



Red Hat Decision Manager 7.13

Deploying Red Hat Decision Manager on Red Hat OpenShift Container Platform

Red Hat Decision Manager 7.13 Deploying Red Hat Decision Manager on Red Hat OpenShift Container Platform

Legal Notice

Copyright © 2024 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

This document describes how to deploy a variety of Red Hat Decision Manager environments on Red Hat OpenShift Container Platform, such as an authoring environment, a managed server environment, an immutable server environment, and other supported environment options.

Table of Contents

PREFACE	3
MAKING OPEN SOURCE MORE INCLUSIVE	4
PART I. DEPLOYING A RED HAT DECISION MANAGER ENVIRONMENT ON RED HAT OPENSIFT CONTAINER PLATFORM 4 USING OPERATORS	5
CHAPTER 1. OVERVIEW OF RED HAT DECISION MANAGER ON RED HAT OPENSIFT CONTAINER PLATFORM	6
1.1. ARCHITECTURE OF AN AUTHORIZING ENVIRONMENT	6
Single authoring environment	6
Clustering KIE Servers and using multiple KIE Servers	7
Smart Router	7
High-availability authoring environment	7
CHAPTER 2. PREPARATION FOR DEPLOYING RED HAT DECISION MANAGER IN YOUR OPENSIFT ENVIRONMENT	9
2.1. ENSURING YOUR ENVIRONMENT IS AUTHENTICATED TO THE RED HAT REGISTRY	9
2.2. CREATING THE SECRETS FOR KIE SERVER	9
2.3. CREATING THE SECRETS FOR BUSINESS CENTRAL	10
2.4. CREATING THE SECRETS FOR THE AMQ BROKER CONNECTION	11
2.5. PREPARING GIT HOOKS	11
2.6. PROVISIONING PERSISTENT VOLUMES WITH READWRITEMANY ACCESS MODE USING NFS	12
2.7. EXTRACTING THE SOURCE CODE FROM BUSINESS CENTRAL FOR USE IN AN S2I BUILD	12
2.8. PREPARING FOR DEPLOYMENT IN A RESTRICTED NETWORK	13
2.9. PREPARING A MAVEN MIRROR REPOSITORY FOR OFFLINE USE	14
CHAPTER 3. DEPLOYMENT AND MANAGEMENT OF A RED HAT DECISION MANAGER ENVIRONMENT USING OPENSIFT OPERATORS	16
3.1. SUBSCRIBING TO THE BUSINESS AUTOMATION OPERATOR	16
3.2. DEPLOYING A RED HAT DECISION MANAGER ENVIRONMENT USING THE OPERATOR	17
3.2.1. Starting the deployment of a Red Hat Decision Manager environment using the Business Automation operator	17
3.2.2. Setting the basic configuration of the environment	17
3.2.2.1. Configuring the image registry to use specific images	20
3.2.3. Setting the security configuration of the environment	22
3.2.4. Setting the Business Central configuration of the environment	25
3.2.5. Setting custom KIE Server configuration of the environment	28
3.3. MODIFYING AN ENVIRONMENT THAT IS DEPLOYED USING OPERATORS	34
3.4. PROVIDING ELYTRON USER CONFIGURATION OR OTHER POST-CONFIGURATION SETTINGS	35
3.5. JVM CONFIGURATION PARAMETERS	36
3.6. KIE CONFIGURATION AND CONFIGMAPS	38
3.6.1. Using ConfigMaps	41
3.7. CREATING CUSTOM IMAGES FOR KIE SERVER	43
3.7.1. Creating a custom KIE Server image with an additional RPM package	43
3.7.2. Creating a custom KIE Server image with an additional JAR file	45
CHAPTER 4. MIGRATION OF INFORMATION FROM A DEPLOYMENT ON RED HAT OPENSIFT CONTAINER PLATFORM 3	47
4.1. MIGRATING INFORMATION IN BUSINESS CENTRAL	47
APPENDIX A. VERSIONING INFORMATION	49
APPENDIX B. CONTACT INFORMATION	50

PREFACE

As a developer or system administrator, you can deploy a variety of Red Hat Decision Manager environments on Red Hat OpenShift Container Platform, such as an authoring environment, a managed server environment, an immutable server environment, and other supported environment options.

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

PART I. DEPLOYING A RED HAT DECISION MANAGER ENVIRONMENT ON RED HAT OPENSIFT CONTAINER PLATFORM 4 USING OPERATORS

As a system engineer, you can deploy a Red Hat Decision Manager environment on Red Hat OpenShift Container Platform 4 to provide an infrastructure to develop or execute services and other business assets. You can use OpenShift Operators to deploy the environment defined in a structured YAML file and to maintain and modify this environment as necessary.

Prerequisites

- A Red Hat OpenShift Container Platform 4 environment is available. For the exact versions of Red Hat OpenShift Container Platform that the current release supports, see [Red Hat Decision Manager 7 Supported Configurations](#).
- The OpenShift project for the deployment is created.
- You are logged into the project using the OpenShift web console.
- The following resources are available on the OpenShift cluster. Depending on the application load, higher resource allocation might be necessary for acceptable performance.
 - For an authoring environment, 4 gigabytes of memory and 2 virtual CPU cores for the Business Central pod. In a high-availability deployment, these resources are required for each replica and two replicas are created by default.
 - 2 gigabytes of memory and 1 virtual CPU core for each replica of each KIE Server pod.
 - In a high-availability authoring deployment, additional resources according to the configured defaults are required for the Red Hat AMQ, and Red Hat Data Grid pods.



NOTE

The default values for **MaxMetaspaceSize** are:

- Business Central images: 1024m
 - KIE Server images: 512m
 - For other images: 256m
- Dynamic persistent volume (PV) provisioning is enabled. Alternatively, if dynamic PV provisioning is not enabled, enough persistent volumes must be available. By default, the deployed components require the following PV sizes:
 - By default, Business Central requires one 1Gi PV. You can change the PV size for Business Central persistent storage.
 - If you intend to deploy a high-availability authoring environment, your OpenShift environment supports persistent volumes with **ReadWriteMany** mode. If your environment does not support this mode, you can use NFS to provision the volumes. For information about access mode support in OpenShift public and dedicated clouds, see [Access Modes](#) in Red Hat OpenShift Container Platform documentation.

CHAPTER 1. OVERVIEW OF RED HAT DECISION MANAGER ON RED HAT OPENSIFT CONTAINER PLATFORM

You can deploy Red Hat Decision Manager into a Red Hat OpenShift Container Platform environment.

In this solution, components of Red Hat Decision Manager are deployed as separate OpenShift pods. You can scale each of the pods up and down individually to provide as few or as many containers as required for a particular component. You can use standard OpenShift methods to manage the pods and balance the load.

The following key components of Red Hat Decision Manager are available on OpenShift:

- KIE Server, also known as *Execution Server*, is the infrastructure element that runs decision services and other deployable assets (collectively referred to as *services*). All logic of the services runs on execution servers.

In some templates, you can scale up a KIE Server pod to provide as many copies as required, running on the same host or different hosts. As you scale a pod up or down, all of its copies run the same services. OpenShift provides load balancing and a request can be handled by any of the pods.

You can deploy a separate KIE Server pod to run a different group of services. That pod can also be scaled up or down. You can have as many separate replicated KIE Server pods as required.

- Business Central is a web-based interactive environment used for authoring services. It also provides a management console. You can use Business Central to develop services and deploy them to KIE Servers.

Business Central is a centralized application. However, you can configure it for high availability, where multiple pods run and share the same data.

Business Central includes a Git repository that holds the source for the services that you develop on it. It also includes a built-in Maven repository. Depending on configuration, Business Central can place the compiled services (KJAR files) into the built-in Maven repository or (if configured) into an external Maven repository.

You can arrange these and other components into various environment configurations within OpenShift.

1.1. ARCHITECTURE OF AN AUTHORIZING ENVIRONMENT

In Red Hat Decision Manager, the Business Central component provides a web-based interactive user interface for authoring services. The KIE Server component runs the services.

You can also use Business Central to deploy services onto a KIE Server. You can use several KIE Servers to run different services and control the servers from the same Business Central.

Single authoring environment

In a single authoring environment, only one instance of Business Central is running. Multiple users can access its web interface at the same time, however the performance can be limited and there is no failover capability.

Business Central includes a built-in Maven repository that stores the built versions of the services that you develop (KJAR files/artifacts). You can use your continuous integration and continuous deployment (CI/CD) tools to retrieve these artifacts from the repository and move them as necessary.

Business Central saves the source code in a built-in Git repository, stored in the **.niogit** directory. It uses a built-in indexing mechanism to index the assets in your services.

Business Central uses persistent storage for the Maven repository and for the Git repository.

A single authoring environment, by default, includes one KIE Server instance.

A single authoring environment can use the *controller strategy*. Business Central includes the *Controller*, a component that can manage KIE Servers. When you configure KIE Server to connect to Business Central, KIE Server uses a REST API to connect to the Controller. This connection opens a persistent WebSocket. In an OpenShift deployment that uses the controller strategy, each KIE Server instance is initially configured to connect to the Business Central Controller.

When you use the Business Central user interface to deploy or manage a service on KIE Server, KIE Server receives the request through the Controller connection WebSocket. To deploy a service, KIE Server requests the necessary artifact from the Maven repository that is a part of Business Central.

Client applications use a REST API to use services that run on KIE Server.

Figure 1.1. Architecture diagram for a single authoring environment



Clustering KIE Servers and using multiple KIE Servers

You can scale a KIE Server pod to run a clustered KIE Server environment.

In a clustered deployment, several instances of KIE Server run the same services. These servers can connect to the Business Central Controller using the same server ID, so they can receive the same requests from the controller. Red Hat OpenShift Container Platform provides load-balancing between the servers. The services that run on a clustered KIE Server instance must be stateless, because requests from the same client might be processed by different instances.

You can also deploy several independent KIE Servers to run different services. In this case, the servers connect to the Business Central Controller with different server ID values. You can use the Business Central UI to deploy services to each of the servers.

Smart Router

The optional Smart Router component provides a layer between client applications and KIE Server instances. It can be useful if you are using several independent KIE Server instances.

The client application can use services running on different KIE Server instances, but always connects to the Smart Router. The Smart Router automatically passes the request to the KIE Server instances that runs the required service. The Smart Router also enables management of service versions and provides an additional load-balancing layer.

High-availability authoring environment

In a high-availability (HA) authoring environment, the Business Central pod is scaled, so several

instances of Business Central are running. Red Hat OpenShift Container Platform provides load balancing for user requests. This environment provides optimal performance for multiple users and supports failover.

Each instance of Business Central includes the Maven repository for the built artifacts and uses the **.niogit** Git repository for source code. The instances use shared persistent storage for the repositories. A persistent volume with **ReadWriteMany** access is required for this storage.

An instance of Red Hat DataGrid provides indexing of all projects and assets developed in Business Central.

An instance of Red Hat AMQ propagates Java CDI messages between all instances of Business Central. For example, when a new project is created or when an asset is locked or modified on one of the instances, this information is immediately reflected in all other instances.

The controller strategy is not suitable for clustered deployment. In an OpenShift deployment, a high-availability Business Central must manage KIE Servers using the *OpenShift startup strategy*.

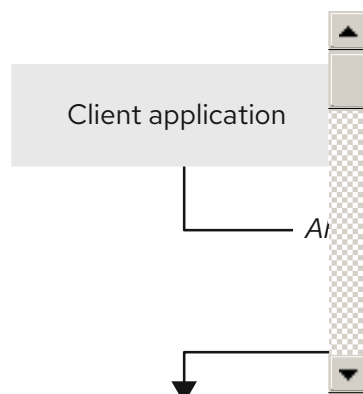
Each KIE Server deployment (which can be scaled) creates a ConfigMap that reflects its current state. The Business Central discovers all KIE Servers by reading their ConfigMaps.

When the user requests a change in the KIE Server configuration (for example, deploys or undeploys a service), Business Central initiates a connection to KIE Server and sends a REST API request. KIE Server changes the ConfigMap to reflect the new configuration state and then triggers its redeployment, so that all instances are redeployed and reflect the new configuration. For more information about ConfigMaps, see [KIE configuration and ConfigMaps](#).

You can deploy several independent KIE Servers in your OpenShift environment. Each of the KIE Servers has a separate ConfigMap with the necessary configuration. You can scale each of the KIE Servers separately.

You can include Smart Router in the OpenShift deployment.

Figure 1.2. Architecture diagram for a high-availability authoring environment



CHAPTER 2. PREPARATION FOR DEPLOYING RED HAT DECISION MANAGER IN YOUR OPENSIFT ENVIRONMENT

Before deploying Red Hat Decision Manager in your OpenShift environment, you must complete several procedures. You do not need to repeat these procedures if you want to deploy additional images, for example, for new versions of decision services or for other decision services



NOTE

If you are deploying a trial environment, complete the procedure described in [Section 2.1, “Ensuring your environment is authenticated to the Red Hat registry”](#) and do not complete any other preparation procedures.

2.1. ENSURING YOUR ENVIRONMENT IS AUTHENTICATED TO THE RED HAT REGISTRY

To deploy Red Hat Decision Manager components of Red Hat OpenShift Container Platform, you must ensure that OpenShift can download the correct images from the Red Hat registry.

OpenShift must be configured to authenticate with the Red Hat registry using your service account user name and password. This configuration is specific for a namespace, and if operators work, the configuration is already completed for the **openshift** namespace.

However, if the image streams for Red Hat Decision Manager are not found in the **openshift** namespace or if the operator is configured to update Red Hat Decision Manager to a new version automatically, the operator needs to download images into the namespace of your project. You must complete the authentication configuration for this namespace.

Procedure

1. Ensure you are logged in to OpenShift with the **oc** command and that your project is active.
2. Complete the steps documented in [Registry Service Accounts for Shared Environments](#). You must log in to Red Hat Customer Portal to access the document and to complete the steps to create a registry service account.
3. Select the **OpenShift Secret** tab and click the link under **Download secret** to download the YAML secret file.
4. View the downloaded file and note the name that is listed in the **name:** entry.
5. Run the following commands:

```
oc create -f <file_name>.yaml
oc secrets link default <secret_name> --for=pull
oc secrets link builder <secret_name> --for=pull
```

Replace **<file_name>** with the name of the downloaded file and **<secret_name>** with the name that is listed in the **name:** entry of the file.

2.2. CREATING THE SECRETS FOR KIE SERVER

OpenShift uses objects called *secrets* to hold sensitive information such as passwords or keystores. For more information about OpenShift secrets, see [What is a secret](#) in the Red Hat OpenShift Container Platform documentation.

In order to provide HTTPS access, KIE Server uses an SSL certificate. The deployment can create a sample secret automatically. However, in production environments you must create an SSL certificate for KIE Server and provide it to your OpenShift environment as a secret.

Procedure

1. Generate an SSL keystore named **keystore.jks** with a private and public key for SSL encryption for KIE Server. For more information about creating keystores and using certificates, see [How to Configure Server Security](#).



NOTE

In a production environment, generate a valid signed certificate that matches the expected URL for KIE Server.

2. Record the name of the certificate. The default value for this name in Red Hat Decision Manager configuration is **jboss**.
3. Record the password of the keystore file. The default value for this name in Red Hat Decision Manager configuration is **mykeystorepass**.
4. Use the **oc** command to generate a secret named **kieserver-app-secret** from the new keystore file:

```
$ oc create secret generic kieserver-app-secret --from-file=keystore.jks
```

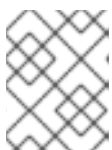
2.3. CREATING THE SECRETS FOR BUSINESS CENTRAL

In order to provide HTTPS access, Business Central uses an SSL certificate. The deployment can create a sample secret automatically. However, in production environments you must create an SSL certificate for Business Central and provide it to your OpenShift environment as a secret.

Do not use the same certificate and keystore for Business Central and KIE Server.

Procedure

1. Generate an SSL keystore named **keystore.jks** with a private and public key for SSL encryption for KIE Server. For more information about creating keystores and using certificates, see [How to Configure Server Security](#).



NOTE

In a production environment, generate a valid signed certificate that matches the expected URL for Business Central.

2. Record the name of the certificate. The default value for this name in Red Hat Decision Manager configuration is **jboss**.
3. Record the password of the keystore file. The default value for this name in Red Hat Decision Manager configuration is **mykeystorepass**.

4. Use the **oc** command to generate a secret named **businesscentral-app-secret** from the new keystore file:

```
$ oc create secret generic businesscentral-app-secret --from-file=keystore.jks
```

2.4. CREATING THE SECRETS FOR THE AMQ BROKER CONNECTION

If you want to connect any KIE Server to an AMQ broker and to use SSL for the AMQ broker connection, you must create an SSL certificate for the connection and provide it to your OpenShift environment as a secret.

Procedure

1. Generate an SSL keystore named **keystore.jks** with a private and public key for SSL encryption for KIE Server. For more information about creating keystores and using certificates, see [How to Configure Server Security](#).



NOTE

In a production environment, generate a valid signed certificate that matches the expected URL for the AMQ broker connection.

2. Record the name of the certificate. The default value for this name in Red Hat Decision Manager configuration is **jboss**.
3. Record the password of the keystore file. The default value for this name in Red Hat Decision Manager configuration is **mykeystorepass**.
4. Use the **oc** command to generate a secret named **broker-app-secret** from the new keystore file:

```
$ oc create secret generic broker-app-secret --from-file=keystore.jks
```

2.5. PREPARING GIT HOOKS

In an authoring environment you can use Git hooks to execute custom operations when the source code of a project in Business Central is changed. The typical use of Git hooks is for interaction with an upstream repository.

To enable Git hooks to interact with an upstream repository using SSH authentication, you must also provide a secret key and a known hosts file for authentication with the repository.

Skip this procedure if you do not want to configure Git hooks.

Procedure

1. Create the Git hooks files. For instructions, see the [Git hooks reference documentation](#).



NOTE

A **pre-commit** script is not supported in Business Central. Use a **post-commit** script.

2. Create a configuration map (ConfigMap) or persistent volume with the files. For more information about ConfigMaps, see [KIE configuration and ConfigMaps](#).
 - If the Git hooks consist of one or several fixed script files, use the **oc** command to create a configuration map. For example:


```
oc create configmap git-hooks --from-file=post-commit=post-commit
```
 - If the Git hooks consist of long files or depend on binaries, such as executable or JAR files, use a persistent volume. You must create a persistent volume, create a persistent volume claim and associate the volume with the claim, and transfer files to the volume. For instructions about persistent volumes and persistent volume claims, see [Storage](#) in the Red Hat OpenShift Container Platform documentation. For instructions about copying files onto a persistent volume, see [Transferring files in and out of containers](#).
3. If the Git hooks scripts must interact with an upstream repository using SSH authentication, prepare a secret with the necessary files:
 - a. Prepare the **id_rsa** file with a private key that matches a public key stored in the repository.
 - b. Prepare the **known_hosts** file with the correct name, address, and public key for the repository.
 - c. Create a secret with the two files using the **oc** command, for example:

```
oc create secret git-hooks-secret --from-file=id_rsa=id_rsa --from-file=known_hosts=known_hosts
```



NOTE

When the deployment uses this secret, it mounts the **id_rsa** and **known_hosts** files into the **/home/jboss/.ssh** directory on Business Central pods.

2.6. PROVISIONING PERSISTENT VOLUMES WITH READWRITEMANY ACCESS MODE USING NFS

If you want to deploy high-availability Business Central, your environment must provision persistent volumes with **ReadWriteMany** access mode. If you want to deploy high-availability Business Central, your environment must provision persistent volumes with **ReadWriteMany** access mode.

If your configuration requires provisioning persistent volumes with **ReadWriteMany** access mode but your environment does not support such provisioning, use NFS to provision the volumes. Otherwise, skip this procedure.

Procedure

Deploy an NFS server and provision the persistent volumes using NFS. For information about provisioning persistent volumes using NFS, see the "Persistent storage using NFS" section of the [OpenShift Container Platform Storage](#) guide.

2.7. EXTRACTING THE SOURCE CODE FROM BUSINESS CENTRAL FOR USE IN AN S2I BUILD

If you are planning to create immutable KIE servers using the source-to-image (S2I) process, you must provide the source code for your services in a Git repository. If you are using Business Central for authoring services, you can extract the source code for your service and place it into a separate Git repository, such as GitHub or an on-premise installation of GitLab, for use in the S2I build.

Skip this procedure if you are not planning to use the S2I process or if you are not using Business Central for authoring services.

Procedure

1. Use the following command to extract the source code:

```
git clone https://<business-central-host>:443/git/<MySpace>/<MyProject>
```

In this command, replace the following variables:

- **<business-central-host>** with the host on which Business Central is running
- **<MySpace>** with the name of the Business Central space in which the project is located
- **<MyProject>** with the name of the project



NOTE

To view the full Git URL for a project in Business Central, click **Menu** → **Design** → **<MyProject>** → **Settings**.



NOTE

If you are using self-signed certificates for HTTPS communication, the command might fail with an **SSL certificate problem** error message. In this case, disable SSL certificate verification in **git**, for example, using the **GIT_SSL_NO_VERIFY** environment variable:

```
env GIT_SSL_NO_VERIFY=true git clone https://<business-central-host>:443/git/<MySpace>/<MyProject>
```

2. Upload the source code to another Git repository, such as GitHub or GitLab, for the S2I build.

2.8. PREPARING FOR DEPLOYMENT IN A RESTRICTED NETWORK

You can deploy Red Hat Decision Manager in a restricted network that is not connected to the public Internet. For instructions about operator deployment in a restricted network, see [Using Operator Lifecycle Manager on restricted networks](#) in Red Hat OpenShift Container Platform documentation.



IMPORTANT

In Red Hat Decision Manager 7.13, deployment on restricted networks is for Technology Preview only. For more information on Red Hat Technology Preview features, see [Technology Preview Features Scope](#).

In order to use a deployment that does not have outgoing access to the public Internet, you must also prepare a Maven repository with a mirror of all the necessary artifacts. For instructions about creating this repository, see [Section 2.9, "Preparing a Maven mirror repository for offline use"](#) .

2.9. PREPARING A MAVEN MIRROR REPOSITORY FOR OFFLINE USE

If your Red Hat OpenShift Container Platform environment does not have outgoing access to the public Internet, you must prepare a Maven repository with a mirror of all the necessary artifacts and make this repository available to your environment.



NOTE

You do not need to complete this procedure if your Red Hat OpenShift Container Platform environment is connected to the Internet.

Prerequisites

- A computer that has outgoing access to the public Internet is available.

Procedure

1. Configure a Maven release repository to which you have write access. The repository must allow read access without authentication and your OpenShift environment must have network access to this repository.

You can deploy a Nexus repository manager in the OpenShift environment. For instructions about setting up Nexus on OpenShift, see [Setting up Nexus](#) in the Red Hat OpenShift Container Platform 3.11 documentation. The documented procedure is applicable to Red Hat OpenShift Container Platform 4.

Use this repository as a mirror to host the publicly available Maven artifacts. You can also provide your own services in this repository in order to deploy these services on immutable servers.

2. On the computer that has an outgoing connection to the public Internet, complete the following steps:
3. Navigate to the [Software Downloads](#) page in the Red Hat Customer Portal (login required), and select the product and version from the drop-down options:

- **Product:** Process Automation Manager
- **Version:** 7.13.5
 - a. Download and extract the **Red Hat Process Automation Manager 7.13.5 Offliner Content List (rhcam-7.13.5-offliner.zip)** product deliverable file.
 - b. Extract the contents of the **rhcam-7.13.5-offliner.zip** file into any directory.
 - c. Change to the directory and enter the following command:

```
./offline-repo-builder.sh offliner.txt
```

This command creates the **repository** subdirectory and downloads the necessary artifacts into this subdirectory. This is the mirror repository.

If a message reports that some downloads have failed, run the same command again. If downloads fail again, contact Red Hat support.

- d. Upload all artifacts from the **repository** subdirectory to the Maven mirror repository that you prepared. You can use the Maven Repository Provisioner utility, available from the [Maven repository tools](#) Git repository, to upload the artifacts.
4. If you developed services outside of Business Central and they have additional dependencies, add the dependencies to the mirror repository. If you developed the services as Maven projects, you can use the following steps to prepare these dependencies automatically. Complete the steps on the computer that has an outgoing connection to the public Internet.
 - a. Create a backup of the local Maven cache directory (`~/.m2/repository`) and then clear the directory.
 - b. Build the source of your projects using the **mvn clean install** command.
 - c. For every project, enter the following command to ensure that Maven downloads all runtime dependencies for all the artifacts generated by the project:

```
mvn -e -DskipTests dependency:go-offline -f /path/to/project/pom.xml --batch-mode -Djava.net.preferIPv4Stack=true
```

Replace `/path/to/project/pom.xml` with the path of the **pom.xml** file of the project.

- d. Upload all artifacts from the local Maven cache directory (`~/.m2/repository`) to the Maven mirror repository that you prepared. You can use the Maven Repository Provisioner utility, available from the [Maven repository tools](#) Git repository, to upload the artifacts.

CHAPTER 3. DEPLOYMENT AND MANAGEMENT OF A RED HAT DECISION MANAGER ENVIRONMENT USING OPENSIFT OPERATORS

To deploy a Red Hat Decision Manager environment, the OpenShift operator uses a YAML source that describes the environment. Red Hat Decision Manager provides an installer that you can use to form the YAML source and deploy the environment.

When the Business Automation operator deploys the environment, it creates a YAML description of the environment, and then ensures that the environment is consistent with the description at all times. You can edit the description to modify the environment.

You can remove the environment by deleting the operator application in Red Hat OpenShift Container Platform.



NOTE

When you remove an environment with a high-availability Business Central, the operator does not delete Persistent Volume Claims that were created as part of the JBoss Datagrid and JBoss AMQ StatefulSet creation. This behaviour is a part of Kubernetes design, as deletion of the Persistent Volume Claims could cause data loss. For more information about handling persistent volumes during deletion of a StatefulSet, see the [Kubernetes documentation](#).

If you create a new environment using the same namespace and the same application name, the environment reuses the persistent volumes for increased performance.

To ensure that new deployments do not use any old data, you can delete the Persistent Volume Claims manually.

3.1. SUBSCRIBING TO THE BUSINESS AUTOMATION OPERATOR

To be able to deploy Red Hat Decision Manager using operators, you must subscribe to the Business Automation operator in OpenShift.

Procedure

1. Enter your project in the OpenShift Web cluster console.
2. In the OpenShift Web console navigation panel, select **Catalog → OperatorHub** or **Operators → OperatorHub**.
3. Search for **Business Automation**, select it and click **Install**.
4. On the **Create Operator Subscription** page, select your target namespace and approval strategy.
Optional: Set **Approval strategy** to **Automatic** to enable automatic operator updates. An operator update does not immediately update the product, but is required before you update the product. Configure automatic or manual product updates using the settings in every particular product deployment.
5. Click **Subscribe** to create a subscription.

3.2. DEPLOYING A RED HAT DECISION MANAGER ENVIRONMENT USING THE OPERATOR

After you subscribe to the Business Automation operator, you can use the installer wizard to configure and deploy a Red Hat Decision Manager environment.



IMPORTANT

In Red Hat Decision Manager 7.13, the operator installer wizard is for Technology Preview only. For more information on Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

3.2.1. Starting the deployment of a Red Hat Decision Manager environment using the Business Automation operator

To start deploying a Red Hat Decision Manager environment using the Business Automation operator, access the installer wizard. The installer wizard is deployed when you subscribe to the operator.

Prerequisites

- You subscribed to the Business Automation operator. For instructions about subscribing to the operator, see [Section 3.1, "Subscribing to the Business Automation operator"](#).

Procedure

1. In the Red Hat OpenShift Container Platform web cluster console menu, select **Catalog** → **Installed operators** or **Operators** → **Installed operators**.
2. Click the name of the operator that contains **businessautomation**. Information about this operator is displayed.
3. Click the **Installer** link located on the right side of the window.
4. If prompted, log in with your OpenShift credentials.

Result

The **Installation** tab of the wizard is displayed.

3.2.2. Setting the basic configuration of the environment

After you start to deploy a Red Hat Decision Manager environment using the Business Automation operator, you must select the type of the environment and set other basic configuration.

Prerequisites

- You started to deploy a Red Hat Decision Manager environment using the Business Automation operator and accessed the installer wizard according to the instructions in [Section 3.2.1, "Starting the deployment of a Red Hat Decision Manager environment using the Business Automation operator"](#).

Procedure

1. In the **Application Name** field, enter a name for the OpenShift application. This name is used in the default URLs for all components.
2. In the **Environment** list, select the type of environment. This type determines the default configuration; you can modify this configuration as necessary. The following types are available for Red Hat Decision Manager:
 - **rhdm-trial**: A trial environment that you can set up quickly and use to evaluate or demonstrate developing and running assets. Includes Business Central and a KIE Server. This environment does not use any persistent storage, and any work you do in the environment is not saved.
 - **rhdm-authoring**: An environment for creating and modifying services using Business Central. It consists of pods that provide Business Central for the authoring work and a KIE Server for test execution of the services. You can also use this environment to run services for staging and production purposes. You can add KIE Servers to the environment and they are managed by the same Business Central.
 - **rhdm-authoring-ha**: An environment for creating and modifying services using Business Central. It consists of pods that provide Business Central for the authoring work and a KIE Server for test execution of the services. This version of the authoring environment supports scaling the Business Central pod to ensure high availability.



IMPORTANT

In Red Hat Decision Manager 7.13, high-availability Business Central functionality is for Technology Preview only. For more information about Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

- **rhpm-production-immutable**: An alternate environment for running existing services for staging and production purposes. You can configure one or more KIE Server pods that build services from source or pull them from a Maven repository. You can then replicate each pod as necessary.
You cannot remove any service from the pod or add any new service to the pod. If you want to use another version of a service or to modify the configuration in any other way, deploy a new server image to replace the old one. You can use any container-based integration workflows to manage the pods.

When configuring this environment, in the **KIE Servers** tab you must customize KIE Server and either click the **Set immutable server configuration** button or set the **KIE_SERVER_CONTAINER_DEPLOYMENT** environment variable. For instructions about configuring KIE Server, see [Section 3.2.5, "Setting custom KIE Server configuration of the environment"](#).

3. If you want to enable automatic upgrades to new versions, select the **Enable Upgrades** box. If this box is selected, when a new patch version of Red Hat Decision Manager 7.13 becomes available, the operator automatically upgrades your deployment to this version. All services are preserved and normally remain available throughout the upgrade process.
If you also want to enable the same automatic upgrade process when a new minor version of Red Hat Decision Manager 7.x becomes available, select the **Include minor version upgrades** box.



NOTE

Disable automatic updates if you want to use a custom image for any component of Red Hat Decision Manager.

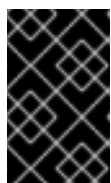
4. If you want to use image tags for downloading images, select the **Use Image Tags** box. This setting is useful if you use a custom registry or if you are directed by Red Hat support.
5. If you want to disable SSL connections to your deployment, select the **Disable SSL Routes** box. In this case, all routes that are externally exposed use clear-text (HTTP) connections.



NOTE

If this box is not selected, only secure (HTTPS) routes are exposed externally.

6. If you want to use a custom image registry, under **Custom registry**, enter the URL of the registry in the **Image registry** field. If this registry does not have a properly signed and recognized SSL certificate, select the **Insecure** box.
For instructions about configuring the image registry to use specific images, see [Section 3.2.2.1, "Configuring the image registry to use specific images"](#).
7. Under **Admin user**, enter the user name and password for the administrative user for Red Hat Decision Manager in the **Username** and **Password** fields.



IMPORTANT

If you use RH-SSO or LDAP authentication, the same user must be configured in your authentication system with the **kie-server,rest-all,admin** roles for Red Hat Decision Manager.

8. Optional: Select the startup strategy. The **OpenShiftStartupStrategy** setting is enabled by default.
In some authoring environments, you might need to ensure that several users can deploy services on the same KIE Server at the same time. By default, after deploying a service onto a KIE Server using Business Central, the user must wait a few seconds before more services can be deployed. The **OpenShiftStartupStrategy** setting is enabled by default and causes this limitation. To remove the limitation, select the **ControllerBasedStartupStrategy** setting from the **Startup Strategy** list.



NOTE

Do not enable the controller-based startup strategy in an environment with a high-availability Business Central.

9. Optional: If you want to use the OpenShift CA bundle as the trust store for HTTPS communication, select the **Use OpenShift CA Bundle** box.
10. Optional: If you want to use a secret that contains the credentials for the Admin user, complete the following tasks:
 - a. From the **Admin user configuration** list, select **Secret configuration**.

- b. Under **OpenShift admin credentials secret**, in the **Secret** field, enter the name of the secret. If the **Secret** field is left blank, the default secret from **kie-admin-credentials** is used.
- c. Under **OpenShift admin credentials secret**, in the **Username** field, enter the username for the admin user to use in the secret. If the **Username** field is left blank, the default username from **kie-admin-credentials** is used.
- d. Under **OpenShift admin credentials secret**, in the **Password** field, enter the password for the admin user to use in the secret. If the **Password** field is left blank, the default password from **kie-admin-credentials** is used.



NOTE

If **kie-admin-credentials** is missing, **kie-admin-credentials** is generated with a default username and password.

Next steps

If you want to deploy the environment with the default configuration, click **Finish** and then click **Deploy** to deploy the environment. Otherwise, continue to set other configuration parameters.

3.2.2.1. Configuring the image registry to use specific images

During the basic configuration of the environment you can configure the custom registry to use specific images. For more information about configuring the basic environment, see [Section 3.2.2, "Setting the basic configuration of the environment"](#).

Prerequisites

- You started to configure a Red Hat Decision Manager environment using the Business Automation operator and accessed the installer wizard according to the instructions in [Section 3.2.2, "Setting the basic configuration of the environment"](#).
- You have set the URL of the image registry in the **Image registry** field under the **Custom registry** of the Red Hat Decision Manager environment configuration according to the instructions in [Section 3.2.2, "Setting the basic configuration of the environment"](#).

Procedure

- During the Red Hat Decision Manager environment configuration, to use a specific image from the custom registry, complete one of the following steps:
 - If you want to specify the image using the installer wizard, set the image context, image name, and image tag parameters in the **Console** and **KIE Server** tabs:

Table 3.1. Parameters

Name	Description
Image context	The context of the image in the registry.
Image	The name of the image.

Name	Description
Image tag	The tag of the image. If you do not set this field, the installation uses the latest tag.

- If you want to specify the image by using a KieApp CR YAML file, add the image registry and image details to the file, for example:

Example

```
apiVersion: app.kiegroup.org/v2
kind: KieApp
spec:
  ...
  useImageTags: true
  imageRegistry:
    registry: registry.example.com:5000
  ...
  objects:
    ...
    servers:
      - id: ...
        ...
        image: YOUR_IMAGE_NAME
        imageContext: YOUR_IMAGE_CONTEXT
        imageTag: YOUR_IMAGE_TAG
    ...
```



NOTE

If an ImageStream with the configured name exists in the namespace that you are using or in the Red Hat OpenShift Container Platform namespace, the operator uses this ImageStream and does not create a new ImageStream.

To configure an ImageStream to update automatically, the ImageStream property **scheduledImportPolicy** must be set to **true** in the KieApp CR YAML file. For example:

Example of **scheduledImportPolicy**

```
apiVersion: app.kiegroup.org/v2
kind: KieApp
spec:
  ...
  useImageTags: true
  scheduledImportPolicy: true
  imageRegistry:
    registry: registry.example.com:5000
  ...
  objects:
    ...
  servers:
    - id: ...
      ...
      image: YOUR_IMAGE_NAME
      imageContext: YOUR_IMAGE_CONTEXT
      imageTag: YOUR_IMAGE_TAG
    ...
```

3.2.3. Setting the security configuration of the environment

After you set the basic configuration of a Red Hat Decision Manager environment using the Business Automation operator, you can optionally configure authentication (security) settings for the environment.

Prerequisites

- You completed basic configuration of a Red Hat Decision Manager environment using the Business Automation operator in the installer wizard according to the instructions in [Section 3.2.2, "Setting the basic configuration of the environment"](#).
- If you want to use RH-SSO or LDAP for authentication, you created users with the correct roles in your authentication system. You must create at least one administrative user (for example, **adminUser**) with the **kie-server,rest-all,admin** roles. This user must have the user name and password that you configured on the **Installation** tab.
- If you want to use RH-SSO authentication, you created the clients in your RH-SSO system for all components of your environment, specifying the correct URLs. This action ensures maximum control. Alternatively, the deployment can create the clients.

Procedure

1. If the **Installation** tab is open, click **Next** to view the **Security** tab.

2. In the **Authentication mode** list, select one of the following modes:
 - **Internal:** You configure the initial administration user when deploying the environment. You can create a post-configuration script to add users in the Elytron security subsystem. For instructions about creating a post-configuration script, see [Section 3.4, "Providing Elytron user configuration or other post-configuration settings"](#).
 - **RH-SSO:** Red Hat Decision Manager uses Red Hat Single Sign-On for authentication.
 - **LDAP:** Red Hat Decision Manager uses LDAP for authentication
3. Complete the security configuration based on the **Authentication mode** that you selected:
 - If you selected **RH-SSO**, configure RH-SSO authentication:
 - a. In the **RH-SSO URL** field, enter the RH-SSO URL.
 - b. In the **Realm** field, enter the RH-SSO realm name.
 - c. If you did not create RH-SSO clients for components of your environment enter the credentials of an administrative user for your RH-SSO system in the **SSO admin user** and **SSO admin password** fields.
 - d. If your RH-SSO system does not have a proper signed SSL certificate, select the **Disable SSL cert validation** box.
 - e. If you want to change the RH-SSO principal attribute used for the user name, in the **Principal attribute** field enter the name of the new attribute.

If you selected **LDAP**, configure LDAP authentication:



IMPORTANT

The **baseFilter** field is a legacy LDAP search filter used to locate the context of the user to authenticate. The input **username** or **userDN** obtained from the login module callback is substituted into the filter anywhere a **{0}** expression is used. A common example for the search filter is **(uid={0})**. For Elytron based subsystems, this property should be configured only with the search filter parameter, without any search expression. For example, search for **uid** instead of **(uid={0})**.

- a. In the **LDAP URL** field, enter the LDAP URL.
- b. Set LDAP parameters. These parameters configure LDAP authentication using the Elytron subsystem of Red Hat JBoss EAP. For more information about using the Elytron subsystem of Red Hat JBoss EAP with LDAP, see [Configure Authentication with an LDAP-Based Identity Store](#).



NOTE

If you want to enable LDAP failover, you can set two or more LDAP server addresses in the **AUTH_LDAP_URL** parameter, separated by a space.

4. If you selected **RH-SSO** or **LDAP**, if your RH-SSO or LDAP system does not define all the roles required for your deployment, you can map authentication system roles to Red Hat Decision Manager roles.

To enable role mapping, you must provide role mapping either as a single configuration string or as a role mapping configuration file. If you use a file for role mapping configuration, you must provide the file in an OpenShift configuration map or secret object in the project namespace.

The string must use the **role=role1,role2;another-role=role2** pattern, for example **admins=kie-server,rest-all,admin;developers=kie-server,rest-all**.

The file must contain entries in the following format:

```
ldap_role=product_role1, product_role2...
```

For example:

```
admins=kie-server,rest-all,admin
```

To enable the use of this string or file, make the following changes:

- a. Under **RoleMapper**, in the **Roles properties file** field, enter the role configuration string or the fully qualified path name of the role mapping configuration file, for example **/opt/eap/standalone/configuration/rolemapping/rolemapping.properties**.
 - b. Optional: Select the **Roles keep mapped** box or the **Roles keep non mapped** box. If you define role mapping, by default only the roles that you define in the mapping are available. If you want to keep the original roles that are defined in the authentication system and that your mapping maps to other roles, select the **Roles keep mapped** box. If you want to keep the original roles that are defined in the authentication system and not mentioned in your mapping, select the **Roles keep non mapped** box.
 - c. If you are using a role configuration file, configure the fields under **RoleMapper Configuration object**:
 - Under the **Kind** label, select the kind of the object that provides the file (**ConfigMap** or **Secret**).
 - In the **Name** field, enter the name of the object. This object is automatically mounted on Business Central and KIE Server pods in the path that you specified for the role mapping configuration file.
5. Configure other passwords, if necessary:
- **AMQ password** and **AMQ cluster password** are passwords for interaction with ActiveMQ using the JMS API.
 - **Keystore password** is the password for the keystore files used in secrets for HTTPS communication. Set this password if you created secrets according to instructions in [Section 2.2, "Creating the secrets for KIE Server"](#) or [Section 2.3, "Creating the secrets for Business Central"](#).
 - **Database password** is the password for database server pods that are a part of the environments.

Next steps

If you want to deploy the environment with the default configuration of all components, click **Finish** and then click **Deploy** to deploy the environment. Otherwise, continue to set configuration parameters for Business Central and KIE Servers.

3.2.4. Setting the Business Central configuration of the environment

After you set the basic and security configuration of a Red Hat Decision Manager environment using the Business Automation operator, you can optionally configure settings for the Business Central component of the environment.

All environment types except **rhpm-production-immutable** include this component.

Do not change these settings for the **rhpm-production-immutable** environment, as this environment does not include Business Central or Business Central Monitoring.

Prerequisites

- You completed basic configuration of a Red Hat Decision Manager environment using the Business Automation operator in the installer wizard according to the instructions in [Section 3.2.2, "Setting the basic configuration of the environment"](#).
- If you want to use RH-SSO or LDAP for authentication, you completed security configuration according to the instructions in [Section 3.2.3, "Setting the security configuration of the environment"](#).

Procedure

1. If the **Installation** or **Security** tab is open, click **Next** until you view the **Console** tab.
2. If you created the secret for Business Central according to the instructions in [Section 2.3, "Creating the secrets for Business Central"](#), enter the name of the secret in the **Keystore secret** field.
3. Optional: If you want to use a custom image for the Business Central deployment, complete the following additional steps:
 - a. Set the custom registry in the **Installation** tab. If you do not set the custom registry, the installation uses the default Red Hat registry. For more information about setting the custom registry value, see [Section 3.2.2, "Setting the basic configuration of the environment"](#).
 - b. In the **Console** tab, set the following fields:
 - **Image context:** The context of the image in the registry.
 - **Image:** The name of the image.
 - **Image tag:** The tag of the image. If you do not set this field, the installation uses the **latest** tag.
For example, if the full address of the image is **registry.example.com/mycontext/mycentral:1.0-SNAPSHOT**, set the custom registry to **registry.example.com**, the **Image context** field to **mycontext**, the **Image** field to **mycentral**, and the **Image tag** field to **1.0-SNAPSHOT**.
4. Optional: To set a custom hostname for the external route, enter a domain in the **Custom hostname to be used on the Business Central external Route** field, formatted as in the following example:

```
`businesscentral.example.com`
```

**NOTE**

The custom hostname must be valid and resolvable.

To change the custom hostname, you can modify the **routeHostname** property.

5. Optional: To enable and set the Edge termination route, complete the following steps:
 - a. Under **Change route termination**, select **Enable Edge termination**.
 - b. Optional: In the **Key** field, enter the private key.
 - c. Optional: In the **Certificate** field, enter the certificate.
 - d. Optional: In the **CaCertificate** field, enter the CaCertificate.
6. Optional: Configure Git hooks.
 In an authoring environment, you can use Git hooks to facilitate interaction between the internal Git repository of Business Central and an external Git repository. If you want to use Git hooks, you must prepare a Git hooks directory in an OpenShift configuration map, secret, or persistent volume claim object in the project namespace. You can also prepare a secret with the SSH key and known hosts files for Git SSH authentication. For instructions about preparing Git hooks, see [Section 2.5, "Preparing Git hooks"](#).

To use a Git hooks directory, make the following changes:

- a. Under **GitHooks**, in the **Mount path** field, enter a fully qualified path for the directory, for example, **/opt/kie/data/git/hooks**.
- b. In the fields under **GitHooks Configuration object**, select the **Kind** of the object that provides the file (**ConfigMap**, **Secret**, or **PersistentVolumeClaim**) and enter the **Name** of the object. This object is automatically mounted on the Business Central pods in the path that you specified for the Git hooks directory.
7. Optional: In the **SSH secret** field enter the name of the secret with the SSH key and known hosts files.
8. Optional: To configure KIE Server for decision management only capabilities so that jBPM and case management features are disabled, select **Execute Kie Server only with Decisions capabilities**.
9. Optional: Enter the number of replicas for Business Central or Business Central monitoring in the **Replicas** field. Do not change this number in a **rhpm-authoring** environment.
10. Optional: To set the Business Central persistent volume size **pvSize**, on the **Console component** page, enter the desired size in the **Persistent Volume Size** field. The default size is 1Gi for Business Central and 64Mb for Business Central Monitoring.
11. Optional: Enter requested and maximum CPU and memory limits in the fields under **Resource quotas**.
12. If you want to customize the configuration of the Java virtual machine on the Business Central pods, select the **Enable JVM configuration** box and then enter information in any of the fields under **Enable JVM configuration**. All fields are optional. For the JVM parameters that you can configure, see [Section 3.5, "JVM configuration parameters"](#).
13. If you selected RH-SSO authentication, configure RH-SSO for Business Central:

- a. Enter the client name in the **Client name** field and the client secret in the **Client secret** field. If a client with this name does not exist, the deployment attempts to create a new client with this name and secret.
 - b. If the deployment is to create a new client, enter the HTTP and HTTPS URLs that will be used for accessing Business Central into the **SSO HTTP URL** and **SSO HTTPS URL** fields. This information is recorded in the client.
14. Optional: If you are configuring a high-availability environment, set the user name and password for the DataGrid component in the **DataGrid username** and **DataGrid password** fields. By default, the user name is **infinispan** and the password is generated automatically.
 15. Optional: Depending on your needs, set environment variables. To set an environment variable, click **Add new Environment variable**, then enter the name and value for the variable in the **Name** and **Value** fields.
 - Optional: If you want to configure the proxy settings, use the following environment variables:
 - **https_proxy**: The location of the https proxy. This takes precedence over **HTTPS_PROXY**, **http_proxy**, and **HTTP_PROXY**, and is used for both Maven builds and Java runtime. For example: **myuser:mypass@127.0.0.1:8080**.
 - **HTTPS_PROXY**: The location of the https proxy. This takes precedence over **http_proxy** and **HTTP_PROXY**, and is used for both Maven builds and Java runtime. For example: **myuser@127.0.0.1:8080**.
 - **http_proxy**: The location of the http proxy. This takes precedence over **HTTP_PROXY** and is used for both Maven builds and Java runtime. For example: **http://127.0.0.1:8080**.
 - **HTTP_PROXY**: The location of the http proxy. This is used for both Maven builds and Java runtime. For example: **127.0.0.1:8080**.
 - **no_proxy**: A comma separated lists of hosts, IP addresses, or domains that can be accessed directly. This takes precedence over **NO_PROXY** and is used for both Maven builds and Java runtime. For example: ***.example.com**.
 - **NO_PROXY**: A comma separated lists of hosts, IP addresses, or domains that can be accessed directly. This is used for both Maven builds and Java runtime. For example: **foo.example.com,bar.example.com**.
 - If you want to use an external Maven repository, set the following variables:
 - **MAVEN_REPO_URL**: The URL for the Maven repository
 - **MAVEN_REPO_ID**: An identifier for the Maven repository, for example, **repo-custom**
 - **MAVEN_REPO_USERNAME**: The user name for the Maven repository
 - **MAVEN_REPO_PASSWORD** The password for the Maven repository



IMPORTANT

In an authoring environment, if you want Business Central to push a project into an external Maven repository, you must configure this repository during deployment and also configure exporting to the repository in every project. For information about exporting Business Central projects to an external Maven repository, see [Packaging and deploying an Red Hat Decision Manager project](#).

- If your OpenShift environment does not have a connection to the public Internet, configure access to a Maven mirror that you set up according to [Section 2.9, "Preparing a Maven mirror repository for offline use"](#). Set the following variables:

- **MAVEN_MIRROR_URL**: The URL for the Maven mirror repository that you set up in [Section 2.9, "Preparing a Maven mirror repository for offline use"](#). This URL must be accessible from a pod in your OpenShift environment.
- **MAVEN_MIRROR_OF**: The value that determines which artifacts are to be retrieved from the mirror. For instructions about setting the **mirrorOf** value, see [Mirror Settings](#) in the Apache Maven documentation. The default value is **external:***. With this value, Maven retrieves every required artifact from the mirror and does not query any other repositories.

If you configure an external Maven repository (**MAVEN_REPO_URL**), change **MAVEN_MIRROR_OF** to exclude the artifacts in this repository from the mirror, for example, **external:*,!repo-custom**. Replace **repo-custom** with the ID that you configured in **MAVEN_REPO_ID**.

If your authoring environment uses a built-in Business Central Maven repository, change **MAVEN_MIRROR_OF** to exclude the artifacts in this repository from the mirror: **external:*,!repo-rhpamcentr**.

- In some cases, you might want to persist the Maven repository cache for Business Central. By default, the cache is not persisted, so when you restart or scale a Business Central pod, all Maven artifacts are downloaded again and all projects within Business Central must be built again. If you enable persistence for the cache, the download is not necessary and startup time can improve in some situations. However, significant additional space on the Business Central persistence volume is required.

To enable persistence for the Maven repository cache, set the **KIE_PERSIST_MAVEN_REPO** environment variable to **true**.

If you set **KIE_PERSIST_MAVEN_REPO** to **true**, you can optionally set a custom path for the cache using the **KIE_M2_REPO_DIR** variable. The default path is **/opt/kie/data/m2**. Files in the **/opt/kie/data** directory tree are persisted.

Next steps

If you want to deploy the environment with the default configuration of KIE Servers, click **Finish** and then click **Deploy** to deploy the environment. Otherwise, continue to set configuration parameters for KIE Servers.

3.2.5. Setting custom KIE Server configuration of the environment

Every environment type in the Business Automation operator includes one or several KIE Servers by default.

Optionally, you can set custom configuration for KIE Servers. In this case, default KIE Servers are not created and only the KIE Servers that you configure are deployed.

Prerequisites

- You completed basic configuration of a Red Hat Decision Manager environment using the Business Automation operator in the installer wizard according to the instructions in [Section 3.2.2, "Setting the basic configuration of the environment"](#).

Procedure

- If the **Installation**, **Security**, or **Console** tab is open, click **Next** until you view the **KIE Servers** tab.
- Click **Add new KIE Server** to add a new KIE Server configuration.
- In the **Id** field, enter an identifier for this KIE Server instance. If the KIE Server instance connects to a Business Central or Business Central Monitoring instance, this identifier determines which server group the server joins.
- In the **Name** field, enter a name for the KIE Server.
- In the **Deployments** field, enter the number of similar KIE Servers that are to be deployed. The installer can deploy several KIE Servers with the same configuration. The identifiers and names of the KIE Servers are modified automatically and remain unique.
- If you created the secret for KIE Server according to the instructions in [Section 2.2, "Creating the secrets for KIE Server"](#), enter the name of the secret in the **Keystore secret** field.
- Optional: To configure KIE Server for decision management only capabilities so that jBPM and case management features are disabled, select **Execute Kie Server only with Decisions capabilities**.
- Optional: Enter the number of replicas for the KIE Server deployment in the **Replicas** field.
- Optional: To set a custom hostname for the external route, enter a domain in the **Custom hostname to be used on the KIE Server external Route** field, formatted as in the following example:

```
`kieserver.example.com`
```



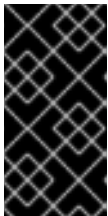
NOTE

The custom hostname must be valid and resolvable.

To change the custom hostname, you can modify the **routeHostname** property.

- Optional: To enable and set the Edge termination route, complete the following steps:
 - Under **Change route termination**, select **Enable Edge termination**.
 - Optional: In the **Key** field, enter the private key.
 - Optional: In the **Certificate** field, enter the certificate.
 - Optional: In the **CaCertificate** field, enter the CaCertificate.

11. Optional: If you want to use a custom KIE Server image, complete the following additional steps:
 - a. Click **Set KIE Server image**
 - b. From the **Kind** list, select **ImageStreamTag** if you want to pull the image from an OpenShift image stream or select **DockerImage** if you want to pull the image from any Docker registry.
 - c. Set the image name by completing one of the following steps:
 - If you selected the **ImageStreamTag** kind, enter the image stream tag name of the image in the **Name** field, for example, **my-custom-is-tag:1.0**.
If a corresponding image stream does not exist in your environment, the operator creates this image stream using the default Red Hat registry and tags. If you configured a custom registry in the **Installation** tab, the operator uses this registry for creating the image stream.
 - If you selected the **DockerImage** kind, enter the fully qualified image name for the image in the **Name** field, for example, **registry.io/test/testing:1.0**.
You can configure an image from any registry that your environment can access.
 - d. If you want to use an image stream that is not in the **openshift** namespace, enter the namespace in the **Namespace** field.
For instructions about creating custom images, see [Section 3.7, "Creating custom images for KIE Server"](#).
12. Optional: If you want to configure an immutable KIE Server using a Source to Image (S2I) build, complete the following additional steps:



IMPORTANT

If you want to configure an immutable KIE Server that pulls services from the Maven repository, do not click **Set Immutable server configuration** and do not complete these steps. Instead, set the **KIE_SERVER_CONTAINER_DEPLOYMENT** environment variable.

- a. Click **Set Immutable server configuration**
- b. In the **KIE Server container deployment** field, enter the identifying information of the services (KJAR files) that the deployment must extract from the result of a Source to Image (S2I) build. The format is **<containerId>=<groupId>:<artifactId>:<version>** or, if you want to specify an alias name for the container, **<containerId>(<aliasId>)=<groupId>:<artifactId>:<version>**. You can provide two or more KJAR files using the **|** separator, as illustrated in the following example:
containerId=groupId:artifactId:version|c2(alias2)=g2:a2:v2.
- c. If your OpenShift environment does not have a connection to the public Internet, enter the URL of the Maven mirror that you set up according to [Section 2.9, "Preparing a Maven mirror repository for offline use"](#) in the **Maven mirror URL** field.
- d. In the **Artifact directory** field, enter the path within the project that contains the required binary files (KJAR files and any other necessary files) after a successful Maven build. Normally this directory is the target directory of the build. However, you can provide prebuilt binaries in this directory in the Git repository.
- e. If you want to use a custom base KIE Server image for the S2I build, click **Set Base build**

image and then enter the name of the image stream in the **Name** field. If the image stream is not in the **openshift** namespace, enter the namespace in the **Namespace** field. If you want to use a Docker image name and not an OpenShift image stream tag, change the **Kind** value to **DockerImage**.

f. Click **Set Git source** and enter information in the following fields:

- **S2I Git URI**The URI for the Git repository that contains the source for your services.
- **Reference**: The branch in the Git repository.
- **Context directory**: (Optional) The path to the source within the project downloaded from the Git repository. By default, the root directory of the downloaded project is the source directory.



NOTE

If you do not configure a Git source, the immutable KIE Server does not use an S2I build. Instead, it pulls the artifacts that you define in the **KIE Server container deployment** field from the configured Maven repository.

- g. If you are using S2I and want to set a Git Webhook so that changes in the Git repository cause an automatic rebuild of the KIE Server, click **Add new Webhook**. Then select the type of the Webhook in the **Type** field and enter the secret string for the Webhook in the **Secret** field.
- h. If you want to set a build environment variable for the S2I build, click **Add new Build Config Environment variable** and then enter the name and value for the variable in the **Name** and **Value** fields.
13. Optional: Enter requested and maximum CPU and memory limits in the fields under **Resource quotas**. If you are configuring several KIE Servers, the limits apply to each server separately.
14. If you selected RH-SSO authentication, configure RH-SSO for the KIE Server:
- a. Enter the client name in the **Client name** field and the client secret in the **Client secret** field. If a client with this name does not exist, the deployment attempts to create a new client with this name and secret.
 - b. If the deployment is to create a new client, enter the HTTP and HTTPS URLs that will be used for accessing this KIE Server instance into the **SSO HTTP URL** and **SSO HTTPS URL** fields. This information is recorded in the client.
15. If you want to interact with the KIE Server through JMS API using an external AMQ message broker, enable the **Enable JMS Integration** setting. Additional fields for configuring JMS Integration are displayed and you must enter the values as necessary:
- **User name, Password**: The user name and password of a standard broker user, if user authentication in the broker is required in your environment.
 - **Executor**: Select this setting to disable the JMS executor. The executor is enabled by default.
 - **Executor transacted**: Select this setting to enable JMS transactions on the executor queue.

- **Enable signal:** Select this setting to enable signal configuration through JMS.
 - **Enable audit:** Select this setting to enable audit logging through JMS.
 - **Audit transacted:** Select this setting to enable JMS transactions on the audit queue.
 - **Queue executor, Queue request, Queue response, Queue signal, Queue audit:** Custom JNDI names of the queues to use. If you set any of these values, you must also set the **AMQ queues** parameter.
 - **AMQ Queues:** AMQ queue names, separated by commas. These queues are automatically created when the broker starts and are accessible as JNDI resources in the JBoss EAP server. If you are using any custom queue names, you must enter the names of all the queues used by the server in this field.
 - **Enable SSL integration:** Select this setting if you want to use an SSL connection to the AMQ broker. In this case you must also provide the name of the secret that you created in [Section 2.4, "Creating the secrets for the AMQ broker connection"](#) and the names and passwords of the key store and trust store that you used for the secret.
16. If you want to customize the configuration of the Java virtual machine on the KIE Server pods, select the **Enable JVM configuration** box and then enter information in any of the fields under **Enable JVM configuration**. All fields are optional. For the JVM parameters that you can configure, see [Section 3.5, "JVM configuration parameters"](#).
17. Optional: If you want to configure the proxy settings, use the following environment variables:
- **https_proxy:** The location of the https proxy. This takes precedence over **HTTPS_PROXY**, **http_proxy**, and **HTTP_PROXY**, and is used for both Maven builds and Java runtime. For example: **myuser:mypass@127.0.0.1:8080**.
 - **HTTPS_PROXY:** The location of the https proxy. This takes precedence over **http_proxy** and **HTTP_PROXY**, and is used for both Maven builds and Java runtime. For example: **myuser@127.0.0.1:8080**.
 - **http_proxy:** The location of the http proxy. This takes precedence over **HTTP_PROXY** and is used for both Maven builds and Java runtime. For example: **http://127.0.0.1:8080**.
 - **HTTP_PROXY:** The location of the http proxy. This is used for both Maven builds and Java runtime. For example: **127.0.0.1:8080**.
 - **no_proxy:** A comma separated lists of hosts, IP addresses, or domains that can be accessed directly. This takes precedence over **NO_PROXY** and is used for both Maven builds and Java runtime. For example: ***.example.com**.
 - **NO_PROXY:** A comma separated lists of hosts, IP addresses, or domains that can be accessed directly. This is used for both Maven builds and Java runtime. For example: **foo.example.com,bar.example.com**.
18. Optional: Depending on your needs, set environment variables. To set an environment variable, click **Add new Environment variable**, then enter the name and value for the variable in the **Name** and **Value** fields.
- If you want to configure an immutable KIE server that pulls services from the configured Maven repository, enter the following settings:
 - i. Set the **KIE_SERVER_CONTAINER_DEPLOYMENT** environment variable. The variable must contain the identifying information of the services (KJAR files) that the

deployment must pull from the Maven repository. The format is **<containerId>=<groupId>:<artifactId>:<version>** or, if you want to specify an alias name for the container, **<containerId>(<aliasId>)=<groupId>:<artifactId>:<version>**. You can provide two or more KJAR files using the | separator, as illustrated in the following example: **containerId=groupId:artifactId:version|c2(alias2)=g2:a2:v2**.

- ii. Configure an external Maven repository.
- If you want to configure an external Maven repository, set the following variables:
 - **MAVEN_REPO_URL**: The URL for the Maven repository
 - **MAVEN_REPO_ID**: An identifier for the Maven repository, for example, **repo-custom**
 - **MAVEN_REPO_USERNAME**: The user name for the Maven repository
 - **MAVEN_REPO_PASSWORD**: The password for the Maven repository
 - If your OpenShift environment does not have a connection to the public Internet, configure access to a Maven mirror that you set up according to [Section 2.9, "Preparing a Maven mirror repository for offline use"](#). Set the following variables:
 - **MAVEN_MIRROR_URL**: The URL for the Maven mirror repository that you set up in [Section 2.9, "Preparing a Maven mirror repository for offline use"](#). This URL must be accessible from a pod in your OpenShift environment. If you configured this KIE Server as S2I, you already entered this URL.
 - **MAVEN_MIRROR_OF**: The value that determines which artifacts are to be retrieved from the mirror. If you configured this KIE Server as S2I, do not set this value. For instructions about setting the **mirrorOf** value, see [Mirror Settings](#) in the Apache Maven documentation. The default value is **external:***. With this value, Maven retrieves every required artifact from the mirror and does not query any other repositories. If you configure an external Maven repository (**MAVEN_REPO_URL**), change **MAVEN_MIRROR_OF** to exclude the artifacts in this repository from the mirror, for example, **external:*,!repo-custom**. Replace **repo-custom** with the ID that you configured in **MAVEN_REPO_ID**.

If your authoring environment uses a built-in Business Central Maven repository, change **MAVEN_MIRROR_OF** to exclude the artifacts in this repository from the mirror: **external:*,!repo-rhpamcentr**.
 - If you want to configure your KIE Server deployment to use Prometheus to collect and store metrics, set the **PROMETHEUS_SERVER_EXT_DISABLED** environment variable to **false**. For instructions about configuring Prometheus metrics collection, see [Managing and monitoring KIE Server](#).
 - If you are using Red Hat Single Sign-On authentication and the interaction of your application with Red Hat Single Sign-On requires support for cross-origin resource sharing (CORS), configure **CORS Filters configuration**:
 - To use CORS with the default configuration, ensure **Default configuration** is selected from the **CORS Filters configuration** list and select **Enable CORS with Default values**
 - To use CORS with a custom configuration, select **Custom configuration** from the **CORS Filters configuration** list and enter the relevant values for the CORS filters.

Next steps

To configure additional KIE Servers, click **Add new KIE Server** again and repeat the procedure for the new server configuration.

Click **Finish** and then click **Deploy** to deploy the environment.

3.3. MODIFYING AN ENVIRONMENT THAT IS DEPLOYED USING OPERATORS

If an environment is deployed using operators, you cannot modify it using typical OpenShift methods. For example, if you delete a deployment configuration or a service, it is re-created automatically with the same parameters.

To modify the environment, you must modify the YAML description of the environment. You can change common settings such as passwords, add new KIE Servers, and scale KIE Servers.

Procedure

1. Enter your project in the OpenShift web cluster console.
2. In the OpenShift Web console navigation panel, select **Catalog → Installed operators** or **Operators → Installed operators**.
3. Find the **Business Automation** operator line in the table and click **KieApp** in the line. Information about the environments that you deployed using this operator is displayed.
4. Click the name of a deployed environment.
5. Select the **YAML** tab.
A YAML source is displayed. In this YAML source, you can edit the content under **spec:** to change the configuration of the environment.
6. If you want to change the deployed version of Red Hat Decision Manager, add the following line under **spec:**

```
version: 7.13.5
```

You can replace **7.13.5** with another required version. Use this setting to upgrade Red Hat Decision Manager to a new version if automatic updates are disabled, for example, if you use a custom image.

7. If you want to change common settings, such as passwords, edit the values under **commonConfig:**.
8. If you want to add new KIE Servers, add their descriptions at the end of the block under **servers:**, as shown in the following examples:

- To add two servers named **server-a** and **server-a-2**, add the following lines:

```
- deployments: 2
  name: server-a
```

- To add an immutable KIE Server that includes services built from source in an S2I process, add the following lines:

```
- build:
```

```

kieServerContainerDeployment: <deployment>
gitSource:
  uri: <url>
  reference: <branch>
  contextDir: <directory>
    
```

Replace the following values:

- **<deployment>**: The identifying information of the decision service (KJAR file) that is built from your source. The format is **<containerId>=<groupId>:<artifactId>:<version>**. You can provide two or more KJAR files using the `|` separator, for example **containerId=groupId:artifactId:version|c2=g2:a2:v2**. The Maven build process must produce all these files from the source in the Git repository.
 - **<url>**: The URL for the Git repository that contains the source for your decision service.
 - **<branch>**: The branch in the Git repository.
 - **<directory>**: The path to the source within the project downloaded from the Git repository.
9. If you want to scale a KIE Server, find the description of the server in the block under **servers:** and add a **replicas:** setting under that description. For example, **replicas: 3** scales the server to three pods.
 10. If you want to make other changes, review the CRD source for the available settings. To view the CRD source, log in to the Red Hat OpenShift Container Platform environment with the **oc** command as an administrative user and then enter the following command:

```
oc get crd kieapps.app.kiegroup.org -o yaml
```

11. Click **Save** and then wait for a **has been updated** pop-up message.
12. Click **Reload** to view the new YAML description of the environment.

3.4. PROVIDING ELYTRON USER CONFIGURATION OR OTHER POST-CONFIGURATION SETTINGS

If you do not use LDAP or RH-SSO authentication, Red Hat Decision Manager relies on internal users in the Elytron subsystem of Red Hat JBoss EAP. By default, only the administrative user is created. You might need to add other users to the Elytron security subsystem of Red Hat JBoss EAP. To do so, you must run an Red Hat JBoss EAP post-configuration script.

You can configure this post-configuration script, or any other Red Hat JBoss EAP post-configuration script, in a deployment of Red Hat Decision Manager on Red Hat OpenShift Container Platform.

Procedure

1. Download sample files from the [GitHub repository](#).
2. Prepare the following files based on the sample files:
 - **postconfigure.sh**: The post-configuration shell script that Red Hat JBoss EAP must run. In the example, this script uses the **add-users.cli** script to add Elytron users. If you want to complete post-configuration tasks outside of the CLI script, modify this script.

- **delayedpostconfigure.sh**: An empty file, required in Red Hat Decision Manager version 7.13.5.
- **add-users.cli**: The Red Hat JBoss EAP command line interface script for configuring Elytron users or for any other CLI tasks. Add your commands between the following lines:

```
embed-server --std-out=echo --server-config=standalone-openshift.xml batch
<your jboss-cli commands>
run-batch quit
```

3. Log in to your Red Hat OpenShift Container Platform cluster with the **oc** command and change to the namespace of your deployment.
4. Create a ConfigMap with the files that you prepared by using the following command:

```
oc create configmap postconfigure \
  --from-file=add-users.cli=add-users.cli \
  --from-file=delayedpostconfigure.sh=delayedpostconfigure.sh \
  --from-file=postconfigure.sh=postconfigure.sh
```

5. Enter the following command to edit the **kieconfigs-7.13.5** config map:

```
oc edit cm kieconfigs-7.13.5
```

6. In the file, modify the deployment configuration under the **console**: section to add the configuration to Business Central and modify all deployment configurations under the **servers**: section to add the configuration to KIE Server instances.
In each deployment configuration, make the following changes:

- Under **deploymentConfigs.metadata.spec.template.spec.containers.volumeMounts**, add the following lines:

```
- name: postconfigure-mount
  mountPath: /opt/eap/extensions
```

- Under **deploymentConfigs.metadata.spec.template.spec.containers.volumeMounts**, add the following lines:

```
- name: "postconfigure-mount"
  configMap:
    name: "postconfigure"
    defaultMode: 0555
```

7. Save the file. After this point, new operator deployments contain the post-configuration settings.

In existing deployments, if the the post-configuration settings are not added automatically, you can delete the Business Central and KIE Server pods. The operator automatically starts updated versions with the post-configuration settings.

3.5. JVM CONFIGURATION PARAMETERS

When deploying Red Hat Decision Manager using the operator, you can optionally set a number of JVM configuration parameters for Business Central and KIE Servers. These parameters set environment variables for the corresponding containers.

The following table lists all JVM configuration parameters that you can set when deploying Red Hat Decision Manager using the operator.

The default settings are optimal for most use cases. Make any changes only when they are required.

Table 3.2. JVM configuration parameters

Configuration field	Environment variable	Description	Example
Java Opts append	JAVA_OPTS_APPEND	User specified Java options to be appended to generated options in JAVA_OPTS.	-Dsome.property=foo
Java max memory ratio	JAVA_MAX_MEM_RATIO	The maximum percentage of container memory that can be used for the Java Virtual Machine. The remaining memory is used for the operating system. The default value is 50 , for a limit of 50%. Sets the -Xmx JVM option. If you enter a value of 0 , the -Xmx option is not set.	40
Java initial memory ratio	JAVA_INITIAL_MEM_RATIO	The percentage of container memory that is initially used for the Java Virtual Machine. The default value is 25 , so 25% of the pod memory is initially allocated for the JVM if this value does not exceed the Java Max Initial Memory value. Sets the -Xms JVM option. If you enter a value of 0 , the -Xms option is not set.	25
Java max initial memory	JAVA_MAX_INITIAL_MEMORY	The maximum amount of memory, in megabytes, that can be initially used for the Java Virtual Machine. If the initial allocated memory, as set in the Java initial memory ratio parameter, would otherwise be greater than this value, the amount of memory set in this value is allocated using the -Xms JVM option. The default value is 4096 .	4096
Java diagnostics	JAVA_DIAGNOSTICS	Enable this setting to enable output of additional JVM diagnostic information to the standard output. Disabled by default.	true

Configuration field	Environment variable	Description	Example
Java debug	JAVA_DEBUG	Enable this setting to switch on remote debugging. Disabled by default. Adds the -agentlib:jdwp=transport=dt_socket,server=y,suspend=n,address=\${debug_port} JVM option, where \${debug_port} defaults to 5005 .	true
Java debug port	JAVA_DEBUG_PORT	The port that is used for remote debugging. The default value is 5005 .	8787
GC min heap free ratio	GC_MIN_HEAP_FREE_RATIO	Minimum percentage of heap free after garbage collection (GC) to avoid expansion. Sets the -XX:MinHeapFreeRatio JVM option.	20
GC max heap free ratio	GC_MAX_HEAP_FREE_RATIO	Maximum percentage of heap free after GC to avoid shrinking. Sets the -XX:MaxHeapFreeRatio JVM option.	40
GC time ratio	GC_TIME_RATIO	Specifies the ratio of the time spent outside the garbage collection (for example, the time spent for application execution) to the time spent in the garbage collection. Sets the -XX:GCTimeRatio JVM option.	4
GC adaptive size policy weight	GC_ADAPTIVE_SIZE_POLICY_WEIGHT	The weighting given to the current GC time versus previous GC times. Sets the -XX:AdaptiveSizePolicyWeight JVM option.	90
GC max metaspace size	GC_MAX_METASPACE_SIZE	The maximum metaspace size. Sets the -XX:MaxMetaspaceSize JVM option.	100

3.6. KIE CONFIGURATION AND CONFIGMAPS

When installing the Red Hat Decision Manager operator, the operator creates ConfigMaps, a YAML file, prefixed as **kieconfig-\$VERSION**, for the current namespace. A ConfigMap is a YAML file containing the configuration for functionality such as **DeploymentConfigs**, secrets, routes, and services for the Red Hat Decision Manager components for that namespace.

Example ConfigMap:

```
# Please edit the object below. Lines beginning with a '#' will be ignored,
# and an empty file will abort the edit. If an error occurs while saving this file will be
```

```
# reopened with the relevant failures.
#
apiVersion: v1
data:
  mysql.yaml: |
    ## KIE Databases BEGIN
    databases:
      ## RANGE BEGINS
      #[[ range $index, $Map := .Databases ]]
  ...
```

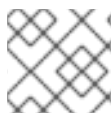
The operator uses ConfigMaps to configure and deploy the components. This includes all supported Red Hat Decision Manager components such as KIE Server, Smart Router, Business Central and service-related configuration such as for persistent volumes, build configuration, and routes. When ConfigMaps are edited manually, the operator uses the new values to create the deployments when the environment is reconciled.

The Red Hat Decision Manager operator can use the current version and the previous version of Red Hat Decision Manager components concurrently, with ConfigMaps for each version, for example, 7.13.0 and 7.12.1.

kieconfigs-7.13.5

This contains the **common.yaml** configuration file. For more information, see [common.yaml](#) on GitHub. You can configure the following components using this configuration file:

- KIE Server, the server object, is identified by the **## KIE Servers BEGIN** placeholder.
- Business Central and Business Central Monitoring, the console object, is defined on the first line of the **common.yaml**.
- Smartrouter, the smartrouter object, is identified by the **## KIE smartrouter BEGIN** placeholder.



NOTE

kieconfigs-7.13.5 also holds the routes and services relating to these three components.

kieconfigs-7.13.5-dashbuilder

This contains the configuration YAML file for the Dashbuilder component. For more information, see [rhcam-standalone-dashbuilder.yaml](#) on GitHub.

kieconfigs-7.13.5-dbs

This contains the base **DeploymentConfig** for the MySQL and PostgreSQL databases. For more information about the MySQL configuration, see [mysql.yaml](#) on GitHub.

For more information about the PostgreSQL configuration, see [postgresql.yaml](#) on GitHub.

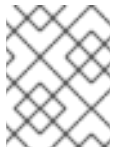
kieconfigs-7.13.5-dbs-pim

This contains the snippet configuration for the Process Instance Migration (PIM) supported databases; external, MySQL, and PostgreSQL.

For more information about the PIM external configuration, see [external.yaml](#) on GitHub.

For more information about the PIM MySQL configuration, see [mysql.yaml](#) on GitHub.

For more information about the PIM PostgreSQL configuration, see [postgresql.yaml](#) on GitHub.



NOTE

These YAML files only hold the specific configuration for the PIM database configured using **application.properties** in this configMap.

kieconfigs-7.13.5-dbs-servers

This contains the snippet configuration for the supported database configurations; external, h2, MySQL, and PostgreSQL.

For more information about the external configuration, see [external.yaml](#) on GitHub.

For more information about the h2 configuration, see [h2.yaml](#) on GitHub. Note that the h2 configuration is not supported on production environments.

For more information about the MySQL configuration, see [mysql.yaml](#) on GitHub.

For more information about the PostgreSQL configuration, see [postgresql.yaml](#) on GitHub.

kieconfigs-7.13.5-envs

This contains the specific configurations for each Red Hat Decision Manager environment such as the authoring or trial environments. This ConfigMap contains the following YAML files:

- rhdm-authoring-ha.yaml
- rhdm-authoring.yaml
- rhdm-production-immutable.yaml
- rhdm-trial.yaml
- rhpam-authoring-ha.yaml
- rhpam-authoring.yaml
- rhpam-production-immutable.yaml
- rhpam-production.yaml
- rhpam-standalone-dashbuilder.yaml
- rhpam-trial.yaml

For more information about specific configurations for each Red Hat Decision Manager environment, see [ConfigMaps for Red Hat Decision Manager environments](#) on GitHub.

kieconfigs-7.13.5-jms

This contains the ActiveMQ configuration for KIE Server when the JMS Executor is enabled. For more information about the JMS Executor configuration, see [activemq-jms-config](#) on GitHub.

kieconfigs-7.13.5-pim

This contains the process instance migration (PIM) DeploymentConfig and related PIM configuration. If you are using MySQL or any other database with PIM you must use the **kieconfigs-7.13.5-dbs-pim** configMap and edit the **mysql.yaml** file.

3.6.1. Using ConfigMaps

You can use ConfigMaps to customize the Red Hat Decision Manager Operator and apply the related configuration. To make changes to ConfigMaps, you can use the **oc** command tool or the Red Hat OpenShift Container Platform console.

Prerequisites

- Operator is installed in your current namespace.
- A KieApp is available.
- **kieconfig-\$VERSION-*** are available.
To check if **kieconfig-\$VERSION-*** are available, run the following command:

```
$ oc get cm | grep kieconfigs
```

Example output of `oc get cm | grep kieconfigs`:

```
kieconfigs-7.13.0          1    64m
kieconfigs-7.13.0-dashbuilder 1    64m
kieconfigs-7.13.0-dbs      2    64m
kieconfigs-7.13.0-dbs-pim  3    64m
kieconfigs-7.13.0-dbs-servers 4    64m
kieconfigs-7.13.0-envs    10   64m
kieconfigs-7.13.0-jms      1    64m
kieconfigs-7.13.0-pim      1    64m
kieconfigs-7.12.1         1    64m
kieconfigs-7.12.1-dashbuilder 1    64m
kieconfigs-7.12.1-dbs      2    64m
kieconfigs-7.12.1-dbs-pim  3    64m
kieconfigs-7.12.1-dbs-servers 4    64m
kieconfigs-7.12.1-envs    10   64m
kieconfigs-7.12.1-jms      1    64m
kieconfigs-7.12.1-pim      1    64m
```

Procedure

1. Create an **rhpm-authoring environment** with one replica for Business Central and KIE Server.

Example **rhpm-authoring**:

```
apiVersion: app.kiegroup.org/v2
kind: KieApp
metadata:
  name: rhpm-authoring
annotations:
  consoleName: rhpm-authoring
  consoleTitle: PAM Authoring
  consoleDesc: Deploys a PAM Authoring environment
spec:
  environment: rhpm-authoring
objects:
  servers:
```

```
- replicas: 1
console:
  replicas: 1
```

2. Complete one of the following steps:

- To open a specific ConfigMap with the **oc** tool, run the following command:

```
$ oc edit cm/<CONFIGMAP_NAME>
```



NOTE

The **oc** tool is similar to the **vim** text editing tool.

- To open a specific ConfigMap using the Red Hat OpenShift Container Platform console, navigate to **kieconfigs-7.13.5** on the **ConfigMaps** page and edit it by opening its YAML version.

3. To modify the YAML file, add **annotations** fields that contain your changes, for example: Add the following to **Console.deploymentConfigs.metadata**:

```
annotations:
  my.custom.annotation/v1: v1-rhpam-app-console
```

To update KIE Server, add the following to the **## KIE Servers Start** placeholder identifier:

```
annotations:
  my.custom.annotation/v1: v1-rhpam-app-kieserver
```



NOTE

You can also edit the Smart Router configuration using this ConfigMap.

4. To save the changes, complete one of the following steps:

- In the **oc** tool editor, to save the change the change and exit, enter **:wq!**.
- In the the Red Hat OpenShift Container Platform console, to save the change using the Red Hat OpenShift Container Platform console, click **Save**.

5. If an environment is running and the operator didn't automatically start redeployment for the component that was changed, you must manually delete the DeploymentConfig of the target component by using **oc** command tool by completing the following steps:

a. To return the DeploymentConfig, run the following command:

```
$ oc get dc
```

The output of **oc get dc** is returned.

Example output of oc get dc:

```
$ oc get dc
```

NAME	REVISION	DESIRED	CURRENT	TRIGGERED BY
rhpm-authoring-kieserver	1	1	1	config
rhpm-authoring-rhpmcentr	1	1	1	config

- b. To delete the DeploymentConfig of the target components, run the following commands:

```
$ oc delete dc/rhpm-authoring-kieserver
```

```
$ oc delete dc/rhpm-authoring-rhpmcentr
```

The deployment is redeployed with the changes applied in the ConfigMap.

6. To verify that the change was applied by checking the annotations for the KIE Server DeploymentConfig, run the following command:

```
$ oc describe dc/rhpm-authoring-kieserver
```

Example output of `oc describe dc/rhpm-authoring-kieserver`:

```
Name: rhpm-authoring-kieserver
Namespace: examplnamespace
Created: 15 minutes ago
Labels: app=rhpm-authoring
        application=rhpm-authoring
        service=rhpm-authoring-kieserver
        services.server.kie.org/kie-server-id=rhpm-authoring-kieserver
Annotations: my.custom.annotation/v1=v1-rhpm-app-kieserver
```

3.7. CREATING CUSTOM IMAGES FOR KIE SERVER

You can create custom images to add files to KIE Server deployments. You must push the images to your own container registry. When deploying Red Hat Decision Manager, you can configure the operator to use the custom images.

If you use a custom image, you must disable automatic version updates. When you want to install a new version, build the image with the same name as before and the new version tag and push the image into your registry. You can then change the version and the operator automatically pulls the new image. For instructions about changing the product version in the operator, see [Section 3.3, “Modifying an environment that is deployed using operators”](#).

In particular, you can create the following types of custom images:

- A custom image of KIE Server that includes an additional RPM package
- A custom image of KIE Server that includes an additional JAR class library

3.7.1. Creating a custom KIE Server image with an additional RPM package

You can create a custom KIE Server image where an additional RPM package is installed. You can push this image into your custom registry and then use it to deploy KIE Server.

You can install any package from the Red Hat Enterprise Linux 8 repository. This example installs the **procps-ng** package, which provides the **ps** utility, but you can modify it to install other packages.

Procedure

1. Authenticate to the **registry.redhat.io** registry using the **podman login** command. For instructions about authenticating to the registry, see [Red Hat Container Registry Authentication](#).

2. To download the supported KIE Server base image, enter the following command:

```
podman pull registry.redhat.io/rhpam-7/rhpam-kieserver-rhel8:7.13.5
```

3. Create a **Dockerfile** that defines a custom image based on the base image. The file must change the current user to **root**, install the RPM package using the **yum** command, and then revert to **USER 185**, the Red Hat JBoss EAP user. The following example shows the content of the **Dockerfile** file:

```
FROM registry.redhat.io/rhpam-7/rhpam-kieserver-rhel8:7.13.5
USER root
RUN yum -y install procps-ng
USER 185
```

Replace the name of the RPM file as necessary. The **yum** command automatically installs all dependencies from the Red Hat Enterprise Linux 8 repository. You might need to install several RPM files, in this case, use several **RUN** commands.

4. Build the custom image using the **Dockerfile**. Supply the fully qualified name for the image, including the registry name. You must use the same version tag as the version of the base image. To build the image, enter the following command:

```
podman build . --tag registry_address/image_name:7.13.5
```

For example:

```
podman build . --tag registry.example.com/custom/rhpam-kieserver-rhel8:7.13.5
```

5. After the build completes, run the image, log in to it, and verify that the customization was successful. Enter the following command:

```
podman run -it --rm registry_address/image_name:7.13.5 /bin/bash
```

For example:

```
podman run -it --rm registry.example.com/custom/rhpam-kieserver-rhel8:7.13.5 /bin/bash
```

In the shell prompt for the image, enter the command to test that the RPM is installed, then enter **exit**. For example, for **procps-ng**, run the **ps** command:

```
[jboss@c2fab36b778e ~]$ ps
PID TTY      TIME CMD
  1 pts/0    00:00:00 bash
 13 pts/0    00:00:00 ps
[jboss@c2fab36b778e ~]$ exit
```

6. To push the custom image into your registry, enter the following command:

■


```
podman push registry_address/image_name:7.13.5
docker://registry_address/image_name:7.13.5
```

For example:

```
podman push registry.example.com/custom/rhpam-kieserver-rhel8:7.13.5
docker://registry.example.com/custom/rhpam-kieserver-rhel8:7.13.5
```

Next steps

When deploying the KIE Server, set the image name and namespace to specify the custom image in your registry. Click **Set KIE Server image**, change the **Kind** value to **DockerImage**, and then provide the image name including the registry name, but without the version tag, for example:

```
registry.example.com/custom/rhpam-kieserver-rhel8
```

For instructions about deploying KIE Server using the operator, see [Section 3.2.5, "Setting custom KIE Server configuration of the environment"](#).

3.7.2. Creating a custom KIE Server image with an additional JAR file

You can create a custom KIE Server image where an additional JAR file (or several JAR files) is installed to extend the capabilities of the server. You can push this image into your custom registry and then use it to deploy KIE Server.

For example, you can create a custom class JAR to provide custom Prometheus metrics in KIE Server. For instructions about creating the custom class, see [Extending Prometheus metrics monitoring in KIE Server with custom metrics](#) in *Managing and monitoring KIE Server*.

Procedure

1. Develop a custom library that works with KIE Server. You can use the following documentation and examples to develop the library:
 - [KIE Server capabilities and extensions](#) in *Managing and monitoring KIE Server*.
 - [Domain-specific Prometheus metrics with Red Hat Process Automation Manager and Decision Manager](#)
 - [Extend KIE Server with additional transport](#)
2. Build the library using Maven, so that the JAR file is placed in the **target** directory. This example uses the **custom-kieserver-ext-1.0.0.Final.jar** file name.
3. Authenticate to the **registry.redhat.io** registry using the **podman login** command. For instructions about authenticating to the registry, see [Red Hat Container Registry Authentication](#).
4. To download the supported KIE Server base image, enter the following command:

```
podman pull registry.redhat.io/rhpam-7/rhpam-kieserver-rhel8:7.13.5
```

5. Create a **Dockerfile** that defines a custom image based on the base image. The file must copy the JAR file (or several JAR files) into the **/opt/eap/standalone/deployments/ROOT.war/WEB-INF/lib/** directory. The following example

shows the content of the **Dockerfile** file:

```
FROM registry.redhat.io/rhpam-7/rhpam-kieserver-rhel8:7.13.5
COPY target/custom-kieserver-ext-1.0.0.Final.jar
/opt/eap/standalone/deployments/ROOT.war/WEB-INF/lib/
```

6. Build the custom image using the **Dockerfile**. Supply the fully qualified name for the image, including the registry name. You must use the same version tag as the version of the base image. To build the image, enter the following command:

```
podman build . --tag registry_address/image_name:7.13.5
```

For example:

```
podman build . --tag registry.example.com/custom/rhpam-kieserver-rhel8:7.13.5
```

7. To push the custom image into your registry, enter the following command:

```
podman push registry_address/image_name:7.13.5
docker://registry_address/image_name:7.13.5
```

For example:

```
podman push registry.example.com/custom/rhpam-kieserver-rhel8:7.13.5
docker://registry.example.com/custom/rhpam-kieserver-rhel8:7.13.5
```

Next steps

When deploying the KIE Server, set the image name and namespace to specify the custom image in your registry. Click **Set KIE Server image**, change the **Kind** value to **DockerImage**, and then provide the image name including the registry name, but without the version tag, for example:

```
registry.example.com/custom/rhpam-kieserver-rhel8
```

For instructions about deploying KIE Server using the operator, see [Section 3.2.5, "Setting custom KIE Server configuration of the environment"](#).

CHAPTER 4. MIGRATION OF INFORMATION FROM A DEPLOYMENT ON RED HAT OPENSIFT CONTAINER PLATFORM 3

If you previously used a Red Hat Decision Manager deployment on Red Hat OpenShift Container Platform 3, you can migrate the information from that deployment to a new deployment on Red Hat OpenShift Container Platform 4.

Before migrating information, you must deploy a new Red Hat Decision Manager infrastructure on Red Hat OpenShift Container Platform 4 using the operator. Include the same elements in the new infrastructure as those present in the old deployment. For example:

- For any existing authoring deployment, create a new authoring infrastructure, including Business Central and at least one KIE Server.
- For any existing immutable KIE Server, deploy a new immutable KIE Server with the same artifacts.

4.1. MIGRATING INFORMATION IN BUSINESS CENTRAL

If you have an existing authoring environment in Red Hat OpenShift Container Platform 3, you can copy the **.niogit** repository and the Maven repository from Business Central in this environment to Business Central in a new deployment on Red Hat OpenShift Container Platform 4. This action makes all the same projects and artifacts available in the new deployment.

Prerequisites

- You must have a machine that has network access to both the Red Hat OpenShift Container Platform 3 and Red Hat OpenShift Container Platform 4 infrastructures.
- The **oc** command-line client from Red Hat OpenShift Container Platform 4 must be installed on the machine. For instructions about installing the command-line client, see [CLI tools](#) in Red Hat OpenShift Container Platform documentation.

Procedure

1. Ensure that no web clients and no client applications are connected to any elements of the old and new deployment, including Business Central and KIE Servers.
2. Create an empty temporary directory and change into it.
3. Using the **oc** command, log in to the Red Hat OpenShift Container Platform 3 infrastructure and switch to the project containing the old deployment.
4. To view the pod names in the old deployment, run the following command:

```
oc get pods
```

Find the Business Central pod. The name of this pod includes **rhpamcentr**. In a high-availability deployment, you can use any of the Business Central pods.

5. Use the **oc** command to copy the **.niogit** repository and the Maven repository from the pod to the local machine, for example:

```
oc cp myapp-rhpamcentr-5-689mw:/opt/kie/data/.niogit .niogit
oc cp myapp-rhpamcentr-5-689mw:/opt/kie/data/maven-repository maven-repository
```

6. Using the **oc** command, log in to the Red Hat OpenShift Container Platform 4 infrastructure and switch to the project containing the new deployment.

7. To view the pod names in the new deployment, run the following command:

```
oc get pods
```

Find the Business Central pod. The name of this pod includes **rhpamcentr**. In a high-availability deployment, you can use any of the Business Central pods.

8. Use the **oc** command to copy the **.niogit** repository and the Maven repository from the local machine to the pod, for example:

```
oc cp .niogit myappnew-rhpamcentr-abd24:/opt/kie/data/.niogit
oc cp maven-repository myappnew-rhpamcentr-abd24:/opt/kie/data/maven-repository
```

APPENDIX A. VERSIONING INFORMATION

Documentation last updated on Thursday, March 14th, 2024.

APPENDIX B. CONTACT INFORMATION

Red Hat Decision Manager documentation team: brms-docs@redhat.com