



Red Hat Integration 2019-12

Data Virtualization

Data Virtualization

Legal Notice

Copyright © 2019 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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, the Red Hat 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 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

Combine data from multiple sources so that applications can connect to a single, virtual data model

Table of Contents

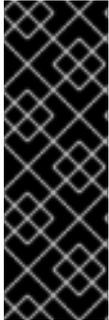
CHAPTER 1. HIGH-LEVEL OVERVIEW OF DATA VIRTUALIZATION	4
CHAPTER 2. VIRTUAL DATABASE CREATION	5
CHAPTER 3. DATA VIRTUALIZATION OPERATOR	6
3.1. INSTALLING THE DATA VIRTUALIZATION OPERATOR ON OPENSIFT	6
3.2. RUNNING THE DATA VIRTUALIZATION OPERATOR TO DEPLOY A VIRTUAL DATABASE	8
3.3. CREATING A VIRTUAL DATABASE FROM THE OPENSIFT WEB CONSOLE	10
3.4. CUSTOM RESOURCES TO SUPPORT DATA SOURCES	11
3.4.1. Settings to configure relational databases as data sources	12
3.4.2. Settings to configure MongoDB as a data source	13
3.4.3. Settings to configure REST, OData, and OpenAPI data sources	14
3.4.4. Settings to configure Salesforce as a data source	16
3.4.5. Setting up an OAuth connection to Salesforce	17
CHAPTER 4. USING MAVEN SPRING BOOT TO CREATE VIRTUAL DATABASES	19
4.1. USING MAVEN TO CREATE A JAVA SHELL PROJECT FOR YOUR VIRTUAL DATABASE	19
4.2. CREATING AND DEPLOYING SECRETS	20
4.3. SECRET OBJECTS FOR STORING DATA SOURCE INFORMATION	21
4.4. SPECIFYING PROJECT DEPENDENCIES IN THE POM.XML FILE	22
4.5. DEFINING THE STRUCTURE FOR VIRTUAL DATABASES IN A DDL FILE	23
4.6. DDL FILES	24
4.7. ADDING JAVA APPLICATION AND CLASS FILES	25
4.8. SAMPLE DATASOURCES.JAVA FILE	26
4.9. SPECIFYING APPLICATION PROPERTIES	26
4.10. SAMPLE APPLICATION.PROPERTIES FILE	27
4.11. DEPLOYMENT CONFIGURATION FILES (DEPLOYMENTCONFIG.YML)	28
4.12. SETTING THE DEPLOYMENT CONFIGURATION	29
4.13. CONNECTION SETTINGS FOR OTHER DATA SOURCES	30
4.13.1. Settings to connect to Salesforce as a data source	31
4.13.2. Settings to connect to Google Sheets as a data source	33
4.14. BUILDING A DATA VIRTUALIZATION PROJECT	34
CHAPTER 5. MAKING VIRTUAL DATABASES AVAILABLE TO API CONSUMERS	36
5.1. CONFIGURING ACCESS FOR ODATA CLIENTS	36
5.2. CONFIGURING ACCESS FOR JDBC CLIENTS	37
5.3. IDENTIFYING THE ODATA ENDPOINT OF A VIRTUAL DATABASE	38
CHAPTER 6. SECURING ODATA APIS FOR A VIRTUAL DATABASE BY USING 3SCALE AND RH-SSO ...	40
6.1. UPDATING THE VIRTUAL DATABASE CONFIGURATION	40
6.1.1. Adding SSO dependencies to pom.xml	41
6.1.2. Adding SSO settings to application.properties	41
6.1.3. Creating a configuration map	42
6.1.4. Updating SSO environment variables in the deploymentconfig.yml file	42
6.1.5. Defining user roles in the DDL file	43
6.2. UPDATING THE 3SCALE CONFIGURATION	43
CHAPTER 7. CREATING AND WORKING WITH VIRTUAL DATABASES IN FUSE ONLINE	45
7.1. CREATING VIRTUAL DATABASES IN FUSE ONLINE	45
7.2. ADDING A VIEW TO A VIRTUAL DATABASE IN FUSE ONLINE	46
7.3. USING THE VIEW EDITOR IN FUSE ONLINE TO MODIFY THE DDL THAT DEFINES A VIRTUAL DATABASE	47
7.4. PREVIEWING A VIRTUAL DATABASE IN FUSE ONLINE BY SUBMITTING SQL TEST QUERIES	48

7.5. PUBLISHING VIRTUAL DATABASES IN FUSE ONLINE TO MAKE THEM AVAILABLE FOR ACCESS	49
7.6. DELETING A VIRTUAL DATABASE IN FUSE ONLINE	49
CHAPTER 8. MIGRATING LEGACY VIRTUAL DATABASE FILES TO DDL FORMAT	51
8.1. VALIDATING A LEGACY VIRTUAL DATABASE XML FILE AND VIEWING IT IN DDL FORMAT	51
8.2. CONVERTING A LEGACY VIRTUAL DATABASE XML FILE AND SAVING IT AS A DDL FILE	52

CHAPTER 1. HIGH-LEVEL OVERVIEW OF DATA VIRTUALIZATION

Data virtualization is a container-native service that provides integrated access to multiple diverse data sources, including relational and noSQL databases, files, web services, and SaaS repositories through a single uniform API. Applications and users connect to a virtual database over standard interfaces (OData REST, or JDBC/ODBC) and can interact with data from all configured data sources as if the data were served from a single relational database.

The Red Hat data virtualization technology is based on Teiid, the open source data virtualization project. For more information about Teiid, see the [Teiid community documentation](#).



IMPORTANT

Data virtualization is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information about the support scope of Red Hat Technology Preview features, see <https://access.redhat.com/support/offerings/techpreview/>.

CHAPTER 2. VIRTUAL DATABASE CREATION

You design and create a virtual database and then deploy it to an OpenShift container. After you create the virtual database, you can make it available to API consumers, which can connect to it as if it were a single relational database.

You can create virtual databases by using a data virtualization operator, by using Fabric8 and Maven Spring Boot, or in Fuse Online.

In this Technology Preview release, the preferred method for deploying virtual databases to OpenShift is to use the data-virtualization operator. The method of using creating and deploying virtual databases with Fabric8 and the Spring Boot Maven plugin is deprecated.

You can create a virtual database in Fuse Online. However, in the current Technology Preview, virtual database created in Fuse Online provide a more limited set of features. Most developers are expected to use the data virtualization operator to deploy virtual databases directly on OpenShift.

Additional resources

- [Chapter 3, Data virtualization operator](#)
- [Chapter 4, Using Maven Spring Boot to create virtual databases](#)
- [Chapter 7, Creating and working with virtual databases in Fuse Online](#)

CHAPTER 3. DATA VIRTUALIZATION OPERATOR

The data-virtualization operator helps to automate the configuration and deployment of virtual databases. After you add the operator to your OpenShift cluster, you can use it to build and deploy virtual database images from a range of data sources.

In this Technology Preview release, the preferred method for deploying virtual databases to OpenShift is to use the data-virtualization operator. The method of using creating and deploying virtual databases with Fabric8 and the Spring Boot Maven plugin is deprecated. Currently, virtual databases that you create from the data virtualization operator are not available in Fuse Online.

You can install the data virtualization operator on OpenShift 3.11 or greater. On OpenShift 4.2 and later the Operator is available in the OperatorHub.

Prerequisites

- You are a cluster administrator. You must have cluster-admin rights to add operators to the OpenShift cluster.
- You have the 2019-12 release of Red Hat Integration.
- You have Developer access to an OpenShift server and you are familiar with using the OpenShift console and CLI.
- You have a DDL file for the virtual database that you want to create, or you know how to write SQL code and create DDL files.
- You have Maven 3.0 installed.
- You have JDK 8 (Java Platform, Standard Edition 8 Development Kit) or greater installed.
- You can connect to one of the following types of data sources:
 - Relational database
 - mongoDB
 - Salesforce
 - REST
 - OData
 - OpenAPI
- You have administrator access to the configured data source.

3.1. INSTALLING THE DATA VIRTUALIZATION OPERATOR ON OPENSIFT

Install the data virtualization operator so that you can use it to deploy virtual database images to OpenShift from YAML-based custom resources (CRs).

Prerequisites

- You have cluster-admin access to an OpenShift 3.11 or 4.2 or greater cluster.

- You can use the **oc** command-line tool to connect to interact with your OpenShift 3.11 cluster, or you have access to the OpenShift 4.2 or greater web console.

Procedure

- Install the operator using one of the following methods, depending on the version of OpenShift that you are running.

Installing on OpenShift 3.11

1. From a terminal window, log in to the OpenShift cluster as a cluster administrator.

```
oc login
```

2. Create or open a project where you want to deploy a virtual database.
3. Type the following commands:

```
export OP_ROOT=https://raw.githubusercontent.com/teiid/teiid-operator/7.5-0.0.x/deploy
oc create -f $OP_ROOT/crds/virtualdatabase.crd.yaml 1
oc create -f $OP_ROOT/service_account.yaml
oc create -f $OP_ROOT/role.yaml
oc create -f $OP_ROOT/role_binding.yaml
oc create -f $OP_ROOT/operator.yaml
```

- 1 If you previously created a CRD in the cluster, the command returns an error, reporting that the CRD already exists. You can ignore the message.

4. Type the following commands to create a pull secret that you can use to access the Red Hat image registry:

```
oc create secret docker-registry dv-pull-secret /
--docker-server=registry.redhat.io /
--docker-username=$username / 1
--docker-password=$password /
--docker-email=$email_address
oc secrets link builder dv-pull-secret
oc secrets link builder dv-pull-secret --for=pull
```

- 1 Substitute the user name and password that you use to log in to the Red Hat Customer Portal.

If the command completes with no errors, the operator is deployed to your OpenShift instance within the current OpenShift project.

5. To enable the data virtualization operator to retrieve images from the Red Hat registry so that you can create virtual databases, link the secret that you created in Step 4 to the service account for the operator.

```
oc secrets link dv-operator dv-pull-secret --for=pull
```

Installing on OpenShift 4.2 or greater

1. From a terminal window, type the following commands to log in to the OpenShift cluster and create a pull secret that you can use to access the Red Hat image registry:

```
oc login
oc create secret docker-registry dv-pull-secret /
--docker-server=registry.redhat.io /
--docker-username=$username / 1
--docker-password=$password /
--docker-email=$email_address
oc secrets link builder dv-pull-secret
oc secrets link builder dv-pull-secret --for=pull
```

1

Use your Red Hat Customer Portal login credentials.

2. Log in to the OpenShift web console as a cluster administrator.
3. From the OpenShift menu, expand **Operators** and click **OperatorHub**.
4. Click **Data Virtualization Operator 7.5.0 provided by Red Hat, Inc**, and then click **Install**.
5. From the **Create Operator Subscription** page, verify that the selected namespace matches the name of the project where you want to install the operator, and then click **Subscribe**. The **Installed Operators** page lists the **Data Virtualization Operator** and reports the status of the installation.
6. From the OpenShift menu, expand **Workloads** and click **Pods** to check the status of the operator pod. After a few minutes, the pod for the operator service begins to run.
7. To enable the data virtualization operator to retrieve images from the Red Hat registry so that you can create virtual databases, link the secret that you created in Step 1 to the service account for the operator.

```
oc secrets link dv-operator dv-pull-secret --for=pull
```

Additional resources

- [Section 3.2, “Running the data virtualization operator to deploy a virtual database”](#) .

3.2. RUNNING THE DATA VIRTUALIZATION OPERATOR TO DEPLOY A VIRTUAL DATABASE

After a cluster administrator adds the data virtualization operator to an OpenShift cluster, other users can run the operator to create a virtual database.

The data virtualization operator processes a virtual database custom resource (CR) to create and deploy a virtual database object on OpenShift. You specify the configuration properties for the data source in the CR. Each type of data source requires a specific set of configuration properties. When you run the operator, you provide it with a **.yaml** file that contains the CR. By running the operator with different CRs, you can create virtual databases from a range of data sources. For more information about specifying data source properties, see [Section 3.4, “Custom resources to support data sources”](#)



NOTE

In this Technology Preview, the data virtualization operator can create virtual databases from the following data sources only:

- Relational databases
- Salesforce databases
- MongoDB
- Open API
- OData

Prerequisites

- You have access to an OpenShift cluster in which the data virtualization operator is installed.
- The operator has access to the Maven repositories that contain the dependencies that the build requires.
- OpenShift can access a supported data source that runs on your network.
- You have the login credentials to access the data source.
- You have a CR in **.yaml** format that provides information about how to create the virtual database.

Procedure

1. From a terminal window, log in to OpenShift and open the project where you want to create the virtual database.
2. On your computer, change to the directory that contains the **.yaml** file that contains the CR.
3. Type the following command to run the operator to create the virtual database:

```
oc create -f <cr_filename.yaml>
```

Replace **<cr_filename.yaml>** with the name of the CR file for your data source. For example,

```
oc create -f dv-customer.yaml
```

After the deployment completes, a service is added to the OpenShift cluster. The name of the service matches the name of the custom resource.

4. Type the following command to verify that the virtual database is created:

```
oc get vdbs
```

OpenShift returns the list of virtual databases in the project.

The deployed service supports connections from the following clients:

- JDBC clients through port 31000.

- PostgreSQL clients, including ODBC clients, through port 35432.
- OData clients, through an HTTP endpoint and route.



NOTE

For OpenShift to create an HTTP endpoint, the value of the property **spec/exposeVia3scale** must be set to **false** in the CR. If the value is set to **true** it is assumed that 3scale manages the endpoint, and no HTTP endpoint is created.

3.3. CREATING A VIRTUAL DATABASE FROM THE OPENSIFT WEB CONSOLE

After a cluster administrator adds the data virtualization operator from the OperatorHub, other users can run the operator from web console to create a virtual database.

When you run the operator from the web console, you can use the embedded editor to create the custom resource (CR) that defines the virtual database. The CR includes properties that specify how to connect to the data source, and a DDL section that determines how information in the data source is used in your virtual database. You use YAML or JSON to define connection properties, and you use SQL-MED in the DDL section to specify the virtual database schema.

The editor includes a sample CR that you can edit in-place, or download to edit it off-line. You can also use drag and drop to upload local copies of properties definitions or DDL files. For information about the properties to use for different data sources, see [Section 3.4, "Custom resources to support data sources"](#)



NOTE

In this Technology Preview, the data virtualization operator can create virtual databases from the following data sources only:

- Relational databases
- Salesforce databases
- MongoDB
- Open API
- OData

Prerequisites

- You have web console access to an OpenShift cluster in which the data virtualization operator is installed.
- The operator has access to the Maven repositories that contain the dependencies that the build requires.
- OpenShift can access a supported data source that runs on your network.
- You have the login credentials to access the data source.

- You have a CR in **.yaml** format that provides information about how to create the virtual database.

Procedure

1. From the OpenShift menu, expand **Operators** and click **Installed Operators**.
2. Click **Data Virtualization Operator**.
3. Click the **VirtualDatabase** tab.
4. Click **Create Virtual Database**
OpenShift displays a sample custom resource.
5. Edit the sample CR as needed to specify the details of the virtualization, and then click **Create**.
OpenShift lists the virtual database service on the **VirtualDatabase** tab.
6. From the OpenShift menu, expand **Workloads** and click **Pods** to check the status of the pod that runs the new virtualization. After it is running, you are ready to use the new virtualization.
7. To view the protocols that the virtualization exposes, from the OpenShift menu, expand **Networking**, click **Services**, and then click the service with the name of the virtualization. The deployed service automatically supports the following connection types:
 - JDBC - Provides access for SQL clients over port 31000.
 - postgresSQL - Provides access for ODBC clients and other postgresSQL clients over port 35432.
 - HTTP - Provides access for OData and REST clients over port 8080.

3.4. CUSTOM RESOURCES TO SUPPORT DATA SOURCES

Before you can use the data virtualization operator to create a virtual database, you must specify properties for the data source in a custom resource (CR) file.

By providing configuration information in a CR, you provide the operator with directions for how to create virtual databases from any of multiple types of data sources. When you run the data virtualization operator, it reads information from the CR to determine the type of the data source, its schema structure, and how to connect to it and authenticate with it.

The CR uses SQL Data Definition Language (DDL) commands to describe the schema of the virtual database and the data source, the data that you want to import into the virtual database, and the mapping between the source schema and the virtual schema. The CR also specifies the *translator* to use to convert the format of commands and data that pass between the virtual database and the data source.

You can specify the values of configuration properties directly in the CR file, or you can reference values that you define in OpenShift *secrets*. For more information about creating secrets, see [Section 4.2, "Creating and deploying secrets"](#).



NOTE

Period characters (.) are not valid for use in environment variables. When you add variables to the CR, replace period characters with underscore characters (_). At runtime, underscores in the variables are converted automatically to periods.

3.4.1. Settings to configure relational databases as data sources

The custom resource (CR) for a relational database management system (RDBMS), such as PostgreSQL or Microsoft SQL Server, must contain specific information so that the data virtualization operator can create a virtual database from the database.

The following tables list some of the properties that are required in a CR to create a virtual database that is based on a relational database. The values that you assign to these properties vary according to your environment and the specific database technology that you use.

Table 3.1. Connection properties in the `spec/env` section of an RDBMS CR

Property Name	Description	Required	Default value
SPRING_DATASOURCE_<NAME>_JDBCURL	URL For the connection	Yes	n/a
SPRING_DATASOURCE_<NAME>_USERNAME	User Name	Yes	n/a
SPRING_DATASOURCE_<NAME>_PASSWORD	Password	Yes	n/a
SPRING_DATASOURCE_<NAME>_DRIVER_CLASS_NAME*	Driver Name	No	n/a
SPRING_DATASOURCE_<NAME>_IMPORTER_SCHEMA_NAME	Schema Name for import	Yes	n/a

In the preceding table, *NAME* specifies a name, in upper case, that refers to the data source. This same name string is used in the DDL that defines the virtual database to represent the data source server.

*The JDBC driver class that you reference must be available as a Maven artifact from a repository that is listed in the **spec/dependencies** section of the CR file. The virtualization operator must have access to download artifacts from the specified repository.

Along with the connection properties You can define properties to control the behavior of the JDBC translator and to provide more control over how you import data from the source database.

Table 3.2. Foreign data wrapper settings in the `build/source/ddl` section of an RDBMS CR

Property Name	Description	Required	Value
FOREIGN DATA WRAPPER	Translator	Yes	<translator-name>*

*The translator value depends on the type of database. Translators are available for the following relational databases.

- *db2*

- *derby*
- *h2*
- *hana* (Connects to SAP HANA through a JDBC driver)
- *hive*jdbc* (Connects to Apache Hive through the Hive JDBC driver)
- *hsql*
- *impala* (Connects to Apache Impala)
- *informix*
- *ingres*
- *jtds* (Connects to Microsoft SQL Server through Java tabular data stream (JTDS) drivers)
- *mysql5*
- *netezza*
- *oracle*
- *osisoft-pi* (Connects to an OSIsoft PI Data Archive through the OSIsoft JDBC driver)
- *phoenix* (Connects to Apache Phoenix using the Phoenix JDBC driver)
- *postgresql*
- *presto-jdbc* (Connects to PrestoDB through a JDBC driver)
- *sqlserver* (Connects to Microsoft SQL Server)
- *sybase*
- *teradata*
- *teiid* (Connecting to a Teiid virtual database through a JDBC driver)
- *vertica*

For a complete list of the translator properties that you can define for JDBC data sources, see [JDBC Translators](#) in the Teiid Reference Guide.

3.4.2. Settings to configure MongoDB as a data source

The custom resource (CR) for a MongoDB database must contain specific information so that the data virtualization operator can create a virtual database from the database.



NOTE

The current release supports MongoDB release 3.

The following tables list the properties that are required in the CR to create a virtual database that is based on a MongoDB database:

Table 3.3. Connection properties in the `spec/env` section of the MongoDB CR

Property Name	Description	Required	Default value
SPRING_TEIID_DATA_MONGODB_{NAME}_REMOTE_SERVER_LIST	List of MongoDB servers ex: (localhost:27012)	Yes	n/a
SPRING_TEIID_DATA_MONGODB_{NAME}_USER	User Name	Yes	n/a
SPRING_TEIID_DATA_MONGODB_{NAME}_PASSWORD	Password	Yes	n/a
SPRING_TEIID_DATA_MONGODB_{NAME}_DATABASE	Database name to connect to	Yes	n/a
SPRING_TEIID_DATA_MONGODB_{NAME}_AUTH_DATABASE	Database name for authorization	No	n/a
SPRING_TEIID_DATA_MONGODB_{NAME}_SSL	Use SSL Connection	No	n/a

In the preceding table, *NAME* specifies a name, in upper case, that refers to the data source. This same name string is used in the DDL that defines the virtual database to represent the data source server.

Table 3.4. Foreign data wrapper settings in the `build/source/ddl` section of the MongoDB CR

Property Name	Description	Required	Value
FOREIGN DATA WRAPPER	Translator	Yes	mongodb

For a complete list of the properties that you can set to control how data is translated between MongoDB and a virtual database, see the [Teiid Reference Guide](#).

**NOTE**

You are not required to list any build dependencies in the CR for a MongoDB virtual database.

3.4.3. Settings to configure REST, OData, and OpenAPI data sources

The custom resource (CR) for a REST-based data source must contain specific information so that the data virtualization operator can create a virtual database from the source.

The set of connection properties that are required for all REST-based data sources is the same. However, services that are based on standard specifications, such as OData or OpenAPI, require specific translators. For information about translator properties, see [Table 3.6, “Foreign data wrapper settings in the `build/source/ddl` section of a REST-based CR”](#).

For generic REST web services that are not based on a standard specification, the data virtualization service cannot automatically convert query criteria into a query parameter. Because these services lack built-in mechanisms to pass SQL query conditions to a REST API endpoint, you must use the `invokeHttp` procedure to pass the query as an XML or JSON payload, and you must specify all query strings and headers.

By default, translators are unable to parse the security configuration of a secured API. To enable translators to access data for a secured API, the CR must specify the security properties for the API.

Table 3.5. Properties in the `spec/env` section of REST-based CRs

Property Name	Description	Required	Default value
SPRING_TEIID_REST_<NAME>_ENDPOINT	Endpoint for the service	Yes	n/a
SPRING_TEIID_REST_<NAME>_SECURITY_TYPE	Security type used. Available options are <i>http-basic</i> , <i>openid-connect</i> or empty	No	no security
SPRING_TEIID_REST_<NAME>_USER_NAME	User Name	Yes	n/a
SPRING_TEIID_REST_<NAME>_PASSWORD	Password	Yes	n/a
SPRING_TEIID_REST_<NAME>_CLIENT_ID	ClientId from connected app	Yes, when the security type is defined as <i>openid-connect</i>	n/a
SPRING_TEIID_REST_<NAME>_CLIENT_SECRET	clientSecret from connected app	Yes, when the security type is defined as <i>openid-connect</i>	n/a
SPRING_TEIID_REST_<NAME>_AUTHORIZE_URL	clientSecret from connected app	Yes, when the security type is defined as <i>openid-connect</i>	n/a
SPRING_TEIID_REST_<NAME>_ACCESS_TOKEN_URL	clientSecret from connected app	Yes, when the security type is defined as <i>openid-connect</i>	n/a
SPRING_TEIID_REST_<NAME>_SCOPE	clientSecret from connected app	No. Applies when the security type is defined as <i>openid-connect</i>	n/a

Property Name	Description	Required	Default value
SPRING_TEIID_REST_<NAME>_REFRESH_TOKEN*	Refresh token for the virtual database service	No	n/a

* When security is set to *openid-connect*, you can use refresh tokens to authenticate, instead of using name and password authentication. Information about obtaining refresh tokens is beyond the scope of this document.

In the [Table 3.5, "Properties in the `spec/env` section of REST-based CRs"](#), *NAME* indicates the name, in upper case, that represents the data source. The *NAME* value is also used in lower case in the **build/source/ddl** section of the CR to represent the data source server.

Table 3.6. Foreign data wrapper settings in the `build/source/ddl` section of a REST-based CR

Data source	Description	Required	Value
REST	Translator	Yes	rest
OData	Translator	Yes	odata
OData4	Translator	Yes	odata4
OpenAPI	Translator	Yes	openapi

For a complete list of the properties that you can set to control how data is translated between REST-based services and a virtual database, see the OData, OData V4, OpenAPI, and Web Services Translator sections in the [Translators section of the Teiid Reference Guide](#).

The OpenAPI translator assumes that the endpoint in the API document is set to the target location **/openapi**, and it builds a source model that is based on that assumption. If the API endpoint is set to a different target, a configuration setting must be specified to enable the translator to locate the endpoint and import data correctly. The following examples show a DDL schema statement and an environment variable that you can set to specify a non-standard endpoint.

Sample DDL statement to specify a non-standard OpenAPI endpoint

```
CREATE SCHEMA sourceModel SERVER oService OPTIONS ("importer.metadataUrl"
'/swagger.json');
```

Sample environment property to specify a non-standard OpenAPI endpoint

```
SPRING_TEIID_REST_OSERVICE_IMPORTER_METADATAURL=/swagger.json
```

3.4.4. Settings to configure Salesforce as a data source

The custom resource (CR) for a Salesforce database must contain specific information so that the data virtualization operator can create a virtual database from the database.

Salesforce uses OAuth 2.0 for authentication and authorization. Before a virtual database can import and query Salesforce data, you must obtain OAuth credentials for the virtual database from Salesforce. For information about how to set up OAuth, see [Section 3.4.5, "Setting up an OAuth connection to Salesforce"](#)

The following tables list the properties that are required in the CR to create a virtual database that is based on a Salesforce database:

Table 3.7. Properties in the `spec/env` section of the Salesforce CR

Property Name	Description	Required	Default value
SPRING_TEIID_DATA_SALESFORCE_<NAME>_USER_NAME*	User for salesforce.com	Yes	n/a
SPRING_TEIID_DATA_SALESFORCE_<NAME>_PASSWORD*	Password for salesforce.com	Yes	n/a
SPRING_TEIID_DATA_SALESFORCE_<NAME>_CLIENT_ID	ClientId from connected app	Yes	n/a
SPRING_TEIID_DATA_SALESFORCE_<NAME>_CLIENT_SECRET	clientSecret from connected app	No	n/a
SPRING_TEIID_DATA_SALESFORCE_<NAME>_REFRESH_TOKEN*	Refresh token for connected app	No	n/a

*If your connected app uses refresh tokens to authenticate, rather than name and password, your CR must specify the refresh token property, in place of the user name and password properties. Information about obtaining refresh tokens is beyond the scope of this document. For information about how to obtain a refresh token for your connected app, see the Salesforce documentation.

In Table 1, *NAME* indicates the name, in upper case, that refers to the data source. The *NAME* value is also used in lower case in the **build/source/ddl** section of the CR to represent the data source server.

Table 3.8. Properties in the `build/source/ddl` section of the Salesforce CR

Property Name	Description	Required	Value
FOREIGN DATA WRAPPER	Translator	Yes	salesforce

For a complete list of the properties that you can set to control how data is translated between Salesforce and a virtual database, see the [Teiid Reference Guide](#).

3.4.5. Setting up an OAuth connection to Salesforce

Before the data virtualization service can retrieve data from a Salesforce database, you must enable configure it as a connected app in Salesforce that is OAuth-enabled. After you configure OAuth, Salesforce generates a client ID and client secret that you must add to the CR file that defines the connection from the virtual database to Salesforce.

To configure OAuth you create a connected app in Salesforce that can request access to Salesforce data on behalf of the data virtualization service. In the settings for the connected app, you enable integration with the Salesforce API by using the OAuth 2.0.

Prerequisites

- You have a Salesforce.com account that has access to the data that you want to integrate in a virtual database.



NOTE

The following steps are based on Salesforce Classic. If you use a different version of Salesforce, you might use a different procedure. For more information about creating connected apps in Salesforce, see [the Salesforce documentation](#).

Procedure

1. From Salesforce, log into your account.
2. Click **SetUp** in the profile menu.
3. In the **Build** section of the navigation sidebar, expand **Create**, and then click **Apps**.
4. In the **Connected Apps** section, click **New**.
5. Complete the required fields.
6. In the section **API (Enable OAuth Settings)**, select **Enable OAuth Settings** to display the OAuth settings.
7. Complete the required OAuth fields. In the **OAuth Scopes** field, you must select the following scopes:
 - Access and manage your data (api).
 - Access your basic information (id, profile, email, address, phone).
 - Allow access to your unique identifier (openid).
 - Full access (full).
 - Perform requests on your behalf at any time (refresh_token, offline_access).
8. Select **Require Secret for Web Server Flow**
9. Click **Save** and then click **Continue**.
10. Make a note of the values in the **Consumer Key** and **Consumer Secret** fields. These values are required for properties in the CR that specifies how the virtual database connects to Salesforce.

CHAPTER 4. USING MAVEN SPRING BOOT TO CREATE VIRTUAL DATABASES

You can use Maven Spring Boot to create, build, and deploy a virtual database on OpenShift.



NOTE

The method of using Fabric8 and Maven Spring Boot plugins to build and deploy virtual databases is deprecated in this release.

The Spring Boot Maven plugin converts the data virtualization library into an executable Spring Boot JAR file. The Fabric8 Maven plugin helps to build a container image that is based on the Spring Boot executable and optionally deploy that container to OpenShift. Both plugins are available from the public Red Hat Maven repository, and they are downloaded on demand at build time after you specify dependencies in your **pom.xml** file. For information about specifying build dependencies, see [Section 4.4, "Specifying project dependencies in the pom.xml file"](#).

Prerequisites

- You have the 2019-07 release of Red Hat Integration, and you are running Fuse 7.5.
- You have a DDL file for the virtual database that you want to create, or you know how to write SQL code and create DDL files.
- You have Maven 3.0 installed.
- You have JDK 8 (Java Platform, Standard Edition 8 Development Kit) or greater installed.
- You can connect to one of the following types of data sources:
 - Relational database
 - mongoDB
 - Salesforce
 - REST
 - OData
 - OpenAPI
 - Google Sheets*
- You have administrator access to the configured data source.

4.1. USING MAVEN TO CREATE A JAVA SHELL PROJECT FOR YOUR VIRTUAL DATABASE

Use Maven to generate a Java shell project that you can then modify.

Prerequisites

- You have Maven 3.0 installed.

- You have Java JDK 11 or greater installed.

Procedure

1. On your local workstation, change to a directory where you want to create the Java project for your virtual database.
2. Run the following Maven command to generate a plain Java shell project:

```
mvn archetype:generate -DgroupId=<domainSuffix>.<domainName> -DartifactId=  
<virtualDbName> -DarchetypeArtifactId=maven-archetype-quickstart -  
DinteractiveMode=false
```

For example:

```
mvn archetype:generate -DgroupId=com.example -DartifactId=sampledbr -  
DarchetypeArtifactId=maven-archetype-quickstart -DinteractiveMode=false
```

All of the artifacts that you need to create a Java project are saved to the project directory.

4.2. CREATING AND DEPLOYING SECRETS

Create and deploy secret objects to store values for your environment variables.

Although secrets exist primarily to protect sensitive data by obscuring the value of a property, you can use them to store the value of any property.

Prerequisites

- You have the login credentials and other information that are required to access the data source.

Procedure

1. Create a secrets file to contain the credentials for your data source, and save it locally as a **.yml** file. For example,

Sample secrets.yml file

```
apiVersion: v1  
kind: Secret  
metadata:  
  name: postgresql  
type: Opaque  
stringData:  
  database-user: bob  
  database-name: sampledbr  
  database-password: bob_password
```

2. Deploy the secret object on OpenShift.
 - a. Log in to OpenShift, and open the project that you want to use for your virtual database. For example,
oc login --token=<token> --server=https://<server>oc project <projectName>

b. Run the following command to deploy the secret file:

```
oc create -f ./secret.yaml
```

3. Set an environment variable to retrieve its value from the secret.

a. In the environment variable, use the format **valueFrom:/secretKeyRef** to specify that the variable retrieves its value from a key in the secret that you created in Step 1.

For example, in the following excerpt, the

SPRING_DATASOURCE_SAMPLEDB_PASSWORD retrieves its value from a reference to the **database-password** key of the **postgresql** secret:

```
- name: SPRING_DATASOURCE_SAMPLEDB_PASSWORD
  valueFrom:
    secretKeyRef:
      name: postgresql
      key: database-password
```

Additional resources

- For more information about how to use secrets on OpenShift, see [Providing sensitive data to pods](#) in the OpenShift documentation.

4.3. SECRET OBJECTS FOR STORING DATA SOURCE INFORMATION

Use secret objects to securely store sensitive information about how your virtual database connects to its data sources.

For OpenShift to open a connection to a data source on behalf your virtual database application, it must pass login credentials and other information to the data source. Typically, to maintain security administrators limit access to database credentials, and do not expose credentials to developers directly. To enable developers indirect access to credentials, it's possible to deploy *secret objects* on OpenShift to securely store and pass credentials.

Use secrets in combination with environment variables. Rather than specify static values directly in your configuration files, you can configure OpenShift to retrieve values for environment variables from secret objects. When a key in an environment variable refers to a secret object, to obtain the key value, OpenShift examines the secret to find a token that has a name that matches the key name. It extracts the token value and then passes it to the environment variable.

For example, the following environment variable is set to retrieve the value for the **database-user** key from a secret object that has the name **postgresql**.

```
- name: SPRING_DATASOURCE_SAMPLEDB_USERNAME
  valueFrom:
    secretKeyRef:
      name: postgresql
      key: database-user
```

When the data virtualization service needs to retrieve the value for the preceding environment variable, it accesses the secret object with the name **postgresql** and reads the value of the **database-user** key.

In the following secret object, the **database-user** token is assigned the value **bob**. OpenShift passes that value to the environment variable.

Sample secrets.yml file

```

apiVersion: v1
kind: Secret
metadata:
  name: postgresql
type: Opaque
stringData:
  database-user: bob

```

4.4. SPECIFYING PROJECT DEPENDENCIES IN THE POM.XML FILE

To provide your Java shell project with the details that are required to build the project, edit the **pom.xml** file to specify project dependencies, plugins to run, and build profiles.

Some properties in **pom.xml** are common to all data source types. For example, the Teiid Spring Boot starter and the Spring Boot Maven plugin are required to connect to both a PostgreSQL database and a Salesforce database. Other properties, such as the drivers that a data source requires, are specific to individual data sources.

Use the sample **pom.xml** file in the Teiid OpenShift repository as the basis for your own file. The sample file contains settings for a PostgreSQL database, but the settings in the **<dependencyManagement>**, **<build>**, and **<profiles>** elements apply to any data source. The build resource must be set to the Spring Boot Maven plugin. This plugin, **spring-boot-maven-plugin**, converts the virtual database schema project into a Spring Boot executable Uber JAR file that becomes the basis for an OpenShift container image. The OpenShift profile element must be set to use the Fabric8 Maven plugin (**<artifactId>fabric8-maven-plugin</artifactId>**), which helps to build a container image from the executable JAR, and optionally deploy it into OpenShift.

Modify the values in the **<dependencies>** element as needed to enable connectivity to the data sources that you want to use.



NOTE

Driver modules for most databases are included in the Red Hat Maven repository. At build time, the drivers are downloaded automatically based on the dependency statements in **pom.xml**. For some proprietary data sources, drivers might not be publicly available. If a driver is not available in the repository, download the driver from a third-party site, and deploy it to your local Maven repository.

Prerequisites

- You have a Java shell project for your virtual database application.
- You are familiar with editing Maven **pom.xml** files.
- Download the sample **pom.xml** file from the Teiid OpenShift repository.
- If the driver for your database is not available from the public Maven repository, you have downloaded the driver and deployed it to your local Maven repository.

Procedure

1. Replace the default **pom.xml** file that you created in your Java shell project with the file that you download from the Teiid OpenShift repository.

2. Edit the **pom.xml** to specify the name of the OpenShift project in the **<fabric8 namespace>** element.
3. Set the value of the properties version element to the data virtualization version that you are using. For example,

```
<properties>
  <version.teiid.spring.boot>${version.teiid.spring-boot} </version.teiid.spring.boot>
</properties>
```

\${version.teiid.spring-boot} represents the build version of the code that is available in the Maven repository. Substitute the value of the build version for the product that you are working with.

4. Specify your data source dependencies in the **<dependencies>** element. If you want to connect a PostgreSQL database, you can use the values in the **<dependencies>** element as they are.

Additional resources

- For information about changes that you must make to the **<dependencies>** element of the **pom.xml** file to support Google Sheets or Salesforce, see [Section 4.13, “Connection settings for other data sources”](#).

4.5. DEFINING THE STRUCTURE FOR VIRTUAL DATABASES IN A DDL FILE

After you complete changes to the **pom.xml** file, you’re ready to define the structure of your virtual database. You define a virtual database through a text-based DDL file. You can supply an existing DDL file, if you have one, or you can create one.



NOTE

If you have a **.vdb** or **.xml** file from an earlier data-virtualization product that you want to reuse, you must convert the file into DDL format before you can use it. For more information about how to convert legacy virtual database files to DDL format, see [Chapter 8, Migrating legacy virtual database files to DDL format](#).

If you want to create your own DDL file you can use the [sample DDL file](#) in the Teiid OpenShift repository as a guide. But be mindful that the entries in your DDL file are unique to your environment and must include settings that are specific to the data sources that you use.

The design of a DDL file can be complex and is beyond the scope of this documentation. For more information about using SQL in DDL files that support data virtualization, see the [Teiid Reference Guide](#).

Prerequisites

- You have a DDL file for the virtual database that you want to create.
- You know how to write SQL code and you are familiar with creating DDL files.

Procedure

1. Add your DDL file to the **src/main/resources** directory. You can create a new file from scratch, modify the [sample DDL file](#), or use an existing file.
2. Use SQL statements to specify the structural elements of your virtual database.

4.6. DDL FILES

DDL files contain SQL commands that describe and define the structure of the virtual database, such as its views or schema, procedures, functions, tables, records, columns, and other fields.

When you are ready to build your virtual database, the data virtualization service reads the information in the DDL file and uses that information to generate the virtual database image.

Structures that you define in the DDL file are converted to relational database tables. Translators in the data virtualization service import and convert data from your data sources, and use the data to create views in the virtual database.

A typical DDL file define database elements by using statements similar to those in the following list:

- The name of the virtual database. For example:

```
CREATE DATABASE customer OPTIONS (ANNOTATION 'Customer VDB');
USE DATABASE customer;
```

- The name of the translator, or foreign data wrapper, that is needed to interpret data from the data source. For example,

```
CREATE FOREIGN DATA WRAPPER postgresql;
```

- The name of the data source server, and the name of the resource adapter that provides the connections details for the external data source. For example,

```
CREATE SERVER sampledb TYPE 'NONE' FOREIGN DATA WRAPPER postgresql
OPTIONS ("resource-name" 'sampledb');
```



NOTE

Several of the files that you use to configure your virtual database refer to the name of the data source that is defined in the **CREATE SERVER** statement of the DDL file. For example, the name of the data source appears in the **Datasources.java** file, in the **application.properties** file, and in the names of environment variables in the **deploymentconfig.yml** file. To wire the various configuration files together, it's important to use the data source name consistently.

- The names of the foreign schema from which you want to import metadata, and the virtual schema into which you want to import that metadata. For example:

```
CREATE SCHEMA accounts SERVER sampledb;
CREATE VIRTUAL SCHEMA portfolio;

SET SCHEMA accounts;
IMPORT FOREIGN SCHEMA public FROM SERVER sampledb INTO accounts
OPTIONS("importer.useFullSchemaName" 'false');
```

- Views in the virtual database and the mapping between data in the data source and in the virtual database view. For example,

```
CREATE VIEW CustomerZip(id bigint PRIMARY KEY, name string, ssn string, zip string) AS
SELECT c.ID as id, c.NAME as name, c.SSN as ssn, a.ZIP as zip
FROM accounts.CUSTOMER c LEFT OUTER JOIN accounts.ADDRESS a
ON c.ID = a.CUSTOMER_ID;
```

4.7. ADDING JAVA APPLICATION AND CLASS FILES

In the application library of your virtual database project, create a Java application file, **Application.java**, and a data source class file, **DataSources.java**.

Application.java is the main Spring Boot application file that bootstraps the virtual database application. **DataSources.java** adds @Bean methods for each data source that you want to use as an input to your virtual database. The Java class serves as a resource adapter, which provides access to the data source.

When the virtual database starts, the virtualization service reads the metadata and generates an internal model from it. The service then uses that model to read and write to the virtual database. Entities within a data source become available in the virtual database for users to access as tables. For example, if you use a Salesforce database as a data source, then **Subjects** in Salesforce become available as tables in the virtual database.

Prerequisites

- You have Java 11 or greater installed.
- You have a Java shell project that you generated with Maven.

Procedure

1. Create the following **Application.java** file in your Java class folder (for example, **src/main/java/com/example**).

```
package com.example;

import org.springframework.boot.SpringApplication;
import org.springframework.boot.autoconfigure.SpringBootApplication;

@SpringBootApplication
public class Application {
    public static void main(String[] args) {
        SpringApplication.run(Application.class, args);
    }
}
```

You can remove the default **app.java** file that Maven generates when you create the project.

2. Create a **Datasources.java** file in the class folder, and add a bean method for each data source that you want to connect to your virtual database. For an example of a **Datasources.java** file that is designed to work with a PostgreSQL database, see the [Section 4.8, "Sample Datasources.java file"](#).

4.8. SAMPLE DATASOURCES.JAVA FILE

The **Datasources.java** file adds a class to represent a connection to a data source. The file also establishes a prefix in the **ConfigurationProperties** argument (**spring.datasource.sampledb**). This prefix must be used in the names of data source properties that you specify in the **application.properties** file.

You can define multiple data sources in **Datasources.java** by adding multiple classes, each with its own prefix designation. In each case you must add corresponding entries to the DDL file and to the properties or deployment configuration files.

To associate the Java bean with the data source that is defined in your DDL file, the bean name must be the same as the name of the **SERVER** and **resource-name** properties in the DDL file. For example, the following sample file establishes a connection to a PostgreSQL database called **sampledb**, which is the name that is assigned in the DDL file to the data source **SERVER** object and to its **resource-name** definition.

```
package com.example;

import javax.sql.DataSource;

import org.springframework.boot.jdbc.DataSourceBuilder;
import org.springframework.boot.context.properties.ConfigurationProperties;
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.Configuration;

@Configuration
public class DataSources {

    @ConfigurationProperties(prefix = "spring.datasource.sampledb") 1
    @Bean
    public DataSource sampledb() { 2
        return DataSourceBuilder.create().build();
    }
}
```

- 1 The prefix must match the prefix that you assign to properties that you define in the **application.properties** file.
- 2 The name **sampledb** in the prefix definition and in the method name must match the name in the **SERVER** and **resource-name** objects that are defined in the virtual database DDL file. The Spring Boot framework automatically associates the names of methods in the **Datasources.java** file with the names of data sources in the DDL file.



NOTE

The preceding sample file is designed to work with a PostgreSQL database. For information about how to adapt the file for use with other data sources, see [Section 4.13, "Connection settings for other data sources"](#).

4.9. SPECIFYING APPLICATION PROPERTIES

You define static properties for your virtual database application in an **application.properties** file in the

`/src/main/resource` directory. Static properties are configuration settings that remain constant across different environments. After you deploy a virtual database on OpenShift, any modifications that you make to the `application.properties` file are not effective unless you rebuild and redeploy your virtual database.

At minimum the `application.properties` file must contain a value for the `teiid.vdb-file` property, which names the DDL file that defines the structure of the the virtual database. For example, `teiid.vdb-file=customer-vdb.ddl`.

You can also use the `application.properties` file to define other properties, such as data source properties, including their names, their drivers, and the URLs, user names, and passwords that are required to connect to them. Using static properties to assign values for these can be useful in a test environment. But if you deploy your virtual database in multiple OpenShift environments, it's best to use environment variables to dynamically assign unique values for each environment. For more information about using environment variables in your virtual database configuration, see [Section 4.12, "Setting the deployment configuration"](#).

If you define data source properties in the `application.properties` file, you must prefix the configuration properties string that you specified in the `Datasources.java` file. The prefix establishes a connection between the properties and the Java class. For example, if you establish the configuration properties prefix `spring.datasource.sampledb` in the `Datasources.java` file, then you must precede the names of the properties that you define in your `application.properties` file with that string. For example,

```
spring.datasource.sampledb.username=<username>
spring.datasource.sampledb.password=<password>
```

Prerequisites

- You have a DDL file in `/src/main/resources` that defines your virtual database structure.
- You have a `Datasources.java` file in your Java class folder that specifies an application prefix.

Procedure

1. Add the file `application.properties` to the `src/main/resources` folder of your Java project.
2. In the file, add the property `teiid.vdb-file` and set its value to the name of the DDL file in the `src/main/resources` folder of your Java project, for example, `customer-vdb.ddl`.
3. (Optional) Add properties to specify connection information for your data source, such as its name, URL, login credentials, and drivers. For an example of an `application.properties` file for a PostgreSQL data source, see [Section 4.10, "Sample application.properties file"](#)

4.10. SAMPLE APPLICATION.PROPERTIES FILE

The following `application.properties` file includes settings for connecting to a PostgreSQL database. You might use a file similar to this for testing purposes, but for an OpenShift deployment, it's best to specify data source properties in a `deploymentconfig.yml` file.

You can specify source-specific properties for other types of data sources. For more information, see [Section 4.13, "Connection settings for other data sources"](#).

```
spring.datasource.sampledb.jdbc-url=jdbc:postgresql://localhost/sampledb 1 2
spring.datasource.sampledb.username=user 3
```

```

spring.datasource.sampledb.password=user
spring.datasource.sampledb.driver-class-name=org.postgresql.Driver 4
spring.datasource.sampledb.platform=sampledb 5

# VDB location
teiid.vdb-file=customer-vdb.ddl 6

# true to allow JDBC connections
teiid.jdbc-enable=true

#logging.level.org.teiid=DEBUG 7

```

- 1 The JDBC URL that the virtual database uses to connect to a local PostgreSQL database as its data source.
- 2 The prefix that is used in each of these properties matches the prefix that is defined in the **Datasources.java** file.
- 3 The user name and password values listed here are displayed in plain text. To secure these credentials in a production deployment on OpenShift, use environment variables in a deployment configuration file to reference the secret object that defines these values.
- 4 The driver that is required to connect to the data source. This driver is defined in the **pom.xml** file.
- 5 The name of the data source.
- 6 The name of the DDL file.
- 7 Uncomment this statement to enable debug logging. You can view the logs for a pod by running the following command: **oc logs <podname>** Logs are also available on the **Logs** tab of the **Pod Overview** page in the OpenShift web console.

4.11. DEPLOYMENT CONFIGURATION FILES (DEPLOYMENTCONFIG.YML)

A deployment configuration file stores settings that govern how the Fabric8 Maven plugin builds and deploys the container image for your virtual database.

The **deploymentconfig.yml** file can also define environment variables for the properties that are required to configure data sources for your virtual databases. The environment variables that you define in the **deploymentconfig.yml** file map to properties in the **application.properties** file. But unlike settings in the properties file, the settings that you define in the **deploymentconfig.yml** file are dynamic. That is, if you add an environment variable or change its value, you do not have to rebuild the virtual database service to put the change into effect.

Environment variables and their corresponding properties have similar names, but they are formatted differently. Separators in the property names are converted from dots or dashes to underscores, and alphabetic characters are converted to uppercase.

Table 4.1. Format of properties versus environment variables

Property in application.properties	Environment variable in deploymentconfig.yml
spring.datasource.sampledb.jdbc-url	SPRING_DATASOURCE_SAMPL EDB_JDBCURL

Because you commonly deploy virtual databases to multiple OpenShift environments, for example a staging and a production environment, you typically specify different data source properties in each environment. For example, the login credentials for accessing a data source in your staging environment probably differ from the credentials that you need to access the data source in the production environment. To define unique values in each environment, you can use environment variables.

The environment variables in **deploymentconfig.yml** replace any static properties that you might set in the **application.properties** file. If a property is defined in both files, the value in the **deploymentconfig.yml** file takes precedence.

You use a single version of the file across environments and use secret objects to isolate the unique details of each environment. Instead of specifying static values for environment variables directly in the file, you can store the values for each deployment environment in secret objects that are unique to each environment. The value of each environment variable in the file contains only a key reference, which specifies the name of a secret object, and the name of a token in the secret. The token stores the actual value. At runtime, environment variables retrieve their values from the tokens.

By using secrets to store the values of your environment variables, you can use a single version of the **deploymentconfig.yml** across environments. The secret objects that you deploy in each environment must have the same name, but each secret object contains token values that are specific to its environment.

Additional resources

- For more information about using secrets, see [Section 4.3, “Secret objects for storing data source information”](#).
- For information about adding a **deploymentconfig.yml** file, see [Section 4.12, “Setting the deployment configuration”](#).

4.12. SETTING THE DEPLOYMENT CONFIGURATION

You set the deployment configuration by editing a **deploymentconfiguration.yml** file. You can define environment variables in the file for each data source that the virtual database uses.

Prerequisites

- You have a copy of the sample **deploymentconfiguration.yml** file from the Teiid OpenShift repository.
- You have information about the connection settings for your data sources.
- If you want to use secrets to store values for your environment variables, you have information about the name of the secret for your virtual database and the names of the tokens that you want to refer to in your environment variables.

Procedure

1. In the `/src/main/fabric8` folder of your Java project, create a **deploymentconfiguration.yml** file.
2. Add environment variables and other settings that are consistent with your environment.

Additional resources

- For more information about the **deploymentconfiguration.yml** file, see [Section 4.11, “Deployment configuration files \(deploymentconfig.yml\)”](#).
- For more information about secrets, see [Section 4.3, “Secret objects for storing data source information”](#).
- For information about deployment configuration settings for other data sources, see [Section 4.13, “Connection settings for other data sources”](#).

4.13. CONNECTION SETTINGS FOR OTHER DATA SOURCES

To enable a virtual database to connect to a data source, you must provide connection details such as the name of the data source, the driver to use, the user credentials and so forth. You specify these settings across several files.

The sample files in the Teiid OpenShift repository, or elsewhere in this documentation provide configuration information that is consistent with using a PostgreSQL database as the data source for your virtual database. If you want to a different data source, you must modify settings in the PostgreSQL versions of the following files:

pom.xml file

Specifies the dependencies for a data source, such as the drivers that are required to connect to the data source. At build time, if the driver for your database type is publically available, Teiid Spring Boot downloads the required drivers automatically.

application.properties file

Defines static application properties that cannot be changed after you deploy the application to OpenShift, for example the name of the DDL file.

Deploymentconfig.yml

Defines application properties through dynamic environment variables so that you can specify values to correspond to a particular deployment environment

Datasources.java

Specifies a Java class to represent the connection to the data source. The service name that you specify in the annotation and as the name of the method must match exactly the name of the SERVER that is listed in the DDL file.

DDL file

Defines the virtual database structure including specific mapping from the source schema to the virtual schema.

When the virtual database service starts, the data virtualization service scans the application’s packages for dependency annotations and uses the information to build the metadata to create the virtual database and deploy it to the OpenShift server.

Prerequisites

- You have reviewed the PostgreSQL versions of the files in the preceding list.

4.13.1. Settings to connect to Salesforce as a data source

If you want to use a Salesforce database as a data source for your virtual database, you must add some source-specific settings. You can use the PostgreSQL files in the Teiid OpenShift repository as a starting point, but you'll have to modify the files to adapt them for use with Salesforce.

The following files contain information that you must modify to use a Salesforce database as a data source:

- **pom.xml**
- **application.properties**
- **deploymentconfig.yml**
- **datasources.java**
- DDL file

pom.xml settings for using Salesforce as a data source

To support Salesforce as a data source, you must add the following dependencies in the **pom.xml** file:

- `teiid-spring-boot-starter`
- `spring-data-salesforce`

For example:

```
<dependency>
  <groupId>org.teiid</groupId>
  <artifactId>teiid-spring-boot-starter</artifactId>
</dependency>
<dependency>
  <groupId>org.teiid.spring</groupId>
  <artifactId>spring-data-salesforce</artifactId>
</dependency>
```

application.properties settings to add to use Salesforce as a data source

You can specify values in **application.properties** to configure Salesforce as a data source, as shown in the following table. Each of the properties uses the prefix **spring.teiid.data.salesforce**

Property Name	Description	Default
<code>url</code>	Login URL	https://login.salesforce.com/services/Soap/u/45.0
<code>requestTimeout</code>	Request timeout	<code>Integer.MAX_VALUE</code>
<code>connectTimeout</code>	Connection timeout	<code>Integer.MAX_VALUE</code>
<code>pollingInterval</code>	Polling interval for bulk results	500

Property Name	Description	Default
clientId	OAuth2 client ID	N/A
clientSecret	OAuth2 client secret	N/A
refreshToken	OAuth2 refresh token	N/A
userName	User name	N/A
password	Password	N/A

datasources.java file for connecting to Salesforce as a data source

The **Datasources.java** file creates a class that Teiid can recognize as a Salesforce data source. The class acts as a resource adapter to enable the virtual database to access information in the data source.

```
package org.teiid.spring.example;

import org.springframework.boot.context.properties.ConfigurationProperties;
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.Configuration;
import org.teiid.spring.data.salesforce.SalesforceConfiguration;
import org.teiid.spring.data.salesforce.SalesforceConnectionFactory;

@Configuration
public class DataSources {

    @Bean
    public SalesforceConnectionFactory accounts(SalesforceConfiguration config) {
        return new SalesforceConnectionFactory(config);
    }

    @Bean
    @ConfigurationProperties("spring.teiid.data.salesforce")
    public SalesforceConfiguration salesforceConfig() {
        return new SalesforceConfiguration();
    }
}
```

The preceding class defines an **accounts** bean. When you create the virtual database, the data virtualization service recognizes the class as a data source, reads its metadata, and generates an internal model from it. The data virtualization service can then read from and write to it.

DDL file changes to connect a virtual database to Salesforce as a data source

Update the name of the **SERVER** object and the **FOREIGN DATA WRAPPER** to reflect the data source. For example, assign the following values to objects in the DDL file:

- Set the name of the **SERVER** object to **salesforce**
- Set the name of the **FOREIGN DATA WRAPPER** to **salesforce**

4.13.2. Settings to connect to Google Sheets as a data source

If you want to use a Google Sheets as a data source for your virtual database, you must add some source-specific settings. You can use the PostgreSQL files in the Teiid OpenShift repository as a starting point, but you'll have to modify the files to adapt them for use with Google Sheets.

The following files contain information that you must modify to use Google Sheets as a data source:

- **pom.xml**
- **application.properties**
- **Deploymentconfig.yml**
- **Datasources.java**
- DDL file

Pom.xml settings for using Google Sheets as a data source

To support Google Sheets as a data source, you must add the following dependencies in the **pom.xml** file:

- `teiid-spring-boot-starter`
- `spring-data-google`

For example:

```
<dependency>
  <groupId>org.teiid</groupId>
  <artifactId>teiid-spring-boot-starter</artifactId>
</dependency>
<dependency>
  <groupId>org.teiid.spring</groupId>
  <artifactId>spring-data-google</artifactId>
</dependency>
```

Application.properties settings to add to use Google Sheets as a data source.

You can specify values in **application.properties** to configure Google Sheets as a data source, as shown in the following table. Each of the properties uses the prefix **spring.teiid.data.google-sheets**

Property Name	Description	Default
<code>spread-sheet-name</code>	Name of the Google Spreadsheet	N/A
<code>spread-sheet-id</code>	Spreadsheet ID	Sheet Id, Look in URL of the spreadsheet. For more info see https://developers.google.com/sheets/api/guides/concepts#spreadsheet_id
<code>refresh-token</code>	OAuth2 refresh token	N/A

Property Name	Description	Default
client-id	Client ID	OAuth2 client ID
client-secret	Client secret	OAuth2 client secret
refresh-token	OAuth2 refresh token	N/A

Datasources.java file for connecting to Google Sheets as a data source

The **Datasources.java** file creates a class that Teiid can recognize as a Google Sheets data source. The class acts as a resource adapter to enable the virtual database to access information in the data source.

```
@Configuration
public class DataSources {

    @Bean
    public SpreadsheetConnectionFactory accounts(SpreadSheetConfiguration config) {
        return new SpreadsheetConnectionFactory(config);
    }

    @Bean
    @ConfigurationProperties("spring.teiid.data.google.sheets")
    public SpreadSheetConfiguration sheetsConfig() {
        return new SpreadSheetConfiguration();
    }
}
```

The preceding class defines an **accounts** bean. When you create the virtual database, the data virtualization service recognizes the class as a data source, reads its metadata, and generates an internal model from it. The data virtualization service read from and write to it.

DDL file changes to connect a virtual database to Google Sheets as a data source

Update the name of the **SERVER** object and the **FOREIGN DATA WRAPPER** to reflect the data source. For example, assign the following values to objects in the DDL file:

- Set the name of the **SERVER** object to **google-spreadsheet**.
- Set the name of the **FOREIGN DATA WRAPPER** to **google-spreadsheet**.

4.14. BUILDING A DATA VIRTUALIZATION PROJECT

After you complete the configuration tasks for your virtual database, you are ready to run Maven to build the project.

The Spring Boot Maven plugin creates a self-contained Uber JAR that includes all of the application code and dependencies in a single JAR file. The resulting JAR file serves as the basis for an OpenShift image. The OpenShift profile includes a Fabric8 Maven plugin that compiles the current build with the Uber JAR to generate an image that you can deploy to OpenShift.

After you make certain changes to the virtual database configuration, for example adding protocol services or routes for the virtual database, you must re-compile the project to update the deployed image.

Prerequisites

- You have completed all of the configuration tasks listed in [Chapter 2, Virtual database creation](#).

Procedure

- Log in to OpenShift and run the following command:

```
mvn clean install -Popenshift -Dfabric8.namespace=`oc project -q`
```

CHAPTER 5. MAKING VIRTUAL DATABASES AVAILABLE TO API CONSUMERS

To enable API consumers to access the virtual database, define services and routes for JDBC or OData protocols for the virtual database service on OpenShift.

After you define a protocol service for a virtual database, any client in the same OpenShift cluster can access the virtual database. Only applications that are in the same cluster have access. Remote clients do not have access. To enable OData access for remote clients, you must define an OData route to the virtual database service.

External JDBC clients do not use routes to access virtual database services. Instead, JDBC clients depend on the OpenShift load balancer service to allocate external IP addresses that external clients can use to access services in the cluster.

You configure services and routes by adding configuration files in **/src/main/fabric8**. Download the sample Fabric8 configuration files from the Teiid OpenShift repository. The following table lists the configuration files to add in **/src/main/fabric8** to configure services and routes for the virtual database.

Table 5.1. Fabric8 configuration files

Name	Purpose	Sample file
Deployment configuration	Controls overall deployment	deploymentconfig.yml
JDBC service configuration	Specifies JDBC service	jdbc-svc.yml
OData service configuration	Specifies OData service	odata-svc.yml
OData route configuration	Specifies OData route	odata-route.yml

If you choose not to create a service for one of the protocols, or to expose a route for that protocol, omit the corresponding file from the **/fabric8** directory.

The service and route configuration files specify default values for port numbers and timeout values. Unless you have a specific reason for modifying those settings, you can retain the default values.

Prerequisites

- You downloaded the sample configuration files that you need from the Teiid OpenShift repository.

5.1. CONFIGURING ACCESS FOR ODATA CLIENTS

The relational model of the data in a virtual database is automatically mapped to JSON or XML to allow applications to consume the data through OData APIs. However, if you want the virtual database to be available to OData clients, you must explicitly define an OData service and route. The default settings in the **odata-svc.yml** and **odata-route.yml** files that are available from the Teiid OpenShift sample repository are configured to enable an OData service and route.

Do not edit instances of the variable **\${project.artifactId}** that appear in the configuration files. At build time, these variables are replaced automatically with information from elements in the **pom.xml** file.

**NOTE**

OData routes are created automatically for virtual databases that you create in Fuse Online.

Prerequisites

- You have the sample template **odata-svc.yml** and **odata-route.yml** files from the Teiid OpenShift repository.
- You have completed the configuration tasks that are summarized in [Chapter 2, Virtual database creation](#).

Procedure

1. Add the following dependency in your **pom.xml** file:

```
<dependency>
  <groupId>org.teiid</groupId>
  <artifactId>spring-odata</artifactId>
</dependency>
```

2. In the **/src/main/fabric8** folder of your Java project, add an **odata-svc.yml** file to create an OData service.
3. If you want to make the virtual database available to OData clients outside of the OpenShift cluster, add an **odata-route.yml** file to create and OData route.

5.2. CONFIGURING ACCESS FOR JDBC CLIENTS

To enable JDBC client applications to access your virtual database, you must configure a JDBC service. After the JDBC service is enabled, OpenShift applications that share the cluster with the virtual database can access the database over JDBC. Third-party client applications that are outside of the OpenShift cluster have no JDBC access to the virtual database.

To provide JDBC access to external clients, you do not create a JDBC route as you do to enable OData access to external clients. Rather, you must configure an OpenShift load balancer service to configure ingress cluster traffic. After that you must provide external applications with the IP address that the load balancer service assigns to the virtual database.

Do not edit instances of the variable **`\${project.artifactId}** that appear in the configuration files. At build time, these variables are replaced automatically with information from elements in the **pom.xml** file.

Prerequisites

- You have the sample **jdbc-svc.yml** file from the Teiid OpenShift repository.
- You have completed the configuration tasks that are summarized in [Chapter 2, Virtual database creation](#).

Procedure

1. To create a JDBC service, add the file **jdbc-svc.yml** to the **/src/main/fabric8** folder of your Java project.

**NOTE**

If you add the **jdbc-svc.yml** file from the Teiid OpenShift sample repository, the JDBC service is enabled by default.

1. Create a file with the name **ingress** and add the following contents to it:

```
apiVersion: v1
kind: Service
metadata:
  name: rdbms-example-ingress
spec:
  ports:
    - name: teiid
      port: 31000
  type: LoadBalancer
  selector:
    app: rdbms-example
  sessionAffinity: ClientIP
```

2. Log in to OpenShift and run the following command to deploy the file to OpenShift:

```
$ oc create -f -ingress
```

3. Run the following command to determine the IP port:

```
$ oc get svc rdbms-example-ingress
```

4. Share the port number that is returned with your API clients.

Additional resources

- For more information about [Configuring ingress cluster traffic using a load balancer](#), see the OpenShift documentation.

5.3. IDENTIFYING THE ODATA ENDPOINT OF A VIRTUAL DATABASE

After you deploy a virtual database, you can share the OData URL with application developers so that they can use REST APIs to query and retrieve data from the virtual database.

**NOTE**

For virtual databases that you create in Fuse Online, OData routes are exposed automatically on the **Data Virtualizations** page.

To retrieve the OData URL for virtual databases that you create using the Spring Boot Maven plugin, you run an OpenShift command. You then append **/odata** to the URL that command returns to establish the API endpoint for the service. .



NOTE

In this Technology Preview release, there is no relationship between virtual databases that you create by running the Maven Spring Boot plugin and those that you create in Fuse Online. As a result, the **Data virtualization** page in Fuse Online does not show virtual databases that you build and create outside of Fuse Online.

Prerequisites

- You used Maven Spring Boot to deploy a virtual database service on OpenShift, and you enabled OData access to the service.

Procedure

- After you deploy the virtual database, log in to OpenShift and run the following command to obtain the OData URL for the service:

```
$ oc describe route <virtual-database-service-name>
```

For example, **oc describe route rdbms-example**

The command returns the root URL for the service.

Additional resources

- For information about how to enable OData access to a virtual database service, see see [Section 5.1, "Configuring access for OData clients"](#).

CHAPTER 6. SECURING ODATA APIS FOR A VIRTUAL DATABASE BY USING 3SCALE AND RH-SSO

You can integrate data virtualization with 3scale and Red Hat's single sign-on technology to apply advanced authorization and authentication controls to the OData endpoints for your virtual database services. By using 3Scale as a gateway to your API, you ensure that only authorized users have access, and you can control the level of access a user has based on who they are (role-based access).

By default, after you create a virtual database, the OData interface to it is discoverable by the 3Scale API management system, as long as 3Scale system is defined to same cluster and namespace. Using 3scale's API management features, you can control access to the OData API, and track usage. By further securing access to the API through SSO, you can define user roles and implement role-based access to the API endpoints. After you complete the configuration, you can control access in the virtual database at the level of the view, column, or data source.

The Red Hat SSO technology uses OpenID-Connect as the authentication mechanism to secure the API, and uses OAuth2 as the authorization mechanism. When a user logs in, 3scale negotiates authentication with the Red Hat SSO package. If the authentication succeeds, 3scale passes a security token to the OData API for verification. The OData API then reads permissions from the token and applies them to the data roles that are defined for the virtual database.

Prerequisites

- You have installed and configured Red Hat's single sign-on technology. For more information about how to configure SSO to secure OData APIs for a virtual database, see the [Teiid OpenShift documentation](#).
- You have 3scale API Management installed in the OpenShift cluster that hosts your virtual database.
- You have configured integration between 3scale and SSO. For more information, see [Configuring Red Hat Single Sign-On integration](#) in Using the Developer Portal.
 - You have configured the Red Hat SSO technology for OpenID Connect.
 - You have created SSO security realms.
 - You have specified an Admin client ID, protocol, and access type.
 - You have assigned the *realm-management* and *manage-clients* roles.
 - You created API users and specified credentials.
 - You configured 3scale to use OpenID-Connect as the authentication mechanism and OAuth2 as the authorization mechanism.

6.1. UPDATING THE VIRTUAL DATABASE CONFIGURATION

After you finish setting up single sign-on in OpenShift, you must modify virtual database settings before you can use SSO to secure your OData APIs. You must update the following artifacts:

- **pom.xml**
- **application.properties**

- Configuration map
- **deploymentconfig.yml**
- DDL file

After you update the virtual database configuration, update the 3scale configuration to complete the integration.

Prerequisites

- You completed the SSO configuration.

6.1.1. Adding SSO dependencies to pom.xml

Prerequisites

- You completed the configuration to integrate 3scale with SSO.

Procedure

- Edit the **pom.xml** file to add the following SSO dependencies:

```
<dependency>  
  <groupId>org.teiid</groupId>  
  <artifactId>spring-keycloak</artifactId>  
</dependency>
```

6.1.2. Adding SSO settings to application.properties

Update properties in the **application.properties** file to support integration with SSO.

Procedure

1. Log in to the Red Hat Single Sign-On Admin Console to find the values for the following authentication properties:
 - keycloak.realm
 - keycloak.auth-server-url
 - keycloak.resource
 - keycloak.ssl-required
 - keycloak.public-client
 - keycloak.principal-attribute
2. Open the **application.properties** file from **src/main/resources/** and add values for the preceding properties. For example:

```
keycloak.realm = 3scale-sso  
keycloak.auth-server-url = http://keycloak-staging.dev.openshiftapps.com/auth
```

```
keycloak.resource = di
keycloak.ssl-required = external
keycloak.public-client = true
keycloak.principal-attribute=preferred_username
```

6.1.3. Creating a configuration map

In **application.properties** you assign static properties that do not change between deployment environments. For settings that can change between environments, such as authentication URLs and credentials, you might want to apply different values depending on the environment. To make it easy to apply different values for environment-specific settings, you can specify their values as properties in a configuration map. In each environment in which you deploy a virtual database, use a configuration map with the same name, but vary the map contents (the authentication URL and secret) to match the environment. Then, in each environment where you deploy the virtual database, the virtual database retrieves its version of the authentication values from the map.

You run an OpenShift command to create a configuration map and add values for the following properties:

- **keycloak.auth-server-url**
- **keycloak.credentials.secret**

Procedure

1. Run the following command:

```
oc create configmap my-config --from-literal=keycloak.auth-server-url=http://<keycloakURL>
--from-literal=keycloak.credentials.secret=<secretValue>
```

For example,

```
oc create configmap my-config --from-literal=keycloak.auth-server-url=http://keycloak-
staging.dev.example.com/auth --from-literal=keycloak.credentials.secret=xxxxxxxx-xxxx-
xxxx-xxxx-xxxxxxxxxxxx
```

The preceding command creates a configuration map in OpenShift with the name *my-config*, which you can then reference from the deployment configuration file (**deploymentconfig.yml**).

6.1.4. Updating SSO environment variables in the deploymentconfig.yml file

In the **deploymentconfig.yml** file you can specify environment variables to apply different settings in each environment where you deploy a virtual database. For example, if you are deploying a virtual database to both a staging environment and a production environment, you can specify different environment variables in each to customize the settings for that environment.

To support SSO, you must add environment variables to the **deploymentconfig.yml** file.

Procedure

- From **src/main/fabric8**, open the **deploymentconfig.yml** file and add the following environment variables:

```
- name: KEYCLOAK_AUTHSERVERURL
```

```

valueFrom:
  configMapKeyRef:
    name: <config-name>
    key: keycloak.auth-server-url
- name: KEYCLOAK_CREDENTIALS_SECRET
valueFrom:
  configMapKeyRef:
    name: <config-name>
    key: keycloak.credentials.secret

```

For example,

```

- name: KEYCLOAK_AUTHSERVERURL
valueFrom:
  configMapKeyRef:
    name: my-config
    key: keycloak.auth-server-url
- name: KEYCLOAK_CREDENTIALS_SECRET
valueFrom:
  configMapKeyRef:
    name: my-config
    key: keycloak.credentials.secret

```

If you deploy a virtual database in multiple environments that use different realm or client settings, be sure to specify unique realm and client properties in the SSO configuration that you establish for each environment.

6.1.5. Defining user roles in the DDL file

To apply unique data roles to different users, you must define those roles in the DDL file.

Procedure

1. Add the following lines to the DDL file in **src/main/resources**:

```

CREATE ROLE ReadRole WITH JAAS ROLE ReadRole;
GRANT SELECT ON TABLE "<tableName.fieldName>" TO ReadRole

```

For example,

```

CREATE ROLE ReadRole WITH JAAS ROLE ReadRole;
GRANT SELECT ON TABLE "portfolio.CustomerZip" TO ReadRole

```

In the preceding example, the first line creates a role with the name **ReadRole**. Map this role to a role with the same name that you created in the SSO configuration, as described in [Securing an OData API with Keycloak](#) in the Teiid OpenShift repository. You can use a different role name, but for simplicity, it's best to use the same name. The second line grants SELECT permission to the **portfolio.CustomerZip** View to users who are assigned the **ReadRole**.

After you make changes to the project to integrate 3scale and enable SSO, rebuild the virtual database image and deploy it to OpenShift. For information about how to build and deploy the image, see [Section 4.14, "Building a data virtualization project"](#).

6.2. UPDATING THE 3SCALE CONFIGURATION

Prerequisites

- You have configured SSO.
- You have made updates to the virtual database service to support SSO.

Procedure

1. Log in to the 3scale admin portal, and from the page header, click the name of the API for which you want to enable OpenId Connect integration with SSO.
2. From the **Overview** page, click **Configuration** in the **Configuration, Methods and Settings** section.
3. Click **edit integration settings**.
4. In the **Authentication** section, click **OpenID Connect**, and click **Update Service**.
5. From the **Configuration** page, click the **edit APIcast configuration**.
6. Expand the **Authentication Settings** section to view the authentication options.
7. In the **OpenID Connect Issuer** field, type the URL of your SSO server, and the client ID and client secret that you previously defined in Red Hat SSO Admin Console.
8. Click **Update the Staging Environment**
9. Create a new application so that 3scale can synchronize it with SSO.
 - a. From the page header, click **Audience**, and then select a Developer account.
 - b. From the submenu, click **Applications**.
 - c. Click **Create Application**, and then type a name and description for the application.
 - d. For more information about creating applications, see [Set up applications](#).
 - e. Select an application plan for the virtual database API that you want to secure. You can select an existing application plan or create a new one. For information about how to create an application plan, see [How to create an application plan](#) in the 3scale Admin Portal Guide.
 - f. From the **API Credentials** section, record the values of the client ID and client secret. You must supply these credentials to access the API. If no value is displayed for the client secret, click **Add random key** to populate the field.
 - g. Specify a redirect URL. You use the redirect URL to test the integration in Postman or other REST clients.
 - h. In the Redirect URL field click **Edit**, type a callback link, for example, <https://openidconnect.net/callback>, and then click **Update**.

CHAPTER 7. CREATING AND WORKING WITH VIRTUAL DATABASES IN FUSE ONLINE

In Fuse Online, you can create a virtual database that integrates data from multiple data sources that you choose. After you deploy the resulting virtual database service, application developers and other database users can connect to as if it were a single physical database.

After you create a virtual database, you can use Fuse Online tools to:

- Add or remove data sources.
- Add or edit views of data from different tables or sources.
- Submit SQL queries to test that views return the expected results.
- Modify the schema that defines the virtual database.
- Publish the virtual database to make it available on OpenShift.
- Delete the virtual database.

Prerequisites

- You have the 2019-07 release of Red Hat Integration, and you are running Fuse 7.4.
- The data virtualization UI for Fuse 7.4 was enabled during installation. For more information, see [Enabling data virtualization in Fuse Online on OCP](#) in the Installing and Operating Fuse Online on OpenShift Container Platform.

7.1. CREATING VIRTUAL DATABASES IN FUSE ONLINE

In Fuse Online, you can create virtual databases that import views from applications or services that are available from the **Connections** page.

For each virtual database that you create, you must import data sources, and select the tables from each data source that you want to include. The views in the resulting virtual database map directly to the database tables that you import. After the initial creation, you can add views to a virtual database that join data from more than one table.



NOTE

In this Technology Preview, you can create virtual databases in Fuse Online only from relational databases, MongoDB, and Salesforce.

Prerequisites

- Your Fuse Online environment has a connection to one or more of the following data sources:
 - Relational database, such as PostgreSQL or MySQL.
 - MongoDB database
 - Salesforce database

Procedure

1. From the navigation sidebar in Fuse Online, click **Data**.
2. Click **Create Data Virtualization**.
3. On the **Create New Data Virtualization** page, type a name for the virtual database and click **Create**.
 - Provide a name that informs people about the database contents or purpose, and that is short enough for application developers and SQL users to easily insert in their code.
 - Names can include only alphanumeric ([a-z][A-Z], [0-9]), and hyphen (-) characters.
4. On the page for your virtualization, click **Import Data Source**.
5. On the **Import Data Source** page, click the tile for an active data source, and then click **Next**.
6. From the list of tables that are available, select one or more tables to include in your virtual database and then click **Done**.

A confirmation message reports when the import completes. The **Views** tab for the draft virtualization lists a view for each table that you imported.

You can now edit the existing views, create another view, or publish the virtual database to make it available for use.

7.2. ADDING A VIEW TO A VIRTUAL DATABASE IN FUSE ONLINE

Add a view to a virtual database to provide a view of the data in a new table.

After you first create a virtual database, it contains only the views that you imported from the initial data source. Add views to the virtual database if you want to incorporate data from other tables. You can add views based on tables in the original data source, or from other data sources.



NOTE

In this Technology Preview release, you can only add one table to each view.

Prerequisites

- The virtual database that you want to add a view to is available in Fuse Online in a *Draft* or *Published* state. You cannot use Fuse Online to add views to virtual databases that were created outside of Fuse Online.
- A Fuse Online connection exists to the data source that contains the table that you want integrate.
- You know the name of the table that you want to use in the view.

Procedure

1. From the navigation sidebar in Fuse Online, click **Data**.
2. From the list on the **Data Virtualizations** page, find the virtual database that you want to modify and click **Edit**.

3. Click **Create a View**.
4. Expand a data source to view the tables that it contains.
5. Select the table that you want to add to the virtual database, and then click **Next**.
6. On the **Create a View** page, type a name in the **View Name** field, and then click **Done**.
The **View Editor** displays the SQL for the view that you created. The **Preview** panel displays the data in the view.
7. If no data displays, click **Refresh**.
8. Click **Done** to close the view.
If the virtual database was previously published, you must republish it to make the new view available.

Additional resources

- Experienced SQL programmers can also add views by directly editing the default SQL statements for the virtual database. For more information, see [Section 7.3, “Using the View Editor in Fuse Online to modify the DDL that defines a virtual database”](#).
- [Section 7.5, “Publishing virtual databases in Fuse Online to make them available for access”](#)
- [Section 7.4, “Previewing a virtual database in Fuse Online by submitting SQL test queries”](#)
- [Section 7.3, “Using the View Editor in Fuse Online to modify the DDL that defines a virtual database”](#)

7.3. USING THE VIEW EDITOR IN FUSE ONLINE TO MODIFY THE DDL THAT DEFINES A VIRTUAL DATABASE

The process of creating a virtual database in Fuse Online is designed to automate many tasks and hide the complexities of the underlying SQL code.

When you create a view for a virtual database, Fuse Online automatically generates the data definition language (DDL) that defines the view. The DDL is a set of SQL statements that describe the view's schema, tables, columns, and other fields.

Fuse Online provides tools to add basic views for a virtual database, but if you know SQL and you want greater control in designing a view, you can directly edit the DDL for the view. In Fuse Online, developers can use the embedded View Editor to modify these SQL statements. To assist you, this SQL editor includes a code-completion feature that provides a list of SQL keywords.

After you save your changes, a built-in validation tool runs to ensure that the SQL code does not contain syntax errors.

Prerequisites

- You have experience using a data definition language (DDL) that is based on the SQL-MED specification to define database structures and to integrate externally stored data.

Procedure

1. From the navigation sidebar in Fuse Online, click **Data**.

2. On the **Data Virtualizations** page, find the virtual database that you want to modify and click **Edit**.
3. In the **Views** tab, find the view that you want to edit, and then click **Edit**.
4. Update the SQL as needed. As you edit, press Ctrl+Space to open the code completion tool.
5. After you complete your changes, click **Save**.
Fuse Online validates the SQL and returns an error if the view contains invalid code.

After the SQL validates, the preview panel shows the result of the updates that you made to the view. The preview displays the first fifteen rows of the results set.

6. Click **Done** to close the **View Editor** and return to the list of views.
If the virtual database was previously published, you must republish it to put your changes into effect.

Additional resources

- [Section 7.5, “Publishing virtual databases in Fuse Online to make them available for access”](#)
- For more information about using SQL in data virtualization DDL files, see the [Teiid Reference Guide](#).
- You can modify the results set by altering the default query to specify different row limits or row offsets. For more information, see [Section 7.4, “Previewing a virtual database in Fuse Online by submitting SQL test queries”](#)

7.4. PREVIEWING A VIRTUAL DATABASE IN FUSE ONLINE BY SUBMITTING SQL TEST QUERIES

Before you publish a virtual database and make it available to applications, you can run test queries against its views to verify that it returns the information that you expect.

Although the default preview shows you the first 15 results returned when a SQL **SELECT * FROM** statement is submitted to a virtual database view, you can use the embedded SQL client in Fuse Online to send modified test queries to your views. You can adjust the default results set by specifying the row limits and row offsets.

If the view that you query originates from a non-SQL data source, the data virtualization engine converts the SQL query into a format that the data source can interpret.

Prerequisites

- You have a valid virtual database that was created in Fuse Online.

Procedure

1. From the navigation sidebar in Fuse Online, click **Data**.
2. On the **Data Virtualizations** page, click **Edit** in the entry for the virtual database that contains the view that you want to test.
3. Click the **SQL Client** tab.

4. From the **View** field, select the view that you want to test.
5. In the **Row Limit** field, specify the number of rows to display.
6. In the **Row Offset** field, specify the number of rows to skip.
7. Click **Submit**. The **Query Results** table displays the result set.

7.5. PUBLISHING VIRTUAL DATABASES IN FUSE ONLINE TO MAKE THEM AVAILABLE FOR ACCESS

After you define a virtual database in Fuse Online, you must publish it to make it available for users and applications to access.

Publishing a virtual database builds the schema definition that you implemented by importing data sources and views into a runtime image. Fuse Online deploys the runtime image to OpenShift as a virtual database container image that you can scale independently.

After you publish the virtual database, it becomes available as a service and is represented on the Fuse Online **Connections** page. The service behaves like any relational database, and clients can connect to it over standard interfaces. It can be incorporated into Fuse Online integration workflows, and it is available to JDBC and OData clients.

Prerequisites

- You created a virtual database in Fuse Online.
- You added any views that you want to the virtual database.

Procedure

1. From the navigation sidebar in Fuse Online, click **Data**.
2. On the **Data Virtualizations** page, find a virtual database that you want to publish, and from the overflow menu, click **Publish**.
A confirmation message notifies you that the virtual database was submitted for publishing, and a progress bar reports the status of the process.

- If the publishing process succeeds, Fuse Online makes the following updates:
 - The status label of the virtual database entry on the **Data virtualizations** page changes from **Draft** to **Published**.
 - The virtual database entry displays a URL link to the OData endpoint for the virtual database.
 - The virtual database service is added to the **Connections** page, and a JDBC connection to it is created.
You can verify the JDBC URL by opening the entry for the virtual database service from the **Connections** page.
- If the publishing process fails, the entry is flagged with the label **Error**.

7.6. DELETING A VIRTUAL DATABASE IN FUSE ONLINE

You can permanently delete virtual databases that you create in Fuse Online. You can delete virtual databases whether they are published or in draft.

The data sources that a virtual database consumes are not affected by the deletion. Connections between Fuse Online and the data sources remain in place.

Prerequisites

- You have a virtual database that was created in Fuse Online and you want to remove it.

Procedure

1. From the navigation sidebar in Fuse Online, click **Data**.
2. On the **Data Virtualizations** page, click the overflow menu for the virtual database that you want to delete, and then click **Delete**.
3. When prompted, click **Delete** to confirm that you want to delete the virtual database. A confirmation message reports when the virtualization is deleted.

CHAPTER 8. MIGRATING LEGACY VIRTUAL DATABASE FILES TO DDL FORMAT

The data virtualization Technology Preview requires that you define the structure of virtual databases in SQL-MED DDL (data definition language) format. By contrast, the structure of legacy Teiid virtual databases, such as those that run on Wildfly, or on the Red Hat JBoss Data Virtualization offering, are defined by using files that are in **.xml** or **.vdb** format.

You can reuse the virtual database designs that you developed for a legacy deployment, but you must first update the format of the files. A migration tool is available to convert your files. After your convert the files you can rebuild the virtual databases as container images and deploy them to OpenShift.

You can use the migration utility in the following two ways:

To validate a VDB file only

Use this method to check whether the utility can a successfully convert a VDB file. The utility converts the VDB file and reports validation errors to the terminal. If there are no validation errors, the utility displays the resulting DDL, but it does not save the converted DDL to a file.

To validate and a VDB file and save it to a DDL file

The file is saved only if there are no validation errors.

The migration tool works on **.xml** files only. Files with a **.vdb** file extension are file archives that contain multiple folders. If you have legacy files in **.vdb** format, use Teiid Designer to export the files to **.xml** format, and then run the migration tool to convert the resulting **.xml** files.

Prerequisites

- You have a legacy virtual database file in **.xml** format.
- You download the [settings.xml](#) file from the Teiid OpenShift repository. Maven uses the information in the file to run the migration tool.

8.1. VALIDATING A LEGACY VIRTUAL DATABASE XML FILE AND VIEWING IT IN DDL FORMAT

You can run a test conversion on a legacy virtual database to check for validation errors and view the resulting DDL file. When you run the migration tool in this way, the converted DDL file is not saved.

Procedure

1. Open the directory that contains the **settings.xml** file that you downloaded from the Teiid OpenShift repository, and type the following command:

```
$ mvn -s settings.xml exec:java -Dvdb=<path_to_vdb_xml_file>
```

For example:

```
$ mvn -s settings.xml exec:java -Dvdb=rdbms-example/src/main/resources/vdb.xml
```

The migration tool checks the specified **.xml** file, and reports any validation errors. If there are no validation errors, the migration tool displays a **.ddl** version of the virtual database on the screen.

8.2. CONVERTING A LEGACY VIRTUAL DATABASE XML FILE AND SAVING IT AS A DDL FILE

You can run the migration tool to convert a legacy virtual database file to **.ddl** format, and then save the **.ddl** file to a specified directory. The migration tool checks the **.xml** file that you provide for validation errors. If there are no validation errors, the migration tool converts the file to **.ddl** format and saves it to the file name and directory that you specify.

Procedure

- Open the directory that contains the **settings.xml** file that you downloaded from the Teiid OpenShift repository, and type the following command:

```
$ mvn -s settings.xml exec:java -Dvdb=<path_to_vdb_xml_file> -Doutput=  
<path_to_save_ddl_file>
```

For example:

```
$ mvn -s settings.xml exec:java -Dvdb=rdbms-example/src/main/resources/vdb.xml -  
Doutput=rdbms-example/src/main/resources/vdb.ddl
```