



Red Hat Decision Manager 7.8

Deploying a Red Hat Decision Manager
environment on Red Hat OpenShift Container
Platform using Operators

Red Hat Decision Manager 7.8 Deploying a Red Hat Decision Manager environment on Red Hat OpenShift Container Platform using Operators

Red Hat Customer Content Services
brms-docs@redhat.com

Legal Notice

Copyright © 2021 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 Red Hat Decision Manager 7.8 environment on Red Hat OpenShift Container Platform using Operators.

Table of Contents

| | |
|--|-----------|
| PREFACE | 3 |
| CHAPTER 1. OVERVIEW OF RED HAT DECISION MANAGER ON RED HAT OPENSIFT CONTAINER PLATFORM | 4 |
| CHAPTER 2. ARCHITECTURE OF AN AUTHORIZING ENVIRONMENT | 5 |
| Single authoring environment | 5 |
| Clustering KIE Servers and using multiple KIE Servers | 6 |
| Smart Router | 6 |
| High-availability authoring environment | 6 |
| CHAPTER 3. PREPARING TO DEPLOY RED HAT DECISION MANAGER IN YOUR OPENSIFT ENVIRONMENT | 9 |
| 3.1. ENSURING YOUR ENVIRONMENT IS AUTHENTICATED TO THE RED HAT REGISTRY | 9 |
| 3.2. CREATING THE SECRETS FOR KIE SERVER | 9 |
| 3.3. CREATING THE SECRETS FOR BUSINESS CENTRAL | 10 |
| 3.4. CREATING THE SECRETS FOR THE AMQ BROKER CONNECTION | 11 |
| 3.5. PREPARING GIT HOOKS | 11 |
| 3.6. PROVISIONING PERSISTENT VOLUMES WITH READWRITEMANY ACCESS MODE USING NFS | 12 |
| 3.7. EXTRACTING THE SOURCE CODE FROM BUSINESS CENTRAL FOR USE IN AN S2I BUILD | 12 |
| 3.8. PREPARING FOR DEPLOYMENT IN A RESTRICTED NETWORK | 13 |
| 3.9. PREPARING A MAVEN MIRROR REPOSITORY FOR OFFLINE USE | 14 |
| CHAPTER 4. DEPLOYMENT AND MANAGEMENT OF A RED HAT DECISION MANAGER ENVIRONMENT USING OPENSIFT OPERATORS | 16 |
| 4.1. SUBSCRIBING TO THE BUSINESS AUTOMATION OPERATOR | 16 |
| 4.2. DEPLOYING A RED HAT DECISION MANAGER ENVIRONMENT USING THE OPERATOR | 17 |
| 4.2.1. Starting the deployment of a Red Hat Decision Manager environment using the Business Automation operator | 17 |
| 4.2.2. Setting the basic configuration of the environment | 17 |
| 4.2.3. Setting the security configuration of the environment | 19 |
| 4.2.4. Setting the Business Central configuration of the environment | 21 |
| 4.2.5. Setting custom KIE Server configuration of the environment | 24 |
| 4.3. MODIFYING AN ENVIRONMENT THAT IS DEPLOYED USING OPERATORS | 28 |
| 4.4. JVM CONFIGURATION PARAMETERS | 29 |
| 4.5. CREATING CUSTOM IMAGES FOR KIE SERVER | 31 |
| 4.5.1. Creating a custom KIE Server image with an additional RPM package | 31 |
| 4.5.2. Creating a custom KIE Server image with an additional JAR file | 33 |
| CHAPTER 5. MIGRATING INFORMATION FROM A DEPLOYMENT ON RED HAT OPENSIFT CONTAINER PLATFORM VERSION 3 | 35 |
| 5.1. MIGRATING INFORMATION IN BUSINESS CENTRAL | 35 |
| APPENDIX A. VERSIONING INFORMATION | 37 |

PREFACE

As a system engineer, you can deploy a Red Hat Decision Manager environment on Red Hat OpenShift Container Platform to provide an infrastructure to develop or execute processes 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 version 4 environment is available. For the exact versions of Red Hat OpenShift Container Platform that the current release supports, see [Red Hat Process Automation Manager 7 Supported Configurations](#).
- At least four gigabytes of memory are available in the OpenShift environment.
- The OpenShift project for the deployment is created.
- You are logged in to the project using the OpenShift web console.
- 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 scale any of the Business Central pods, 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](#).

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.

CHAPTER 2. 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.

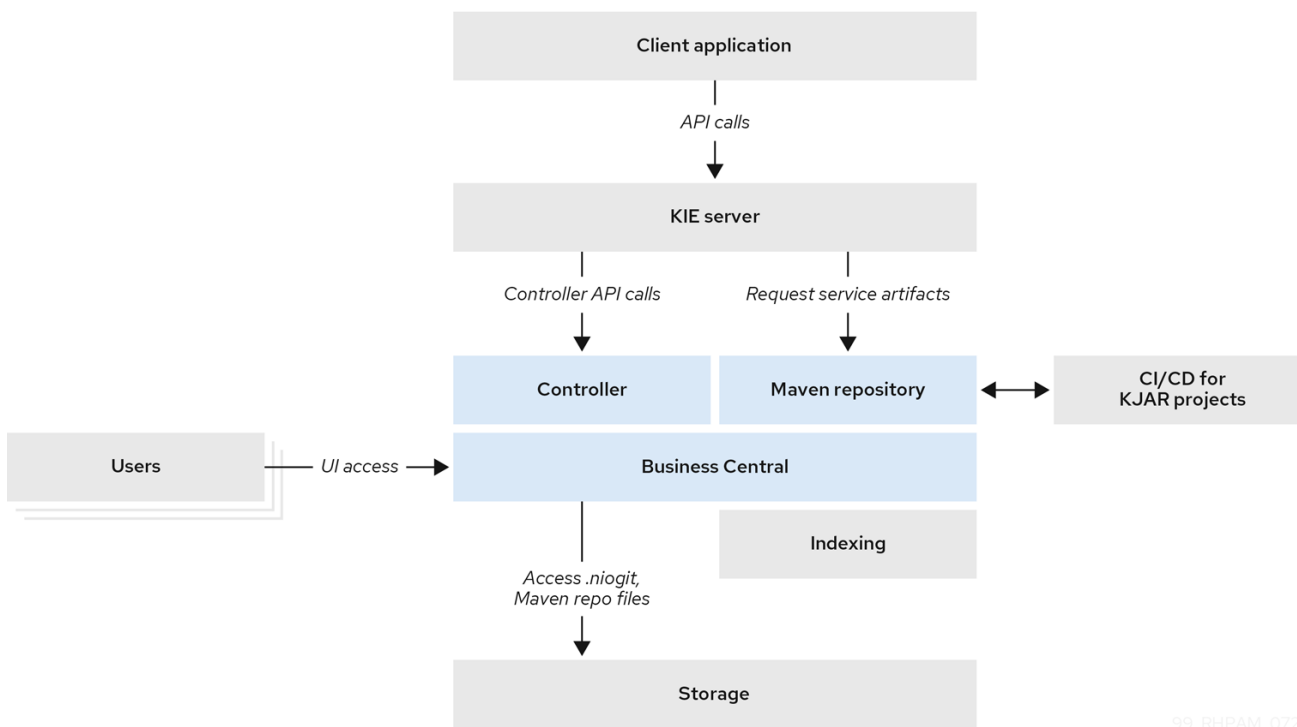
A single authoring environment, by default, includes one KIE Server.

A single authoring environment can use the *controller strategy*. Business Central includes the *Controller*, a component that can manage KIE Servers. When you configure a KIE Server to connect to Business Central, the 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 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 the KIE Server, the KIE Server receives the request through the Controller connection WebSocket. To deploy a service, the 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 the KIE Server.

Figure 2.1. Architecture diagram for a single authoring environment



99_RHPAM_0720

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 the 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 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 Servers. It can be useful if you are using several independent KIE Servers.

The client application can use services running on different KIE Servers, but always connects to the Smart Router. The Smart Router automatically passes the request to the KIE Servers 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 KIE Server configuration (for example, deploys or undeploys a service), Business Central initiates a connection to the KIE Server and sends a REST API request. The KIE Server changes the ConfigMap to reflect the new configuration state and then triggers its own redeployment, so that all instances are redeployed and reflect the new configuration.

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 2.2. Architecture diagram for a high-availability authoring environment



99_RHPAM_0720

CHAPTER 3. PREPARING TO DEPLOY RED HAT DECISION MANAGER IN YOUR OPENSIFT ENVIRONMENT

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

3.1. 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.

3.2. 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 with a private and public key for SSL encryption for KIE Server. For more information on how to create a keystore with self-signed or purchased SSL certificates, see [Generate a SSL Encryption Key and Certificate](#).



NOTE

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

2. Save the keystore in a file named **keystore.jks**.
3. Record the name of the certificate. The default value for this name in Red Hat Decision Manager configuration is **jboss**.
4. Record the password of the keystore file. The default value for this name in Red Hat Decision Manager configuration is **mykeystorepass**.
5. 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
```

3.3. 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 with a private and public key for SSL encryption for Business Central. For more information on how to create a keystore with self-signed or purchased SSL certificates, see [Generate a SSL Encryption Key and Certificate](#).



NOTE

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

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

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

3.4. 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 with a private and public key for SSL encryption for the AMQ broker connection. For more information on how to create a keystore with self-signed or purchased SSL certificates, see [Generate a SSL Encryption Key and Certificate](#).



NOTE

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

2. Save the keystore in a file named **keystore.jks**.
3. Record the name of the certificate. The default value for this name in Red Hat Decision Manager configuration is **jboss**.
4. Record the password of the keystore file. The default value for this name in Red Hat Decision Manager configuration is **mykeystorepass**.
5. 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
```

3.5. 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.

- 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:

- Prepare the **id_rsa** file with a private key that matches a public key stored in the repository.
- Prepare the **known_hosts** file with the correct name, address, and public key for the repository.
- 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.

3.6. 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 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 4.3 Storage](#) guide.

3.7. 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://<decision-central-host>:443/git/<MySpace>/<MyProject>
```

In this command, replace the following variables:

- **<decision-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://<decision-central-host>:443/git/<MySpace>/<MyProject>
```

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

3.8. 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 7.8, 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 3.9, “Preparing a Maven mirror repository for offline use”](#).

3.9. 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 version 4. Use this repository as a separate mirror repository.

Alternatively, if you use a custom external repository (for example, Nexus) for your services, you can use the same repository as a mirror repository.

2. On the computer that has an outgoing connection to the public Internet, complete the following steps:
 - a. Click [Red Hat Process Automation Manager 7.8.0 Offliner Content List](#) to download the **rhdm-7.8.0-offliner.zip** product deliverable file from the [Software Downloads](#) page of the Red Hat Customer Portal.
 - b. Extract the contents of the **rhdm-7.8.0-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 a **repository** subdirectory and downloads the necessary artifacts into this subdirectory.

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.
3. If you developed services outside 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 correct path to 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 4. 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.

You can delete the Persistent Volume Claims manually.

4.1. 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.

4.2. 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 7.8, the operator installer wizard is for Technology Preview only. For more information on Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

4.2.1. 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 4.1, "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.

4.2.2. 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 4.2.1, "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 7.8, 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](#).

- **rhdm-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 the KIE Server and either click the **Set immutable server configuration** button or set the **KIE_SERVER_CONTAINER_DEPLOYMENT** environment variable. For instructions about configuring the KIE Server, see [Section 4.2.5, "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 7.8 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 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.
5. 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.

6. If you want to use a custom version tag for images, complete the following steps:
 - a. Click **Next** to access the **Security** tab.
 - b. Scroll to the bottom of the window.
 - c. Enter the image tag in the **Image tag** field.

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.

4.2.3. 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 4.2.2, "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. The user can use Business Central to set up other users as necessary.
 - **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:

- a. In the **LDAP URL** field, enter the LDAP URL.
 - b. Configure LDAP parameters that correspond to the settings of the LdapExtended Login module of Red Hat JBoss EAP. For instructions about using these settings, see [LdapExtended Login Module](#).
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 a role mapping configuration file in an OpenShift configuration map or secret object in the project namespace. 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 file, make the following changes:

- a. Under **RoleMapper**, in the **Roles properties file** field, enter the fully qualified path name of the role mapping configuration file, for example, **/opt/eap/standalone/configuration/rolemapping/rolemapping.properties**.
- b. If you want to replace roles defined in the authentication system with roles that you define in the mapping file, select the **Replace roles** box. Otherwise, both the roles defined in RH-SSO or LDAP and the roles defined in the configuration file are available.

- c. In the fields under **RoleMapper Configuration object**, select the **Kind** of the object that provides the file (**ConfigMap** or **Secret**) and 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 3.2, "Creating the secrets for KIE Server"](#) or [Section 3.3, "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.

4.2.4. 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 or Business Central Monitoring component of the environment.

Do not change these settings for an immutable server 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 4.2.2, "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 4.2.3, "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 3.3, "Creating the secrets for Business Central"](#), enter the name of the secret in the **Secret** field.
3. Optionally, 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 3.5, “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.
 - c. Optionally, in the **SSH secret** field enter the name of the secret with the SSH key and known hosts files.
4. Optionally, enter the number of replicas for Business Central or Business Central monitoring in the **Replicas** field. Do not change this number in a **rhdm-authoring** environment.
 5. Optionally, enter requested and maximum CPU and memory limits in the fields under **Resource quotas**.
 6. 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 4.4, “JVM configuration parameters”](#).
 7. 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 the KIE Server into the **SSO HTTP URL** and **SSO HTTPS URL** fields. This information is recorded in the client.
 8. Optionally, 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 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 a 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 3.9, "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 3.9, "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-rhdmcentr**.
- 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 needs to wait for some seconds before more services can be deployed. The **OpenShiftStartupStrategy** setting is enabled by default and causes this limitation. To remove the limitation, you can configure an **rhdm-authoring** environment to use the *controller strategy*. Do not make this change unless a specific need for it exists; if you decide to enable controller strategy, make this change on Business Central and on all KIE Servers in the same environment.



NOTE

Do not enable the controller strategy in an environment with a high-availability Business Central. In such environments the controller strategy does not function correctly.

To enable the controller strategy on Business Central, set the **KIE_SERVER_CONTROLLER_OPENSIFT_ENABLED** environment variable to **false**.

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.

4.2.5. 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 4.2.2, "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 4.2.3, "Setting the security configuration of the environment"](#).

Procedure

1. If the **Installation**, **Security**, or **Console** tab is open, click **Next** until you view the **KIE Servers** tab.
2. Click **Add new KIE Server** to add a new KIE Server configuration.
3. In the **Id** field, enter an identifier for the KIE Server. If the KIE Server connects to a Business Central or Business Central Monitoring instance, this identifier determines which server group the server joins.
4. In the **Name** field, enter a name for the KIE Server.
5. 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.
6. If you created the secret for KIE Server according to the instructions in [Section 3.2, "Creating the secrets for KIE Server"](#), enter the name of the secret in the **Keystore secret** field.
7. Optionally, enter the number of replicas for the KIE Server in the **Replicas** field.
8. If you want to use a custom image for the KIE Server, complete the following additional steps:
 - a. Click **Set KIE Server image**
 - b. If you want to use a Docker image name and not an OpenShift image stream tag, change the **Kind** value to **DockerImage**.
 - c. Enter the name of the image stream in the **Name** field.
 - d. If the image stream is not in the **openshift** namespace, enter the namespace in the **Namespace** field.
For instructions about creating custom images, see [Section 4.5, "Creating custom images for KIE Server"](#).
9. 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_REPLOYMENT** 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 3.9, "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.
 - g. If you want to set a Git Webhook so changes in the Git repository cause an automatic rebuild of the KIE Server, click **Add new Webhook**. Select the type of the Webhook from the **Type** list and enter the secret string for the Webhook in the **Secret** field.
10. Optionally, 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.
 11. 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 the KIE Server into the **SSO HTTP URL** and **SSO HTTPS URL** fields. This information is recorded in the client.
12. 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 uses 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 3.4, "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.
13. 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 4.4, "JVM configuration parameters"](#).
14. Optionally, 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 3.9, “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 3.9, “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-rhdmcentr**.

- 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 `{URL_MANAGING_SETTINGS}#prometheus-monitoring-ocp-proc-execution-server` [*Managing and monitoring KIE Server*].

- If you are using `{RH-SSO}` authentication and the interaction of your application with `{RH-SSO}` requires support for CORS, set the **SSO_ENABLE_CORS** variable to **true**.

- 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 needs to wait for some seconds before more services can be deployed. The **OpenShiftStartupStrategy** setting is enabled by default and causes this limitation. To remove the limitation, you can configure an **rhdm-authoring** environment to use the *controller strategy*. Do not make this change unless a specific need for it exists; if you decide to enable controller strategy, make this change on Business Central and on all KIE Servers in the same environment.



NOTE

Do not enable the controller strategy in an environment with a high-availability Business Central. In such environments the controller strategy does not function correctly.

To enable controller strategy on a KIE Server, set the **KIE_SERVER_STARTUP_STRATEGY** environment variable to **ControllerBasedStartupStrategy** and the **KIE_SERVER_CONTROLLER_OPENSHIFT_ENABLED** environment variable to **false**.

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.

4.3. 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: 7.8.0
```

You can replace **7.8.0** 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.

4.4. 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 4.1. 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 |
| 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 |

| Configuration field | Environment variable | Description | Example |
|--------------------------------|--------------------------------|--|------------|
| 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 |

4.5. 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 4.3, “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 library

4.5.1. 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 the 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/rhdm-7/rhdm-kieserver-rhel8:7.8.0
```

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/rhdm-7/rhdm-kieserver-rhel8:7.8.0
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:7.8.0
```

For example:

```
podman build . --tag registry.example.com/custom/rhdm-kieserver-rhel8:7.8.0
```

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:7.8.0 /bin/bash
```

For example:

```
podman run -it --rm registry.example.com/custom/rhdm-kieserver-rhel8:7.8.0 /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:7.8.0
docker://registry_address/image_name:7.8.0
```

For example:

```
podman push registry.example.com/custom/rhdm-kieserver-rhel8:7.8.0
docker://registry.example.com/custom/rhdm-kieserver-rhel8:7.8.0
```

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/rhdm-kieserver-rhel8
```

For instructions about deploying the KIE Server using the operator, see [Section 4.2.5, “Setting custom KIE Server configuration of the environment”](#).

4.5.2. 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 the KIE Server.

Procedure

1. Develop a custom library that works with the 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/rhdm-7/rhdm-kieserver-rhel8:7.8.0
```

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/rhdm-7/rhdm-kieserver-rhel8:7.8.0
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:7.8.0
```

For example:

```
podman build . --tag registry.example.com/custom/rhdm-kieserver-rhel8:7.8.0
```

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

```
podman push registry_address/image_name:7.8.0
docker://registry_address/image_name:7.8.0
```

For example:

```
podman push registry.example.com/custom/rhdm-kieserver-rhel8:7.8.0
docker://registry.example.com/custom/rhdm-kieserver-rhel8:7.8.0
```

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/rhdm-kieserver-rhel8
```

For instructions about deploying the KIE Server using the operator, see [Section 4.2.5, "Setting custom KIE Server configuration of the environment"](#).

CHAPTER 5. MIGRATING INFORMATION FROM A DEPLOYMENT ON RED HAT OPENSIFT CONTAINER PLATFORM VERSION 3

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

Before migrating information, you must deploy a new Red Hat Decision Manager infrastructure on Red Hat OpenShift Container Platform version 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.

5.1. MIGRATING INFORMATION IN BUSINESS CENTRAL

If you have an existing authoring environment in Red Hat OpenShift Container Platform version 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 version 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 version 3 and Red Hat OpenShift Container Platform version 4 infrastructures.
- The **oc** command-line client from Red Hat OpenShift Container Platform version 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 version 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 **rhdmcentr**. 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-rhdmcentr-5-689mw:/opt/kie/data/.niogit .niogit
oc cp myapp-rhdmcentr-5-689mw:/opt/kie/data/maven-repository maven-repository
```

- Using the **oc** command, log in to the Red Hat OpenShift Container Platform version 4 infrastructure and switch to the project containing the new deployment.
- 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 **rhdmcentr**. In a high-availability deployment, you can use any of the Business Central pods.

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


APPENDIX A. VERSIONING INFORMATION

Documentation last updated on Friday, June 25, 2021.