



Red Hat JBoss Data Virtualization

6.3

Development Guide Volume 1: Client Development

This guide is for developers wanting to interface to Red Hat JBoss Data Virtualization from within client applications.

Red Hat Customer Content
Services

Red Hat JBoss Data Virtualization 6.3 Development Guide Volume 1: Client Development

This guide is for developers wanting to interface to Red Hat JBoss Data Virtualization from within client applications.

Red Hat Customer Content Services

Legal Notice

Copyright © 2016 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This document provides information on concepts and tasks relating to interfacing to Red Hat JBoss Data Virtualization from within client applications.

Table of Contents

| | |
|--|-----------|
| Chapter 1. Read Me | 3 |
| 1.1. Back Up Your Data | 3 |
| 1.2. Variable Name: EAP_HOME | 3 |
| 1.3. Variable Name: MODE | 3 |
| 1.4. Red Hat Documentation Site | 3 |
| Chapter 2. Connecting to a Virtual Database | 4 |
| 2.1. Java Database Connectivity (JDBC) | 4 |
| 2.2. JBoss Data Virtualization and JDBC | 4 |
| 2.3. JBoss Data Virtualization and ODBC | 4 |
| 2.4. Configure JDBC for SAP Business Objects on Microsoft Windows | 4 |
| 2.5. Generated Keys | 5 |
| 2.6. Connection Methods | 5 |
| 2.7. Connecting to a VDB Using the Driver Class | 5 |
| 2.8. Create a Connection to a VDB Using the Driver Class | 6 |
| 2.9. Driver Connection URL Format | 6 |
| 2.10. Connection Properties for the Driver and Data Source Classes | 7 |
| 2.11. Connecting to a VDB Using the Data Source Class | 10 |
| 2.12. Create a Connection to a VDB Using the Data Source Class | 10 |
| 2.13. Additional Data Source Connection Properties | 11 |
| 2.14. Connecting to a VDB as a Data Source | 12 |
| Chapter 3. Connecting with node.js | 14 |
| 3.1. Access Red Hat JBoss Data Virtualization from node.js | 14 |
| Chapter 4. Multiple Hosts | 15 |
| 4.1. Using Multiple Hosts | 15 |
| 4.2. Failover | 15 |
| 4.3. Load Balancing | 15 |
| 4.4. Increase the Maximum Number of Cached Instances | 16 |
| 4.5. Advanced Configuration | 16 |
| 4.6. Reauthentication | 17 |
| Chapter 5. Extensions to JDBC | 18 |
| 5.1. Prepared Statements | 18 |
| 5.2. JDBC Statement Extensions | 18 |
| 5.3. Non-blocking Statement Execution | 19 |
| 5.4. Continuous Execution | 20 |
| 5.5. Execution Properties | 20 |
| 5.6. XML Extensions | 21 |
| 5.7. XML Document Formatting | 21 |
| 5.8. XML Schema Validation | 22 |
| 5.9. The SET Statement | 22 |
| 5.10. The SHOW Statement | 23 |
| 5.11. Transaction Statements | 23 |
| 5.12. Partial Results Mode | 24 |
| 5.13. Setting Partial Results Mode | 24 |
| 5.14. Partial Results Warnings | 25 |
| Chapter 6. JDBC Transactions | 26 |
| 6.1. JDBC Transaction Types | 26 |
| 6.2. Local Transactions | 26 |
| 6.3. Ending Local Transactions | 27 |

| | |
|--|-----------|
| 6.4. Turning Off Local Transactions | 27 |
| 6.5. Request Level Transactions | 27 |
| 6.6. Transaction Wrapping Modes | 28 |
| 6.7. Set the Transaction Wrapping Mode | 28 |
| 6.8. Multiple Insert Batches | 28 |
| 6.9. Global Transactions | 28 |
| 6.10. Enterprise Information System Support | 30 |
| Chapter 7. Client SSL Connections | 31 |
| 7.1. SSL Client Connections | 31 |
| Chapter 8. Using Hibernate with JBoss Data Virtualization | 34 |
| 8.1. Configure Hibernate for Use with JBoss Data Virtualization | 34 |
| 8.2. Limitations of Using Hibernate with JBoss Data Virtualization | 35 |
| Chapter 9. OData Support | 36 |
| 9.1. OData Support | 36 |
| 9.2. OData Version 4.0 Support | 42 |
| Appendix A. Unsupported JDBC Methods | 48 |
| A.1. Unsupported JDBC Methods | 48 |
| A.2. ResultSet Limitations | 48 |
| A.3. Unsupported Classes and Methods in java.sql | 48 |
| A.4. Unsupported Classes and Methods in javax.sql | 50 |
| Appendix B. Keytool | 52 |
| B.1. Keytool | 52 |
| B.2. Using Keytool with JBoss Data Virtualization | 52 |
| B.3. Create a Private/Public Key Pair with Keytool | 52 |
| B.4. Extract a Self-signed Certificate from the Keystore | 53 |
| B.5. Add a Certificate to a Truststore Using Keytool | 53 |
| Appendix C. Revision History | 55 |

Chapter 1. Read Me

1.1. Back Up Your Data



Warning

Red Hat recommends that you back up your system settings and data before undertaking any of the configuration tasks mentioned in this book.

1.2. Variable Name: EAP_HOME

EAP_HOME refers to the root directory of the Red Hat JBoss Enterprise Application Platform installation on which JBoss Data Virtualization has been deployed.

1.3. Variable Name: MODE

MODE will either be **standalone** or **domain** depending on whether JBoss Data Virtualization is running in standalone or domain mode. Substitute one of these whenever you see **MODE** in a file path in this documentation. (You need to set this variable yourself, based on where the product has been installed in your directory structure.)

1.4. Red Hat Documentation Site

Red Hat's official documentation site is available at <https://access.redhat.com/site/documentation/>. There you will find the latest version of every book, including this one.

Chapter 2. Connecting to a Virtual Database

2.1. Java Database Connectivity (JDBC)

Java Database Connectivity (JDBC) is an application to database connectivity tool. This Application Program Interface (API) enables communication between applications written in Java and data stored in databases providing methods for data querying and updating.

JDBC is very similar to Open Database Connectivity (ODBC).

2.2. JBoss Data Virtualization and JDBC

JBoss Data Virtualization provides an API that builds on Java Database Connectivity (JDBC), allowing client applications to issue SQL queries against deployed virtual databases (VDBs).



Note

The Java 1.6 or 1.7 JDK is required by client applications connecting to an JBoss Data Virtualization VDB.



Note

The JBoss Data Virtualization JDBC API is compatible with the JDBC 4.0 specification but does not fully support all methods. Advanced features, such as *updatable* result sets and SQL3 data types are also not supported.

See [Section A.3, “Unsupported Classes and Methods in java.sql”](#) and [Section A.4, “Unsupported Classes and Methods in javax.sql”](#) for more information about unsupported classes and methods.

2.3. JBoss Data Virtualization and ODBC

To learn how to configure ODBC for Red Hat JBoss Data Virtualization for Red Hat Enterprise Linux and Microsoft Windows, please refer to the Installation Guide.

2.4. Configure JDBC for SAP Business Objects on Microsoft Windows

Follow these steps if you wish to use JDBC drivers in SAP Business Objects 3.1 on Microsoft Windows to create a connection in order to query a VDB.

Procedure 2.1. Configure JDBC Drivers for SAP BO and Create a Connection

1. Download and save the JDBC drivers for the database to which you wish to connect.
2. Navigate to the directory in which BusinessObjects is installed.
3. Open the [Database_Name].sbo file in your text editor for the database with which you wish to connect. For instance, if you wish to connect to an Oracle database, open oracle.sbo.

- Find these lines in the file:

```
<ClassPath>
  <Path>${ROOT$}/drivers/java/dbd_jdbcwrapper.jar</Path>
</ClassPath>
```

- Modify the code so that it looks like this:

```
<ClassPath>
  <Path>${ROOT$}/drivers/java/dbd_jdbcwrapper.jar</Path>
  <Path>C:\JDBCDrivers\ojdbc6.jar</Path>
</ClassPath>
```

- Save the file and exit.
- Create a database connection in the Information Design Tool. Click New->Relational Connection and then give the connection a name and click Next.
- Select the type of connection (JDBC Drivers) and click Next.
- Enter your connection parameters and click Test.
- If the test is successful, click Finish.

2.5. Generated Keys

JBoss Data Virtualization supports returning generated keys for JDBC sources and from JBoss Data Virtualization temp tables with SERIAL primary key columns. However the current implementation returns only the last set of keys generated and returns the key results directly from the source - no view projection of other intermediate handling is performed. For most scenarios (single source inserts) this handling is sufficient. A custom solution may need to be developed if you are using a FOR EACH ROW instead of trigger to process your inserts and target multiple tables that each return generated keys. It is possible to develop a UDF that also manipulates the returned generated keys - see the **org.teiid.CommandContext** methods dealing with generated keys for more.

2.6. Connection Methods

The following are methods of creating a JDBC connection to an JBoss Data Virtualization virtual database (VDB):

- » Using the **org.teiid.jdbc.TeiidDriver** driver class
- » Using the **org.teiid.jdbc.TeiidDataSource** data source class
- » Configuring a VDB as a JBoss data source

2.7. Connecting to a VDB Using the Driver Class

The **org.teiid.jdbc.TeiidDriver** class, found in **EAP_HOME/dataVirtualization/jdbc/teiid-VERSION-redhat-N-jdbc.jar**, should be used to

create a connection using the Java `java.sql.DriverManager` class.

2.8. Create a Connection to a VDB Using the Driver Class

Prerequisites

- ✦ The client application must have the JBoss Data Virtualization `EAP_HOME/dataVirtualization/jdbc/teiid-VERSION-redhat-N-jdbc.jar` JAR file in its classpath. (If the application is running on the same application server as the JBoss Data Virtualization instance to which the connection is desired, then this will already be the case.)
- ✦ JBoss Data Virtualization must be installed and running, with the relevant virtual databases (VDBs) deployed.

Procedure 2.2. Create a Connection to a VDB Using the Driver Class

✦ Create a Connection to the VDB

Within your client application code, create a **Connection** to the VDB using the `DriverManager.getConnection()` method:

```
public class TeiidClient
{
    public Connection getConnection(String user, String password) throws
    Exception
    {
        String url =
        "jdbc:teiid:myVDB@mm://localhost:31000;ApplicationName=myApp";
        return DriverManager.getConnection(url, user, password);
    }
}
```

See Also:

- ✦ [Section 2.9, "Driver Connection URL Format"](#)

2.9. Driver Connection URL Format


URLs used when establishing a connection using the driver class have the following format:

```
jdbc:teiid:VDB-NAME@mm[s]://HOSTNAME:PORT;[prop-name=prop-value;]*
```

Given this format, the following table describes the variable parts of the URL:

Table 2.1. URL Entities

| Variable Name | Description |
|---------------|-------------|
|---------------|-------------|

| Variable Name | Description |
|------------------------|--|
| VDB-NAME | The name of the virtual database (VDB) to which the application is connected. |
| |  Important VDB names can contain version information; for example, myvdb.2 . If such a name is used in the URL, this has the same effect as supplying a version=2 connection property. Note that if the VDB name contains version information, you cannot also use the version property in the same request. |
| mm[s] | The JBoss Data Virtualization JDBC protocol. mm is the default for normal connections. mms uses SSL for encryption and is the default for the AdminAPI tools. |
| HOSTNAME | The server where JBoss Data Virtualization is installed. |
| PORT | The port on which JBoss Data Virtualization is listening for incoming JDBC connections. |
| [prop-name=prop-value] | Any number of additional name-value pairs can be supplied in the URL, separated by semi-colons. Property values must be URL encoded if they contain reserved characters, for example, ? , = , and ; . |

2.10. Connection Properties for the Driver and Data Source Classes

The following table shows all the connection properties that can be used with the JBoss Data Virtualization JDBC driver URL connection string and the JBoss Data Virtualization JDBC data source class.

These properties are defined in **org.teiid.net.TeiidURL** and **org.teiid.jdbc.ExecutionProperties** (**waitForLoad** is defined in **org.teiid.jdbc.EmbeddedProfile**), and the corresponding set methods are defined for the data source class in **org.teiid.jdbc.TeiidDataSource** and its superclass, **org.teiid.jdbc.BaseDataSource**.

Property names that can be used in the driver URL connection string are listed in the *Property Name* column, and the corresponding set methods for use with the data source class are listed in the *Method Name* column.

Table 2.2. Connection Properties

| Property Name | Method Name | Type | Description |
|---------------|-------------|------|-------------|
|---------------|-------------|------|-------------|

| Property Name | Method Name | Type | Description |
|------------------------------|---------------------------------|----------------|--|
| ansiQuotedIdentifiers | setAnsiQuotedIdentifiers | boolean | Sets the parsing behavior for double quoted entries in SQL. If true, then parses doubled quoted entries as identifiers. If false, then double quoted values that are valid string literals are parsed as string literals. Default is true. |
| ApplicationName | setApplicationName | String | Name of the client application; allows the administrator to distinguish between connections. |
| autoCommitTxn | setAutoCommitTxn | String | <p>Only applies only when autoCommit is set to "true". This determines how an executed command needs to be wrapped as a transaction inside the JBoss Data Virtualization engine to maintain the data integrity.</p> <ul style="list-style-type: none"> ✦ ON - Always wrap command in distributed transaction ✦ OFF - Never wrap command in distributed transaction ✦ DETECT (default)- If the executed command is spanning more than one source it automatically uses distributed transaction. <p>See Section 6.1, "JDBC Transaction Types" for more information.</p> |
| autoFailover | setAutoFailover | boolean | If true, automatically selects a new server instance after a communication exception. This is typically not needed when connections are managed, as the connection can be purged from the pool. Default is false. |
| disableLocalTxn | setDisableLocalTxn | boolean | If true, the autoCommit setting, commit and rollback is ignored for local transactions. Default is false. |
| fetchSize | setFetchSize | int | Size of the resultset. Default is 2048. <=0 indicates that the default should be used. |
| NOEXEC | setNoExec | String | (Typically not set as a connection property.) Can be ON or OFF. ON prevents query execution, but parsing and planning still occurs. Default is OFF. |
| partialResultsMode | setPartialResultsMode | boolean | Enable/disable support partial results mode. Default false. See Section 5.12, "Partial Results Mode" for more details. |

| Property Name | Method Name | Type | Description |
|--|--|----------------|--|
| Passthrough Authentication | setPassthrough Authentication | boolean | Only applies to local connections. When this option is set to true, JBoss Data Virtualization looks for an authenticated security context on the calling thread. If one is found, it uses that user's credentials to create a session. JBoss Data Virtualization also verifies that the same user is using this connection during the life of the connection. If it finds a different security context on the calling thread and the new user is eligible to login to JBoss Data Virtualization, it switches the identity on the connection. |
| useCallingThread | Not applicable. Must be set using setAdditionalProperties method if connecting via the data source class. | boolean | Only applies to local connections. When this option is set to true, the calling thread is used to process the query. If false, then an engine thread is used. Default is true. |
| QUERYTIMEOUT | setQueryTimeout | int | Default query timeout in seconds. Must be ≥ 0 . 0 indicates no timeout. Can be overridden by Statement.setQueryTimeout . Default is 0. |
| useJDBC4ColumnNameAndLabelSemantics | setUseJDBC4ColumnNameAndLabelSemantics | boolean | A change was made in JDBC4 to return unaliased column names as the ResultSetMetadata column name. Prior to this, if a column alias was used it was returned as the column name. Setting this property to false enables backwards compatibility when JDBC3 and older support is still required. Defaults to true. |
| password | setPassword | String | Credential for user |
| resultSetCacheMode | setResultSetCacheMode | boolean | ResultSet caching is turned on/off. Default is false. |
| SHOWPLAN | setShowPlan | String | (Typically not set as a connection property.) Can be ON, OFF or DEBUG. ON returns the query plan along with the results and DEBUG additionally prints the query planner debug information in the log and returns it with the results. Both the plan and the log are available through JDBC API extensions. Default is OFF. |
| user | setUser | String | User name. |
| version | setDatabaseVersion | String | Version number of the VDB. |

| Property Name | Method Name | Type | Description |
|-------------------------------------|--|----------------|--|
| jaasName | setJaasName | String | JAAS configuration name. Only applies when configuring GSS authentication. See the <i>JBoss Data Virtualization Administration Guide</i> for more information about GSS configuration. |
| kerberosServicePrincipleName | setKerberosServicePrincipleName | String | Kerberos authenticated principle name. Only applies when configuring GSS authentication. See the <i>JBoss Data Virtualization Administration Guide</i> for more information about GSS configuration. |
| encryptRequests | setEncryptRequests | boolean | Only applies to non-SSL socket connections. When set to true, the request message and any associated payload is encrypted using the connection cryptor. Default is false. |
| waitForLoad | Not applicable. Needs to be specified using setAdditionalProperties if set using the data source class. | String | Only applies to local connections. When this option is set to a non-negative value, the connection will wait that number of milliseconds for the VDB to become active. Setting to a negative number uses the system default setting. |

2.11. Connecting to a VDB Using the Data Source Class

The `org.teiid.jdbc.TeiidDataSource` class is based on the `javax.sql.DataSource` connection factory. It can be used to create **ManagedConnections** and **XAConnections** to both **DataSources** and **XADataSources**. XA transactions are extended to JBoss Data Virtualization sources that support XA.

The JBoss Data Virtualization **DataSource** class is *serializable* and can be used with JNDI naming services.

See `EAP_HOME/quickstarts/simpleclient/` example for more information.

2.12. Create a Connection to a VDB Using the Data Source Class

Prerequisites

- ✦ The client application must have the JBoss Data Virtualization `EAP_HOME/dataVirtualization/jdbc/teiid-VERSION-redhat-N-jdbc.jar` JAR file in its classpath. (If the application is running on the same application server as the JBoss Data Virtualization instance to which the connection is desired, then this will already be the case.)
- ✦ JBoss Data Virtualization must be installed and running, with the relevant virtual databases (VDBs) deployed.

Procedure 2.3. Create a Connection to a VDB Using the Data Source Class

- ✦ **Create a Connection Object**

Create a `org.teiid.jdbc.TeiidDataSource` object, set the required properties, and use the `TeiidDataSource.getConnection()` method to obtain a Connection object. For example:

```
public class TeiidClient
{
    public Connection getConnection(String user, String password) throws
    Exception
    {
        TeiidDataSource ds = new TeiidDataSource();
        ds.setUser(user);
        ds.setPassword(password);
        ds.setServerName("localhost");
        ds.setPortNumber(31000);
        ds.setDatabaseName("myVDB");
        return ds.getConnection();
    }
}
```


See Also:

- [Section 2.10, “Connection Properties for the Driver and Data Source Classes”](#)
- [Section 2.13, “Additional Data Source Connection Properties”](#)

2.13. Additional Data Source Connection Properties

When using the driver class, various properties are derived from the URL. For the data source class, these properties are set using the following additional methods:

Table 2.3. Data Source Connection Properties

| Method Name | Type | Description |
|--|---------------------|---|
| <code>setAlternateServers</code> | <code>String</code> | Optional delimited list of host:port entries. Refer to Section 4.1, “Using Multiple Hosts” for more information. |
| <code>setAdditionalProperties</code> | <code>String</code> | Optional setting of properties that has the same format as the property string in a driver connection URL. Refer to Section 2.9, “Driver Connection URL Format” |
| <code>setDatabaseName</code> | <code>String</code> | The name of a virtual database (VDB) deployed to JBoss Data Virtualization. |
|  Important | | |
| <p>VDB names can contain version information; for example, <code>myvdb.2</code>. If such a name is used in the URL, this has the same effect as supplying a <code>version=2</code> connection property. Note that if the VDB name contains version information, you cannot also use the <code>version</code> property in the same request.</p> | | |
| <code>setDatabaseVersion</code> | <code>String</code> | The VDB version. |

| Method Name | Type | Description |
|--------------------------|----------------|--|
| setDataSourceName | String | The name given to this data source |
| setPortNumber | int | The port number on which the server process is listening. |
| setServerName | String | The server hostname where the JBoss Data Virtualization runtime is installed. |
| setSecure | boolean | Secure connection. Flag to indicate to use SSL (mms) based connection between client and server. |



Note

All the properties from URL Connection Properties can be used on DataSource using the `AdditionalProperties` setter method if the corresponding setter method is not already available. For example, you can add **useCallingThread** property as

```
<xa-datasource-property
name="AdditionalProperties">useCallingThread=false</xa-datasource-
property>
```

2.14. Connecting to a VDB as a Data Source

JBoss Data Virtualization virtual databases (VDBs) can be configured as an JBoss Enterprise Application Platform (EAP) data source. The data source can then be accessed from JNDI or injected into your Java EE applications. A JBoss Data Virtualization data source is deployed in the same way as any other database resource.



Note

The recommended approach for configuring data sources is to use JBoss CLI or Management Console, not directly editing the **standalone.xml** configuration file. See the Red Hat JBoss Enterprise Application Platform *Administration and Configuration Guide* for more information on how to configure data sources in JBoss EAP.

Chapter 3. Connecting with node.js

3.1. Access Red Hat JBoss Data Virtualization from node.js

If you are writing a "node.js" application and would like to access Red Hat JBoss Data Virtualization from with it, you can do so by using an NPM package called "pg". Red Hat JBoss Data Virtualization supports the PG transport, so you can use this PostgreSQL client for "node.js" to access it.

For example, if you have a virtual database called "northwind" deployed on your Red Hat JBoss Data Virtualization server, and it has a table called "customers", your default configuration will look like this:

```
user = 'user'
password = 'user'
host = 127.0.0.1
port = 35432
```

You can then write a short access program like this:

```
var pg = require('pg');
var connectionString = "pg://user:user@localhost:35432/northwind"
pg.connect(connectionString, function(err, client) {
  client.query('SELECT CustomerID, ContactName, ContactTitle FROM
Customers', function(err, result) {
    console.log(result.rows)
  });
});
```

If you want to access one row at a time, you can also use the event mechanism by writing a piece of code like this:

```
var pg = require('pg');
var connectionString = "pg://user:user@localhost:35432/northwind"
pg.connect(connectionString, function(err, client) {
  var query = client.query('SELECT CustomerID, ContactName,
ContactTitle FROM Customers');
  query.on('row', function(row) {
    console.log(row);
  });
  query.on('error', function(error) {
    console.log("failed to query");
  });
  query.on('end', function(error) {
    console.log("closing client");
    client.end();
  });
});
```

For more information, please refer to the `node-postgres` package documentation. You can also find the source code on [GitHub](#).

Chapter 4. Multiple Hosts

4.1. Using Multiple Hosts

JBoss Data Virtualization may be clustered over several servers utilizing *failover* and *load balancing*.

The easiest way to enable these features is for the client to specify multiple hostname and port number combinations in the URL connection string as a comma separated list of **host:port** combinations:

```
jdbc:teiid:<vdb-name>@mm://host1:31000,host1:31001,host2:31000;version=2
```

If you are connecting with the data source class, the **setAlternateServers** method can be used to specify the failover servers. The format is also a comma separated list of **host:port** combinations.

The client randomly selects one of the JBoss Data Virtualization servers from the list and establishes a session with that server. If a connection cannot be established, then each of the remaining servers will be tried in random order. This allows for both connection time failover and random server selection load balancing.

4.2. Failover

Post connection failover will be used if you are using an administration connection (such as what is used by AdminShell) or if the **autoFailover** connection property is set to true. Post connection failover works by sending a ping, at most every second, to test the connection prior to use. If the ping fails, a new instance will be selected prior to the operation being attempted.

This is not considered to be true transparent application failover because the client does not restart transactions or queries, nor will it recreate session scoped temporary tables.



Warning

Extreme caution should be exercised if using this with non-admin connections.

4.3. Load Balancing

Post connection load balancing can be utilized in one of two ways.

If you are using the **TeiidDataSource** class and the **PooledConnection** returned by **getPooledConnection** is terminated using the **close()** method, then a new server instance will be selected automatically. (When using driver based connections or when using the **TeiidDataSource** class in a connection pool, the automatic load balancing will not happen.)

Alternatively, you can explicitly trigger load balancing through the use of the set statement: **SET NEWINSTANCE TRUE**. Typically you will not issue this statement manually, but you can use it as the connection test query on your data source configuration:

```
<datasources>
  <datasource jndi-name="java://teiidDS" pool-name="teiidDS">
    <connection-url>jdbc:teiid:{vdb}@mm://{host}:31000</connection-
url>
```

```

    <driver>teiid</driver>
    <pool>
      <prefill>>false</prefill>
      <use-strict-min>>false</use-strict-min>
      <flush-strategy>FailingConnectionOnly</flush-strategy>
      <check-valid-connection-sql>SET NEWINSTANCE TRUE</check-
valid-connection-sql>
    </pool>
    <security>
      <user-name>{user}</user-name>
      <password>{password}</password>
    </security>
  </datasource>
</drivers>
  <driver name="teiid" module="org.jboss.teiid.client">
    <driver-class>org.teiid.jdbc.TeiidDriver</driver-class>
    <xa-datasource-class>org.teiid.jdbc.TeiidDataSource</xa-
datasource-class>
  </driver>
</drivers>
</datasources>

```



Important

Session level temporary tables, currently running transactions, session level cache entries, and PreparedPlans for a given session will not be available on other cluster members. Therefore, it is recommended that post connection time load balancing is only used when the logical connection could have been closed, but the actual connection is reused. This is the typical connection pool pattern.

4.4. Increase the Maximum Number of Cached Instances

JBoss Data Virtualization maintains a pool of extra socket connections that are reused. For load balancing, this reduces the potential cost of switching a connection to another server instance. The default setting is to maintain 16 connections. If the client application is connecting to a large number of JBoss Data Virtualization instances and post connection time load balancing is used, then consider increasing the number of cached instances.

Procedure 4.1. Increase the Maximum Number of Cached Instances

- ✳ Increase the **org.teiid.sockets.maxCachedInstances** property in the **teiid-client-settings.properties** file located in **EAP_HOME/dataVirtualization/jdbc/teiid-VERSION-redhat-N-jdbc.jar**.



Note

An example file (**teiid-client-settings.orig.properties**) is packaged in the client JAR file.

4.5. Advanced Configuration

Features such as server discovery, load balancing, failover, retry, and retry delay, may be customized if the default policy is not appropriate. Refer to the `org.teiid.net.socket.ServerDiscovery` interface and default implementation `org.teiid.net.socket.UrlServerDiscovery` for more information on customization.

The `UrlServerDiscovery` implementation provides the following features:

- ✦ discovery of servers from URL hosts (including the data source server and alternative servers)
- ✦ random selection for load balancing and failover
- ✦ one connection attempt per host
- ✦ no biasing
- ✦ black listing

Typically you'll want to extend the `UrlServerDiscovery` so that it can be used as the fall-back strategy and to only implement the necessary changed methods. It is important to consider that one `ServerDiscovery` instance will be created for each connection. Any sharing of information between instances should be done through static state or some other shared lookup.

Your customized server discovery class must be referenced by the `discoveryStrategy` connection/DataSource property using its full class name.

4.6. Reauthentication

JBoss Data Virtualization connections (defined by the `org.teiid.jdbc.TeiidConnection` interface) support the `changeUser` method to reauthenticate a given connection. If reauthentication is successful, the current connection may be used with the given identity. Existing statements and resultsets are still available for use under the old identity.

Chapter 5. Extensions to JDBC

5.1. Prepared Statements

JBoss Data Virtualization provides `org.teiid.jdbc.TeiidPreparedStatement`, a custom interface for the standard `java.sql.PreparedStatement`, and implementations `org.teiid.jdbc.CallableStatementImpl` and `org.teiid.jdbc.PreparedStatementImpl`. Prepared statements can be important in speeding up common statement execution, since they allow the server to skip parsing, resolving, and planning of the statement.

The following points should be considered when using prepared statements:


- It is not necessary to pool client-side JBoss Data Virtualization prepared statements, because JBoss Data Virtualization performs plan caching on the server side.
- The number of cached plans is configurable. The plans are purged in order of least recently used (LRU).
- Cached plans are not distributed through a cluster. A new plan must be created for each cluster member.
- Plans are cached for the entire VDB or for just a particular session. The scope of a plan is detected automatically based upon the functions evaluated during its planning process.
- Runtime updates of costing information do not yet cause re-planning. At this time only session-scoped temporary table or internally materialized tables update their costing information.
- Stored procedures executed through a callable statement have their plans cached in the same way as a prepared statement.
- Bind variable types in function signatures, for example `where t.col = abs(?)`, can be determined if the function has only one signature or if the function is used in a predicate where the return type can be determined. In more complex situations it may be necessary to add a type hint with a cast or convert, for example `upper(convert(?, string))`.

5.2. JDBC Statement Extensions

The JBoss Data Virtualization statement interface, `org.teiid.jdbc.TeiidStatement`, provides functionality beyond the JDBC standard. To use this interface, cast or unwrap the statement returned by the connection. The interface provides the following methods for setting and retrieving statement properties:

Table 5.1. Statement Properties

| Method Name | Description |
|-----------------------------------|--|
| <code>getAnnotations</code> | This method has been deprecated. Use the SHOW statement. |
| <code>getDebugLog</code> | This method has been deprecated. Use the SHOW statement. |
| <code>getExecutionProperty</code> | This method has been deprecated. Use the SHOW statement. |
| <code>getPlanDescription</code> | Get the query plan description if the statement was last executed with SHOWPLAN ON DEBUG. The plan is a tree made up of <code>org.teiid.client.plan.PlanNode</code> objects. Typically <code>PlanNode.toString()</code> or <code>PlanNode.toXml()</code> will be used to convert the plan into a textual form. |
| <code>getRequestIdentifier</code> | Get an identifier for the last command executed on this statement. If no command has been executed yet, null is returned. |

| Method Name | Description |
|-----------------------------|---|
| setExecutionProperty | <div style="border: 1px solid gray; padding: 5px; margin: 5px 0;">  <p>Note</p> <p>This method has been deprecated. Use the SET statement.</p> </div> |
| setPayload | Set a per-command payload to pass to translators. Currently the only built-in use is for sending hints for an Oracle data source. |

5.3. Non-blocking Statement Execution

JDBC query execution can indefinitely block the calling thread when a statement is executed or a resultset is being iterated. In some situations you may wish to have your calling threads held in these blocked states. When using embedded connections, you may optionally use the **org.teiid.jdbc.TeiidStatement** and **org.teiid.jdbc.TeiidPreparedStatement** interfaces to execute queries with a callback **org.teiid.jdbc.StatementCallback** that will be notified of statement events, such as an available row, an exception, or completion. Your calling thread will be free to perform other work. The callback will be executed by an engine processing thread as needed. If your results processing is blocking and you want query processing to run concurrently with results processing, then your callback should implement `onRow` handling in a multi-threaded manner to allow the engine thread to continue.

```

PreparedStatement stmt = connection.prepareStatement(sql);
TeiidPreparedStatement tStmt = stmt.unwrap(TeiidPreparedStatement.class);
tStmt.submitExecute(new StatementCallback() {
    @Override
    public void onRow(Statement s, ResultSet rs) {
        //any logic that accesses the current row ...
        System.out.println(rs.getString(1));
    }

    @Override
    public void onException(Statement s, Exception e) throws Exception {
        s.close();
    }

    @Override
    public void onComplete(Statement s) throws Exception {
        s.close();
    }, new RequestOptions()
});

```



Note

The non-blocking logic is limited to statement execution only. Other JDBC operations, such as connection creation or batched executions do not yet have non-blocking options.

If you access forward positions in the `onRow` method (calling `next`, `isLast`, `isAfterLast`, `absolute`), they may not yet be valid and a **org.teiid.jdbc.AsynchPositioningException** will be thrown. That exception

is recoverable if caught or can be avoided by calling `TeiidResultSet.available()` to determine if your desired positioning will be valid.

5.4. Continuous Execution

The `RequestOptions` object may be used to specify a special type of continuous async execution via the `continuous` or `setContinuous` methods. In continuous mode the statement will be continuously re-executed. This is intended for consuming real-time or other data streams processed through a SQL plan. A continuous query will only terminate on an error or when the statement is explicitly closed. The SQL for a continuous query is no different than any other statement. Care should be taken to ensure that retrievals from non-continuous sources is appropriately cached for reuse, such as by using materialized views or session scoped temp tables. A continuous query must return a result set, must be executed with a forward-only resultset, and cannot be used in the scope of a transaction. Since resource consumption is expected to be different in a continuous plan, it does not count against the server max active plan limit. Typically custom sources will be used to provide data streams.

When the client wishes to end the continuous query, the `Statement.close()` or `Statement.cancel()` method should be called. Typically your callback will close whenever it no longer needs to process results.

See also the `ContinuousStatementCallback` for use as the `StatementCallback` for additional methods related to continuous processing.

5.5. Execution Properties

The following table provides a list of execution properties as defined in `org.teiid.jdbc.ExecutionProperties`. These can be modified using the SET statement.

Table 5.2. Execution Properties

| Constant Identifier | String Value | Description |
|---|--|---|
| <code>ANSI_QUOTED_IDENTIFIERIFIERS</code> | <code>ansiQuotedIdentifiers</code> | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . |
| <code>DISABLE_LOCAL_TRANSACTIONS</code> | <code>disableLocalTxn</code> | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . |
| <code>JDBC4COLUMNNAMEANDLABELSEMANTICS</code> | <code>useJDBC4ColumnNameAndLabelSemantics</code> | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . |
| <code>NOEXEC</code> | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . | |
| <code>PROP_FETCH_SIZE</code> | <code>fetchSize</code> | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . |
| <code>PROP_PARTIAL_RESULTS_MODE</code> | <code>partialResultsMode</code> | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” and Section 5.12, “Partial Results Mode” . |
| <code>PROP_TXN_AUTOWRAP</code> | <code>autoCommitTxn</code> | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . |

| Constant Identifier | String Value | Description |
|------------------------------|---------------------------|---|
| PROP_XML_FORMAT | XMLFormat | Determines the formatting of XML documents returned by XML document models. Can be one of XML_COMPACT_FORMAT or XML_TREE_FORMAT . See Section 5.7, “XML Document Formatting” for more information. |
| PROP_XML_VALIDATION | XMLValidation | Determines whether XML documents returned by XML document models will be validated against their schema after processing. See Section 5.8, “XML Schema Validation” and topics on “XML SELECT” in the <i>JBoss Data Virtualization Development Guide: Reference Material</i> for more information. |
| QUERYTIMEOUT | QUERYTIMEOUT | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . |
| RESULT_SET_CACHE_MODE | resultSetCacheMode | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . |
| SQL_OPTION_SHOWPLAN | SHOWPLAN | See Section 2.10, “Connection Properties for the Driver and Data Source Classes” . |

See Also:

- ✦ [Section 5.9, “The SET Statement”](#)

5.6. XML Extensions

JBoss Data Virtualization provides some XML extensions defined in `org.teiid.jdbc.ExecutionProperties`:

- ✦ **PROP_XML_FORMAT** (defined as XMLFormat)
- ✦ **PROP_XML_VALIDATION** (defined as XMLValidation)



Note

These extensions apply to XML results from queries to XML document models, but not to XML produced by SQL/XML or read from some other source.

5.7. XML Document Formatting

The **PROP_XML_FORMAT** execution property, defined in `org.teiid.jdbc.ExecutionProperties`, can be set using the SET statement to modify the way that XML documents are formatted from XML document models. The following valid values are also defined:

XML_TREE_FORMAT

XML_TREE_FORMAT (defined as Tree) returns a version of the XML formatted for display. The XML will use line breaks and tabs as appropriate to format the XML as a tree. This format is slower due to the formatting time and the larger document size.

XML_COMPACT_FORMAT

XML_COMPACT_FORMAT (defined as Compact) returns a version of the XML formatted for optimal performance. The XML is a single long string without any unnecessary white space.



Important

If the **XML_COMPACT_FORMAT** execution property is not set, the formatting flag of the XML document in the original model is honored. This may produce either the tree or compact form of the document depending on the document setting.

5.8. XML Schema Validation

The **PROP_XML_VALIDATION** execution property, defined in **org.teiid.jdbc.ExecutionProperties**, can be set using the SET statement to indicate that the server should validate XML document model documents against their schema before returning them to the client. If schema validation is on, then the server sends an SQLWarning if the document does not conform to the schema it is associated with.



Note

Using schema validation will reduce the performance of your XML queries.

5.9. The SET Statement

Execution properties are set on the connection using the SET statement. The SET statement is not yet a language feature of JBoss Data Virtualization and is handled only in the JDBC client.

SET Syntax:

- ✦ SET [PAYLOAD] (parameter|SESSION AUTHORIZATION) value
- ✦ SET SESSION CHARACTERISTICS AS TRANSACTION ISOLATION LEVEL (READ UNCOMMITTED|READ COMMITTED|REPEATABLE READ|SERIALIZABLE)

Syntax Rules:

- ✦ The parameter must be an identifier - it cannot contain spaces or other special characters only if quoted.
- ✦ The value may be either a non-quoted identifier or a quoted string literal value.
- ✦ If payload is specified, for example, **SET PAYLOAD x y**, then a session scoped payload properties object will have the corresponding name value pair set. The payload object is not fully session scoped. It will be removed from the session when the XAConnection handle is closed/returned to the pool (assumes the use of TeiidDataSource). The session scoped payload is superseded by usage of the **TeiidStatement.setPayload**.

Using SET SESSION CHARACTERISTICS AS TRANSACTION ISOLATION LEVEL is equivalent to calling Connection.setTransactionIsolation with the corresponding level.

The SET statement is most commonly used to control planning and execution.

- ✦ SET SHOWPLAN (ON|DEBUG|OFF)
- ✦ SET NOEXEC (ON|OFF)

The following is an example of how to use the SET statement to enable a debug plan:

```
Statement s = connection.createStatement();
s.execute("SET SHOWPLAN DEBUG");

Statement s1 = connection.createStatement();
ResultSet rs = s1.executeQuery("select col from table");

ResultSet planRs = s1.executeQuery("SHOW PLAN");
planRs.next();
String debugLog = planRs.getString("DEBUG_LOG");
Query Plan without executing the query
s.execute("SET NOEXEC ON");
s.execute("SET SHOWPLAN DEBUG");
...
e.execute("SET NOEXEC OFF");
```

The SET statement may also be used to control authorization. A SET SESSION AUTHORIZATION statement will perform a reauthentication (see [Section 4.6, "Reauthentication"](#)) given the credentials currently set on the connection.

The connection credentials may be changed by issuing a SET PASSWORD statement.

```
Statement s = connection.createStatement();
s.execute("SET PASSWORD 'someval'");
s.execute("SET SESSION AUTHORIZATION 'newuser'");
```

5.10. The SHOW Statement

The SHOW statement can be used to see a variety of information. The SHOW statement is not yet a language feature of JBoss Data Virtualization and is handled only in the JDBC client.

SHOW PLAN

SHOW PLAN returns a resultset with a CLOB column PLAN_TEXT, an xml column PLAN_XML, and a CLOB column DEBUG_LOG with a row containing the values from the previously executed query. If SHOWPLAN is OFF or no plan is available, no rows are returned. If SHOWPLAN is not set to DEBUG, then DEBUG_LOG will return a null value.

SHOW ANNOTATIONS

SHOW ANNOTATIONS returns a resultset with string columns CATEGORY, PRIORITY, ANNOTATION, RESOLUTION and a row for each annotation on the previously executed query. If SHOWPLAN is OFF or no plan is available, no rows are returned.

SHOW <property>

SHOW <property> is the inverse of SET and shows the property value for the property supplied. It returns a resultset with a single string column with a name matching the property key.

SHOW ALL

SHOW ALL returns a resultset with a NAME string column and a VALUE string column with a row entry for every property value.

5.11. Transaction Statements

5.11. Transaction Statements

In situations where direct use of the JDBC connection is not possible, transaction statements can be used to control a local transaction.

START TRANSACTION

START TRANSACTION is a synonym for `connection.setAutoCommit(false)`

COMMIT

COMMIT is a synonym for `connection.setAutoCommit(true)`

ROLLBACK

ROLLBACK is a synonym for `connection.rollback()` and returning to auto commit mode.

5.12. Partial Results Mode

JBoss Data Virtualization supports a partial results query mode. This mode changes the behavior of the query processor so the server returns results even when some data sources are unavailable.

For example, suppose that two data sources exist for different suppliers and your data designers have created a virtual group that creates a union between the information from the two suppliers. If your application submits a query without using partial results query mode and one of the suppliers' databases is down, the query against the virtual group returns an exception. However, if your application runs the same query in partial results query mode, the server returns data from the running data source and no data from the data source that is down.

When using partial results mode, if a source throws an exception during processing it does not cause the user's query to fail. Rather, that source is treated as returning no more rows after the failure point. Most commonly, that source will return 0 rows.

This behavior is most useful when using **UNION** or **OUTER JOIN** queries as these operations handle missing information in a useful way. Most other kinds of queries will simply return 0 rows to the user when used in partial results mode and the source is unavailable.



Note

In some instances, (typically when you are using JDBC sources), if the source is not available initially, its absence will prevent Teiid from automatically determining the appropriate set of source capabilities. If you see an exception indicating that the capabilities for an unavailable source are not valid in partial results mode, then it may be necessary to manually set the database version or similar property on the translator to ensure that the capabilities are known even if the source is not available.

5.13. Setting Partial Results Mode

Partial results mode is off by default but you can turn it on for all queries in a connection by using either `setPartialResultsMode(true)` on a `DataSource` or `partialResultsMode=true` on a JDBC URL. In either case, you can toggle partial results mode on or off later with a SET statement.

This is how you configure the partial results mode using the SET statement:

```
Statement statement = ...obtain statement from Connection...
statement.execute("set partialResultsMode true");
```

5.14. Partial Results Warnings

For each source that is excluded from the query, a warning will be generated describing the source and the failure. These warnings can be obtained from the **Statement.getWarnings()** method. This method returns a **SQLWarning** object, but in the case of partial results warnings this object will be an instance of the **org.teiid.jdbc.PartialResultsWarning** class. This class can be used to obtain a list of all the failed sources by name and to obtain the specific exception thrown by each resource adaptor.



Note

Since JBoss Data Virtualization supports cursoring before the entire result is formed, it is possible that a data source failure will not be determined until after the first batch of results have been returned to the client. This can happen in the case of unions, but not joins. To ensure that all warnings have been accumulated, the statement should be checked after the entire result set has been read.

The following is an example of how to obtain partial results warnings:

```
statement.execute("set partialResultsMode true");
ResultSet results = statement.executeQuery("SELECT Name FROM Accounts");
while (results.next())
{
    //process the result set
}

SQLWarning warning = statement.getWarnings();

if(warning instanceof PartialResultsWarning)
{
    PartialResultsWarning partialWarning = (PartialResultsWarning)warning;
    Collection failedConnectors = partialWarning.getFailedConnectors();
    Iterator iter = failedConnectors.iterator();
    while(iter.hasNext())
    {
        String connectorName = (String) iter.next();
        SQLException connectorException =
partialWarning.getConnectorException(connectorName);
        System.out.println(connectorName + ": " +
connectorException.getMessage());
    }
}
```

Chapter 6. JDBC Transactions

6.1. JDBC Transaction Types

The JBoss Data Virtualization JDBC API supports three types of transactions from a client perspective:

- ✧ global transactions,
- ✧ local transactions, and
- ✧ request level transactions.

All are implemented by JBoss Data Virtualization as XA transactions. Refer to the JTA specification at <http://www.oracle.com/technetwork/java/javaee/tech/jta-138684.html> for more information on XA Transactions.



Warning

The use of global, local, and request level transactions are all mutually exclusive. Request level transactions only apply when not in a global or local transaction. Any attempt to mix global and local transactions concurrently will result in an exception.

6.2. Local Transactions

A connection uses the **autoCommit** flag to explicitly control local transactions. By default, autoCommit is set to **true**, which indicates request level or implicit transaction control:

```
// Set auto commit to false and start a transaction
connection.setAutoCommit(false);

try
{
    // Execute multiple updates
    Statement statement = connection.createStatement();
    statement.executeUpdate("INSERT INTO Accounts (ID, Name) VALUES (10,
'Mike')");
    statement.executeUpdate("INSERT INTO Accounts (ID, Name) VALUES (15,
'John')");
    statement.close();

    // Commit the transaction
    connection.commit();
}
catch(SQLException e)
{
    // If an error occurs, rollback the transaction
    connection.rollback();
}
```

This example demonstrates several things:

1. Setting **autoCommit** flag to false. This will start a transaction bound to the connection.

2. Executing multiple updates within the context of the transaction.
3. When the statements are complete, the transaction is committed by calling `commit()`.
4. If an error occurs, the transaction is rolled back using the `rollback()` method.

6.3. Ending Local Transactions

Any of the following operations will end a local transaction:

1. `Connection.setAutoCommit(true)` if previously set to false
2. `Connection.commit()`
3. `Connection.rollback()`
4. A transaction will be rolled back automatically if it times out.

6.4. Turning Off Local Transactions

In some cases, tools or frameworks above JBoss Data Virtualization will call `setAutoCommit(false)`, `commit()` and `rollback()` even when all access is read-only and no transactions are necessary. In the scope of a local transaction, JBoss Data Virtualization will start and attempt to commit an XA transaction, possibly complicating configuration or causing performance degradation.

In these cases, you can override the default JDBC behavior to indicate that these methods should perform no action regardless of the commands being executed. To turn off the use of local transactions, add the following property to the JDBC connection URL:

```
disableLocalTxn=true
```



Warning

Turning off local transactions can be dangerous and can result in inconsistent results when reading or inconsistent data in data stores when writing. For safety, this mode should be used only if you are certain that the calling application does not need local transactions.

6.5. Request Level Transactions

Request level transactions are used when the request is not in the scope of a global or local transaction, which implies `autoCommit` is `true`. In a request level transaction, your application does not need to explicitly call `commit` or `rollback`, rather every command is assumed to be its own transaction that will automatically be committed or rolled back by the server.

JBoss Data Virtualization can perform updates through virtual tables. These updates might result in an update against multiple physical systems, even though the application issues the update command against a single virtual table. Often, a user might not know whether the queried tables actually update multiple sources and require a transaction.

For that reason, JBoss Data Virtualization allows your application to automatically wrap commands in transactions when necessary. Because this wrapping incurs a performance penalty for your queries, you can choose from a number of available wrapping modes to suit your environment. You need to choose between

the highest degree of integrity and performance your application needs. For example, if your data sources are not transaction-compliant, you might turn transaction wrapping off to maximize performance.

6.6. Transaction Wrapping Modes

You can set your transaction wrapping to one of the following modes:

1. ON

This mode always wraps every command in a transaction without checking whether it is required. This is the safest mode.

2. OFF

This mode never automatically wraps a command in a transaction or checks whether it needs to wrap a command. This mode can be dangerous as it will allow multiple source updates outside of a transaction without an error. This mode has best performance for applications that do not use updates or transactions.

3. DETECT

This mode assumes that the user does not know how to execute multiple source updates in a transaction. JBoss Data Virtualization checks every command to see whether it is a multiple source update and wraps it in a transaction. If it is single source then it uses the source level command transaction.

6.7. Set the Transaction Wrapping Mode

You can set the transaction mode as a property when you establish the connection using:

- the `autoCommitTxn` property in the connection URL (see [Section 2.10, “Connection Properties for the Driver and Data Source Classes”](#)),
- the `setAutoCommitTxn` method (see [Section 2.10, “Connection Properties for the Driver and Data Source Classes”](#)),
- or on a per-query basis, using the SET statement with the `PROP_TXN_AUTO_WRAP` property (see [Section 5.5, “Execution Properties”](#)).

6.8. Multiple Insert Batches

If you work with request-level transactions and issue INSERTs with a query expression (or the deprecated SELECT INTO), you may find that multiple insert batches handled by separate source INSERTs are processed by the Red Hat JBoss Data Virtualization server. Take care to ensure that targeted sources support XA or that compensating actions are taken in the event of a failure.

6.9. Global Transactions

Global or client XA transactions allow the JBoss Data Virtualization JDBC API to participate in transactions that are beyond the scope of a single client resource. For this, use the `org.teiid.jdbc.TeiidDataSource` class for establishing connections.

When the data source class is used in the context of a user transaction in an application server, such as JBoss, WebSphere, or Weblogic, the resulting connection will already be associated with the current XA transaction. No additional client JDBC code is necessary to interact with the XA transaction.

The following code demonstrates usage of UserTransactions.

```
UserTransaction ut = context.getUserTransaction();
try {
    ut.begin();
    Datasource oracle = lookup(...)
    Datasource teiid = lookup(...)
    Connection c1 = oracle.getConnection();
    Connection c2 = teiid.getConnection();
    // do something with Oracle connection
    // do something with Teiid connection
    c1.close();
    c2.close();
    ut.commit();
} catch (Exception ex) {
    ut.rollback();
}
```

The following code demonstrates manual usage of XA transactions.

```
XAConnection xaConn = null;
XAResource xaRes = null;
Connection conn = null;
Statement stmt = null;

try
{
    xaConn = <XADataSource instance>.getXAConnection();
    xaRes = xaConn.getXAResource();
    Xid xid = <new Xid instance>;
    conn = xaConn.getConnection();
    stmt = conn.createStatement();

    xaRes.start(xid, XAResource.TMNOFLAGS);
    stmt.executeUpdate("insert into â!");
    // other statements on this connection or other resources enlisted in
    this transaction
    // ...
    xaRes.end(xid, XAResource.TMSUCCESS);

    if (xaRes.prepare(xid) == XAResource.XA_OK)
    {
        xaRes.commit(xid, false);
    }
}
catch (XAException e)
{
    xaRes.rollback(xid);
}
finally
```



```
{  
    // clean up code  
    // ...  
}
```

With the use of global transactions, multiple XAConnections may participate in the same transaction. It is important to note that the JDBC XAResource `isSameRM()` method only returns `true` if connections are made to the same server instance in a cluster. If the JBoss Data Virtualization connections are to different server instances then transactional behavior may not be the same as if they were to the same cluster member. For example, if the client transaction manager uses the same **XID** for each connection, duplicate XID exceptions may arise from the same physical source accessed through different cluster members. If the client transaction manager uses a different branch identifier for each connection, issues may arise with sources that lock or isolate changes based upon branch identifiers.

6.10. Enterprise Information System Support

The underlying resource adaptors that represent the Enterprise Information System (EIS) and the EIS itself must support XA transactions to be able to participate in distributed XA transaction using JBoss Data Virtualization. If a source system does not support the XA, then it can not participate in the distributed transaction. However, the source is still eligible to participate in data integration without XA support.

Participation in an XA transaction is automatically determined based on the resource adaptor's XA capability. It is the developer's responsibility to ensure they configure an XA resource when they require them to participate in distributed transactions.

Chapter 7. Client SSL Connections

7.1. SSL Client Connections

You need to define certain properties for each SSL mode.



Important

When connecting to the Red Hat JBoss Data Virtualization server with SSL enabled, you must use the "mms" protocol, instead of "mm" in the JDBC connection URL:

```
jdbc:teiid:<myVdb>@mms://<host>:<port>
```

There are two different sets of properties that a client can configure to enable 1-way or 2-way SSL.

The first option is to use Java SSL properties. These are standard Java defined system properties to configure the SSL under any JVM, Teiid is not unique in its use of SSL. Provide the following system properties to the client VM process.

1-way SSL:

```
-Djavax.net.ssl.trustStore=<dir>/server.truststore (required)
-Djavax.net.ssl.trustStorePassword=<password> (optional)
-Djavax.net.ssl.keyStoreType (optional)
```

2-way SSL:

```
-Djavax.net.ssl.keyStore=<dir>/client.keystore (required)
-Djavax.net.ssl.keyStorePassword=<password> (optional)
-Djavax.net.ssl.trustStore=<dir>/server.truststore (required)
-Djavax.net.ssl.trustStorePassword=<password> (optional)
-Djavax.net.ssl.keyStoreType=<keystore type> (optional)
```

The second option is to use Teiid-specific properties. Use this option when the above "javax" based properties are already in use by the host process. For example if your client application is a Tomcat process that is configured for https protocol and the above Java based properties are already in use, and importing Teiid-specific certificate keys into those https certificate keystores is not allowed.

In this scenario, a different set of Teiid-specific SSL properties can be set as system properties or defined inside the a "teiid-client-settings.properties" file. A sample "teiid-client-settings.properties" file can be found inside the "teiid-version-client.jar" file at the root called "teiid-client-settings.orig.properties". Extract this file, make a copy, change the property values required for the chosen SSL mode, and place this file in the client application's classpath before the "teiid-version-client.jar" file.

Here are the SSL properties and definitions that can be set in a **teiid-client-settings.properties** file:

```
#####
# SSL Settings
#####
```

```
#
# The key store type. Defaults to JKS
#

org.teiid.ssl.keyStoreType=JKS

#
# The key store algorithm, defaults to
# the system property "ssl.TrustManagerFactory.algorithm"
#

#org.teiid.ssl.algorithm=

#
# The classpath or filesystem location of the
# key store.
#
# This property is required only if performing 2-way
# authentication that requires a specific private
# key.
#

#org.teiid.ssl.keyStore=

#
# The key store password (not required)
#

#org.teiid.ssl.keyStorePassword=

#
# The key alias(not required, if given named certificate is used)
#

#org.teiid.ssl.keyAlias=

#
# The key password(not required, used if the key password is different than
# the keystore password)
#

#org.teiid.ssl.keyPassword=

#
# The classpath or filesystem location of the
# trust store.
#
# This property is required if performing 1-way
# authentication that requires trust not provided
# by the system defaults.
#

#org.teiid.ssl.trustStore=

#
# The trust store password (not required)
```

```
#  
  
#org.teiid.ssl.trustStorePassword=  
  
#  
# The cipher protocol, defaults to TLSv3  
#  
  
org.teiid.ssl.protocol=TLSv1  
  
#  
# Whether to allow anonymous SSL  
# (the TLS_DH_anon_WITH_AES_128_CBC_SHA cipher suite)  
# defaults to true  
#  
  
org.teiid.ssl.allowAnon=true  
1-way SSL  
org.teiid.ssl.trustStore=<dir>/server.truststore (required)  
2-way SSL  
org.teiid.ssl.keyStore=<dir>/client.keystore (required)  
org.teiid.ssl.trustStore=<dir>/server.truststore (required)
```

Chapter 8. Using Hibernate with JBoss Data Virtualization

8.1. Configure Hibernate for Use with JBoss Data Virtualization

Prerequisites

- ✦ You must have the JBoss Data Virtualization JDBC API client JAR file (**teiid-client.jar**) and the JBoss Data Virtualization hibernate dialect JAR file (**teiid-hibernate-dialect-VERSION.jar**) in Hibernate's classpath. These files are found in **EAP_HOME/modules/system/layers/dv/org/jboss/teiid/client/main/**.

These are required for the **org.teiid.dialect.TeiidDialect**, **org.teiid.jdbc.TeiidDriver** and **org.teiid.jdbc.TeiidDataSource** classes.

Procedure 8.1. Configure Hibernate for Use with JBoss Data Virtualization

1. Open the Hibernate configuration file

Open the **hibernate.cfg.xml** file.

2. Specify the JBoss Data Virtualization driver class

Specify the JBoss Data Virtualization driver class in the **connection.driver_class** property:

```
<property name="connection.driver_class">
    org.teiid.jdbc.TeiidDriver
</property>
```

3. Set the Connection URL

Specify the URL for the VDB in the **connection.url** property:

```
<property name="connection.url">
    jdbc:teiid:VDB-NAME@mm://HOST:POST;user=USERNAME;password=PASSWORD
</property>
```



Note

Be sure to use a local connection if Hibernate is in the same VM as the application server. .

4. Specify the dialect class

Specify the JBoss Data Virtualization dialect class in the **dialect** property:

```
<property name="dialect">
    org.teiid.dialect.TeiidDialect
</property>
```



Note

Alternatively, the connection properties can be added to the **hibernate.properties** file instead of **hibernate.cfg.xml**:

```
hibernate.connection.driver_class=org.teiid.jdbc.TeiidDriver
hibernate.connection.url=jdbc:teiid:VDB-NAME@mm://HOST:PORT
hibernate.connection.username=USERNAME
hibernate.connection.password=PASSWORD
hibernate.dialect=org.teiid.dialect.TeiidDialect
```



Note

Since your VDBs will likely contain multiple source and view models with identical table names, you will need to fully qualify table names specified in Hibernate mapping files:

```
<class name="CLASSNAME" table="SOURCE/VIEW_MODEL_NAME .
[SCHEMA_NAME.]TABLENAME">
    ...
</class>
```

For example:

```
<class name="org.teiid.example.Publisher"
table="BOOKS.BOOKS.PUBLISHERS">
    ...
</class>
```

8.2. Limitations of Using Hibernate with JBoss Data Virtualization

Many Hibernate use cases assume a data source has the ability (with proper user permissions) to process Data Definition Language (DDL) statements like CREATE TABLE and DROP TABLE as well as Data Manipulation Language (DML) statements like SELECT, UPDATE, INSERT and DELETE. JBoss Data Virtualization can handle a broad range of DML, but does not directly support DDL against a particular source.

Sequence and Identity generation are not supported. Identifier generation based upon table values, such as the hilo generator, require that the identifier table(s) be exposed through JBoss Data Virtualization. The GUID identifier generation strategy is directly supported.

Chapter 9. OData Support

9.1. OData Support

9.1.1. OData

OData (Open Data Protocol) is a web protocol that utilizes URIs to identify and interact with resources allowing for a great level of data integration and interoperability across clients, servers and services.

9.1.2. Use Case

REST is commonly used by companies that have implemented service-oriented architectures. However, it does not provide unified calling semantics nor does it provide a data model, meaning that each company has needed to define its own. OData is a solution to this problem because it provides a specification that defines standard ways to define data source operations and also standardizes the way in which you define your data schema.

By standardising your service and providing metadata that applications can use, you can avoid vendor lock-in.

9.1.3. Implementation in Red Hat JBoss Data Virtualization

When you deploy a virtual database, the OData protocol is supported by default. OData support is implemented and deployed through two WAR files that you can find at `EAP_HOME/dataVirtualization/vdb/teiid-odata.war` and `EAP_HOME/dataVirtualization/vdb/teiid-olingo-odata4.war`.

9.1.4. Accessing Information via OData

If you have a virtual database called "northwind" deployed and you wish to access a customers table in an NW model, use HTTP GET via this URL:

```
http://localhost:8080/odata/northwind.1/NW.customers
```

The equivalent SQL command via a JDBC/ODBC connection is this:

```
SELECT * FROM NW.customers
```



Note

Make sure that you fully qualify the table name along with the model name. Also, use the correct case (upper or lower) to match what you have in the virtual database.

You can output your query results in AtomPub XML (the default), Atom or JSON format.

**Note**

Sometimes you may see this message:

```
<error
  xmlns="http://schemas.microsoft.com/ado/2007/08/dataservices/metadata">
  <code>NotFoundException</code>
  <message lang="en-US">EdmEntitySet NW.customer is not found</message>
</error>
```

This means you either supplied an incorrect model-name.table-name combination so please check the spelling and case. Alternatively, it may mean that your table did not have any PRIMARY KEY or UNIQUE KEY(s) on them. Since OData access is key oriented, every table must have a primary key or at least one unique key. (Note that Red Hat JBoss Data Virtualization does support composite primary keys.)

If you do not see all the rows, issue another call with the same URL, but with the \$skiptoken query option specified from the previous result:

```
http://localhost:8080/odata/northwind.1/NW.customers?$skiptoken=xxx
```

You can submit criteria with your query to filter the results:

```
http://localhost:8080/odata/northwind.1/NW.customers?$filter=name eq 'bob'
```

This is the equivalent of making a JDBC/ODBC connection and issuing the SQL

```
SELECT * FROM NW.customers where name = 'bob'
```

To request the result in JSON format, add this option: \$format=json. Here is an example:

```
http://localhost:8080/odata/northwind.1/NW.customers?$format=JSON
```

You can combine query options as needed. Here is an example:

```
http://localhost:8080/odata/northwind.1/NW.customers?$filter=name eq
'bob'&$format=JSON
```

You can navigate from one entity to another like this:

```
http://localhost:8080/odata/northwind.1/NW.customers(1234)/NW.orders?
$filter=orderdate gt datetime'2012-12-31T21:23:38Z'
```

That is the equivalent to this SQL query:

```
SELECT o.* FROM NW.orders o join NW.customers c on o.customer_id = c.id
where c.id=1234 and o.orderdate > {ts '2012-12-31 21:23:38'}
```

You can also use OData to perform CREATE/UPDATE/DELETE operations along with READ operations. Here are the HTTP methods involved:

INSERT/CREATE is accomplished through the "POST" HTTP method:

```
POST /service.svc/Customers HTTP/1.1
Host: host
Content-Type: application/atom+xml
Accept: application/atom+xml
Content-Length: nnn
<?xml version="1.0" encoding="utf-8"?>
<entry xmlns:d="http://schemas.microsoft.com/ado/2007/08/dataservices"
xmlns:m="http://schemas.microsoft.com/ado/2007/08/dataservices/metadata"
xmlns="http://www.w3.org/2005/Atom">
  <id>http://host/service.svc/Customers('ASDFG')</id>
  <title type="text" />
  <updated>2008-12-07T8:00:00Z</updated>
  <author>
    <name />
  </author>
  <link
rel="http://schemas.microsoft.com/ado/2007/08/dataservices/related/Orders"
href="Orders(1)" />
  <link
rel="http://schemas.microsoft.com/ado/2007/08/dataservices/related/Orders"
href="Orders(2)" />
  <content type="application/xml">
    <m:properties>
      <d:CustomerID>ASDFG</d:CustomerID>
      <d:CompanyName>Contoso Widgets</d:CompanyName>
      <d:Address>
        <d:Street>58 Contoso St</d:Street>
        <d:City>Seattle</d:City>
      </d:Address>
    </m:properties>
  </content>
</entry>
```

An UPDATE is performed with an HTTP "PUT" command. Here is an example:

```
PUT /service.svc/Customers('ALFKI') HTTP/1.1
Host: host
Content-Type: application/atom+xml
Accept: application/atom+xml
Content-Length: nnn
DataServiceVersion: 1.0
MaxDataServiceVersion: 3.0
Prefer: return-content
<?xml version="1.0" encoding="utf-8"?>
<entry xmlns:d="http://schemas.microsoft.com/ado/2007/08/dataservices"
xmlns:m="http://schemas.microsoft.com/ado/2007/08/dataservices/metadata"
xmlns="http://www.w3.org/2005/Atom">
  <content type="application/xml">
    <m:properties>
      <d:CustomerID>ALFKI</d:CustomerID>
      <d:CompanyName>Updated Company Name</d:CompanyName>
      <d:Address>
        <d:Street>Updated Street</d:Street>
```

```

        </d:Address>
    </m:properties>
</content>
</entry>

```

Here is an example delete command:

```

DELETE /service.svc/Customers('ALFKI') HTTP/1.1
Host: host
Content-Type: application/atom+xml
Accept: application/atom+xml
Content-Length: nnn
DataServiceVersion: 1.0

```

9.1.5. Security

OData access is restricted by HTTPBasic authentication which is authenticated against the "teiid-security" domain. You must assign users the odata role.

In order to change the default security domain, use a JBoss AS deployment-overlay to override the web.xml file.

9.1.6. Proxy and Cloud Environments

If the Red Hat JBoss Data Virtualization server is configured behind a proxy server or deployed in cloud environment, or using a load-balancer then the URI of the server which is handling the OData request is different from URI of the proxy. To generate valid links in the OData responses, you must configure the "proxy-base-uri" property in the web.xml file. If it is a system-wide property, configure it like this:

```

<context-param>
    <param-name>proxy-base-uri</param-name>
    <param-value>${system-property-name}</param-value>
</context-param>

```



Note

When the volume of result rows exceed the configured batch size, cursoring logic is implemented. On every request, only batch-size number of rows are returned. Each such request is considered an active cursor, with a specified amount of idle time specified by the skip-token-cache-time parameter. After the cursor has timed out, it closes and the remaining results and will not be available for any further queries. Since there is no session-based tracking of these cursors, if the request for skiptoken comes after the expiry, the original query will be executed again and tries to reposition the cursor to relative absolute position, however the results are not guaranteed to be same as the underlying sources may have been updated with new information meanwhile.

9.1.7. How Red Hat JBoss Data Virtualization exposes schema for OData

OData defines its schema using Conceptual Schema Definition Language (CSDL). Every VDB, that is deployed in an ACTIVE state in the Red Hat JBoss Data Virtualization server exposes its metadata in CSDL format. For example if you want retrieve metadata for the Northwind database, you need to issue a query like this:

```
http://localhost:8080/odata/northwind/$metadata
```

Since the OData schema model is not a relational schema model, Red Hat JBoss Data Virtualization maps its relational schema model to OData's schema model like this:

Table 9.1. Table Mappings

| Relational Entity | Mapped OData Entity |
|--------------------------|---|
| Model Name | Schema Namespace, EntityContainer Name |
| Table/View | EntityType, EntitySet |
| Table Columns | EntityType's Properties |
| Primary Key | EntityType's Key Properties |
| Foreign Key | FunctionImport |
| Procedure | Navigation Property on EntityType, Association, AssosiationSet |
| Procedure's Table Return | ComplexType. (Note that Red Hat JBoss Data Virtualization does not define any "embedded" ComplexType in the EntityType. It also does not define any one EntityContainer that resulted from different VDB models as a default container. Therefore, you can only access entities by supplying their full paths.) |

9.1.8. Native Queries

Native and direct queries are not supported. However, you can use the web service translator's `invokehttp` method directly to issue a REST-based call and parse the results of this via SQLXML.

9.1.9. Server Use

Red Hat JBoss Data Virtualization can expose any data source as an OData-based web service.

9.1.10. Configuration

The OData WAR file can be configured with the following properties in the `web.xml` file.

Table 9.2. web.xml Settings

| Property Name | Description | Default Value |
|--|---|--|
| <code>batch-size</code> | Number of rows to send back each time, -1 returns all rows | 256 |
| <code>skiptoken-cache-time</code> | Time interval between the results being recycled/expired between \$skiptoken requests | 300000 |
| <code>local-transport-name</code> | Data Virtualization local transport name for connection | odata |
| <code>invalid-xml10-character-replacement</code> | Replacement string if an invalid XML 1.0 character appears in the data - note that this replacement will occur even if JSON is requested. | No value (the default) means that an exception will be thrown with XML results if such a character is encountered. |

| Property Name | Description | Default Value |
|----------------|---|---------------|
| proxy-base-uri | Defines the proxy server's URI to be used in OData responses. | n/a |
| connection.XXX | Sets XXX as an execution property on the local connection. Can be used for example to enable result set cache mode. | n/a |

Red Hat JBoss Data Virtualization OData server implements cursoring logic when the result rows exceed the configured batch size. On every request, only batch-size number of rows are returned. Each such request is considered an active cursor, with a specified amount of idle time specified by skip-token-cache-time. After the cursor is timed out, the cursor will be closed and remaining results will be cleaned up, and will no longer be available for further queries. Since there is no session based tracking of these cursors, if the request for skiptoken comes after the expired time, the original query will be executed again and tries to reposition the cursor to relative absolute position, however the results are not guaranteed to be same as the underlying sources may have been updated with new information meanwhile.

9.1.11. Limitations

The following feature limitations currently apply.

- JBoss Data Virtualization now supports OData Version 4.
- Blob support for media types are not supported.
- \$value construct to retrieve individual column value is supported.
- create/update/delete \$links is not supported.
- \$expand support is limited to a single parent/child.

9.1.12. Client Tools for Access

There are various ways you write your OData access layer. Your choice depends upon your needs and your programming model. Here are some suggestions:

- Your Browser: The OData Explorer is an online tool for browsing an OData data service.
- Microsoft .NET Framework 3.51: the WCF Data Services framework is available as a separate download for .NET 3.x.
- Microsoft .NET Framework 4.0: the WCF Data Services framework built into .NET 4.0.
- Silverlight 3: the Data Services client library for Silverlight is available for download.
- Java: the Restlet 2.0 library for Java (including Java on your Android phone) supports the OData protocol.
- Java: Use a library like OData4J for Java based access, or any Rest based framework
- JavaScript: the XMLHttpRequest object is standard in modern browsers or you can use jQuery, which comes out of the box with .NET 4.0 or is available for download.
- PHP: the Toolkit for PHP provides OData support for PHP clients.
- AJAX: if you're using AJAX for ASP.NET, Microsoft provides the ASP.NET Ajax Library for getting to OData.
- Excel 2010 PowerPivot: PowerPivot comes with OData support built right in.

- ✦ Windows Desktop: LINQPad is a good tool for building OData queries interactively but it has some limitations: firstly, it is not able to handle FunctionImports (procedures). If the model contains only procedures, no tables then also it acts error. Secondly, it does not work with multiple schema, as it does not show all the schemas, only the default one. Since no default schema is set in the VDB's \$metadata, it finds the first one and uses it. OData V2 does allow multiple Schemas and multiple EntityContainers in a single \$metadata so it appears that LINQPad is not yet fully compliant with this specification.
- ✦ Shell Scripts: use CURL tool

9.2. OData Version 4.0 Support

Teiid strives to be compliant with the OData specification.

For example, if you have deployed a VDB named northwind that has a customers table in a NW model, then you can access that table with an HTTP GET via this URL:

<http://localhost:8080/odata4/northwind/NW/customers>. This is akin to making a JDBC/ODBC connection and issuing the SQL:

```
SELECT * FROM NW.customers
```



Note

Use correct case (upper or lower) in the resource path. Unlike SQL, the names used in the URI are case-sensitive.



Note

The returned results from the OData query are output in either Atom/AtomPub XML or JSON format. JSON results are returned by default.

You can submit predicates with your query to filter the results:

[http://localhost:8080/odata4/northwind/NW/customers?\\$filter=name eq 'bob'](http://localhost:8080/odata4/northwind/NW/customers?$filter=name eq 'bob')



Note

The spaces around 'eq' are for readability of the example only; in real URLs they must be percent-encoded as %20. OData mandates percent encoding for all spaces in URLs. <http://docs.oasis-open.org/odata/odata/v4.0/odata-v4.0-part2-url-conventions.html>

This is similar to making a JDBC/ODBC connection and issuing the SQL

```
SELECT * FROM NW.customers where name = 'bob'
```

To request the result to be formatted in a specific format, add the query option \$format like this:

[http://localhost:8080/odata4/northwind/NW/customers?\\$format=JSON](http://localhost:8080/odata4/northwind/NW/customers?$format=JSON)

Query options can be combined as needed. For example here is how you format with a filter:

```
http://localhost:8080/odata4/northwind/NW/customers?$filter=name eq
'bob'&$format=xml
```

OData allows for querying navigations from one entity to another. A navigation is similar to the foreign key relationships in relational databases.

For example, if the customers table has an exported key to the orders table on the customers primary key called the customer_fk, then an OData GET could be issued like this:

```
http://localhost:8080/odata4/northwind/NW/customers(1234)/customer_fk?$filter=orderdate gt 2012-12-31T21:23:38Z
```

This would be akin to making a JDBC/ODBC connection and issuing this SQL:

```
SELECT o.* FROM NW.orders o join NW.customers c on o.customer_id = c.id
where c.id=1234 and o.orderdate > {ts '2012-12-31 21:23:38'}
```



Note

For detailed protocol access you can read the specification at <http://odata.org>. You can also read this very useful web resource for an example of accessing an OData server.



Important

If you are not seeing all the rows, see the configuration section below for more details. Generally batching is being utilized, which tooling should understand automatically, and additional queries with a \$skiptoken query option specified are needed: `http://localhost:8080/odata4/northwind/NW/customers?$skiptoken=xxx`



Important

Sometimes you may encounter an "EntitySet Not Found" error. This happens when you issue the above query and you see a message like this:

```
{"error":{"code":null,"message":"Cannot find EntitySet, Singleton,
ActionImport or FunctionImport with name 'xxx'."}}
```

This message means that either you supplied an incorrect model-name/table-name combination. Check that the spelling and case are correct.

It is possible that the entity is not part of the metadata, such as when a table does not have any PRIMARY KEY or UNIQUE KEY(s).

It is possible to update your data. Using the OData protocol it is possible to perform CREATE/UPDATE/DELETE operations along with the READ operations shown above. These operations use different HTTP methods.

INSERT/CREATE is accomplished through the "POST" HTTP method. Here is an example:

```

POST /service.svc/Customers HTTP/1.1
Host: host
Content-Type: application/json
Accept: application/json
{
  "CustomerID": "AS123X",
  "CompanyName": "Contoso Widgets",
  "Address" : {
    "Street": "58 Contoso St",
    "City": "Seattle"
  }
}

```

An UPDATE is performed with an HTTP "PUT":

```

PUT /service.svc/Customers('ALFKI') HTTP/1.1
Host: host
Content-Type: application/json
Accept: application/json
{
  "CustomerID": "AS123X",
  "CompanyName": "Updated Company Name",
  "Address" : {
    "Street": "Updated Street"
  }
}

```

The DELETE operation uses the HTTP "DELETE" method.

```

DELETE /service.svc/Customers('ALFKI') HTTP/1.1
Host: host
Content-Type: application/json
Accept: application/json

```

9.2.1. Security

By default OData access is secured using HTTPBasic authentication. The user will be authenticated against Red Hat JBoss Data Virtualization's default security domain "teiid-security".



Important

Users are expected to have the odata role. Be sure to create user with this role when you are using add-user.sh script to create a new user.

However, if you wish to change the security domain use a deployment-overlay to override the web.xml file in the odata4 file in the [modules]/org/jboss/teiid/main/deployments directory.

The OData WAR can also support Kerberos, SAML and OAuth2 authentications. To learn about these, please see the Security Guide.

9.2.2. Configuration

The OData WAR file can be configured by configuring the following properties in the web.xml file:

Table 9.3. Configuring OData 4

| Property Name | Description | Default Value |
|-------------------------------------|--|---------------|
| batch-size | Number of rows to send back each time, -1 returns all rows | 256 |
| skiptoken-cache-time | Time interval between the results being recycled/expired between \$skiptoken requests | 300000 |
| local-transport-name | Data Virtualization Local transport name for connection | odata |
| invalid-xml10-character-replacement | Replacement string if an invalid XML 1.0 character appears in the data - note that this replacement will occur even if JSON is requested. No value (the default) means that an exception will be thrown with XML results if such a character is encountered. | NA |
| proxy-base-uri | Defines the proxy server's URI to be used in OData responses. | NA |
| connection.XXX | Sets XXX as an execution property on the local connection. Can be used for example to enable result set cache mode. | NA |



Note

If the Data Virtualization server is configured behind a proxy server or deployed in cloud environment, or using a load-balancer then the URI of the server which is handling the OData request is different from URI of proxy. To generate valid links in the OData responses, configure "proxy-base-uri" property in the web.xml. If this value is available as system property then define the property value as per the below.

```
<init-param>
  <param-name>proxy-base-uri</param-name>
  <param-value>${system-property-name}</param-value>
</init-param>
```

To modify the web.xml file, create a deployment-overlay using the CLI with the modified contents:

```
deployment-overlay add --name=myOverlay --content=/WEB-INF/web.xml=/modified/web.xml --deployments=teiid-odata-odata4.war --redeploy-affected
```

The Red Hat JBoss Data Virtualization OData server implements cursoring logic when the result rows exceed the configured batch size. On every request, only batch-size number of rows are returned. Each such request is considered an active cursor, with a specified amount of idle time specified by skip-token-cache-time. After the cursor is timed out, the cursor will be closed and remaining results will be cleaned up, and will no longer

be available for further queries. Since there is no session based tracking of these cursors, if the request for skiptoken comes after the expired time, the original query will be executed again and tries to reposition the cursor to relative absolute position, however the results are not guaranteed to be same as the underlying sources may have been updated with new information meanwhile.



Important

The following feature limitations currently apply.

- ✦ delta processing is not supported
- ✦ data-aggregation extension to specification is not supported.
- ✦ \$expand is minimally supported, where only one navigation can be expanded.

9.2.3. Client Tools for Access

There are different tools you can use. Depending upon your programming model and needs there are various ways you write your access layer into OData. Here are some suggestions:

- ✦ Your Browser: The OData Explorer is an online tool for browsing an OData data service.
- ✦ Olingo: Is a Java framework that supports OData V4, has both consumer and producer framework.
- ✦ Microsoft has various .Net based libraries. See <http://odata.github.io/>
- ✦ Windows Desktop: LINQPad is a wonderful tool for building OData queries interactively. See <https://www.linqpad.net/>
- ✦ Shell Scripts: use the CURL tool

9.2.4. OData Metadata

OData defines its schema using Conceptual Schema Definition Language (CSDL). Every VDB, that is deployed in an ACTIVE state on the Data Virtualization server exposes its metadata in CSDL format. For example if you want retrieve metadata for your vdb northwind, you need to issue a query like this:
[http://localhost:8080/odata4/northwind/NW/\\$metadata](http://localhost:8080/odata4/northwind/NW/$metadata)

Since OData schema model is not a relational schema model, Red Hat JBoss Data Virtualization uses the following semantics to map its relational schema model to OData schema model.

Table 9.4. Mapping OData to Data Virtualization

| Relational Entity | Mapped OData Entity |
|--------------------------|-----------------------------------|
| Table/View | EntityType, EntitySet |
| Table Columns | EntityType's Properties |
| Primary Key | EntityType's Key Properties |
| Foreign Key | Navigation Property on EntityType |
| Procedure | FunctionImport, Action Import |
| Procedure's Table Return | ComplexType |

By design, Red Hat JBoss Data Virtualization does not define any "embedded" ComplexType in the EntityType.

Since OData access is more key based, it is * MANDATORY* that every table Data Virtualization exposes through OData must have a PK or at least one UNIQUE key. A table which does not either of these will be dropped out of the \$metadata.

Appendix A. Unsupported JDBC Methods

A.1. Unsupported JDBC Methods

This appendix lists those JDBC methods that JBoss Data Virtualization does not support, based on the JDK 1.7 JDBC. Unless specified below, JBoss Data Virtualization supports all other JDBC Methods.

Those methods listed without comments throw an exception of type **SQLException** stating that it is not supported.

Where specified, some listed methods do not throw an exception, but possibly exhibit unexpected behavior. If no arguments are specified, then all related (overridden) methods are not supported. If an argument is listed then only those forms of the method specified are not supported.

A.2. ResultSet Limitations

- ✦ **TYPE_SCROLL_SENSITIVE** is not supported.
- ✦ **UPDATABLE ResultSets** are not supported.
- ✦ Returning multiple **ResultSets** from Procedure execution is not supported.

A.3. Unsupported Classes and Methods in java.sql

Table A.1. Connection Properties

| Class name | Methods |
|--------------------------|--|
| Array | Not Supported |
| Blob | <pre> getBinaryStream(long, long) - throws SQLFeatureNotSupportedException setBinaryStream(long) - - throws SQLFeatureNotSupportedException setBytes - - throws SQLFeatureNotSupportedException truncate(long) - throws SQLFeatureNotSupportedException </pre> |
| CallableStatement | <pre> getObject(int parameterIndex, Map&lt;String, Class&lt;?&gt;&gt; map) - throws SQLFeatureNotSupportedException getRef - throws SQLFeatureNotSupportedException getRowId - throws SQLFeatureNotSupportedException getURL(String parameterName) - throws SQLFeatureNotSupportedException registerOutParameter - ignores registerOutParameter(String parameterName, *) - throws SQLFeatureNotSupportedException setRowId(String parameterName, RowId x) - throws SQLFeatureNotSupportedException setURL(String parameterName, URL val) - throws SQLFeatureNotSupportedException </pre> |

| Class name | Methods |
|-------------------------|--|
| Clob | <pre> getCharacterStream(long arg0, long arg1) - throws SQLFeatureNotSupportedException setAsciiStream(long arg0) - throws SQLFeatureNotSupportedException setCharacterStream(long arg0) - throws SQLFeatureNotSupportedException setString - throws SQLFeatureNotSupportedException truncate - throws SQLFeatureNotSupportedException </pre> |
| Connection | <pre> createArrayOf - throws SQLFeatureNotSupportedException createBlob - throws SQLFeatureNotSupportedException createClob - throws SQLFeatureNotSupportedException createNClob - throws SQLFeatureNotSupportedException createSQLXML - throws SQLFeatureNotSupportedException createStruct(String typeName, Object[] attributes) - throws SQLFeatureNotSupportedException getClientInfo - throws SQLFeatureNotSupportedException releaseSavepoint - throws SQLFeatureNotSupportedException rollback(Savepoint savepoint) - throws SQLFeatureNotSupportedException setHoldability - throws SQLFeatureNotSupportedException setSavepoint - throws SQLFeatureNotSupportedException setTypeMap - throws SQLFeatureNotSupportedException </pre> |
| DatabaseMetaData | <pre> getAttributes - throws SQLFeatureNotSupportedException getClientInfoProperties - throws SQLFeatureNotSupportedException getFunctionColumns - throws SQLFeatureNotSupportedException getFunctions - throws SQLFeatureNotSupportedException getRowIdLifetime - throws SQLFeatureNotSupportedException </pre> |
| NClob | Not Supported |

| Class name | Methods |
|--------------------------|--|
| PreparedStatement | <pre> setArray - throws SQLFeatureNotSupportedException setRef - throws SQLFeatureNotSupportedException setRowId - throws SQLFeatureNotSupportedException setUnicodeStream - throws SQLFeatureNotSupportedException </pre> |
| Ref | Not Implemented |
| ResultSet | <pre> deleteRow - throws SQLFeatureNotSupportedException getHoldability - throws SQLFeatureNotSupportedException getObject(*, Map<String, Class<? >>> map) - throws SQLFeatureNotSupportedException getRef - throws SQLFeatureNotSupportedException getRowId - throws SQLFeatureNotSupportedException getUnicodeStream - throws SQLFeatureNotSupportedException getURL - throws SQLFeatureNotSupportedException insertRow - throws SQLFeatureNotSupportedException moveToInsertRow - throws SQLFeatureNotSupportedException refreshRow - throws SQLFeatureNotSupportedException rowDeleted - throws SQLFeatureNotSupportedException rowInserted - throws SQLFeatureNotSupportedException rowUpdated - throws SQLFeatureNotSupportedException setFetchDirection - throws SQLFeatureNotSupportedException update* - throws SQLFeatureNotSupportedException </pre> |
| RowId | Not Supported |
| Savepoint | not Supported |
| SQLData | Not Supported |
| SQLInput | not Supported |
| SQLOutput | Not Supported |
| Statement | <pre> setCursorName(String) </pre> |
| Struct | Not Supported |

A.4. Unsupported Classes and Methods in javax.sql

Table A.2. Connection Properties

| Class name | Methods |
|----------------|----------------------|
| RowSet* | Not Supported |

Appendix B. Keytool

B.1. Keytool

Keytool is an encryption key and certificate management utility. It enables users to create and manage their own public/private key pairs and associated certificates for use in self-authentication, and also to cache public keys (in the form of certificates) belonging to other parties, for securing communication to those parties.

B.2. Using Keytool with JBoss Data Virtualization

When using the keytool to manage public key cryptography for JBoss Data Virtualization, use the following options:

- » Set the alias to **teiid** using the **-alias teiid** option.
- » Set the algorithm to **RSA** using the **-keyalg RSA** option.
- » Set the validity period to **365** days using the **-validity 365** option.
- » Set the store type to **JKS** using the **-storetype JKS** option.

B.3. Create a Private/Public Key Pair with Keytool

Procedure B.1. Create a Private/Public Key Pair with Keytool

1. Run the **keytool -genkey -alias ALIAS -keyalg ALGORITHM -validity DAYS -keystore server.keystore -storetype TYPE** command:

```
keytool -genkey -alias teiid -keyalg RSA -validity 365 -keystore
server.keystore -storetype JKS
```

2. If the specified keystore already exists, enter the existing password for that keystore, otherwise enter a new password:

```
Enter keystore password: <password>
```

3. Answer the following questions when prompted:

```
What is your first and last name?
[Unknown]: <user's name>
What is the name of your organizational unit?
[Unknown]: <department name>
What is the name of your organization?
[Unknown]: <company name>
What is the name of your City or Locality?
[Unknown]: <city name>
What is the name of your State or Province?
[Unknown]: <state name>
What is the two-letter country code for this unit?
[Unknown]: <country name>
```

4. Enter **yes** to confirm the provided information is correct:

```
Is CN=<user's name>, OU=<department name>, O="<company name>",
L=<city name>, ST=<state name>, C=<country name> correct?
[no]: yes
```

5. Enter your desired keystore password:

```
Enter key password for <server>
(Return if same as keystore password)
```

Result

The `server.keystore` file contains the newly generated public and private key pair.

B.4. Extract a Self-signed Certificate from the Keystore

Procedure B.2. Extract a Self-signed Certificate from the Keystore

1. Run the `keytool -export -alias ALIAS -keystore server.keystore -rfc -file public.cert` command:

```
keytool -export -alias teiid -keystore server.keystore -rfc -file
public.cert
```

2. Enter the keystore password when prompted:

```
Enter keystore password: <password>
```

Result

This creates the `public.cert` file that contains a certificate signed with the private key in the `server.keystore`.

B.5. Add a Certificate to a Truststore Using Keytool

Procedure B.3. Add a Certificate to a Truststore Using Keytool

1. Run the `keytool -import -alias ALIAS -file public.cert -storetype TYPE -keystore server.truststore` command:

```
keytool -import -alias teiid -file public.cert -storetype JKS -keystore
server.truststore
```

2. If the specified truststore already exists, enter the existing password for that truststore, otherwise enter a new password:

```
Enter keystore password: <password>
```

3. Enter **yes** when prompted to trust the certificate:

```
Owner: CN=<user's name>, OU=<dept name>, O=<company name>, L=<city>,
ST=<state>, C=<country>
Issuer: CN=<user's name>, OU=<dept name>, O=<company name>, L=<city>,
ST=<state>, C=<country>
Serial number: 416d8636
Valid from: Fri Jul 31 14:47:02 CDT 2009 until: Sat Jul 31 14:47:02 CDT
2010
Certificate fingerprints:
    MD5:  22:4C:A4:9D:2E:C8:CA:E8:81:5D:81:35:A1:84:78:2F
    SHA1:
05:FE:43:CC:EA:39:DC:1C:1E:40:26:45:B7:12:1C:B9:22:1E:64:63
Trust this certificate? [no]:  yes
```

Result

The certificate in **public.cert** has been added to the new truststore named **server.truststore**.

Appendix C. Revision History

| | | |
|--------------------------|-----------------------|----------------------|
| Revision 6.3.0-10 | Fri Oct 7 2016 | David Le Sage |
| Updates for 6.3. | | |

| | | |
|--------------------------|------------------------|----------------------|
| Revision 6.2.0-56 | Thu Dec 10 2015 | David Le Sage |
| Updates for 6.2. | | |