Kerberos authentication* uses Kerberos, a trusted third-party authentication service, to verify user identities. Kerberos authentication can take advantage of the user name and password maintained by the operating system to authenticate users to the database or use another set of user credentials specified by the application.

This method requires knowledge of how to configure your Kerberos environment and supports Windows Active Directory Kerberos only.

By default, the driver is configured to use user ID/password authentication (`AuthenticationMethod=userIdPassword`).

See also

[AuthenticationMethod](#) on page 83

Configuring the driver for user ID/password authentication

Take the following steps to configure user ID/Password authentication.

1. Set the `AuthenticationMethod` property to `userIdPassword`.
2. Set the `User` property to provide the user ID.
3. Set the `Password` property to provide the password.
4. Specify values for minimum required properties for establishing a connection.
 - a) Set the `ServerName` property to specify either the IP address in IPv4 or IPv6 format, or the server name for your server.
 - b) Set the `PortNumber` property to specify the TCP port of the primary database server that is listening for connections to the database.

For example, the following is a connection string with only the required properties for making a connection using user ID/password authentication.

```
Connection conn = DriverManager.getConnection
    ("jdbc:datadirect:greenplum://server1:5432;
    AuthenticationMethod=userIdPassword;User=test;
    Password=secret");
```

See also

[AuthenticationMethod](#) on page 83

[User](#) on page 119

[Password](#) on page 106

Configuring the driver for Kerberos authentication

Your Kerberos environment should be fully configured before you configure the driver for Kerberos authentication. You should refer to your Greenplum documentation and Java documentation for instructions on configuring Kerberos. For a Windows Active Directory implementation, you should also consult your Windows documentation. For a non-Active Directory implementation (on a Windows or non-Windows operating system), you should consult MIT Kerberos documentation.

Important: A properly configured Kerberos environment must include a means of obtaining a Kerberos Ticket Granting Ticket (TGT). For a Windows Active Directory implementation, Active Directory automatically obtains the TGT. However, for a non-Active Directory implementation, the means of obtaining the TGT must be automated or handled manually.

Once your Kerberos environment has been configured, take the following steps to configure the driver.

1. Use one of the following methods to integrate the JAAS configuration file into your Kerberos environment. (See "The JAAS login configuration file" for details.)

Note: The `install_dir/lib/JDBCdriverLogin.conf` file is the JAAS login configuration file installed with the driver. You can use this file or another file as your JAAS login configuration file.

Note: Regardless of operating system, forward slashes must be used when designating the path of the JAAS login configuration file.

Option 1. Specify a login configuration file directly in your application with the

`java.security.auth.login.config` system property. For example:

```
System.setProperty("java.security.auth.login.config","install_dir/lib/JDBCdriverLogin.conf");
```

Option 2. Set up a default configuration. Modify the Java security properties file to indicate the URL of the login configuration file with the `login.config.url.n` property where *n* is an integer connoting separate, consecutive login configuration files. When more than one login configuration file is specified, then the files are read and concatenated into a single configuration.

- a) Open the Java security properties file. The security properties file is the `java.security` file in the `/jre/lib/security` directory of your Java installation.
- b) Find the line `# Default login configuration file` in the security properties file.
- c) Below the `# Default login configuration file` line, add the URL of the login configuration file as the value for a `login.config.url.n` property. For example:

```
# Default login configuration file
login.config.url.1=file:${user.home}/.java.login.config
login.config.url.2=file:jdbc_driver_install_dir/lib/JDBCdriverLogin.conf
```

2. Ensure your JAAS login configuration file includes an entry with authentication technology that the driver can use to establish a Kerberos connection. (See "The JAAS login configuration file" for details.)

Note: The JAAS login configuration file installed with the driver (`install_dir/lib/JDBCdriverLogin.conf`) includes a default entry with the name `JDBC_DRIVER_01`. This entry specifies the Kerberos authentication technology used with an Oracle JVM.

The following examples show that the authentication technology used in a Kerberos environment depends on your JVM.

Oracle JVM

```
JDBC_DRIVER_01 {
    com.sun.security.auth.module.Krb5LoginModule required useTicketCache=true;
};
```

IBM JVM

```
JDBC_DRIVER_01 {
    com.ibm.security.auth.module.Krb5LoginModule required useDefaultCcache=true;
};
```

3. Set the driver's `AuthenticationMethod` connection property to `kerberos`. (See "AuthenticationMethod" for details.)
4. Set the `ServicePrincipalName` connection property to the service principal name registered with the KDC.

The `ServicePrincipalName` takes the following form.

```
Service_Name/Fully_Qualified_Domain_Name@REALM.COM
```

The value of the `ServicePrincipalName` property can include the Kerberos realm name, but it is optional. If you do not specify the realm name, the default realm is used.

See "ServicePrincipalName" for details on the composition of the service principal name.

- Optionally, set the LoginConfigName connection property if the name of the JAAS login configuration file entry is different from the driver default `JDBC_DRIVER_01`. (See "The JAAS login configuration file" and "LoginConfigName" for details.)

`JDBC_DRIVER_01` is the default entry name for the JAAS login configuration file (`JDBCDriverLogin.conf`) installed with the driver. When configuring your Kerberos environment, your network or system administrator may have used a different entry name. Check with your administrator to verify the correct entry name.

See also

[Kerberos authentication requirements](#) on page 54

[The JAAS login configuration file](#) on page 55

[AuthenticationMethod](#) on page 83

[ServicePrincipalName](#) on page 114

[LoginConfigName](#) on page 101

Kerberos authentication requirements

Verify that your environment meets the requirements listed in the following table before you configure the driver for Kerberos authentication.

Note: For Windows Active Directory, the domain controller must administer both the database server and the client.

Table 4: Kerberos configuration requirements

Component	Requirements
Database server	No restrictions
Kerberos server	<p>The Kerberos server is the machine where the user IDs for authentication are administered. The Kerberos server is also the location of the Kerberos key distribution center (KDC). Network authentication must be provided by one of the following methods.</p> <ul style="list-style-type: none"> Windows Active Directory on one of the following operating systems: <ul style="list-style-type: none"> Windows Server 2003 or higher Windows 2000 Server Service Pack 3 or higher MIT Kerberos 1.5 or higher
Client	Java SE 8 or higher must be installed.

See also

[Configuring the driver for Kerberos authentication](#) on page 52

The JAAS login configuration file

The Java Authentication and Authorization Service (JAAS) login configuration file contains one or more entries that specify authentication technologies to be used by applications. To establish Kerberos connections with the driver, the JAAS login configuration file must include an entry specifically for the driver. In addition, the login configuration file must be referenced either by setting the `java.security.auth.login.config` system property or by setting up a default configuration using the Java security properties file.

Setting up a default configuration

To set up a default configuration, you must modify the Java security properties file to indicate the URL of the login configuration file with the `login.config.url.n` property where `n` is an integer connoting separate, consecutive login configuration files. When more than one login configuration file is specified, then the files are read and concatenated into a single configuration. The following steps summarize how to modify the security properties file.

1. Open the Java security properties file. The security properties file is the `java.security` file in the `/jre/lib/security` directory of your Java installation.
2. Find the line `# Default login configuration file` in the security properties file.
3. Below the `# Default login configuration file` line, add the URL of the login configuration file as the value for a `login.config.url.n` property. For example:

```
# Default login configuration file
login.config.url.1=file:${user.home}/.java.login.config
login.config.url.2=file:jdbc_driver_install_dir/lib/JDBCdriverLogin.conf
```

JAAS login configuration file entry for the driver

You can create your own JAAS login configuration file, or you can use the `JDBCdriverLogin.conf` file installed in the `/lib` directory of the product installation directory. In either case, the login configuration file must include an entry that specifies the Kerberos authentication technology to be used by the driver.

JAAS login configuration file entries begin with an entry name followed by one or more `LoginModule` items. Each `LoginModule` item contains information that is passed to the `LoginModule`. A login configuration file entry takes the following form.

```
entry_name {
    login_module flag_value module_options
};
```

where:

entry_name

is the name of the login configuration file entry. The driver's `LoginConfigName` connection property can be used to specify the name of this entry. `JDBC_DRIVER_01` is the default entry name for the `JDBCDriverLogin.conf` file installed with the driver.

login_module

is the fully qualified class name of the authentication technology used with the driver.

flag_value

specifies whether the success of the module is `required`, `requisite`, `sufficient`, or `optional`.

module_options

specifies available options for the `LoginModule`. These options vary depending on the `LoginModule` being used.

The following examples show that the `LoginModule` used for a Kerberos implementation depends on your JVM.

Oracle JVM

```
JDBC_DRIVER_01 {
    com.sun.security.auth.module.Krb5LoginModule required useTicketCache=true;
};
```

IBM JVM

```
JDBC_DRIVER_01 {
    com.ibm.security.auth.module.Krb5LoginModule required useDefaultCcache=true;
};
```

Refer to Java Authentication and Authorization Service documentation for information about the JAAS login configuration file and implementing authentication technologies.

See also

[Configuring the driver for Kerberos authentication](#) on page 52

[LoginConfigName](#) on page 101

Data Encryption

TLS/SSL works by allowing the client and server to send each other encrypted data that only they can decrypt. TLS/SSL negotiates the terms of the encryption in a sequence of events known as the *handshake*. The handshake involves the following types of authentication:

- *TLS/SSL server authentication* requires the server to authenticate itself to the client.
- *TLS/SSL client authentication* is optional and requires the client to authenticate itself to the server after the server has authenticated itself to the client.

Configuring TLS/SSL encryption

The following steps outline how to configure TLS/SSL encryption.

Note: Connection hangs can occur when the driver is configured for TLS/SSL and the database server does not support TLS/SSL. You may want to set a login timeout using the LoginTimeout property to avoid problems when connecting to a server that does not support TLS/SSL.

To configure SSL encryption:

Important: The driver complies with FIPS when FIPS mode is enabled with the client JVM. See "FIPS (Federal Information Processing Standard)" for more information.

- Set general properties required for connection:
 - Set the ServerName property to the name or the IP address of the server to which you want to connect. For example, `myserver`.
 - Set the PortNumber property to specify the port number of the server listener. The default is 5432.
 - Set properties related to your authentication method. See "Authentication" for details.
- Set the EncryptionMethod property to `SSL`.
- (Optional) Set the CryptoProtocolVersion property to specify acceptable cryptographic protocol versions (for example, TLSv1.3) supported by your server.
- (Optional) Specify the location and password of the truststore file used for SSL server authentication. Either set the TrustStore and TrustStorePassword properties or their corresponding Java system properties (`javax.net.ssl.trustStore` and `javax.net.ssl.trustStorePassword`, respectively).
- (Optional) To validate certificates sent by the database server, set the ValidateServerCertificate property to `true`.
- (Optional) Set the HostNameInCertificate property to a host name to be used to validate the certificate. The HostNameInCertificate property provides additional security against man-in-the-middle (MITM) attacks by ensuring that the server the driver is connecting to is the server that was requested.
- (Optional) If your database server is configured for SSL client authentication, configure your keystore information:
 - Specify the location and password of the keystore file. Either set the KeyStore and KeyStorePassword properties or their corresponding Java system properties (`javax.net.ssl.keyStore` and `javax.net.ssl.keyStorePassword`, respectively).
 - If any key entry in the keystore file is password-protected, set the KeyPassword property to the key password.

The following examples demonstrate the required properties for a session using TLS/SSL encryption with user ID and password authentication.

For a connection URL:

```
Connection conn = DriverManager.getConnection
("jdbc:datadirect:greenplum://myserver:5432;
EncryptionMethod=SSL";User=jsmith;Password=secret);
```

For a data source:

```
GreenplumDataSource mds = new GreenplumDataSource();
mds.setDescription("My Greenplum Data Source");
mds.setEncryptionMethod("SSL");
mds.setServerName("myserver");
mds.setUser("jsmith");
mds.setServerName("secret");
```

See also

[CryptoProtocolVersion](#) on page 89

[KeyStore](#) on page 99

[KeyStorePassword](#) on page 99

[LoginTimeout](#) on page 102

[EncryptionMethod](#) on page 92

[TrustStore](#) on page 118

[TrustStorePassword](#) on page 118

[HostNameInCertificate](#) on page 94

[ValidateServerCertificate](#) on page 120

[KeyPassword](#) on page 98

[FIPS \(Federal Information Processing Standard\)](#) on page 60

[Authentication](#) on page 51

[Connection property descriptions](#) on page 73

[Configuring TLS/SSL Server Authentication](#) on page 58

[Configuring TLS/SSL Client Authentication](#) on page 59

Configuring TLS/SSL Server Authentication

When the client makes a connection request, the server presents its public certificate for the client to accept or deny. The client checks the issuer of the certificate against a list of trusted Certificate Authorities (CAs) that resides in an encrypted file on the client known as a *truststore*. Optionally, the client may check the subject (owner) of the certificate. If the certificate matches a trusted CA in the truststore (and the certificate's subject matches the value that the application expects), an encrypted connection is established between the client and server. If the certificate does not match, the connection fails and the driver throws an exception.

To check the issuer of the certificate against the contents of the truststore, the driver must be able to locate the truststore and unlock the truststore with the appropriate password. You can specify truststore information in either of the following ways:

- Specify values for the Java system properties `javax.net.ssl.trustStore` and `javax.net.ssl.trustStorePassword`. For example:

```
java -Djavax.net.ssl.trustStore=C:\Certificates\MyTruststore
-Djavax.net.ssl.trustStorePassword=MyTruststorePassword
```

This method sets values for all TLS/SSL sockets created in the JVM.

- Specify values for the connection properties `TrustStore` and `TrustStorePassword` in the connection URL. For example:

```
TrustStore=C:\Certificates\MyTruststore
```

and

```
TrustStorePassword=MyTruststorePassword
```

Any values specified by the `TrustStore` and `TrustStorePassword` properties override values specified by the Java system properties. This allows you to choose which truststore file you want to use for a particular connection.

Alternatively, you can configure the drivers to trust any certificate sent by the server, even if the issuer is not a trusted CA. Allowing a driver to trust any certificate sent from the server is useful in test environments because it eliminates the need to specify truststore information on each client in the test environment. If the driver is configured to trust any certificate sent from the server, the issuer information in the certificate is ignored.

Configuring TLS/SSL Client Authentication

If the server is configured for TLS/SSL client authentication, the server asks the client to verify its identity after the server has proved its identity. Similar to TLS/SSL server authentication, the client sends a public certificate to the server to accept or deny. The client stores its public certificate in an encrypted file known as a *keystore*.

The driver must be able to locate the keystore and unlock the keystore with the appropriate keystore password. Depending on the type of keystore used, the driver also may need to unlock the keystore entry with a password to gain access to the certificate and its private key.

The drivers can use the following types of keystores:

- Java Keystore (JKS) contains a collection of certificates. Each entry is identified by an alias. The value of each entry is a certificate and the certificate's private key. Each keystore entry can have the same password as the keystore password or a different password. If a keystore entry has a password different than the keystore password, the driver must provide this password to unlock the entry and gain access to the certificate and its private key.
- PKCS #12 keystores. To gain access to the certificate and its private key, the driver must provide the keystore password. The file extension of the keystore must be `.pfx` or `.p12`.

You can specify this information in either of the following ways:

- Specify values for the Java system properties `javax.net.ssl.keyStore` and `javax.net.ssl.keyStorePassword`. For example:

```
java -Djavax.net.ssl.keyStore=C:\Certificates\MyKeystore
-Djavax.net.ssl.keyStorePassword=MyKeystorePassword
```

This method sets values for all TLS/SSL sockets created in the JVM.

Note: If the keystore specified by the `javax.net.ssl.keyStore` Java system property is a JKS and the keystore entry has a password different than the keystore password, the `KeyPassword` connection property must specify the password of the keystore entry (for example, `KeyPassword=MyKeyPassword`).

- Specify values for the connection properties `KeyStore` and `KeyStorePassword` in the connection URL. For example:

```
KeyStore=C:\Certificates\MyKeyStore
and
KeyStorePassword=MyKeystorePassword
```

Note: If the keystore specified by the `KeyStore` connection property is a JKS and the keystore entry has a password different than the keystore password, the `KeyPassword` connection property must specify the password of the keystore entry (for example, `KeyPassword=MyKeyPassword`).

Any values specified by the `KeyStore` and `KeyStorePassword` properties override values specified by the Java system properties. This allows you to choose which keystore file you want to use for a particular connection.

FIPS (Federal Information Processing Standard)

The Federal Information Processing Standard (or FIPS) is a cryptography standard created by the U.S. government. FIPS specifications require certain secure algorithms, cryptographic modules, and random number generation. The driver is FIPS compliant for data encryption when FIPS is enabled for the JVM on the client machine.

The following applies when the driver is running in a FIPS environment:

- The driver complies with 140-3 and 140-2 standards.
- The driver uses PKCS #11 providers to access keystores.

The driver was tested with FIPS 140-3 enabled using Red Hat OpenJDK 21 on a Red Hat Universal Base Image 9 instance.

See also

[Configuring TLS/SSL Server Authentication](#) on page 58

[Configuring TLS/SSL Client Authentication](#) on page 59

Performance considerations

You can optimize your applications performance by adopting guidelines described in this section.

InsensitiveResultSetBufferSize: To improve performance, result set data can be cached instead of written to disk. If the size of the result set data is greater than the size allocated for the cache, the driver writes the result set to disk. The maximum cache size setting is 2 GB.

MaxPooledStatements: To improve performance, the driver's own internal prepared statement pooling should be enabled when the driver does not run from within an application server or from within another application that does not provide its own prepared statement pooling. When the driver's internal prepared statement pooling is enabled, the driver caches a certain number of prepared statements created by an application. For example, if the `MaxPooledStatements` property is set to 20, the driver caches the last 20 prepared statements created by the application. If the value set for this property is greater than the number of prepared statements used by the application, all prepared statements are cached.

Refer to "Designing JDBC Applications for Performance Optimization" in the *Progress DataDirect for JDBC Drivers Reference* for more information about using prepared statement pooling to optimize performance.

ResultSetMetaDataOptions: The driver's performance may be adversely affected if you set this option to 1. If set to 1 and the `ResultSetMetaData.getTableName` method is called, the driver performs emulations which take additional processing.

VarcharClobThreshold: There are performance penalties when enabling CLOB functionality. To provide the benefits associated with Clobs, data must be cached. Because data is cached, your application will incur a performance penalty, particularly if data is read once sequentially. This performance penalty can be severe if the size of the long data is larger than available memory. If you want to avoid the performance penalties associated with CLOB functionality, you should set this value at a value greater than the maximum *Character varying* column width your application handles.

Additional features and functionality

The following section describes additionally supported features and functionality that are specific to the driver.

For details, see the following topics:

- [Statement pooling](#)
- [Connection pooling](#)
- [Using client information](#)
- [Bulk load support](#)
- [Failover](#)
- [Refcursors](#)
- [Isolation levels](#)
- [Using scrollable cursors](#)
- [JTA support](#)
- [Large object \(LOB\) support](#)
- [Using COPY command](#)
- [Parameter metadata support](#)
- [ResultSet metadata support](#)
- [Rowset support](#)

Statement pooling

Most applications have a certain set of SQL statements that are executed multiple times and a few SQL statements that are executed only once or twice during the life of the application. Similar to connection pooling, *statement pooling* provides performance gains for applications that execute the same SQL statements multiple times over the life of the application.

A *statement pool* is a group of prepared statements that can be reused by an application. If you have an application that repeatedly executes the exact same SQL statements, statement pooling can improve performance because the database server does not have to repeatedly parse and create cursors for the same statement. In addition, the associated network round trips to the database server are avoided.

The drivers have an internal prepared statement pooling mechanism, which allows you to realize the performance benefits of statement pooling when you are not running from within an application server or another application that provides its own statement pooling. You can enable this driver-based internal statement pooling with the `MaxPooledStatements` connection property.

The DataDirect for JDBC drivers also support the DataDirect Statement Pool Monitor. You can use the Statement Pool Monitor to load statements into and remove statements from the statement pool as well as generate information to help you troubleshoot statement pooling performance. The Statement Pool Monitor is an integrated component of the driver, and you can manage statement pooling directly with DataDirect-specific methods. In addition, the Statement Pool Monitor can be enabled as a Java Management Extensions (JMX) MBean. When enabled as a JMX MBean, the Statement Pool Monitor can be used to manage statement pooling with standard JMX API calls, and it can easily be used by JMX-compliant tools, such as JConsole. To enable the Statement Pool Monitor as a JMX MBean, you must register the Statement Pool Monitor MBean with the `RegisterStatementPoolMonitorMBean` connection property.

Refer to "Statement Pool Monitor" in the *Progress DataDirect for JDBC Drivers Reference* for further details.

See also

[MaxPooledStatements](#) on page 104

[RegisterStatementPoolMonitorMBean](#) on page 112

Connection pooling

Progress DataDirect for JDBC drivers support connection pooling using the DataDirect Connection Pool Manager. Typically, creating a connection is the most performance-expensive operations that an application performs. Connection pooling allows you to reuse connections rather than create a new one every time a driver needs to connect to the database. Further, connection pooling manages connection sharing across multiple user requests to maintain performance and reduce the number of new connections to be created.

Refer to "Connection Pool Manager" in the *Progress DataDirect for JDBC Drivers Reference* for more information.

Using client information

Many databases allow applications to store client information associated with a connection, which can be useful for database administration and monitoring purposes. The driver allows applications to store and return the following types of client information.

- Name of the application currently using the connection.

- User ID for whom the application using the connection is performing work. The user ID may be different than the user ID that was used to establish the connection.
- Host name of the client on which the application using the connection is running.
- Product name and version of the driver on the client.
- Additional information that may be used for accounting or troubleshooting purposes, such as an accounting ID.

Refer to "Client information" in the *Progress DataDirect for JDBC Drivers Reference* for more information.

Bulk load support

As Greenplum does not have native bulk load support, the Greenplum driver emulates bulk load using the standard batch mechanism.

The BulkLoadBatchSize connection property affects how bulk load works with the Greenplum driver. It suggests the number of rows to be loaded to the database when bulk loading data. The default is 1000.

See also

[BulkLoadBatchSize](#) on page 84

Failover

The driver provides connection failover support to ensure continuous, uninterrupted access to data.

Note:

For implementation and configuration details, see [Configuring failover](#) on page 63.

Note: For more information on connection failover and different levels of failover protection provided by the driver, refer to "Failover" in the *Progress DataDirect for JDBC Drivers Reference*.

Configuring failover

The driver allows you to specify a list of alternate database servers that are tried at connection time if the primary server is not accepting connections. Connection attempts continue until a connection is successfully established or until all the database servers in the list have been tried the specified number of times.

Specifying primary and alternate servers

Connection information for primary and alternate servers can be specified using either a connection URL through the JDBC Driver Manager or a JDBC data source. For example, the following connection URL for the Greenplum driver specifies connection information for the primary and alternate servers using a connection URL:

```
jdbc:datadirect:greenplum://server1:5432;DatabaseName=TEST;User=test;
Password=secret;AlternateServers=(server2:5432;DatabaseName=TEST2,
server3:5432;DatabaseName=TEST3)
```

In this example:

```
...server1:5432;DatabaseName=TEST...
```

is the part of the connection URL that specifies connection information for the primary server. Alternate servers are specified using the `AlternateServers` property. For example:

```
...;AlternateServers=(server2:5432;DatabaseName=TEST2,server3:5432;
DatabaseName=TEST3)
```

Similarly, the same connection information for the primary and alternate servers specified using a JDBC data source would look like this:

```
GreenplumDataSource mds = new GreenplumDataSource();
mds.setDescription("My GreenplumDataSource");
mds.setServerName("server1");
mds.setPortNumber(5432);
mds.setDatabaseName("TEST");
mds.setUser("test");
mds.setPassword("secret");
mds.setAlternateServers("server2:5432;DatabaseName=TEST2,server3:5432;
DatabaseName=TEST3")
```

In this example, connection information for the primary server is specified using the `ServerName`, `PortNumber`, and `DatabaseName` properties. Connection information for the alternate servers is specified using the `AlternateServers` property.

The value of the `AlternateServers` property is a string that has the format:

```
(servername1[:port1][;property=value[;...]][,servername2[:port2]
[:property=value[;...]])...
```

where:

servername1 is the IP address or server name of the first alternate database server, *servername2* is the IP address or server name of the second alternate database server, and so on. The IP address or server name is required for each alternate server entry.

port1 is the port number on which the first alternate database server is listening, *port2* is the port number on which the second alternate database server is listening, and so on. Port numbers are optional for each alternate server entry. If unspecified, the port number specified for the primary server is used. If a port number is unspecified for the primary server, a default port number of 5432 is used.

property=value is the DatabaseName connection property. This property is optional for each alternate server entry. For example:

```
jdbc:datadirect:greenplum://server1:5432;DatabaseName=TEST;User=test;
Password=secret;AlternateServers=(server2:5432;DatabaseName=TEST2,
server3:5432;DatabaseName=TEST3)
```

If you do not specify the DatabaseName connection property in an alternate server entry, the connection to that alternate server uses the property specified in the URL for the primary server. For example, if you specify DatabaseName=TEST for the primary server, but do not specify a database name in the alternate server entry as shown in the following URL, the driver tries to connect to the TEST database on the alternate server.

```
jdbc:datadirect:greenplum://server1:5432;DatabaseName=TEST;User=test;
Password=secret;AlternateServers=(server2:5432,server3:5432)
```

Specifying connection retry

Connection retry allows the Greenplum driver to retry connections to the primary database server, and if specified, alternate servers until a successful connection is established. You use the ConnectionRetryCount and ConnectionRetryDelay properties to enable and control how connection retry works. For example:

```
jdbc:datadirect:greenplum://server1:5432;DatabaseName=TEST;User=test;
Password=secret;AlternateServers=(server2:5432;DatabaseName=TEST2,
server3:5432;DatabaseName=TEST3);ConnectionRetryCount=2;
ConnectionRetryDelay=5
```

In this example, if a successful connection is not established on the Greenplum driver's first pass through the list of database servers (primary and alternate), the driver retries the list of servers in the same sequence twice (ConnectionRetryCount=2). Because the connection retry delay has been set to five seconds (ConnectionRetryDelay=5), the driver waits five seconds between retry passes.

Failover properties

The following table summarizes the connection properties that control how failover works with the Greenplum driver. See [Connection property descriptions](#) on page 73 for details about configuring each property.

Table 5: Summary: Failover Properties for the Greenplum Driver

Property	Characteristic
AlternateServers	One or multiple alternate database servers. An IP address or server name identifying each server is required. Port number and the connection property DatabaseName are optional. If the port number is unspecified, the port number specified for the primary server is used. If a port number is unspecified for the primary server, the default port number of 5432 is used.
ConnectionRetryCount	Number of times the driver retries the primary database server, and if specified, alternate servers until a successful connection is established. The default is 5.
ConnectionRetryDelay	Wait interval, in seconds, between connection retry attempts when the ConnectionRetryCount property is set to a positive integer. The default is 1.
DatabaseName	Name of the database to which you want to connect.

Property	Characteristic
LoadBalancing	Sets whether the driver will use client load balancing in its attempts to connect to the database servers (primary and alternate). If client load balancing is enabled, the driver uses a random pattern instead of a sequential pattern in its attempts to connect. The default is <code>false</code> (client load balancing is disabled).
PortNumber	Port listening for connections on the primary database server. This property is supported only for DataSource connections. The default port number is <code>5432</code> .
ServerName	IP address or server name for the primary database server. This property is supported only for DataSource connections.

Refcursors

The driver returns the cursor name for refcursors. The application must fetch the actual data from the refcursor using the cursor name and must close the cursor before additional processing can be done on the statement. The application must close the cursor regardless of whether it actually fetches data from the cursor.

Isolation levels

The driver supports the Read Committed, Read Uncommitted, Repeatable Read, and Serializable isolation levels. The default is Read Committed.

Using scrollable cursors

The driver supports scroll-insensitive result sets and updatable result sets.

Note: When the driver cannot support the requested result set type or concurrency, it automatically downgrades the cursor and generates one or more SQLWarnings with detailed information.

JTA support

JDBC distributed transactions through JTA are not supported by the driver.

Large object (LOB) support

The Greenplum driver allows you to retrieve and update long data, specifically LONGVARBINARY and LONGVARCHAR¹³ data, using JDBC methods designed for Blobs and Clobs. When using these methods to update long data as Blobs or Clobs, the updates are made to the local copy of the data contained in the Blob or Clob object.

Retrieving and updating long data using JDBC methods designed for Blobs and Clobs provides some of the same benefits as retrieving and updating Blobs and Clobs, such as:

- Provides random access to data
- Allows searching for patterns in the data, such as retrieving long data that begins with a specific character string

To provide these benefits normally associated with Blobs and Clobs, data must be cached. Because data is cached, your application will incur a performance penalty, particularly if data is read once sequentially. This performance penalty can be severe if the size of the long data is larger than available memory.

Using COPY command

The driver supports a customized version of the Greenplum COPY command. It provides an additional keyword, LOCALFILE, to allow you to copy data from or to standard file-system files that are stored anywhere on your network, not just on the database server. The file types supported by the driver are .txt and .csv.

Note: The customized COPY command supports all the options supported by the Greenplum COPY command.

Syntax

To copy data from a file to an existing table in the database, use the following syntax:

```
COPY tablename[(columnname[,...])] from {LOCALFILE 'filepath'} [[WITH](option[,...])]
```

where:

tablename

is the name of the table that you are copying data to.

columnname

is the name of the column available in the table you are copying data to.

filepath

is the absolute path of the file you are copying data from.

¹³ You may determine whether *Character varying* columns are described as VARCHAR or LONGVARCHAR by setting the `VarcharClobThreshold` on page 120 connection property.

option

can be one of the following: FORMAT, OIDS, DELIMITER, NULL, HEADER, QUOTE, ESCAPE, FORCE_QUOTE, FORCE_NOT_NULL, and ENCODING.

Example:

```
COPY testcopytable(intcol, varcharcol, charcol) from {LOCALFILE 'C:\\Users\\abc\\data.txt'}  
with DELIMITER ','
```

To copy data from an existing table in the database to a file, use the following syntax:

```
COPY tablename[(columnname[,...])] to {LOCALFILE 'filepath'} [[WITH](option[,...])]
```

where:

tablename

is the name of the table that you are copying data from.

columnname

is the name of the column available in the table you are copying data from.

filepath

is the absolute path of the file you are copying data to.

option

can be one of the following: FORMAT, OIDS, DELIMITER, NULL, HEADER, QUOTE, ESCAPE, FORCE_QUOTE, FORCE_NOT_NULL, and ENCODING.

Example:

```
COPY testcopytable(intcol, varcharcol, charcol) to {LOCALFILE  
'/home/users/abc/copydata.csv'} with DELIMITER ','
```

Parameter metadata support

The driver supports returning parameter metadata. The driver returns only the parameter's data type.

User-defined functions' results

Greenplum provides functionality to create user-defined functions. Greenplum does not define a call mechanism for invoking a user-defined function. User-defined functions must be invoked via a SQL statement. For example, given a function defined as:

```
CREATE table foo (intcol int, varcharcol varchar(123))  
CREATE or REPLACE FUNCTION insertFoo  
(IN idVal int, IN nameVal varchar) RETURNS void  
AS $$  
    insert into foo values ($1, $2);
```

```
$$
LANGUAGE SQL;
```

must be invoked natively as:

```
SELECT * FROM insertFoo(100, 'Mark')
```

even though the function does not return a value or results. The Select SQL statement returns a result set that has one column named insertFoo and no row data.

The Greenplum driver supports invoking user-defined functions using the JDBC call Escape. The previously described function can be invoked using:

```
{call insertFoo(100, 'Mark')}
```

Greenplum functions return data from functions as a result set. If multiple output parameters are specified, the values for the output parameters are returned as columns in the result set. For example, the function defined as:

```
CREATE or REPLACE FUNCTION addValues(in v1 int, in v2 int)
  RETURNS int
  AS $$
    SELECT $1 + $2;
  $$
  LANGUAGE SQL;
```

returns a result set with a single column of type INTEGER, whereas the function defined as:

```
CREATE or REPLACE FUNCTION selectFooRow2
  (IN idVal int, OUT id int, OUT name varchar)
  AS $$
    select intcol, varcharcol from foo where intcol = $1;
  $$
  LANGUAGE SQL
```

returns a result set that contains two columns, a INTEGER id column and a VARCHAR name column.

In addition, when calling Greenplum functions that contain output parameters, the native syntax requires that the output parameter values be omitted from the function call. This, in addition to output parameter values being returned as a result set, makes the Greenplum behavior of calling functions different from most other databases.

Note: When using the ?= version of the call escape on a function that returns a set of values, only the first result in the set will be returned.

The Greenplum driver provides a mechanism that makes the invoking of functions more consistent with how other databases behave. In particular, the Greenplum driver allows parameter markers for output parameters to be specified in the function argument list when the Escape call is used. The driver extracts the output parameter values from the result set returned by the server and makes the values available via the getxxx methods on a CallableStatement.

For example, the function selectFooRow2 described previously can be invoked as:

```
sql = "{call selectFooRow2 (?, ?, ?)}";
CallableStatement cSTMT = connection.prepareStatement(sql);
cSTMT.setInt(1, idVal);
cSTMT.registerOutParameter(2, Types.INTEGER);
```

```
cSTMT.registerOutParameter(3, Types.VARCHAR);
cSTMT.execute();
int myID = cSTMT.getInt(2);
String myName = cSTMT.getString(3);
```

The values of the id and name output parameters are returned in the `myID` and `myName` variables.

An error is returned if the number of output parameters registered when the function is executed is less than the number of output parameters defined in the function. If no output parameters are registered to a function call, the driver returns the output parameters as a result set.

Greenplum can also return results from a function as a refcursor. There can be, at most, one refcursor per result; however, a function can return multiple results where each result is a refcursor. See [Refcursors](#) on page 66 for more information.

ResultSet metadata support

If your application requires table name information, the driver can return table name information in `ResultSet` metadata for `Select` statements. If you set the `ResultSetMetaDataOptions` property to 1, the driver performs additional processing to determine the correct table name for each column in the result set when the `ResultSetMetaData.getTableName()` method is called. Otherwise, the `getTableName()` method may return an empty string for each column in the result set.

When the `ResultSetMetaDataOptions` property is set to 1 and the `ResultSetMetaData.getTableName()` method is called, the table name information that is returned by the driver depends on whether the column in a result set maps to a column in a table in the database. For each column in a result set that maps to a column in a table in the database, the driver returns the table name associated with that column. For columns in a result set that do not map to a column in a table (for example, aggregates and literals), the driver returns an empty string.

The `Select` statements for which `ResultSet` metadata is returned may contain aliases, joins, and fully qualified names. The following queries are examples of `Select` statements for which the `ResultSetMetaData.getTableName()` method returns the correct table name for columns in the `Select` list:

```
SELECT id, name FROM Employee
SELECT E.id, E.name FROM Employee E
SELECT E.id, E.name AS EmployeeName FROM Employee E
SELECT E.id, E.name, I.location, I.phone FROM Employee E, EmployeeInfo I
    WHERE E.id = I.id
SELECT id, name, location, phone FROM Employee, EmployeeInfo WHERE id = empId
SELECT Employee.id, Employee.name, EmployeeInfo.location, EmployeeInfo.phone
    FROM Employee, EmployeeInfo WHERE Employee.id = EmployeeInfo.id
```

The table name returned by the driver for generated columns is an empty string. The following query is an example of a `Select` statement that returns a result set that contains a generated column (the column named "upper").

```
SELECT E.id, E.name as EmployeeName, {fn UCASE(E.name)} AS upper FROM Employee E
```

The driver also can return catalog name information when the `ResultSetMetaData.getCatalogName()` method is called if the driver can determine that information. For example, for the following statement, the driver returns "test" for the catalog name and "foo" for the table name:

```
SELECT * FROM test.foo
```

The additional processing required to return table name and catalog name information is only performed if the `ResultSetMetaData.getTableName()` or `ResultSetMetaData.getCatalogName()` methods are called.

Rowset support

The driver supports any JSR 114 implementation of the RowSet interface, including:

- `CachedRowSets`
- `FilteredRowSets`
- `WebRowSets`
- `JoinRowSets`
- `JDBCRowSets`

Visit <https://www.jcp.org/en/jsr/detail?id=114> for more information about JSR 114.

Connection property descriptions

You can use connection properties to customize the driver for your environment. This section organizes connection properties according to functionality. You can use connection properties with either the JDBC `DriverManager` or a JDBC data source. For a `DriverManager` connection, a property is expressed as a key value pair and takes the form `property=value`. For a data source connection, a property is expressed as a JDBC method and takes the form `setProperty(value)`.

Note:

- In a JDBC data source, string values must be enclosed in double quotation marks, for example, `setUser("abc@defcorp.com")`.
- The data type listed for each connection property is the Java data type used for the property value in a JDBC data source.
- Connection property names are case-insensitive. For example, `Password` is the same as `password`.
- For connection properties that support string values, use the following escape sequence to specify values containing leading or trailing spaces and curly brackets: `{value}`. For example: `User={hello }` or `Password={{hello}}`.

The following tables describe the connection properties by functionality.

- [Required properties](#)
- [User ID/password properties](#)
- [Kerberos properties](#)
- [Data encryption properties](#)
- [Bulk load properties](#)

- [Failover properties](#)
- [Proxy server properties](#)
- [Timeout properties](#)
- [Client information properties](#)
- [Statement pooling properties](#)
- [Additional properties](#)

Required properties

The following table summarizes connection properties required to connect to a database.

Property	Data Source Method	Default
PortNumber on page 106	setPortNumber getPortNumber	5432
ServerName on page 113	setServerName getServerName	No default value

User ID/password authentication properties

The following table summarizes the connection properties required for user ID/password authentication.

Property	Data Source Method	Default
User on page 119	setUser getUser	No default value
Password on page 106	setPassword getPassword	No default value

Kerberos authentication properties

The following table summarizes the connection properties required for Kerberos authentication.

Property	Data Source Method	Default
AuthenticationMethod on page 83	setAuthenticationMethod getAuthenticationMethod	userIdPassword
DatabaseName on page 91	setDatabaseName getDatabaseName	No default value

Property	Data Source Method	Default
ServicePrincipalName on page 114	setServicePrincipalName getServicePrincipalName	No default value
LoginConfigName on page 101	setLoginConfigName getLoginConfigName	JDBC_DRIVER_01

Data encryption properties

The following table summarizes connection properties which can be used to enable SSL.

Property	Data Source Method	Default
CryptoProtocolVersion on page 89	setCryptoProtocolVersion getCryptoProtocolVersion	No default value
EncryptionMethod on page 92	setEncryptionMethod getEncryptionMethod	<i>noEncryption</i>
HostNameInCertificate on page 94	setHostNameInCertificate getHostNameInCertificate	Empty string
TrustStore on page 118	setTrustStore getTrustStore	No default value
TrustStorePassword on page 118	setTrustStorePassword getTrustStorePassword	No default value
ValidateServerCertificate on page 120	setValidateServerCertificate getValidateServerCertificate	<i>true</i>

Bulk load properties

The following table contains the only connection property that affects how bulk load works with the Greenplum driver.

Property	Data Source Method	Default
BulkLoadBatchSize on page 84	setBulkLoadBatchSize getBulkLoadBatchSize	1000

Failover properties

The following table summarizes the connection properties used for configuring failover.

Property	Data Source Method	Default
AlternateServers on page 82	setAlternateServers getAlternateServers	No default value
ConnectionRetryCount on page 87	setConnectionRetryCount getConnectionRetryCount	5
ConnectionRetryDelay on page 88	setConnectionRetryDelay getConnectionRetryDelay	1 (second)
DatabaseName on page 91	setDatabaseName getDatabaseName	No default value
LoadBalancing on page 100	setLoadBalancing getLoadBalancing	false

Proxy server properties

The following table summarizes proxy server connection properties.

Property	Data Source Method	Default
ProxyHost on page 109	setProxyHost	Empty string
ProxyPassword on page 109	setProxyPassword	Empty string
ProxyPort on page 110	setProxyPort	0 which means the default is determined by the ProxyHost property. For HTTP URLs: 80 For HTTPS URLs: 443
ProxyUser on page 110	setProxyUser	Empty string

Timeout properties

The following table summarizes timeout connection properties.

Property	Data Source Method	Default
EnableCancelTimeout on page 92	setEnableCancelTimeout getEnableCancelTimeout	false

Property	Data Source Method	Default
LoginTimeout on page 102	setLoginTimeout getLoginTimeout	0
QueryTimeout on page 111	setQueryTimeout getQueryTimeout	0

Client information properties

The following table summarizes connection properties that can be used to return client information.

Property	Data Source Method	Default
AccountingInfo on page 81	setAccountingInfo getAccountingInfo	Empty string
ApplicationName on page 83	setApplicationName getApplicationName	Empty string
ClientHostName on page 85	setClientHostName getClientHostName	Empty string
ClientUser on page 86	setClientUser getClientUser	Empty string
ProgramID on page 108	setProgramID getProgramID	Empty string

Statement pooling properties

The following table summarizes statement pooling connection properties.

Property	Data Source Method	Default
ImportStatementPool on page 95	setImportStatementPool getImportStatementPool	Empty string
MaxPooledStatements on page 104	setMaxPooledStatements getMaxPooledStatements	0
RegisterStatementPoolMonitorMBean on page 112	setRegisterStatementPoolMonitorMBean getRegisterStatementPoolMonitorMBean	false

Additional properties

The following table summarizes additional connection properties.

Property	Data Source Method	Default
CatalogOptions on page 85	setCatalogOptions getCatalogOptions	2
ConvertNull on page 89	setConvertNull getConvertNull	1
Database on page 91	setDatabase getDatabase	No default value
ExtendedColumnMetadata on page 93	setExtendedColumnMetadata getExtendedColumnMetadata	false
InitializationString on page 96	setInitializationString getInitializationString	No default value
InsensitiveResultSetBufferSize on page 96	setInsensitiveResultSetBufferSize getInsensitiveResultSetBufferSize	2048
JavaDoubleToString on page 97	setJavaDoubleToString getJavaDoubleToString	false
KeyPassword on page 98	setKeyPassword getKeyPassword	No default value
KeyStore on page 99	setKeyStore getKeyStore	No default value
KeyStorePassword on page 99	setKeyStorePassword getKeyStorePassword	No default value
MaxLongVarcharSize on page 102	setMaxLongVarcharSize getMaxLongVarcharSize	1, 073, 741, 823
MaxNumericPrecision on page 103	setMaxNumericPrecision getMaxNumericPrecision	1000
MaxNumericScale on page 103	setMaxNumericScale getMaxNumericScale	998

Property	Data Source Method	Default
MaxStatements on page 105	setMaxStatements getMaxStatements	0
MaxVarcharSize on page 105	setMaxVarcharSize getMaxVarcharSize	10,485,760
PrepareThreshold on page 107	setPrepareThreshold getPrepareThreshold	0
ResultSetMetaDataOptions on page 112	resultSetMetaDataOptions getResultSetMetaDataOptions	0
SpyAttributes on page 115	setSpyAttributes getSpyAttributes	No default value
SupportsCatalogs on page 116	setSupportsCatalogs getSupportsCatalogs	true
TransactionErrorBehavior on page 117	setTransactionErrorBehavior getTransactionErrorBehavior	RollbackTransaction
TrustStore on page 118	setTrustStore getTrustStore	No default value
VarcharClobThreshold on page 120	setVarcharClobThreshold getVarcharClobThreshold	32768

For details, see the following topics:

- [AccountingInfo](#)
- [AlternateServers](#)
- [ApplicationName](#)
- [AuthenticationMethod](#)
- [BulkLoadBatchSize](#)
- [CatalogOptions](#)
- [ClientHostName](#)
- [ClientUser](#)
- [ConnectionRetryCount](#)

- [ConnectionRetryDelay](#)
- [ConvertNull](#)
- [CryptoProtocolVersion](#)
- [Database](#)
- [DatabaseName](#)
- [EnableCancelTimeout](#)
- [EncryptionMethod](#)
- [ExtendedColumnMetadata](#)
- [HostNameInCertificate](#)
- [ImportStatementPool](#)
- [InitializationString](#)
- [InsensitiveResultSetBufferSize](#)
- [JavaDoubleToString](#)
- [KeyPassword](#)
- [KeyStore](#)
- [KeyStorePassword](#)
- [LoadBalancing](#)
- [LoginConfigName](#)
- [LoginTimeout](#)
- [MaxLongVarcharSize](#)
- [MaxNumericPrecision](#)
- [MaxNumericScale](#)
- [MaxPooledStatements](#)
- [MaxStatements](#)
- [MaxVarcharSize](#)
- [Password](#)
- [PortNumber](#)
- [PrepareThreshold](#)
- [ProgramID](#)
- [ProxyHost](#)
- [ProxyPassword](#)
- [ProxyPort](#)
- [ProxyUser](#)

- [QueryTimeout](#)
- [RegisterStatementPoolMonitorMBean](#)
- [ResultSetMetaDataOptions](#)
- [ServerName](#)
- [ServicePrincipalName](#)
- [SpyAttributes](#)
- [SupportsCatalogs](#)
- [TransactionErrorBehavior](#)
- [TrustStore](#)
- [TrustStorePassword](#)
- [User](#)
- [ValidateServerCertificate](#)
- [VarcharClobThreshold](#)

AccountingInfo

Purpose

Defines accounting information. This value is stored locally and is used for database administration/monitoring purposes.

Valid Values

string

where:

string

is the accounting information.

Data Source Methods

```
public String getAccountingInfo()  
public void setAccountingInfo(String)
```

Default

empty string

Data Type

String

See also

[Using client information](#) on page 62

AlternateServers

Purpose

Defines a list of alternate database servers that is used to failover new or lost connections, depending on the failover method selected.

Valid Values

```
(servername1[:port1][;DatabaseName=value][,servername2[:port2][;DatabaseName=value]]...)
```

The server name (*servername1*, *servername2*, and so on) is required for each alternate server entry. Port number (*port1*, *port2*, and so on) and connection properties (*property=**value*) are optional for each alternate server entry. If the port is unspecified, the port number of the primary server is used. If the port number of the primary server is unspecified, the default port number of 5432 is used.

DatabaseName is an optional connection property.

Data Source Methods

```
public String getAlternateServers()  
public void setAlternateServers(String)
```

Default

None

Data Type

No default value

Example

The following URL contains alternate server entries for server2 and server3. The alternate server entries contain the optional DatabaseName property.

```
jdbc:datadirect:greenplum://server1:5432;DatabaseName=TEST;User=test;  
Password=secret;AlternateServers=(server2:5432;DatabaseName=TEST2,  
server3:5432;DatabaseName=TEST3)
```

See also

[Configuring failover](#) on page 63

ApplicationName

Purpose

Specifies the name of the application. This value is stored locally and is used for database administration/monitoring purposes.

Valid Values

string

where:

string

is the name of the application.

Data Source Methods

```
public String getApplicationName()  
public void setApplicationName(String)
```

Default

empty string

Data Type

String

See also

[Using client information](#) on page 62

AuthenticationMethod

Purpose

Determines which authentication method the driver uses when establishing a connection. If the specified authentication method is not supported by the database server, the connection fails and the driver throws an exception.

Valid Values

`userIdPassword` | `kerberos`

Behavior

If set to `userIdPassword`, the driver uses user ID/password authentication. If a user ID and password is not specified, the driver throws an exception.

If set to `kerberos`, the driver uses Kerberos authentication. The `ServicePrincipalName` property must be specified.

Notes

- The `User` property provides the user ID. The `Password` property provides the password.

Data Source Methods

```
public String getAuthenticationMethod()  
public void setAuthenticationMethod(String)
```

Default

`userIdPassword`

Data Type

String

See also

- [Authentication](#) on page 51
- [ServicePrincipalName](#) on page 114

BulkLoadBatchSize

Purpose

Provides a suggestion to the driver for the number of rows to load to the database at a time when bulk loading data. Performance can be improved by increasing the number of rows the driver loads at a time because fewer network round trips are required. Be aware that increasing the number of rows that are loaded also causes the driver to consume more memory on the client.

Valid Values

`x`

where

`x`

is a positive integer that represents a number of rows.

Data Source Methods

```
public Long getBulkLoadBatchSize()  
public void setBulkLoadBatchSize(Long)
```

Default

1000

Data Type

long

See also

[Bulk load support](#) on page 63

CatalogOptions

Purpose

Determines which type of metadata information is included in result sets when an application calls DatabaseMetaData methods.

Valid Values

2 | 4

Behavior

If set to 2, the driver queries database catalogs for column information.

If set to 4, a hint is provided to the driver to emulate getColumnns() calls using the ResultSetMetaData object instead of querying database catalogs for column information. Using emulation can improve performance because the SQL statement that is formulated by the emulation is less complex than the SQL statement that is formulated using getColumnns(). The argument to getColumnns() must evaluate to a single table. If it does not, because of a wildcard or null value, for example, the driver reverts to the default behavior for getColumnns() calls.

Data Source Methods

```
public Integer getCatalogOptions()  
public void setCatalogOptions(Integer)
```

Default

2

Data Type

int

ClientHostName

Purpose

Specifies the host name of the client machine. This value is stored locally and is used for database administration/monitoring purposes.

Valid Values

string

where:

string

is the host name of the client machine.

Data Source Methods

```
public String getClientHostName()  
public void setClientHostName(String)
```

Default

empty string

Data Type

String

See also

[Using client information](#) on page 62

ClientUser

Purpose

Specifies the user ID. This value is stored locally and is used for database administration/monitoring purposes.

Valid Values

string

where:

string

is a valid user ID.

Data Source Methods

```
public String getClientUser()  
public void setClientUser(String)
```

Default

empty string

Data Type

String

See also

[Using client information](#) on page 62

ConnectionRetryCount

Purpose

Specifies the number of times the driver retries connection attempts to the primary database server, and if specified, alternate servers until a successful connection is established.

Valid Values

0 | x

where:

x

is a positive integer that represents the number of retries.

Behavior

If set to 0, the driver does not try to reconnect after the initial unsuccessful attempt.

If set to x , the driver retries connection attempts the specified number of times. If a connection is not established during the retry attempts, the driver returns an exception that is generated by the last database server to which it tried to connect.

Example

If this property is set to 2 and alternate servers are specified using the AlternateServers property, the driver retries the list of servers (primary and alternate) twice after the initial retry attempt.

Notes

- If an application sets a login timeout value (for example, using DataSource.loginTimeout or DriverManager.loginTimeout), and the login timeout expires, the driver ceases connection attempts.
- The ConnectionRetryDelay property specifies the wait interval, in seconds, to occur between retry attempts.

Data Source Methods

```
public Integer getConnectionRetryCount()  
public void setConnectionRetryCount(Integer)
```

Default

5

Data Type

int

See also

[Specifying connection retry](#)

ConnectionRetryDelay

Purpose

Specifies the number of seconds the driver waits between connection retry attempts when ConnectionRetryCount is set to a positive integer.

Valid Values

0 | x

where:

x

is a number of seconds.

Behavior

If set to 0, the driver does not delay between retries.

If set to x , the driver waits between connection retry attempts the specified number of seconds.

Example

If ConnectionRetryCount is set to 2, ConnectionRetryDelay is set to 3, and alternate servers are specified using the AlternateServers property, the driver retries the list of servers (primary and alternate) twice after the initial retry attempt. The driver waits 3 seconds between retry attempts.

Data Source Methods

```
public Integer getConnectionRetryDelay()  
public void setConnectionRetryDelay(Integer)
```

Default

1 (second)

Data Type

int

See also

[Specifying connection retry](#)

ConvertNull

Purpose

Controls how data conversions are handled for null values.

Valid Values

0 | 1

Behavior

If set to 0, the driver does not perform the data type check if the value of the column is null. This allows null values to be returned even though a conversion between the requested type and the column type is undefined.

If set to 1, the driver checks the data type that is requested against the data type of the table column that stores the data. If a conversion between the requested type and column type is not defined, the driver generates an "unsupported data conversion" exception regardless of the data type of the column value.

Data Source Methods

```
public Integer getConvertNull()  
public void setConvertNull(Integer)
```

Default

1

Data Type

int

CryptoProtocolVersion

Purpose

Specifies a cryptographic protocol or comma-separated list of cryptographic protocols that can be used when TLS/SSL is enabled using the EncryptionMethod connection property.

Valid Values

cryptographic_protocol [, *cryptographic_protocol*]...

where:

cryptographic_protocol

is one of the following cryptographic protocols:

TLSv1.3 | TLSv1.2 | TLSv1.1 | TLSv1 | SSLv3 | SSLv2

Caution: To avoid vulnerabilities associated with SSLv3 and SSLv2, good security practices recommend using TLSv1 or higher.

Example

If your server supports TLSv1.3 and TLSv1.2, you can specify acceptable cryptographic protocols with the following key-value pair:

```
CryptoProtocolVersion=TLSv1.3,TLSv1.2
```

Notes

- When multiple protocols are specified, the driver uses the highest version supported by the server. If none of the specified protocols are supported by the server, the connection fails and the driver returns an error.
- When no value has been specified for `CryptoProtocolVersion`, the cryptographic protocol used depends on the highest protocol version supported by the server and the highest protocol version supported by the JDK. Refer to the database management system documentation for information on which cryptographic protocols are supported.
- The TLSv1.3 protocol works with Java SE 11 or higher by default.
- To enable the TLSv1.3 protocol when using Java SE 8, set the `jdk.tls.client.protocols` Java system property to `TLSv1.3`. For example, `$ java -Djdk.tls.client.protocols="TLSv1.3" myApp`.
In the following versions of Oracle JDK and OpenJDK, support for the TLSv1.3 protocol is not enabled by default.
 - Oracle JDK 8u261 or later but earlier than Oracle JDK 8u341
 - OpenJDK 8u272 or later but earlier than OpenJDK 8u352

Data Source Methods

```
public String getCryptoProtocolVersion()  
public void setCryptoProtocolVersion(String)
```

Default

No default value

Data Type

String

See also

- [EncryptionMethod](#) on page 92
- [Data Encryption](#) on page 56

Database

Purpose

An alias for the DatabaseName property.

DatabaseName

Purpose

Specifies the name of the database to which you want to connect.

Valid Values

string

where:

string

is the name of a Greenplum database.

Data Source Methods

```
public String getDatabaseName()  
public void setDatabaseName(String)
```

Default

No default value

Data Type

String

Alias

Database property. If both the Database and DatabaseName properties are specified in a connection URL, the last property that is positioned in the connection URL is used. For example, if your application specifies the following connection URL, the value of the Database connection property would be used instead of the value of the DatabaseName connection property.

```
jdbc:datadirect:greenplum://server1:5432;DatabaseName=jdbc;Database=acct;  
User=test;Password=secret
```

EnableCancelTimeout

Purpose

Determines whether a cancel request that is sent by the driver as the result of a query timing out is subject to the same query timeout value as the statement it cancels.

Valid Values

true | false

Behavior

If set to `true`, the cancel request times out using the same timeout value, in seconds, that is set for the statement it cancels. For example, if your application calls `Statement.setQueryTimeout(5)` on a statement and that statement is cancelled because its timeout value was exceeded, the driver sends a cancel request that also will time out if its execution exceeds 5 seconds. If the cancel request times out, because the server is down, for example, the driver throws an exception indicating that the cancel request was timed out and the connection is no longer valid.

If set to `false`, the cancel request does not time out.

Data Source Methods

```
public Boolean getEnableCancelTimeout()  
public void setEnableCancelTimeout(Boolean)
```

Default

false

Data Type

boolean

EncryptionMethod

Purpose

Determines whether data is encrypted and decrypted when transmitted over the network between the driver and database server.

Valid Values

noEncryption | SSL | requestSSL

Behavior

If set to `noEncryption`, data is not encrypted or decrypted.

If set to `SSL`, data is encrypted using SSL. If the database server does not support SSL, the connection fails and the driver throws an exception.

If set to *requestSSL*, the login request and data is encrypted using SSL. If the database server does not support SSL, the driver establishes an unencrypted connection.

Notes

- When SSL is enabled, the following properties also apply:

CryptoProtocolVersion

HostNameInCertificate

KeyStore (for SSL client authentication)

KeyStorePassword (for SSL client authentication)

KeyPassword (for SSL client authentication)

TrustStore

TrustStorePassword

ValidateServerCertificate

Data Source Methods

```
public String getEncryptionMethod()
public void setEncryptionMethod(String)
```

Default

noEncryption

Data Type

String

ExtendedColumnMetadata

Purpose

Determines how the driver returns column metadata when retrieving results with `ResultSetMetaData` methods.

Valid Values

true | false

Behavior

If set to `true`, the driver makes an additional roundtrip to the database to retrieve actual values for column metadata. For example, the driver returns nullability as `IS_NULLABLE`, `NOT_NULLABLE`, or `NULLABILITY_UNKNOWN`, depending on the actual status of the column.

If set to `false`, the driver returns only the values for column metadata hardcoded in the driver. For example, the driver returns nullability as `NULLABILITY_UNKNOWN`.

Notes

Setting `ExtendedColumnMetadata` to enabled may diminish performance because an additional roundtrip to the database is required to retrieve actual values for the column metadata.

Data Source Methods

```
public Boolean getExtendedColumnMetadata()  
public void setExtendedColumnMetadata(Boolean)
```

Default

false

Data Type

Boolean

HostNameInCertificate

Purpose

Specifies a host name for certificate validation when SSL encryption is enabled (`EncryptionMethod=SSL`) and validation is enabled (`ValidateServerCertificate=true`).

This connection property is optional and provides additional security against man-in-the-middle (MITM) attacks by ensuring that the server the driver is connecting to is the server that was requested.

Valid Values

host_name | `#SERVERNAME#`

where:

host_name

is a valid host name.

Behavior

If *host_name* is specified, the driver compares the specified host name to the `DNSName` value of the `SubjectAlternativeName` in the certificate. If the certificate does not have a `SubjectAlternativeName`, the driver compares the host name with the `Common Name (CN)` part of the certificate. If the values do not match, the connection fails and the driver throws an exception.

If `#SERVERNAME#` is specified, the driver compares the server name that is specified in the connection URL or data source of the connection to the `DNSName` value of the `SubjectAlternativeName` in the certificate. If the certificate does not have a `SubjectAlternativeName`, the driver compares the host name to the `CN` part of the certificate's `Subject` name. If the values do not match, the connection fails and the driver throws an exception. If multiple `CN` parts are present, the driver validates the host name against each `CN` part. If any one validation succeeds, a connection is established.

Data Source Methods

```
public String getHostNameInCertificate()  
public void setHostNameInCertificate(String)
```

Default

empty string

Data Type

String

ImportStatementPool

Purpose

Specifies the path and file name of the file to be used to load the contents of the statement pool. When this property is specified, statements are imported into the statement pool from the specified file.

If the driver cannot locate the specified file when establishing the connection, the connection fails and the driver throws an exception.

Valid Values

string

where:

string

is the path and file name of the file to be used to load the contents of the statement pool.

Data Source Methods

```
public String getImportStatementPool()  
public void setImportStatementPool(String)
```

Default

empty string

Data Type

String

See also

Refer to "Statement Pool Monitor" in the *Progress DataDirect for JDBC Drivers Reference* for further details.

InitializationString

Purpose

Specifies one or multiple SQL commands to be executed by the driver after it has established the connection to the database and has performed all initialization for the connection. If the execution of a SQL command fails, the connection attempt also fails and the driver throws an exception indicating which SQL command or commands failed.

You can use this property to override the session `wait_timeout` value that is used by the `InteractiveClient` property for the current connection.

Valid Values

string

where:

string

is one or multiple SQL commands.

Multiple commands must be separated by semicolons. In addition, if this property is specified in a connection URL, the entire value must be enclosed in parentheses when multiple commands are specified.

Example

The following URL sets the session `wait_timeout` variable to 3600 (seconds) and turns on a warning count for simple inserts.

```
jdbc:datadirect:greenplum://server1:5432;DatabaseName=test;  
InitializationString=(command1;command2)
```

Data Source Methods

```
public String getInitializationString()  
public void setInitializationString(String)
```

Default

No default value

Data Type

String

InsensitiveResultSetBufferSize

Purpose

Determines the amount of memory used by the driver to cache insensitive result set data.

Valid Values

-1 | 0 | x

where:

x

is a positive integer that represents the size of the memory buffer.

Behavior

If set to -1, the driver caches insensitive result set data in memory. If the size of the result set exceeds available memory, an `OutOfMemoryException` is generated. With no need to write result set data to disk, the driver processes the data efficiently.

If set to 0, the driver caches insensitive result set data in memory, up to a maximum of 2 GB. If the size of the result set data exceeds available memory, the driver pages the result set data to disk. Because result set data may be written to disk, the driver may have to reformat the data to write it correctly to disk.

If set to x , the driver caches insensitive result set data in memory and uses this value to set the size (in KB) of the memory buffer for caching insensitive result set data. If the size of the result set data exceeds available memory, the driver pages the result set data to disk. Because the result set data may be written to disk, the driver may have to reformat the data to write it correctly to disk. Specifying a buffer size that is a power of 2 results in efficient memory use.

Data Source Methods

```
public Integer getInsensitiveResultSetBufferSize()
```

```
public void setInsensitiveResultSetBufferSize(Integer)
```

Default

2048

Data Type

int

See also

[Performance Considerations](#)

JavaDoubleToString

Purpose

Determines which algorithm the driver uses when converting a double or float value to a string value. By default, the driver uses its own internal conversion algorithm, which improves performance.

Valid Values

true | false

Behavior

If set to `true`, the driver uses the JVM algorithm when converting a double or float value to a string value. If your application cannot accept rounding differences and you are willing to sacrifice performance, set this value to `true` to use the JVM conversion algorithm.

If set to `false`, the driver uses its own internal algorithm when converting a double or float value to a string value. This value improves performance, but slight rounding differences within the allowable error of the double and float data types can occur when compared to the same conversion using the JVM algorithm.

Data Source Methods

```
public Boolean getJavaDoubleToString()  
public void setJavaDoubleToString(Boolean)
```

Default

`false`

Data Type

`boolean`

KeyPassword

Purpose

Specifies the password that is used to access the individual keys in the keystore file when SSL is enabled (EncryptionMethod=SSL) and SSL client authentication is enabled on the database server. This property is useful when individual keys in the keystore file have a different password than the keystore file.

Valid Values

string

where:

string

is a valid password.

Data Source Methods

```
public String getKeyPassword()  
public void setKeyPassword(String)
```

Default

No default value

Data Type

`String`

KeyStore

Purpose

Specifies the directory of the keystore file to be used when SSL is enabled (`EncryptionMethod=SSL`) and SSL client authentication is enabled on the database server. The keystore file contains the certificates that the client sends to the server in response to the server's certificate request.

This value overrides the directory of the keystore file that is specified by the `javax.net.ssl.keyStoreJava` system property. If this property is not specified, the keystore directory is specified by the `javax.net.ssl.keyStore` Java system property.

Valid Values

string

where:

string is a valid directory of a keystore file.

Notes

The keystore and truststore files can be the same file.

Data Source Methods

```
public String getKeyStore()  
public void setKeyStore(String)
```

Default

No default value

Data Type

String

KeyStorePassword

Purpose

Specifies the password that is used to access the keystore file when SSL is enabled (`EncryptionMethod=SSL`) and SSL client authentication is enabled on the database server. The keystore file contains the certificates that the client sends to the server in response to the server's certificate request.

This value overrides the password of the keystore file that is specified by the `javax.net.ssl.keyStorePassword` Java system property. If this property is not specified, the keystore password is specified by the `javax.net.ssl.keyStorePassword` Java system property.

Valid Values

string

where:

string

is a valid password.

Notes

- The keystore and truststore files can be the same file.

Data Source Methods

```
public String getKeyStorePassword()  
public void setKeyStorePassword(String)
```

Default

No default value

Data Type

String

LoadBalancing

Purpose

Determines whether the driver uses client load balancing in its attempts to connect to the database servers (primary and alternate). You can specify one or multiple alternate servers by setting the `AlternateServers` property.

Valid Values

`true` | `false`

Behavior

If set to `true`, the driver uses client load balancing and attempts to connect to the database servers (primary and alternate) in random order. The driver randomly selects from the list of primary and alternate servers which server to connect to first. If that connection fails, the driver again randomly selects from this list of servers until all servers in the list have been tried or a connection is successfully established.

If set to `false`, the driver does not use client load balancing and connects to each server based on their sequential order (primary server first, then, alternate servers in the order they are specified).

Data Source Methods

```
public Boolean getLoadBalancing()  
public void setLoadBalancing(Boolean)
```

Default

`false`

Data Type

boolean

See also

[Failover support](#)

LoginConfigName

Purpose

Specifies the name of the entry in the JAAS login configuration file that contains the authentication technology used by the driver to establish a Kerberos connection. The LoginModule-specific items found in the entry are passed on to the LoginModule.

Valid Values

entry_name

where:

entry_name

is the name of the entry that contains the authentication technology used with the driver.

Example

In the following example, JDBC_DRIVER_01 is the entry name while the authentication technology and related settings are found in the brackets.

```
JDBC_DRIVER_01 {  
    com.sun.security.auth.module.Krb5LoginModule required useTicketCache=true;  
};
```

Data Source Methods

```
public String getLoginConfigName()  
public void setLoginConfigName(String)
```

Default

JDBC_DRIVER_01

Data Type

String

See also

- [Authentication](#) on page 51
- [The JAAS login configuration file](#) on page 55

LoginTimeout

Purpose

Specifies the amount of time, in seconds, that the driver waits for a connection to be established before timing out the connection request.

Valid Values

0 | x

where:

x

is a positive integer that represents a number of seconds.

Behavior

If set to 0, the driver does not time out a connection request.

If set to x , the driver waits for the specified number of seconds before returning control to the application and throwing a timeout exception.

Data Source Methods

```
public Integer getLoginTimeout()  
public void setLoginTimeout(Integer)
```

Default

0

Data Type

int

MaxLongVarcharSize

Purpose

Determines the maximum size of LONGVARCHAR columns when described within the result set metadata.

Valid Values

x

where:

x

is an integer greater than or equal to 1 and less than or equal to 1,073,741,823.

Data Source Methods

```
public Integer getMaxLongVarcharSize()  
public void setMaxLongVarcharSize(Integer)
```

Default

1,073,741,823

Data Type

Integer

MaxNumericPrecision

Purpose

Determines the maximum precision of NUMERIC columns when described within the result set metadata.

Valid Values

x

where:

x

is an integer greater than or equal to 1 and less than or equal to 1000.

Data Source Methods

```
public Integer getMaxNumericPrecision()  
public void setMaxNumericPrecision(Integer)
```

Default

1000

Data Type

Integer

MaxNumericScale

Purpose

Determines the maximum scale of NUMERIC columns when described within the result set metadata.

Valid Values

x

where:

x

is an integer greater than or equal to 0 and less than or equal to 998.

Data Source Methods

```
public Integer getMaxNumericScale()  
public void setMaxNumericScale(Integer)
```

Default

998

Data Type

Integer

MaxPooledStatements

Purpose

Specifies the maximum number of prepared statements to be pooled for each connection and enables the driver's internal prepared statement pooling when set to an integer greater than zero (0). The driver's internal prepared statement pooling provides performance benefits when the driver is not running from within an application server or another application that provides its own statement pooling.

Valid Values

0 | x

where:

x

is a positive integer that represents a number of prepared statements to be cached.

Behavior

If set to 0, the driver's internal prepared statement pooling is not enabled.

If set to x , the driver's internal prepared statement pooling is enabled and the driver uses the specified value to cache up to that many prepared statements created by an application. If the value set for this property is greater than the number of prepared statements that are used by the application, all prepared statements that are created by the application are cached. Because CallableStatement is a sub-class of PreparedStatement, CallableStatements also are cached.

Notes

When you enable statement pooling, your applications can access the Statement Pool Monitor directly with DataDirect-specific methods. However, you can also enable the Statement Pool Monitor as a JMX MBean. To enable the Statement Pool Monitor as an MBean, statement pooling must be enabled with MaxPooledStatements and the Statement Pool Monitor MBean must be registered using the RegisterStatementPoolMonitorMBean connection property.

Example

If the value of this property is set to 20, the driver caches the last 20 prepared statements that are created by the application.

Data Source Methods

```
public Integer getMaxPooledStatements()  
public void setMaxPooledStatements(Integer)
```

Default

0

Data Type

int

See also

- [Statement pooling properties](#)
- [Performance considerations](#) on page 60

MaxStatements

Purpose

An alias for the MaxPooledStatements property.

MaxVarcharSize

Purpose

Determines the maximum size of VARCHAR columns when described within the result set metadata.

Valid Values

x

where:

x

is an integer greater than or equal to 1 and less than or equal to 10,485,760.

Notes

When string functions are used within the Select list of a Select statement, the driver describes the resulting string value with a size of 10,485,760. For instance, concatenating two columns that are each defined as VARCHAR(10) will result in a VARCHAR(10485760) rather than a VARCHAR(20). The unnecessarily long size can result in undesirable behavior in some JDBC applications.

Data Source Methods

```
public Integer getMaxVarcharSize()  
public void setMaxVarcharSize(Integer)
```

Default

10,485,760

Data Type

Integer

Password

Purpose

Specifies a password that is used to connect to your Greenplum database. A password is required if user ID/Password authentication is enabled on your database. Contact your system administrator to obtain your password.

Valid Values

string

where:

string

is a valid password. The password is case-sensitive.

Data Source Methods

```
public String getPassword()  
public void setPassword(String)
```

Default

No default value

Data Type

String

PortNumber

Purpose

The TCP port of the primary database server that is listening for connections to the Greenplum database. This property is supported only for data source connections.

Valid Values

port

where:

port

is the port number.

Data Source Methods

```
public Integer getPortNumber()
public void setPortNumber(Integer)
```

Default

5432

Data Type

int

PrepareThreshold

Purpose

Specifies the number of PreparedStatement executions that must occur before the driver begins using server-side prepared statements.

Valid Values

0 | *x*

where:

x

is a positive integer that represents the number of PreparedStatement executions.

Behavior

If set to 0, the driver uses server-side prepared statements by default.

If set to *x*, the driver waits till it reaches the specified number of PreparedStatement executions before using server-side prepared statements.

Example

If PrepareThreshold is set to 5, the driver will start using server-side prepared statements on the fifth execution of the same PreparedStatement object.

Data Source Methods

```
public Integer getPrepareThreshold()
public void setPrepareThreshold(Integer)
```

Default

0

Data Type

int

See also

[QueryTimeout](#)

ProgramID

Purpose

The driver and version information on the client to be stored in the database. This value is stored locally and is used for database administration/monitoring purposes.

Valid Values

string

where:

string

is a value that identifies the product and version of the driver on the client.

Example

DDJ04200

Data Source Methods

```
public String getProgramID()  
public void setProgramID(String)
```

Default

empty string

Data Type

String

See also

[Using client information](#) on page 62

ProxyHost

Description

Identifies a proxy server to use for the first connection.

Valid Values

server_name | *IP_address*

where:

server_name

is the name of the proxy server, which may be qualified with the domain name.

IP_address

is an IP address, specified in either IPv4 or IPv6 format, or a combination of the two.

Data Source Method

setProxyHost

Default

empty string

See also

- [Proxy server](#) on page 15

ProxyPassword

Purpose

Specifies the password needed to connect to a proxy server for the first connection.

Valid Values

password

where:

password

is a valid password for that server. Contact your system administrator to obtain a valid password.

Data Source Method

setProxyPassword

Default

empty string

See also

[Proxy server](#) on page 15

ProxyPort

Purpose

Specifies the port number where the proxy server is listening for HTTP or HTTPS requests for the first connection.

Valid Values

port

where:

port

is the port number on which the proxy server is listening. Contact your system administrator to obtain the correct port.

Data Source Method

`setProxyPort`

Default

0 which means that the default value is determined by whether the value specified for the ProxyHost property is an HTTP or HTTPS URL.

For HTTP: 80

For HTTPS: 443

See also

[Proxy server](#) on page 15

ProxyUser

Purpose

Specifies the specifies the user name needed to connect to a proxy server for the first connection.

Valid Values

user_name

where:

`user_name`

is a valid user ID for the proxy server.

Data Source Method

`setProxyUser`

Default

empty string

See also

[Proxy server](#) on page 15

QueryTimeout

Purpose

Sets the default query timeout (in seconds) for all statements created by a connection.

Valid Values

`-1` | `0` | `x`

where:

`x`

is a positive integer that represents a number of seconds.

Behavior

If set to `-1`, the query timeout functionality is disabled. The driver silently ignores calls to the `Statement.setQueryTimeout()` method.

If set to `0`, the default query timeout is infinite (the query does not time out).

If set to `x`, the driver uses the value as the default timeout for any statement that is created by the connection. To override the default timeout value set by this connection option, call the `Statement.setQueryTimeout()` method to set a timeout value for a particular statement.

Data Source Methods

```
public Integer getQueryTimeout()  
public void setQueryTimeout(Integer)
```

Default

0

Data Type

int

RegisterStatementPoolMonitorMBean

Purpose

Registers the Statement Pool Monitor as a JMX MBean when statement pooling has been enabled with `MaxPooledStatements`. This allows you to manage statement pooling with standard JMX API calls and to use JMX-compliant tools, such as JConsole.

Valid Values

`true` | `false`

Behavior

If set to `true`, the driver registers an MBean for the statement pool monitor for each statement pool. This gives applications access to the Statement Pool Monitor through JMX when statement pooling is enabled.

If set to `false`, the driver does not register an MBean for the statement pool monitor for any statement pool.

Notes

Registering the MBean exports a reference to the Statement Pool Monitor. The exported reference can prevent garbage collection on connections if the connections are not properly closed. When garbage collection does not take place on these connections, out of memory errors can occur.

Data Source Methods

```
public Boolean getRegisterStatementPoolMonitorMBean()  
public void setRegisterStatementPoolMonitorMBean(Boolean)
```

Default

`false`

Data Type

`boolean`

See also

- [Statement pooling](#) on page 62
- [MaxPooledStatements](#) on page 104

ResultSetMetaDataOptions

Purpose

Determines whether the driver returns table name information in the `ResultSet` metadata for `Select` statements.

Valid Values

0 | 1

Behavior

If set to 0 and the `ResultSetMetaData.getTableName()` method is called, the driver does not perform additional processing to determine the correct table name for each column in the result set. The `getTableName()` method may return an empty string for each column in the result set.

If set to 1 and the `ResultSetMetaData.getTableName()` method is called, the driver performs additional processing to determine the correct table name for each column in the result set. The driver returns schema name and catalog name information when the `ResultSetMetaData.getSchemaName()` and `ResultSetMetaData.getCatalogName()` methods are called if the driver can determine that information.

Data Source Methods

```
public Integer getResultSetMetaDataOptions()  
public void setResultSetMetaDataOptions(Integer)
```

Default

0

Data Type

int

See also

[Performance Considerations](#)

ServerName

Purpose

Specifies either the IP address in IPv4 or IPv6 format, or the server name (if your network supports named servers) of the primary database server.

This property is supported only for data source connections.

Valid Values

string

where:

string

is a valid IP address or server name.

Example

122.23.15.12 or GreenplumServer

Data Source Methods

```
public String getServerName()  
public void setServerName(String)
```

Default

No default value

Data Type

String

ServicePrincipalName

Purpose

Specifies the service principal name to be used for Kerberos authentication.

Valid Values

ServicePrincipalName

where:

ServicePrincipalName

is the three-part service principal name registered with the key distribution center (KDC).

Specify the service principal name using the following format:

Service_Name/Fully_Qualified_Domain_Name@REALM.COM

where:

Service_Name

is the name of the service hosting the instance. It is the same value as the `krb_svrname` configuration parameter on the server. The default is `postgres`.

Fully_Qualified_Domain_Name

is the fully qualified domain name (FQDN) of the host machine. This value must match the FQDN registered with the KDC. The FQDN consists of a host name and a domain name. For the example `myserver.example.com`, `myserver` is the host name and `example.com` is the domain name.

REALM.COM

is the domain name of the host machine. This value is optional. If no value is specified, the default domain is used. The domain must be specified in upper-case characters. For example, `EXAMPLE.COM`. For Windows Active Directory, the Kerberos realm name is the Windows domain name.

Example

The following is an example of a valid service principal name:

```
postgres/myserver.example.com@EXAMPLE.COM
```

Notes

- If `AuthenticationMethod` is set to `userIdPassword`, the value of the `ServicePrincipalName` property is ignored.

Data Source Methods

```
public String getServicePrincipalName()
public void setServicePrincipalName(String)
```

Default

No default value

See also

- [Authentication](#) on page 51
- [AuthenticationMethod](#) on page 83
- [Configuring the driver for Kerberos authentication](#) on page 52

SpyAttributes

Purpose

Enables DataDirect Spy to log detailed information about calls issued by the driver on behalf of the application. DataDirect Spy is not enabled by default.

Valid Values

```
( spy_attribute [ ; spy_attribute ] ... )
```

where:

spy_attribute

is any valid DataDirect Spy attribute. Refer to "Tracking JDBC Calls with DataDirect Spy" in the *Progress DataDirect for JDBC Drivers Reference* for a list of supported attributes.

Example

The following value instructs the driver to log all JDBC activity to a file using a maximum of 80 characters for each line.

```
(log=(file)/tmp/spy.log;linelimit=80)
```

Notes

- If coding a path on Windows to the log file in a Java string, the backslash character (\) must be preceded by the Java escape character, a backslash. For example: `log=(file)C:\\temp\\spy.log`.

Data Source Methods

```
public String getSpyAttributes()  
public void setSpyAttributes(String)
```

Default

No default value

Data Type

String

See also

Refer to "Tracking JDBC Calls with DataDirect Spy" in the *Progress DataDirect for JDBC Drivers Reference* for more information about using DataDirect Spy.

SupportsCatalogs

Purpose

Enables support for catalogs. While Greenplum has the notion of catalogs (or databases), you can only select from tables that reside within the catalog you specified at connect time. Catalogs cannot be changed after connecting. Therefore, most applications behave better if you indicate support for catalogs.

Valid Values

`true` | `false`

Behavior

If set to `true`, the driver returns the database as the catalog for catalog calls, for example, `getTables` and `getColumns`.

If set to `false`, the driver returns NULL for the catalog in catalog calls.

Notes

The `SupportsCatalogs` connection property affects the following catalog methods:

- `getCatalogSeparator`
- `getCatalogTerm`
- `getMaxCatalogNameLength`
- `isCatalogAtStart`
- `supportsCatalogsInDataManipulation`
- `supportsCatalogsInProcedureCalls`

- supportsCatalogsInTableDefinitions
- supportsCatalogsInIndexDefinitions
- supportsCatalogsInPrivilegeDefinitions

Data Source Methods

```
public Boolean getSupportsCatalogs()
public void setSupportsCatalogs(Boolean)
```

Default

true

Data Type

boolean

TransactionErrorBehavior

Purpose

Determines how the driver handles errors that occur within a transaction. When an error occurs in a transaction, the Greenplum server does not allow any operations on the connection except for rolling back the transaction.

Valid Values

none | RollbackTransaction

Behavior

If set to `none`, the driver does not roll back the transaction when an error occurs. The application must handle the error and roll back the transaction. Any operation on the statement other than a rollback results in an error.

If set to `RollbackTransaction`, the driver rolls back the transaction when an error occurs. In addition to the original error message, the driver posts an error message indicating that the transaction has been rolled back.

Data Source Methods

```
public String getTransactionErrorBehavior()
public void setTransactionErrorBehavior(String)
```

Default

RollbackTransaction

Data Type

String

TrustStore

Purpose

Specifies the directory of the truststore file to be used when SSL is enabled (`EncryptionMethod=SSL`) and server authentication is used. The truststore file contains a list of the Certificate Authorities (CAs) that the client trusts.

This value overrides the directory of the truststore file that is specified by the `javax.net.ssl.trustStore` Java system property. If this property is not specified, the truststore directory is specified by the `javax.net.ssl.trustStore` Java system property.

This property is ignored if `ValidateServerCertificate=false`.

Valid Values

string

where:

string

is the directory of the truststore file.

Data Source Methods

```
public String getTrustStore()  
public void setTrustStore(String)
```

Default

No default value

Data Type

String

TrustStorePassword

Purpose

Specifies the password that is used to access the truststore file when SSL is enabled (`EncryptionMethod=SSL`) and server authentication is used. The truststore file contains a list of the Certificate Authorities (CAs) that the client trusts.

This value overrides the password of the truststore file that is specified by the `javax.net.ssl.trustStorePassword` Java system property. If this property is not specified, the truststore password is specified by the `javax.net.ssl.trustStorePassword` Java system property.

This property is ignored if `ValidateServerCertificate=false`.

Valid Values

string

where:

string

is the directory of the truststore file.

Data Source Methods

```
public String getTrustStorePassword()  
public void setTrustStorePassword(String)
```

Default

No default value

Data Type

String

User

Purpose

Specifies the user name that is used to connect to the Greenplum database. A user name is required if user ID/password authentication is enabled on your database. Contact your system administrator to obtain your user name.

Valid Values

string

where:

string

is a valid user name. The user name is case-sensitive.

Data Source Methods

```
public String getUser()  
public void setUser(String)
```

Default

No default value

Data Type

String

ValidateServerCertificate

Purpose

Determines whether the driver validates the certificate that is sent by the database server when SSL encryption is enabled (`EncryptionMethod=SSL`). When using SSL server authentication, any certificate that is sent by the server must be issued by a trusted Certificate Authority (CA). Allowing the driver to trust any certificate that is returned from the server even if the issuer is not a trusted CA is useful in test environments because it eliminates the need to specify truststore information on each client in the test environment.

Valid Values

true | false

Behavior

If set to `true`, the driver validates the certificate that is sent by the database server. Any certificate from the server must be issued by a trusted CA in the truststore file. If the `HostNameInCertificate` property is specified, the driver also validates the certificate using a host name. The `HostNameInCertificate` property is optional and provides additional security against man-in-the-middle (MITM) attacks by ensuring that the server the driver is connecting to is the server that was requested.

If set to `false`, the driver does not validate the certificate that is sent by the database server. The driver ignores any truststore information that is specified by the `TrustStore` and `TrustStorePassword` properties or Java system properties.

Data Source Methods

```
public Boolean getValidateServerCertificate()  
public void setValidateServerCertificate(Boolean)
```

Default

true

Data Type

boolean

VarcharClobThreshold

Purpose

Enables CLOB functionality when handling character data. Determines whether columns of the *Character varying* data type will be described as `VARCHAR` or `LONGVARCHAR` (CLOB).

Valid Values

x

where:

x

is a number of characters.

Behavior

If the width of a *Character varying* column is greater than *x*, the *Character varying* data type is described as LONGVARCHAR and CLOB functionality is enabled. The driver allows you to retrieve and update long data by using JDBC methods designed for Clobs. This functionality provides random access to data and allows searching for patterns in the data. To provide these benefits, data must be cached. Because data is cached, your application will incur a performance penalty, particularly if data is read once sequentially.

If the width of a *Character varying* column is less than or equal to *x*, the *Character varying* data type is described as VARCHAR. If you want to avoid the performance penalties associated with CLOB functionality, you should set this value at a value greater than the maximum *Character varying* column width your application handles.

Data Source Methods

```
public Varchar getVarcharClobThreshold()  
public void setVarcharClobThreshold(Varchar)
```

Default

32768

Data Type

VARCHAR or LONGVARCHAR

See also

[Performance considerations](#) on page 60

[Large object \(LOB\) support](#) on page 67

