



## Red Hat Process Automation Manager 7.8

Deploying a Red Hat Process Automation  
Manager authoring environment on Red Hat  
OpenShift Container Platform



# Red Hat Process Automation Manager 7.8 Deploying a Red Hat Process Automation Manager authoring environment on Red Hat OpenShift Container Platform

---

Red Hat Customer Content Services

[brms-docs@redhat.com](mailto: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<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

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

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

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

Node.js<sup>®</sup> 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<sup>®</sup> 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 Process Automation Manager 7.8 authoring environment on Red Hat OpenShift Container Platform.

# Table of Contents

<b>PREFACE</b> .....	<b>4</b>
<b>CHAPTER 1. OVERVIEW OF RED HAT PROCESS AUTOMATION MANAGER ON RED HAT OPENSIFT CONTAINER PLATFORM</b> .....	<b>6</b>
<b>CHAPTER 2. ARCHITECTURE OF AN AUTHORIZING ENVIRONMENT</b> .....	<b>8</b>
Single authoring environment	8
Clustering KIE Servers and using multiple KIE Servers	9
Smart Router	9
High-availability authoring environment	9
<b>CHAPTER 3. PREPARING TO DEPLOY RED HAT PROCESS AUTOMATION MANAGER IN YOUR OPENSIFT ENVIRONMENT</b> .....	<b>12</b>
3.1. ENSURING THE AVAILABILITY OF IMAGE STREAMS AND THE IMAGE REGISTRY	12
3.2. CREATING THE SECRETS FOR KIE SERVER	13
3.3. CREATING THE SECRETS FOR BUSINESS CENTRAL	14
3.4. CREATING THE SECRET FOR THE ADMINISTRATIVE USER	14
3.5. CHANGING GLUSTERFS CONFIGURATION	15
3.6. PROVISIONING PERSISTENT VOLUMES WITH READWRITEMANY ACCESS MODE USING NFS	16
3.7. PREPARING A MAVEN MIRROR REPOSITORY FOR OFFLINE USE	17
3.8. BUILDING A CUSTOM KIE SERVER EXTENSION IMAGE FOR AN EXTERNAL DATABASE	18
<b>CHAPTER 4. AUTHORIZING ENVIRONMENT</b> .....	<b>22</b>
4.1. DEPLOYING AN AUTHORIZING ENVIRONMENT	23
4.1.1. Starting configuration of the template for an authoring environment	23
4.1.2. Setting required parameters for an authoring environment	24
4.1.3. Configuring the image stream namespace for an authoring environment	25
4.1.4. Setting an optional Maven repository for an authoring environment	25
4.1.5. Configuring access to a Maven mirror in an environment without a connection to the public Internet for an authoring environment	26
4.1.6. Configuring Business Central and KIE Server replicas for a high-availability authoring environment	26
4.1.7. Specifying the Git hooks directory for an authoring environment	27
4.1.8. Configuring resource usage for a high-availability deployment	27
4.1.9. Setting parameters for RH-SSO authentication for an authoring environment	28
4.1.10. Setting parameters for LDAP authentication for an authoring environment	30
4.1.11. Setting parameters for using an external database server for an authoring environment	31
4.1.12. Enabling Prometheus metric collection for an authoring environment	33
4.1.13. Completing deployment of the template for an authoring environment	33
4.2. (OPTIONAL) PROVIDING THE GIT HOOKS DIRECTORY	34
4.3. (OPTIONAL) PROVIDING A TRUSTSTORE FOR ACCESSING HTTPS SERVERS WITH SELF-SIGNED CERTIFICATES	35
4.4. (OPTIONAL) PROVIDING THE LDAP ROLE MAPPING FILE	37
4.5. ENABLING THE OPENSIFTSTARTUPSTRATEGY SETTING TO CONNECT ADDITIONAL KIE SERVERS TO BUSINESS CENTRAL	38
4.6. MODIFYING THE TEMPLATE FOR THE SINGLE AUTHORIZING ENVIRONMENT	39
4.7. MODIFYING THE TEMPLATE FOR THE HIGH AVAILABILITY AUTHORIZING ENVIRONMENT	41
<b>CHAPTER 5. RED HAT PROCESS AUTOMATION MANAGER ROLES AND USERS</b> .....	<b>43</b>
<b>CHAPTER 6. OPENSIFT TEMPLATE REFERENCE INFORMATION</b> .....	<b>45</b>
6.1. RHPAM78-AUTHORIZING.YAML TEMPLATE	45
6.1.1. Parameters	45
6.1.2. Objects	59

6.1.2.1. Services	59
6.1.2.2. Routes	60
6.1.2.3. Deployment Configurations	60
6.1.2.3.1. Triggers	60
6.1.2.3.2. Replicas	60
6.1.2.3.3. Pod Template	61
6.1.2.3.3.1. Service Accounts	61
6.1.2.3.3.2. Image	61
6.1.2.3.3.3. Readiness Probe	61
6.1.2.3.3.4. Liveness Probe	61
6.1.2.3.3.5. Exposed Ports	62
6.1.2.3.3.6. Image Environment Variables	62
6.1.2.3.3.7. Volumes	80
6.1.2.4. External Dependencies	81
6.1.2.4.1. Volume Claims	81
6.1.2.4.2. Secrets	81
6.2. RHPAM78-AUTHORING-HA.YAML TEMPLATE	81
6.2.1. Parameters	81
6.2.2. Objects	99
6.2.2.1. Services	99
6.2.2.2. Routes	100
6.2.2.3. Deployment Configurations	101
6.2.2.3.1. Triggers	101
6.2.2.3.2. Replicas	101
6.2.2.3.3. Pod Template	101
6.2.2.3.3.1. Service Accounts	101
6.2.2.3.3.2. Image	102
6.2.2.3.3.3. Readiness Probe	102
6.2.2.3.3.4. Liveness Probe	102
6.2.2.3.3.5. Exposed Ports	103
6.2.2.3.3.6. Image Environment Variables	103
6.2.2.3.3.7. Volumes	123
6.2.2.4. External Dependencies	123
6.2.2.4.1. Volume Claims	123
6.2.2.4.2. Secrets	123
6.2.2.4.3. Clustering	123
6.3. OPENSIFT USAGE QUICK REFERENCE	125
<b>APPENDIX A. VERSIONING INFORMATION</b>	<b>127</b>



## PREFACE

As a system engineer, you can deploy a Red Hat Process Automation Manager authoring environment on Red Hat OpenShift Container Platform to provide a platform for development of services, process applications, and other business assets.

### Prerequisites

- Red Hat OpenShift Container Platform version 3.11 is deployed.
- At least four gigabytes of memory are available in the OpenShift cluster/namespace.
- For a high-availability deployment, the following resources are available on the OpenShift cluster:
  - For the Business Central replicated pod, 8 gigabytes of memory and 2 CPU cores are required for each replica. Two replicas are created by default.
  - For the KIE Server replicated pod, 1 gigabyte of memory and 1 CPU core are required for each replica. Two replicas are created by default.
  - For the Red Hat Data Grid replicated pod, 2 gigabytes of memory and 1 CPU core are required for each replica. Two replicas are created by default.
  - The Red Hat AMQ replicated pod uses the default resource limits configured on your cluster.
  - The MySQL replicated pod uses the default resource limits configured on your cluster.



### NOTE

For instructions about checking the capacity of your cluster, see [Analyzing cluster capacity](#) in the Red Hat OpenShift Container Platform 3.11 product documentation.

- The OpenShift project for the deployment is created.
- You are logged in to the project using the **oc** command. For more information about the **oc** command-line tool, see the OpenShift [CLI Reference](#). If you want to use the OpenShift Web console to deploy templates, you must also be logged on using the 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:
  - The replicated set of KIE Server pods requires one 1Gi PV for the database by default. You can change the database PV size in the template parameters. This requirement does not apply if you use an external database server.
  - Business Central requires one 1Gi PV by default. You can change the PV size for Business Central persistent storage in the template parameters.
- If you intend to deploy a high-availability authoring environment, which includes high-availability Business Central, 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. However, for best performance and reliability, use GlusterFS to provision persistent volumes for a high-availability authoring environment. For information about access mode support in OpenShift public and dedicated clouds, see [Access Modes](#).

**NOTE**

Since Red Hat Process Automation Manager version 7.5, images and templates for Red Hat OpenShift Container Platform 3.x are deprecated. These images and templates do not get new features, but remain supported until the end of full support for Red Hat OpenShift Container Platform version 3.x. For more information about the full support lifecycle phase for Red Hat OpenShift Container Platform version 3.x, see [Red Hat OpenShift Container Platform Life Cycle Policy \(non-current versions\)](#).

**NOTE**

Do not use Red Hat Process Automation Manager templates with Red Hat OpenShift Container Platform 4.x. To deploy Red Hat Process Automation Manager on Red Hat OpenShift Container Platform 4.x, see the instructions in [Deploying a Red Hat Process Automation Manager environment on Red Hat OpenShift Container Platform using Operators](#).

# CHAPTER 1. OVERVIEW OF RED HAT PROCESS AUTOMATION MANAGER ON RED HAT OPENSIFT CONTAINER PLATFORM

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

In this solution, components of Red Hat Process Automation 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 Process Automation Manager are available on OpenShift:

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

A database server is normally required for KIE Server. You can provide a database server in another OpenShift pod or configure an execution server on OpenShift to use any other database server. Alternatively, KIE Server can use an H2 database; in this case, you cannot scale the pod.

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 use the same database server and 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 and monitoring console. You can use Business Central to develop services and deploy them to KIE Servers. You can also use Business Central to monitor the execution of processes.

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.

- Business Central Monitoring is a web-based management and monitoring console. It can manage the deployment of services to KIE Servers and provide monitoring information, but does not include authoring capabilities. You can use this component to manage staging and production environments.
- Smart Router is an optional layer between KIE Servers and other components that interact with them. When your environment includes many services running on different KIE Servers, Smart Router provides a single endpoint to all client applications. A client application can make a REST API call that requires any service. Smart Router automatically calls the KIE Server that can process a particular request.

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

The following environment types are typical:

- *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. For instructions about deploying this environment, see [Deploying a Red Hat Process Automation Manager authoring environment on Red Hat OpenShift Container Platform](#).
- *Managed deployment*: An environment for running existing services for staging and production purposes. This environment includes several groups of KIE Server pods; you can deploy and undeploy services on every such group and also scale the group up or down as necessary. Use Business Central Monitoring to deploy, run, and stop the services and to monitor their execution. You can deploy two types of managed environment. In a *freeform* server environment, you initially deploy Business Central Monitoring and one KIE Server. You can additionally deploy any number of KIE Servers. Business Central Monitoring can connect to all servers in the same namespace. For instructions about deploying this environment, see [Deploying a Red Hat Process Automation Manager freeform managed server environment on Red Hat OpenShift Container Platform](#).

Alternatively, you can deploy a *fixed* managed server environment. A single deployment includes Business Central Monitoring, Smart Router, and a preset number of KIE Servers (by default, two servers, but you can modify the template to change the number). You cannot easily add or remove servers at a later time. For instructions about deploying this environment, see [Deploying a Red Hat Process Automation Manager fixed managed server environment on Red Hat OpenShift Container Platform](#).

- *Deployment with immutable servers*: An alternate environment for running existing services for staging and production purposes. In this environment, when you deploy a KIE Server pod, it builds an image that loads and starts a service or group of services. You cannot stop any service on the pod or add any new service to the pod. If you want to use another version of a service or modify the configuration in any other way, you deploy a new server image and displace the old one. In this system, the KIE Server runs like any other pod on the OpenShift environment; you can use any container-based integration workflows and do not need to use any other tools to manage the pods. Optionally, you can use Business Central Monitoring to monitor the performance of the environment and to stop and restart some of the service instances, but not to deploy additional services to any KIE Server or undeploy any existing ones (you cannot add or remove containers). For instructions about deploying this environment, see [Deploying a Red Hat Process Automation Manager immutable server environment on Red Hat OpenShift Container Platform](#).

You can also deploy a *trial* or evaluation environment. This environment includes Business Central and a KIE Server. You can set it up quickly and use it to evaluate or demonstrate developing and running assets. However, the environment does not use any persistent storage, and any work you do in the environment is not saved. For instructions about deploying this environment, see [Deploying a Red Hat Process Automation Manager trial environment on Red Hat OpenShift Container Platform](#).

To deploy a Red Hat Process Automation Manager environment on OpenShift, you can use the templates that are provided with Red Hat Process Automation Manager. You can modify the templates to ensure that the configuration suits your environment.

## CHAPTER 2. ARCHITECTURE OF AN AUTHORIZING ENVIRONMENT

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

The KIE Server uses a database server to store the state of process 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.

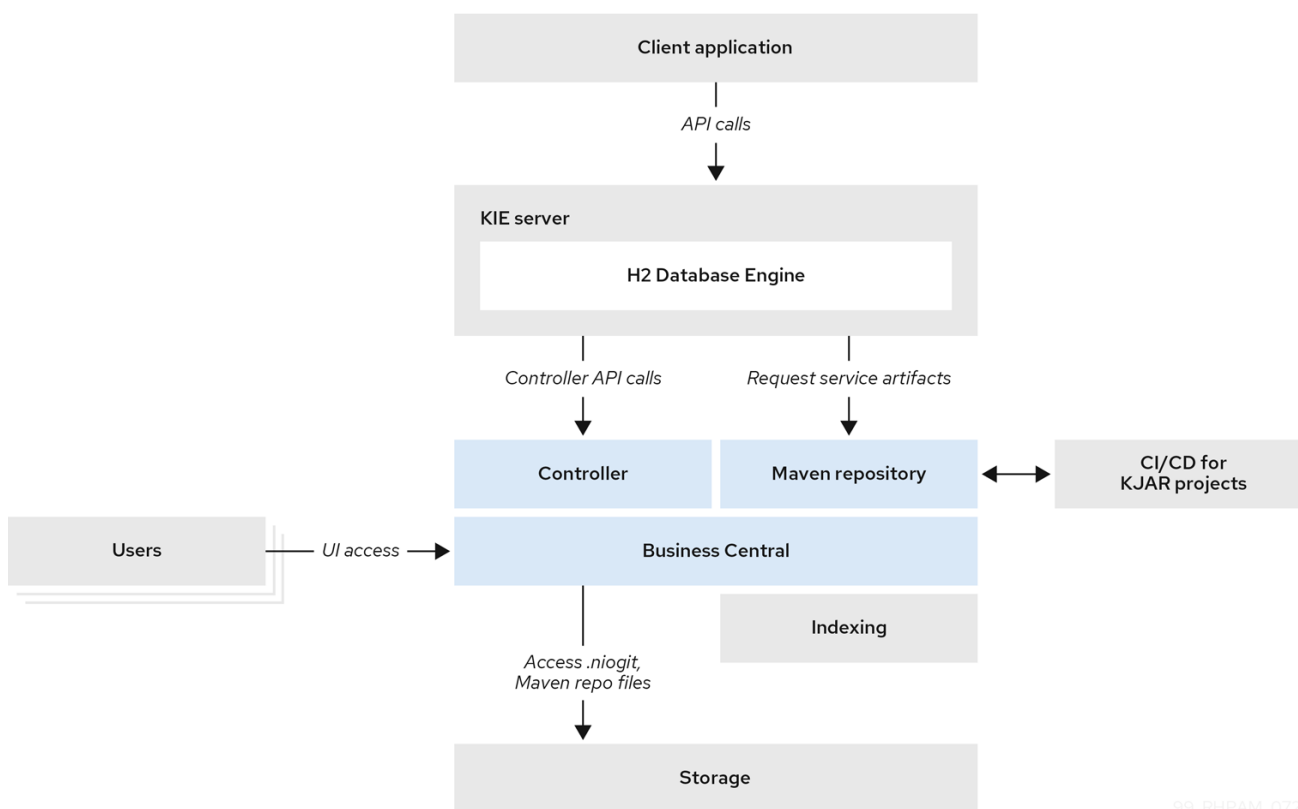
A single authoring environment, by default, includes one KIE Server. This KIE Server uses a built-in H2 database engine to store the state of process services.

A single authoring environment, by default, uses 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. To scale a KIE Server, you must ensure that it uses a database server in a separate pod or an external database server, and not a built-in H2 database engine.

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. Decision services and business optimizer 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 PROCESS AUTOMATION MANAGER IN YOUR OPENSIFT ENVIRONMENT

Before deploying Red Hat Process Automation 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 processes or for other processes.

### 3.1. ENSURING THE AVAILABILITY OF IMAGE STREAMS AND THE IMAGE REGISTRY

To deploy Red Hat Process Automation Manager components on Red Hat OpenShift Container Platform, you must ensure that OpenShift can download the correct images from the Red Hat registry. To download the images, OpenShift requires *image streams*, which contain the information about the location of images. OpenShift also must be configured to authenticate with the Red Hat registry using your service account user name and password.

Some versions of the OpenShift environment include the required image streams. You must check if they are available. If image streams are available in OpenShift by default, you can use them if the OpenShift infrastructure is configured for registry authentication server. The administrator must complete the registry authentication configuration when installing the OpenShift environment.

Otherwise, you can configure registry authentication in your own project and install the image streams in that project.

#### Procedure

1. Determine whether Red Hat OpenShift Container Platform is configured with the user name and password for Red Hat registry access. For details about the required configuration, see [Configuring a Registry Location](#). If you are using an OpenShift Online subscription, it is configured for Red Hat registry access.
2. If Red Hat OpenShift Container Platform is configured with the user name and password for Red Hat registry access, enter the following commands:

```
$ oc get imagestreamtag -n openshift | grep -F rhpam-businesscentral | grep -F 7.8
$ oc get imagestreamtag -n openshift | grep -F rhpam-kieserver | grep -F 7.8
```

If the outputs of both commands are not empty, the required image streams are available in the **openshift** namespace and no further action is required.

3. If the output of one or both of the commands is empty or if OpenShift is not configured with the user name and password for Red Hat registry access, complete the following steps:
  - a. Ensure you are logged in to OpenShift with the **oc** command and that your project is active.
  - b. Complete the steps documented in [Registry Service Accounts for Shared Environments](#). You must log in to the Red Hat Customer Portal to access the document and to complete the steps to create a registry service account.
  - c. Select the **OpenShift Secret** tab and click the link under **Download secret** to download the YAML secret file.
  - d. View the downloaded file and note the name that is listed in the **name:** entry.



- e. Enter 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.

- f. Download the **rhcam-7.8.0-openshift-templates.zip** product deliverable file from the [Software Downloads](#) page and extract the **rhcam78-image-streams.yaml** file.
- g. Enter the following command:

```
$ oc apply -f rhcam78-image-streams.yaml
```



#### NOTE

If you complete these steps, you install the image streams into the namespace of your project. In this case, when you deploy the templates, you must set the **IMAGE\_STREAM\_NAMESPACE** parameter to the name of this project.

## 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 the [Secrets chapter](#) in the Red Hat OpenShift Container Platform documentation.

You must create an SSL certificate for HTTP access to 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 Process Automation Manager configuration is **jboss**.
4. Record the password of the keystore file. The default value for this name in Red Hat Process Automation 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

You must create an SSL certificate for HTTP access to 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 Process Automation Manager configuration is **jboss**.
4. Record the password of the keystore file. The default value for this name in Red Hat Process Automation Manager configuration is **mykeystorepass**.
5. 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
```

### 3.4. CREATING THE SECRET FOR THE ADMINISTRATIVE USER

You must create a generic secret that contains the user name and password for a Red Hat Process Automation Manager administrative user account. This secret is required for deploying Red Hat Process Automation Manager using any template except the trial template.

The secret must contain the user name and password as literals. The key name for the user name is **KIE\_ADMIN\_USER**. The key name for the password is **KIE\_ADMIN\_PWD**.

If you are using multiple templates to deploy components of Red Hat Process Automation Manager, use the same secret for all these deployments. The components utilize this user account to communicate with each other.

You can also use this user account to log in to Business Central.



#### IMPORTANT

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

## Procedure

Use the **oc** command to generate a generic secret named **kie-admin-user-secret** from the user name and password:

```
$ oc create secret generic rhpam-credentials --from-literal=KIE_ADMIN_USER=adminUser --from-literal=KIE_ADMIN_PWD=adminPassword
```

In this command, replace *adminPassword* with the password for the administrative user. Optionally, you can replace *adminUser* with another user name for the administrative user.

## 3.5. CHANGING GLUSTERFS CONFIGURATION

You must check whether your OpenShift environment uses GlusterFS to provide permanent storage volumes. If it uses GlusterFS, to ensure optimal performance of Business Central, you must tune your GlusterFS storage by changing the storage class configuration.

### Procedure

1. To check whether your environment uses GlusterFS, enter the following command:

```
oc get storageclass
```

In the results, check whether the **(default)** marker is on the storage class that lists **glusterfs**. For example, in the following output the default storage class is **gluster-container**, which does list **glusterfs**:

```
NAME             PROVISIONER             AGE
gluster-block    gluster.org/glusterblock 8d
gluster-container (default) kubernetes.io/glusterfs 8d
```

If the result has a default storage class that does not list **glusterfs** or if the result is empty, you do not need to make any changes. In this case, skip the rest of this procedure.

2. To save the configuration of the default storage class into a YAML file, enter the following command:

```
oc get storageclass <class-name> -o yaml >storage_config.yaml
```

Replace **<class-name>** with the name of the default storage class. Example:

```
oc get storageclass gluster-container -o yaml >storage_config.yaml
```

3. Edit the **storage\_config.yaml** file:
  - a. Remove the lines with the following keys:
    - **creationTimestamp**
    - **resourceVersion**
    - **selfLink**
    - **uid**

- b. If you are planning to use Business Central only as a single pod, without high-availability configuration, on the line with the **volumeoptions** key, add the following options:

```
features.cache-invalidation on
performance.nl-cache on
```

For example:

**volumeoptions: client.ssl off, server.ssl off, features.cache-invalidation on, performance.nl-cache on**

- c. If you are planning to use Business Central in a high-availability configuration, on the line with the **volumeoptions** key, add the following options:

```
features.cache-invalidation on
nfs.trusted-write on
nfs.trusted-sync on
performance.nl-cache on
performance.stat-prefetch off
performance.read-ahead off
performance.write-behind off
performance.readdir-ahead off
performance.io-cache off
performance.quick-read off
performance.open-behind off
locks.mandatory-locking off
performance.strict-o-direct on
```

For example:

**volumeoptions: client.ssl off, server.ssl off, features.cache-invalidation on, nfs.trusted-write on, nfs.trusted-sync on, performance.nl-cache on, performance.stat-prefetch off, performance.read-ahead off, performance.write-behind off, performance.readdir-ahead off, performance.io-cache off, performance.quick-read off, performance.open-behind off, locks.mandatory-locking off, performance.strict-o-direct on**

4. To remove the existing default storage class, enter the following command:

```
oc delete storageclass <class-name>
```

Replace **<class-name>** with the name of the default storage class. Example:

```
oc delete storageclass gluster-container
```

5. To re-create the storage class using the new configuration, enter the following command:

```
oc create -f storage_config.yaml
```

### 3.6. PROVISIONING PERSISTENT VOLUMES WITH READWRITE MANY ACCESS MODE USING NFS

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



#### NOTE

If you want to deploy a high-availability authoring environment, for optimal performance and reliability, provision persistent volumes using GlusterFS. Configure the GlusterFS storage class as described in [Section 3.5, "Changing GlusterFS configuration"](#).

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 [Configuring Clusters](#) guide in the Red Hat OpenShift Container Platform 3.11 documentation.

### 3.7. 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. 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 **rhpm-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 **rhpm-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.

### 3.8. BUILDING A CUSTOM KIE SERVER EXTENSION IMAGE FOR AN EXTERNAL DATABASE

If you want to use an external database server for a KIE Server and the database server is not a MySQL or PostgreSQL server, you must build a custom KIE Server extension image with drivers for this server before deploying your environment.

Complete the steps in this build procedure to provide drivers for any of the following database servers:

- Microsoft SQL Server
- IBM DB2
- Oracle Database
- Sybase

Optionally, you can use this procedure to build a new version of drivers for any of the following database servers:

- MySQL
- MariaDB
- PostgreSQL

For the supported versions of the database servers, see [Red Hat Process Automation Manager 7 Supported Configurations](#).

The build procedure creates a custom extension image that extends the existing KIE Server image. You must import this custom extension image into your OpenShift environment and then reference it in the **EXTENSIONS\_IMAGE** parameter.

### Prerequisites

- You are logged in to your OpenShift environment using the **oc** command. Your OpenShift user must have the **registry-editor** role.
- For Oracle Database, IBM DB2, or Sybase, you downloaded the JDBC driver from the database server vendor.
- You have installed the following required software:
  - Docker: For installation instructions, see [Get Docker](#).
  - Cekit version 3.2: For installation instructions, see [Installation](#).
  - The following libraries and extensions for Cekit. For more information, see [Dependencies](#).
    - **docker**, provided by the **python3-docker** package or similar package
    - **docker-squash**, provided by the **python3-docker-squash** package or similar package
    - **behave**, provided by the **python3-behave** package or similar package

### Procedure

1. For IBM DB2, Oracle Database, or Sybase, provide the JDBC driver JAR file in a local directory.
2. Download the **rhcam-7.8.0-openshift-templates.zip** product deliverable file from the [Software Downloads](#) page of the Red Hat Customer Portal.
3. Unzip the file and, using the command line, change to the **templates/contrib/jdbc/cekit** directory of the unzipped file. This directory contains the source code for the custom build.
4. Enter one of the following commands, depending on the database server type:
  - For Microsoft SQL Server:

```
make mssql
```
  - For MySQL:

```
make mysql
```
  - For PostgreSQL:

```
-
```

```
make postgresql
```

- For MariaDB:

```
make mariadb
```

- For IBM DB2:

```
make db2 artifact=/tmp/db2jcc4.jar version=10.2
```

In this command, replace **/tmp/db2jcc4.jar** with the path name of the IBM DB2 driver and **10.2** with the version of the driver.

- For Oracle Database:

```
make oracle artifact=/tmp/ojdbc7.jar version=7.0
```

In this command, replace **/tmp/ojdbc7.jar** with the path name of the Oracle Database driver and **7.0** with the version of the driver.

- For Sybase:

```
make build sybase artifact=/tmp/jconn4-16.0_PL05.jar version=16.0_PL05
```

In this command, replace **/tmp/jconn4-16.0\_PL05.jar** with the path name of the downloaded Sybase driver and **16.0\_PL05** with the version of the driver.

Alternatively, if you need to update the driver class or driver XA class for the Sybase driver, you can set the **DRIVER\_CLASS** or **DRIVER\_XA\_CLASS** variable for this command, for example:

```
export DRIVER_CLASS=another.class.Sybase && make sybase artifact=/tmp/jconn4-16.0_PL05.jar version=16.0_PL05
```

5. Enter the following command to list the Docker images that are available locally:

```
docker images
```

Note the name of the image that was built, for example, **jboss-kie-db2-extension-openshift-image**, and the version tag of the image, for example, **11.1.4.4** (not the **latest** tag).

6. Access the registry of your OpenShift environment directly and push the image to the registry. Depending on your user permissions, you can push the image into the **openshift** namespace or into a project namespace. For instructions about accessing the registry and pushing the images, see [Accessing the Registry Directly](#) in the Red Hat OpenShift Container Platform product documentation.
7. When configuring your KIE Server deployment with a template that supports an external database server, set the following parameters:
  - **Drivers Extension Image (EXTENSIONS\_IMAGE)**: The ImageStreamTag definition of the extension image, for example, **jboss-kie-db2-extension-openshift-image:11.1.4.4**



- **Drivers ImageStream Namespace (`EXTENSIONS_IMAGE_NAMESPACE`):** The namespace to which you uploaded the extension image, for example, **openshift** or your project namespace.

## CHAPTER 4. AUTHORIZING ENVIRONMENT

You can deploy an environment for creating and modifying processes using Business Central. It consists of Business Central for the authoring work and KIE Server for test execution of the processes. If necessary, you can connect additional KIE Servers to the Business Central.

Depending on your needs, you can deploy either a single authoring environment template or a high-availability (HA) authoring environment template.

A single authoring environment contains two pods. One of the pods runs Business Central, the other runs KIE Server. The KIE Server by default includes an embedded H2 database engine. This environment is most suitable for single-user authoring or when your OpenShift infrastructure has limited resources. It does not require persistent volumes that support the **ReadWriteMany** access mode.

In a single authoring environment, you cannot scale Business Central. By default, you also cannot scale KIE Server, as the H2 database engine does not support scaling. However, you can modify the template to use a separate MySQL or PostgreSQL database server pod; in this case, you can scale KIE Server. For instructions about modifying the single authoring environment template, see [Section 4.6, "Modifying the template for the single authoring environment"](#).

In an HA authoring environment, both Business Central and KIE Server are provided in scalable pods. When pods are scaled, persistent storage is shared between the copies. The database is provided by a separate pod.

To enable high-availability functionality in Business Central, additional pods with AMQ and Data Grid are required. These pods are configured and deployed by the high-availability authoring template. Use a high-availability authoring environment to provide maximum reliability and responsiveness, especially if several users are involved in authoring at the same time.

In the current version of Red Hat Process Automation Manager, an HA authoring environment is supported with certain limitations:

- If a Business Central pod crashes while a user works with it, the user can get an error message and then is redirected to another pod. Logging on again is not required.
- If a Business Central pod crashes during a user operation, data that was not committed (saved) might be lost.
- If a Business Central pod crashes during creation of a project, an unusable project might be created.
- If a Business Central pod crashes during creation of an asset, the asset might be created but not indexed, so it cannot be used. The user can open the asset in Business Central and save it again to make it indexed.
- When a user deploys a service to the KIE Server, the KIE Server deployment is rolled out again. Users can not deploy another service to the same KIE Server until the roll-out completes.

In a high-availability authoring environment you can also deploy additional managed or immutable KIE Servers, if required. Business Central can automatically discover any KIE Servers in the same namespace, including immutable KIE Servers and managed KIE Servers.

If you want to deploy additional managed or immutable KIE Servers in a single authoring environment, you must complete an additional manual step to enable the **OpenShiftStartupStrategy** setting in the environment, as described in [Section 4.5, "Enabling the \*\*OpenShiftStartupStrategy\*\* setting to connect additional KIE Servers to Business Central"](#). This setting enables the discovery of other KIE Servers.

For instructions about deploying managed KIE Servers, see [Deploying a Red Hat Process Automation Manager freeform managed server environment on Red Hat OpenShift Container Platform](#). For instructions about deploying immutable KIE Servers, see [Deploying a Red Hat Process Automation Manager immutable server environment on Red Hat OpenShift Container Platform](#).

## 4.1. DEPLOYING AN AUTHORIZING ENVIRONMENT

You can use OpenShift templates to deploy a single or high-availability authoring environment. This environment consists of Business Central and a single KIE Server.

### 4.1.1. Starting configuration of the template for an authoring environment

If you want to deploy a single authoring environment, use the **rhcam78-authoring.yaml** template file. By default, the single authoring template uses the H2 database with permanent storage. If you prefer to create a MySQL or PostgreSQL pod or to use an external database server (outside the OpenShift project), modify the template before deploying the environment. For instructions about modifying the template, see [Section 4.6, “Modifying the template for the single authoring environment”](#).

If you want to deploy a high-availability authoring environment, use the **rhcam78-authoring-ha.yaml** template file. By default, the high-availability authoring template creates a MySQL pod to provide the database server for the KIE Server. If you prefer to use PostgreSQL or to use an external server (outside the OpenShift project) you need to modify the template before deploying the environment. You can also modify the template to change the number of replicas initially created for Business Central. For instructions about modifying the template, see [Section 4.7, “Modifying the template for the High Availability authoring environment”](#).

#### Procedure

1. Download the **rhcam-7.8.0-openshift-templates.zip** product deliverable file from the [Software Downloads](#) page of the Red Hat Customer Portal.
2. Extract the required template file.
3. Use one of the following methods to start deploying the template:
  - To use the OpenShift Web UI, in the OpenShift application console select **Add to Project** → **Import YAML / JSON** and then select or paste the **<template-file-name>.yaml** file. In the **Add Template** window, ensure **Process the template** is selected and click **Continue**.
  - To use the OpenShift command line console, prepare the following command line:

```
oc new-app -f <template-path>/<template-file-name>.yaml -p
BUSINESS_CENTRAL_HTTPS_SECRET=businesscentral-app-secret -p
KIE_SERVER_HTTPS_SECRET=kieserver-app-secret -p PARAMETER=value
```

In this command line, make the following changes:

- Replace **<template-path>** with the path to the downloaded template file.
- Replace **<template-file-name>** with the name of the template file.
- Use as many **-p PARAMETER=value** pairs as needed to set the required parameters.

#### Next steps

Set the parameters for the template. Follow the steps in [Section 4.1.2, “Setting required parameters for an authoring environment”](#) to set common parameters. You can view the template file to see descriptions for all parameters.

## 4.1.2. Setting required parameters for an authoring environment

When configuring the template to deploy an authoring environment, you must set the following parameters in all cases.

### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, “Starting configuration of the template for an authoring environment”](#).

### Procedure

1. Set the following parameters:

- **Credentials secret (CREDENTIALS\_SECRET)**: The name of the secret containing the administrative user credentials, as created in [Section 3.4, “Creating the secret for the administrative user”](#).
- **Business Central Server Keystore Secret Name (BUSINESS\_CENTRAL\_HTTPS\_SECRET)**: The name of the secret for Business Central, as created in [Section 3.3, “Creating the secrets for Business Central”](#).
- **KIE Server Keystore Secret Name (KIE\_SERVER\_HTTPS\_SECRET)**: The name of the secret for KIE Server, as created in [Section 3.2, “Creating the secrets for KIE Server”](#).
- **Business Central Server Certificate Name (BUSINESS\_CENTRAL\_HTTPS\_NAME)**: The name of the certificate in the keystore that you created in [Section 3.3, “Creating the secrets for Business Central”](#).
- **Business Central Server Keystore Password (BUSINESS\_CENTRAL\_HTTPS\_PASSWORD)**: The password for the keystore that you created in [Section 3.3, “Creating the secrets for Business Central”](#).
- **KIE Server Certificate Name (KIE\_SERVER\_HTTPS\_NAME)**: The name of the certificate in the keystore that you created in [Section 3.2, “Creating the secrets for KIE Server”](#).
- **KIE Server Keystore Password (KIE\_SERVER\_HTTPS\_PASSWORD)**: The password for the keystore that you created in [Section 3.2, “Creating the secrets for KIE Server”](#).
- **Application Name (APPLICATION\_NAME)**: The name of the OpenShift application. It is used in the default URLs for Business Central Monitoring and KIE Server. OpenShift uses the application name to create a separate set of deployment configurations, services, routes, labels, and artifacts.
- **ImageStream Namespace (IMAGE\_STREAM\_NAMESPACE)**: The namespace where the image streams are available. If the image streams were already available in your OpenShift environment (see [Section 3.1, “Ensuring the availability of image streams and the image registry”](#)), the namespace is **openshift**. If you have installed the image streams file, the namespace is the name of the OpenShift project.

### Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, “Completing deployment of the template for an authoring environment”](#).

### 4.1.3. Configuring the image stream namespace for an authoring environment

If you created image streams in a namespace that is not **openshift**, you must configure the namespace in the template.

If all image streams were already available in your Red Hat OpenShift Container Platform environment, you can skip this procedure.

#### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, “Starting configuration of the template for an authoring environment”](#).

#### Procedure

If you installed an image streams file according to instructions in [Section 3.1, “Ensuring the availability of image streams and the image registry”](#), set the **ImageStream Namespace (IMAGE\_STREAM\_NAMESPACE)** parameter to the name of your OpenShift project.

### 4.1.4. Setting an optional Maven repository for an authoring environment

When configuring the template to deploy an authoring environment, if you want to place the built KJAR files into an external Maven repository, you must set parameters to access the repository.

#### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, “Starting configuration of the template for an authoring environment”](#).

#### Procedure


To configure access to a custom Maven repository, set the following parameters:

- **Maven repository URL (MAVEN\_REPO\_URL)**: The URL for the Maven repository.
- **Maven repository ID (MAVEN\_REPO\_ID)**: An identifier for the Maven repository. The default value is **repo-custom**.
- **Maven repository username (MAVEN\_REPO\_USERNAME)**: The user name for the Maven repository.
- **Maven repository password (MAVEN\_REPO\_PASSWORD)**: The password for the Maven repository.

#### Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, “Completing deployment of the template for an authoring environment”](#).



## IMPORTANT

To export or push Business Central projects as KJAR artifacts to the external Maven repository, you must also add the repository information in the **pom.xml** file for every project. For information about exporting Business Central projects to an external repository, see [Packaging and deploying a Red Hat Process Automation Manager project](#) .

### 4.1.5. Configuring access to a Maven mirror in an environment without a connection to the public Internet for an authoring environment

When configuring the template to deploy an authoring environment, if your OpenShift environment does not have a connection to the public Internet, you must configure access to a Maven mirror that you set up according to [Section 3.7, "Preparing a Maven mirror repository for offline use"](#) .

#### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, "Starting configuration of the template for an authoring environment"](#).

#### Procedure

To configure access to the Maven mirror, set the following parameters:

- **Maven mirror URL (MAVEN\_MIRROR\_URL)**: The URL for the Maven mirror repository that you set up in [Section 3.7, "Preparing a Maven mirror repository for offline use"](#) . This URL must be accessible from a pod in your OpenShift environment.
- **Maven mirror of (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:\*,!repo-rhpamcentr**; with this value, Maven retrieves artifacts from the built-in Maven repository of Business Central directly and retrieves any other required artifacts from the mirror. If you configure an external Maven repository (**MAVEN\_REPO\_URL**), change **MAVEN\_MIRROR\_OF** to exclude the artifacts in this repository, for example, **external:\*,!repo-custom**. Replace **repo-custom** with the ID that you configured in **MAVEN\_REPO\_ID**.

#### Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, "Completing deployment of the template for an authoring environment"](#).

### 4.1.6. Configuring Business Central and KIE Server replicas for a high-availability authoring environment

If you are deploying a high-availability authoring environment, by default two replicas of Business Central and two replicas of the KIE Server are initially created.

Optionally, you can modify the number of replicas.

Skip this procedure for a single authoring environment.

#### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, “Starting configuration of the template for an authoring environment”](#).

## Procedure

To modify the numbers of initial replicas, set the following parameters:

- **Business Central Container Replicas**(**BUSINESS\_CENTRAL\_CONTAINER\_REPLICAS**): The number of replicas that the deployment initially creates for Business Central.
- **KIE Server Container Replicas**(**KIE\_SERVER\_CONTAINER\_REPLICAS**): The number of replicas that the deployment initially creates for the KIE Server.

## Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, “Completing deployment of the template for an authoring environment”](#).

## 4.1.7. Specifying the Git hooks directory for 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 configure a Git hooks directory.

### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, “Starting configuration of the template for an authoring environment”](#).

## Procedure

To configure a Git hooks directory, set the following parameter:

- **Git hooks directory** (**GIT\_HOOKS\_DIR**): The fully qualified path to a Git hooks directory, for example, **/opt/kie/data/git/hooks**. You must provide the content of this directory and mount it at the specified path. For instructions about providing and mounting the Git hooks directory using a configuration map or a persistent volume, see [Section 4.2, “\(Optional\) Providing the Git hooks directory”](#).

## Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, “Completing deployment of the template for an authoring environment”](#).

## 4.1.8. Configuring resource usage for a high-availability deployment

If you are deploying the high-availability template (**rhcam78-authoring-ha.yaml**), you can optionally configure resource usage to optimize performance for your requirements.

If you are deploying the single authoring environment template (**rhcam78-authoring.yaml**), skip this procedure.

For more information about sizing resources, see the following sections in the Red Hat OpenShift Container Platform 3.11 product documentation:

- [Application memory sizing](#)
- [Compute resources](#)

### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, “Starting configuration of the template for an authoring environment”](#).


### Procedure

Set the following parameters of the template as applicable:

- **Business Central Container Memory Limit(BUSINESS\_CENTRAL\_MEMORY\_LIMIT):** The amount of memory requested in the OpenShift environment for the Business Central container. The default value is **8Gi**.
- **Business Central JVM Max Memory Ratio (BUSINESS\_CENTRAL\_JAVA\_MAX\_MEM\_RATIO):** The percentage of container memory that is used for the Java Virtual Machine for Business Central. The remaining memory is used for the operating system. The default value is **80**, for a limit of 80%.
- **Business Central Container CPU Limit(BUSINESS\_CENTRAL\_CPU\_LIMIT):** The maximum CPU usage for Business Central. The default value is **2000m**.
- **KIE Server Container Memory Limit(KIE\_SERVER\_MEMORY\_LIMIT):** The amount of memory requested in the OpenShift environment for the KIE Server container. The default value is **1Gi**.
- **KIE Server Container CPU Limit(KIE\_SERVER\_CPU\_LIMIT):** The maximum CPU usage for KIE Server. The default value is **1000m**.
- **DataGrid Container Memory Limit(DATAGRID\_MEMORY\_LIMIT):** The amount of memory requested in the OpenShift environment for the Red Hat Data Grid container. The default value is **2Gi**.
- **DataGrid Container CPU Limit(DATAGRID\_CPU\_LIMIT):** The maximum CPU usage for Red Hat Data Grid. The default value is **1000m**.

### 4.1.9. Setting parameters for RH-SSO authentication for an authoring environment

If you want to use RH-SSO authentication, complete the following additional configuration when configuring the template to deploy an authoring environment.



#### IMPORTANT

Do not configure LDAP authentication and RH-SSO authentication in the same deployment.

### Prerequisites

- A realm for Red Hat Process Automation Manager is created in the RH-SSO authentication system.



- User names and passwords for Red Hat Process Automation Manager are created in the RH-SSO authentication system. For a list of the available roles, see [Chapter 5, Red Hat Process Automation Manager roles and users](#).  
You must create a user with the username and password configured in the secret for the administrative user, as described in [Section 3.4, “Creating the secret for the administrative user”](#). This user must have the **kie-server,rest-all,admin** roles.
- Clients are created in the RH-SSO authentication system for all components of the Red Hat Process Automation Manager environment that you are deploying. The client setup contains the URLs for the components. You can review and edit the URLs after deploying the environment. Alternatively, the Red Hat Process Automation Manager deployment can create the clients. However, this option provides less detailed control over the environment.
- You started the configuration of the template, as described in [Section 4.1.1, “Starting configuration of the template for an authoring environment”](#).

## Procedure

1. Set the following parameters:
  - **RH-SSO URL (SSO\_URL)**: The URL for RH-SSO.
  - **RH-SSO Realm name (SSO\_REALM)**: The RH-SSO realm for Red Hat Process Automation Manager.
  - **RH-SSO Disable SSL Certificate Validation (SSO\_DISABLE\_SSL\_CERTIFICATE\_VALIDATION)**: Set to **true** if your RH-SSO installation does not use a valid HTTPS certificate.
2. Complete one of the following procedures:
  - a. If you created the clients for Red Hat Process Automation Manager within RH-SSO, set the following parameters in the template:
    - **Business Central RH-SSO Client name (BUSINESS\_CENTRAL\_SSO\_CLIENT)**: The RH-SSO client name for Business Central.
    - **Business Central RH-SSO Client Secret (BUSINESS\_CENTRAL\_SSO\_SECRET)**: The secret string that is set in RH-SSO for the client for Business Central.
    - **KIE Server RH-SSO Client name (KIE\_SERVER\_SSO\_CLIENT)**: The RH-SSO client name for KIE Server.
    - **KIE Server RH-SSO Client Secret (KIE\_SERVER\_SSO\_SECRET)**: The secret string that is set in RH-SSO for the client for KIE Server.
  - b. To create the clients for Red Hat Process Automation Manager within RH-SSO, set the following parameters in the template:
    - **Business Central RH-SSO Client name (BUSINESS\_CENTRAL\_SSO\_CLIENT)**: The name of the client to create in RH-SSO for Business Central.
    - **Business Central RH-SSO Client Secret (BUSINESS\_CENTRAL\_SSO\_SECRET)**: The secret string to set in RH-SSO for the client for Business Central.
    - **KIE Server RH-SSO Client name (KIE\_SERVER\_SSO\_CLIENT)**: The name of the client to create in RH-SSO for KIE Server.

- **KIE Server RH-SSO Client Secret**(**KIE\_SERVER\_SSO\_SECRET**): The secret string to set in RH-SSO for the client for KIE Server.
- **RH-SSO Realm Admin Username**(**SSO\_USERNAME**) and **RH-SSO Realm Admin Password** (**SSO\_PASSWORD**): The user name and password for the realm administrator user for the RH-SSO realm for Red Hat Process Automation Manager. You must provide this user name and password in order to create the required clients.

## Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, "Completing deployment of the template for an authoring environment"](#).

After completing the deployment, review the URLs for components of Red Hat Process Automation Manager in the RH-SSO authentication system to ensure they are correct.

## 4.1.10. Setting parameters for LDAP authentication for an authoring environment

If you want to use LDAP authentication, complete the following additional configuration when configuring the template to deploy an authoring environment.



### IMPORTANT

Do not configure LDAP authentication and RH-SSO authentication in the same deployment.

## Prerequisites

- You created user names and passwords for Red Hat Process Automation Manager in the LDAP system. For a list of the available roles, see [Chapter 5, Red Hat Process Automation Manager roles and users](#).  
You must create a user with the username and password configured in the secret for the administrative user, as described in [Section 3.4, "Creating the secret for the administrative user"](#). This user must have the **kie-server,rest-all,admin** roles.
- You started the configuration of the template, as described in [Section 4.1.1, "Starting configuration of the template for an authoring environment"](#).

## Procedure

1. Set the **AUTH\_LDAP\*** parameters of the template. These parameters correspond to the settings of the **LdapExtended** Login module of Red Hat JBoss EAP. For instructions about using these settings, see [LdapExtended login module](#).  
If the LDAP server does not define all the roles required for your deployment, you can map LDAP groups to Red Hat Process Automation Manager roles. To enable LDAP role mapping, set the following parameters:
  - **RoleMapping rolesProperties file path** (**AUTH\_ROLE\_MAPPER\_ROLES\_PROPERTIES**): The fully qualified path name of a file that defines role mapping, for example, **/opt/eap/standalone/configuration/rolemapping/rolemapping.properties**. You must provide this file and mount it at this path in all applicable deployment configurations; for instructions, see [Section 4.4, "\(Optional\) Providing the LDAP role mapping file"](#).

- **RoleMapping replaceRole property(AUTH\_ROLE\_MAPPER\_REPLACE\_ROLE)**: If set to **true**, mapped roles replace the roles defined on the LDAP server; if set to **false**, both mapped roles and roles defined on the LDAP server are set as user application roles. The default setting is **false**.

## Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, "Completing deployment of the template for an authoring environment"](#).

## 4.1.11. Setting parameters for using an external database server for an authoring environment

If you modified the template to use an external database server for the KIE Server, as described in [Section 4.6, "Modifying the template for the single authoring environment"](#) or [Section 4.7, "Modifying the template for the High Availability authoring environment"](#), complete the following additional configuration when configuring the template to deploy an authoring environment.

### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, "Starting configuration of the template for an authoring environment"](#).

### Procedure

1. Set the following parameters:

- **KIE Server External Database Driver(KIE\_SERVER\_EXTERNALDB\_DRIVER)**: The driver for the server, depending on the server type:
  - **mysql**
  - **postgresql**
  - **mariadb**
  - **mssql**
  - **db2**
  - **oracle**
  - **sybase**
- **KIE Server External Database User(KIE\_SERVER\_EXTERNALDB\_USER)** and **KIE Server External Database Password (KIE\_SERVER\_EXTERNALDB\_PWD)**: The user name and password for the external database server
- **KIE Server External Database URL(KIE\_SERVER\_EXTERNALDB\_URL)**: The JDBC URL for the external database server



## NOTE

If you are using the EnterpriseDB Postgres database server, use an URL starting with **jdbc:postgresql://** and not with **jdbc:edb://**. Alternatively, do not set the URL and set the host and port parameters instead.

- **KIE Server External Database Host**(**KIE\_SERVER\_EXTERNALDB\_SERVICE\_HOST**) and **KIE Server External Database Port** (**KIE\_SERVER\_EXTERNALDB\_SERVICE\_PORT**): The host name and port number of the external database server. You can set these parameters as an alternative to setting the **KIE\_SERVER\_EXTERNALDB\_URL** parameter.
  - **KIE Server External Database Dialect**(**KIE\_SERVER\_EXTERNALDB\_DIALECT**): The Hibernate dialect for the server, depending on the server type. The common settings are:
    - **org.hibernate.dialect.MySQL5InnoDBDialect**
    - **org.hibernate.dialect.MySQL8Dialect**
    - **org.hibernate.dialect.MariaDB102Dialect**
    - **org.hibernate.dialect.PostgreSQL95Dialect**
    - **org.hibernate.dialect.PostgresPlusDialect** (used for EnterpriseDB Postgres Advanced Server)
    - **org.hibernate.dialect.SQLServer2012Dialect** (used for MS SQL)
    - **org.hibernate.dialect.DB2Dialect**
    - **org.hibernate.dialect.Oracle10gDialect**
    - **org.hibernate.dialect.SybaseASE15Dialect**  
For a complete list of supported dialects, see Table A.7 in [Hibernate properties](#) in the Red Hat JBoss EAP documentation.
  - **KIE Server External Database name**(**KIE\_SERVER\_EXTERNALDB\_DB**): The database name to use on the external database server
  - **JDBC Connection Checker class** (**KIE\_SERVER\_EXTERNALDB\_CONNECTION\_CHECKER**): The name of the JDBC connection checker class for the database server. Without this information, a database server connection cannot be restored after it is lost, for example, if the database server is rebooted.
  - **JDBC Exception Sorter class** (**KIE\_SERVER\_EXTERNALDB\_EXCEPTION\_SORTER**): The name of the JDBC exception sorter class for the database server. Without this information, a database server connection cannot be restored after it is lost, for example, if the database server is rebooted.
2. If you created a custom image for using an external database server, as described in [Section 3.8, “Building a custom KIE Server extension image for an external database”](#), set the following parameters:
- **Drivers Extension Image** (**EXTENSIONS\_IMAGE**): The ImageStreamTag definition of the extension image, for example, **jboss-kie-db2-extension-openshift-image:11.1.4.4**

- **Drivers ImageStream Namespace (EXTENSIONS\_IMAGE\_NAMESPACE)**: The namespace to which you uploaded the extension image, for example, **openshift** or your project namespace.
3. If you are using a MySQL version 8 external database server, enable the **mysql\_native\_password** plugin and use it for authentication. For instructions about this pluding, see [Native Pluggable Authentication](#) in the *MySQL 8.0 Reference Manual*. If you are using a MySQL version 8 image provided by Red Hat on Red Hat OpenShift Container Platform, to enable the plugin, set the **MYSQL\_DEFAULT\_AUTHENTICATION\_PLUGIN** environment variable to **mysql\_native\_password**.

If you created users on the MySQL version 8 server before enabling the **mysql\_native\_password** plugin, you must update the **mysql-user** table after you enable the plugin.

### Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, “Completing deployment of the template for an authoring environment”](#).

## 4.1.12. Enabling Prometheus metric collection for an authoring environment

If you want to configure your KIE Server deployment to use Prometheus to collect and store metrics, enable support for this feature in KIE Server at deployment time.

### Prerequisites

- You started the configuration of the template, as described in [Section 4.1.1, “Starting configuration of the template for an authoring environment”](#).

### Procedure

To enable support for Prometheus metric collection, set the **Prometheus Server Extension Disabled (PROMETHEUS\_SERVER\_EXT\_DISABLED)** parameter to **false**.

### Next steps

If necessary, set additional parameters.

To complete the deployment, follow the procedure in [Section 4.1.13, “Completing deployment of the template for an authoring environment”](#).

For instructions about configuring Prometheus metrics collection, see [Managing and monitoring KIE Server](#).

## 4.1.13. Completing deployment of the template for an authoring environment

After setting all the required parameters in the OpenShift Web UI or in the command line, complete deployment of the template.

### Procedure

Depending on the method that you are using, complete the following steps:

- In the OpenShift Web UI, click **Create**.

- If the **This will create resources that may have security or project behavior implications** message appears, click **Create Anyway**.
- Complete the command line and press Enter.

## 4.2. (OPTIONAL) PROVIDING THE GIT HOOKS DIRECTORY

If you configure the **GIT\_HOOKS\_DIR** parameter, you must provide a directory of Git hooks and must mount this directory on the Business Central deployment.

The typical use of Git hooks is interaction with an upstream repository. To enable Git hooks to push commits into an upstream repository, you must also provide a secret key that corresponds to a public key configured on the upstream repository.

### Procedure

1. If interaction with an upstream repository using SSH authentication is required, complete the following steps to prepare and mount 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
```


- d. Mount the secret in the SSH key path of the Business Central deployment, for example:

```
oc set volume dc/<myapp>-rhpamcentr --add --type secret --secret-name git-hooks-secret --mount-path=/home/jboss/.ssh --name=ssh-key
```

Replace **<myapp>** with the application name that you set when configuring the template.

2. Create the Git hooks directory. For instructions, see the [Git hooks reference documentation](#). For example, a simple Git hooks directory can provide a post-commit hook that pushes the changes upstream. If the project was imported into Business Central from a repository, this repository remains configured as the upstream repository. Create a file named **post-commit** with permission values **755** and the following content:

```
git push
```



### NOTE

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

3. Supply the Git hooks directory to the Business Central deployment. You can use a configuration map or a persistent volume.
  - a. If the Git hooks consist of one or several fixed script files, use a configuration map. Complete the following steps:

- i. Change into the Git hooks directory that you have created.
- ii. Create an OpenShift configuration map from the files in the directory. Run the following command:

```
oc create configmap git-hooks --from-file=<file_1>=<file_1> --from-file=<file_2>=<file_2> ...
```

Replace **file\_1**, **file\_2**, and so on with Git hook script file names. Example:

```
oc create configmap git-hooks --from-file=post-commit=post-commit
```

- iii. Mount the configuration map on the Business Central deployment in the path that you have configured:

```
oc set volume dc/<myapp>-rhpamcentr --add --type configmap --configmap-name git-hooks --mount-path=<git_hooks_dir> --name=git-hooks
```

Replace **<myapp>** with the application name that was set when configuring the template and **<git\_hooks\_dir>** is the value of **GIT\_HOOKS\_DIR** that was set when configuring the template.

- b. If the Git hooks consist of long files or depend on binaries, such as executable or KJAR files, use a persistence volume. You must create a persistent volume, create a persistent volume claim and associate the volume with the claim, transfer files to the volume, and mount the volume in the **myapp-rhpamcentr** deployment configuration (replace *myapp* with the application name). For instructions about creating and mounting persistence volumes, see [Using persistent volumes](#). For instructions about copying files onto a persistent volume, see [Transferring files in and out of containers](#).
4. Wait a few minutes, then review the list and status of pods in your project. Because Business Central does not start until you provide the Git hooks directory, the KIE Server might not start at all. To see if it has started, check the output of the following command:

```
oc get pods
```

If a working KIE Server pod is not present, start it:

```
oc rollout latest dc/<myapp>-kieserver
```

Replace **<myapp>** with the application name that was set when configuring the template.

### 4.3. (OPTIONAL) PROVIDING A TRUSTSTORE FOR ACCESSING HTTPS SERVERS WITH SELF-SIGNED CERTIFICATES

Components of your Red Hat Process Automation Manager infrastructure might need to use HTTPS access to servers that have a self-signed HTTPS certificate. For example, Business Central and KIE Server might need to interact with an internal Nexus repository that uses a self-signed HTTPS server certificate.

In this case, to ensure that HTTPS connections complete successfully, you must provide client certificates for these services using a truststore.

Skip this procedure if you do not need Red Hat Process Automation Manager components to communicate with servers that use self-signed HTTPS server certificates.

## Procedure

1. Prepare a truststore with the certificates. Use the following command to create a truststore or to add a certificate to an existing truststore. Add all the necessary certificates to one truststore.

```
keytool -importcert -file certificate-file -alias alias -keyalg algorithm -keysize size -
trustcacerts -noprompt -storetype JKS -keypass truststore-password -storepass truststore-
password -keystore keystore-file
```

Replace the following values:

- ***certificate-file***: The pathname of the certificate that you want to add to the truststore.
- ***alias***: The alias for the certificate in the truststore. If you are adding more than one certificate to the truststore, every certificate must have a unique alias.
- ***algorithm***: The encryption algorithm used for the certificate, typically **RSA**.
- ***size***: The size of the certificate key in bytes, for example, **2048**.
- ***truststore-password***: The password for the truststore.
- ***keystore-file***: The pathname of the truststore file. If the file does not exist, the command creates a new truststore.

The following example command adds a certificate from the **/var/certs/nexus.cer** file to a truststore in the **/var/keystores/custom-trustore.jks** file. The truststore password is **mykeystorepass**.

```
keytool -importcert -file /var/certs/nexus.cer -alias nexus-cert -keyalg RSA -keysize 2048
-trustcacerts -noprompt -storetype JKS -keypass mykeystorepass -storepass
mykeystorepass -keystore /var/keystores/custom-trustore.jks
```

2. Create a secret with the truststore file using the **oc** command, for example:

```
oc create secret generic truststore-secret --from-file=/var/keystores/custom-trustore.jks
```

3. In the deployment for the necessary components of your infrastructure, mount the secret and then set the **JAVA\_OPTS\_APPEND** option to enable the Java application infrastructure to use the trust store, for example:

```
oc set volume dc/myapp-rhpamcentr --add --overwrite --name=custom-trustore-volume --
mount-path /etc/custom-secret-volume --secret-name=custom-secret
```

```
oc set env dc/myapp-rhpamcentr JAVA_OPTS_APPEND='-
Djavax.net.ssl.trustStore=/etc/custom-secret-volume/custom-trustore.jks -
Djavax.net.ssl.trustStoreType=jks -Djavax.net.ssl.trustStorePassword=mykeystorepass'
```

```
oc set volume dc/myapp-kieserver --add --overwrite --name=custom-trustore-volume --
mount-path /etc/custom-secret-volume --secret-name=custom-secret
```



```
oc set env dc/myapp-kieserver JAVA_OPTS_APPEND='-
Djavax.net.ssl.trustStore=/etc/custom-secret-volume/custom-trustore.jks -
Djavax.net.ssl.trustStoreType=jks -Djavax.net.ssl.trustStorePassword=mykeystorepass'
```

Replace **myapp** with the application name that you set when configuring the template.

## 4.4. (OPTIONAL) PROVIDING THE LDAP ROLE MAPPING FILE

If you configure the **AUTH\_ROLE\_MAPPER\_ROLES\_PROPERTIES** parameter, you must provide a file that defines the role mapping. Mount this file on all affected deployment configurations.

### Procedure

1. Create the role mapping properties file, for example, **my-role-map**. The file must contain entries in the following format:

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

For example:

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

2. Create an OpenShift configuration map from the file by entering the following command:

```
oc create configmap ldap-role-mapping --from-file=<new_name>=<existing_name>
```

Replace **<new\_name>** with the name that the file is to have on the pods (it must be the same as the name specified in the **AUTH\_ROLE\_MAPPER\_ROLES\_PROPERTIES** file) and **<existing\_name>** with the name of the file that you created. Example:

```
oc create configmap ldap-role-mapping --from-file=rolemapping.properties=my-role-map
```

3. Mount the configuration map on every deployment configuration that is configured for role mapping.

The following deployment configurations can be affected in this environment:

- **myapp-rhpamcentr**: Business Central
- **myapp-kieserver**: KIE Server

Replace **myapp** with the application name. Sometimes, several KIE Server deployments can be present under different application names.

For every deployment configuration, run the command:

```
oc set volume dc/<deployment_config_name> --add --type configmap --configmap-name
ldap-role-mapping --mount-path=<mapping_dir> --name=ldap-role-mapping
```

Replace **<mapping\_dir>** with the directory name (without file name) set in the **AUTH\_ROLE\_MAPPER\_ROLES\_PROPERTIES** parameter, for example, **/opt/eap/standalone/configuration/rolemapping**.

## 4.5. ENABLING THE OPENSIFTSTARTUPSTRATEGY SETTING TO CONNECT ADDITIONAL KIE SERVERS TO BUSINESS CENTRAL

In an environment deployed using Red Hat Process Automation Manager authoring templates, Business Central manages one KIE Server. If you use the high-availability authoring template or if you modified the single authoring template to use a database server other than an embedded H2 database, you can scale the KIE Server pod, but all the copies execute the same services.

You can connect additional KIE Servers to Business Central. However, if you deployed a single authoring environment using the **rhcam78-authoring.yaml**, you must enable the **OpenShiftStartupStrategy** setting in the environment. When **OpenShiftStartupStrategy** is enabled, Business Central automatically discovers KIE Servers in the same namespace and these KIE Servers can be configured to connect to the Business Central.

With the **OpenShiftStartupStrategy** setting, when a user deploys a service to the KIE Server, the KIE Server deployment is rolled out again. Users can not deploy another service to the same KIE Server until the roll-out completes. Because the roll-out might take noticeable time, the **OpenShiftStartupStrategy** setting might not be suitable for some authoring environments.

Do not complete this procedure if you deployed a high-availability authoring environment using the **rhcam78-authoring-ha.yaml** template. In this environment, the **OpenShiftStartupStrategy** setting is enabled by default.

Do not complete this procedure unless you want to connect additional KIE Servers to Business Central.

### Prerequisites

- You deployed an authoring environment using the **rhcam78-authoring.yaml** template.
- You are logged in to the OpenShift project where the environment is deployed using the **oc** tool.

### Procedure

1. Enter the following command to view the deployment configurations that are deployed in the project:

```
$ oc get dc
```

2. In the output of the command, find the deployment configuration names for the Business Central and KIE Server pods:
  - The name of the deployment configuration for Business Central is **myapp-rhcamcentr**. Replace **myapp** with the application name of the environment, which is set in the **APPLICATION\_NAME** parameter of the template.
  - The name of the deployment configuration for KIE Server is **myapp-kieserver**. Replace **myapp** with the application name.
3. Enter the following commands to enable the **OpenShiftStartupStrategy** setting on the pods:

```
$ oc env myapp-rhcamcentr KIE_SERVER_CONTROLLER_OPENSIFT_ENABLED=true
$ oc env myapp-kieserver KIE_SERVER_STARTUP_STRATEGY=OpenShiftStartupStrategy
```

In these commands, replace **myapp-rhpamcentr** with the Business Central deployment configuration name and **myapp-kieserver** with the KIE Server deployment configuration name.

4. When you enable the **OpenShiftStartupStrategy** setting, by default Business Central discovers only KIE Servers that are deployed with the same value of the **APPLICATION\_NAME** parameter as the authoring template. If you want to connect KIE Servers with any other application names to the Business Central, enter the following command:

```
$ oc env myapp-rhpamcentr
KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED=true
```

In this command, replace **myapp-rhpamcentr** with the Business Central deployment configuration name.

## 4.6. MODIFYING THE TEMPLATE FOR THE SINGLE AUTHORIZING ENVIRONMENT

By default, the single authoring template uses the H2 database with permanent storage. If you prefer to create a MySQL or PostgreSQL pod or to use an external database server (outside the OpenShift project), modify the template before deploying the environment.

You must use a MySQL or PostgreSQL pod or an external database server if you want to scale the KIE Server pod. An OpenShift template defines a set of objects that can be created by OpenShift. To change an environment configuration, you need to modify, add, or delete these objects. To simplify this task, comments are provided in the Red Hat Process Automation Manager templates.

Some comments mark blocks within the template, starting with **BEGIN** and ending with **END**. For example, the following block is named **Sample block**:

```
## Sample block BEGIN
sample line 1
sample line 2
sample line 3
## Sample block END
```


For some changes, you might need to replace a block in one template file with a block from another template file provided with Red Hat Process Automation Manager. In this case, delete the block, then paste the new block in its exact location.

### Procedure

Edit the **rhpam78-authoring.yaml** template file to make any of the following changes as necessary.

- If you want to use MySQL instead of the H2 database, you need to replace several blocks of the file, marked with comments from **BEGIN** to **END**, with blocks from the **rhpam78-kieserver-mysql.yaml** file that are also marked with comments. You also need to remove several other blocks and to add blocks in designated locations:
  1. Replace the block named **H2 database parameters** with the block named **MySQL database parameters**. (Take this block and all subsequent replacement blocks from the **rhpam78-kieserver-mysql.yaml** file.)
  2. Replace the block named **H2 driver settings** with the block named **MySQL driver settings**.
  3. Replace the block named **H2 persistent volume claim** with the block named **MySQL persistent volume claim**.

4. Remove the blocks named **H2 volume mount** and **H2 volume settings**.
  5. Under the comment **Place to add database service**, add the block named **MySQL service**.
  6. Under the comment **Place to add database deployment config**, add the block named **MySQL deployment config**.
- If you want to use PostgreSQL instead of the H2 database, you need to replace several blocks of the file, marked with comments from **BEGIN** to **END**, with blocks from the **rhcam78-kieserver-postgresql.yaml** file that are also marked with comments. You also need to remove several other blocks and to add blocks in designated locations:
    1. Replace the block named **H2 database parameters** with the block named **PostgreSQL database parameters**. (Take this block and all subsequent replacement blocks from the **rhcam78-kieserver-postgresql.yaml** file.)
    2. Replace the block named **H2 driver settings** with the block named **PostgreSQL driver settings**.
    3. Replace the block named **H2 persistent volume claim** with the block named **PostgreSQL persistent volume claim**.
    4. Remove the blocks named **H2 volume mount** and **H2 volume settings**.
    5. Under the comment **Place to add database service**, add the block named **PostgreSQL service**.
    6. Under the comment **Place to add database deployment config**, add the block named **PostgreSQL deployment config**.
  - If you want to use an external database server, replace several blocks of the file, marked with comments from **BEGIN** to **END**, with blocks from the **rhcam78-kieserver-externaldb.yaml** file, and also remove some blocks:
    1. Replace the block named **H2 database parameters** with the block named **External database parameters**. (Take this block and all subsequent replacement blocks from the **rhcam78-kieserver-externaldb.yaml** file.)
    2. Replace the block named **H2 driver settings** with the block named **External database driver settings**.
    3. Remove the following blocks of the file, marked with comments from **BEGIN** to **END**:
      - **H2 persistent volume claim**
      - **H2 volume mount**
      - **H2 volume settings**



## IMPORTANT

The standard KIE Server image includes drivers for MySQL, MariaDB, and PostgreSQL external database servers. If you want to use another database server, you must build a custom KIE Server image. For instructions, see [Section 3.8, “Building a custom KIE Server extension image for an external database”](#).

## 4.7. MODIFYING THE TEMPLATE FOR THE HIGH AVAILABILITY AUTHORIZING ENVIRONMENT

By default, the high-availability authoring template creates a MySQL pod to provide the database server for the KIE Server. If you prefer to use PostgreSQL or to use an external server (outside the OpenShift project), you need to modify the template before deploying the environment.

You can also modify the High Availability authoring template to change the number of replicas initially created for Business Central.

An OpenShift template defines a set of objects that can be created by OpenShift. To change an environment configuration, you need to modify, add, or delete these objects. To simplify this task, comments are provided in the Red Hat Process Automation Manager templates.

Some comments mark blocks within the template, starting with **BEGIN** and ending with **END**. For example, the following block is named **Sample block**:

```
## Sample block BEGIN
sample line 1
sample line 2
sample line 3
## Sample block END
```

For some changes, you might need to replace a block in one template file with a block from another template file provided with Red Hat Process Automation Manager. In this case, delete the block, then paste the new block in its exact location.

### Procedure

Edit the **rhpm78-authoring-ha.yaml** template file to make any of the following changes as necessary.

- If you want to use PostgreSQL instead of MySQL, replace several blocks of the file, marked with comments from **BEGIN** to **END**, with blocks from the **rhpm78-kieserver-postgresql.yaml** file:
  1. Replace the block named **MySQL database parameters** with the block named **PostgreSQL database parameters**. (Take this block and all subsequent replacement blocks from the **rhpm78-kieserver-postgresql.yaml** file.)
  2. Replace the block named **MySQL service** with the block named **PostgreSQL service**.
  3. Replace the block named **MySQL driver settings** with the block named **PostgreSQL driver settings**.
  4. Replace the block named **MySQL deployment config** with the block named **PostgreSQL deployment config**.
  5. Replace the block named **MySQL persistent volume claim** with the block named **PostgreSQL persistent volume claim**.
- If you want to use an external database server, replace several blocks of the file, marked with comments from **BEGIN** to **END**, with blocks from the **rhpm78-kieserver-externaldb.yaml** file, and also remove some blocks:
  1. Replace the block named **MySQL database parameters** with the block named **External database parameters**. (Take this block and all subsequent replacement blocks from the **rhpm78-kieserver-externaldb.yaml** file.)

2. Replace the block named **MySQL driver settings** with the block named **External database driver settings**.
3. Remove the following blocks of the file, marked with comments from **BEGIN** to **END**:
  - **MySQL service**
  - **MySQL deployment config**
  - **MySQL persistent volume claim**



### IMPORTANT

The standard KIE Server image includes drivers for MySQL, MariaDB, and PostgreSQL external database servers. If you want to use another database server, you must build a custom KIE Server image. For instructions, see [Section 3.8, “Building a custom KIE Server extension image for an external database”](#).

- If you want to change the number of replicas initially created for Business Central, on the line below the comment **## Replicas for Business Central**, change the number of replicas to the desired value.

## CHAPTER 5. RED HAT PROCESS AUTOMATION MANAGER ROLES AND USERS

To access Business Central or KIE Server, you must create users and assign them appropriate roles before the servers are started.

The Business Central and KIE Server use Java Authentication and Authorization Service (JAAS) login module to authenticate the users. If both Business Central and KIE Server are running on a single instance, then they share the same JAAS subject and security domain. Therefore, a user, who is authenticated for Business Central can also access KIE Server.

However, if Business Central and KIE Server are running on different instances, then the JAAS login module is triggered for both individually. Therefore, a user, who is authenticated for Business Central, needs to be authenticated separately to access the KIE Server (for example, to view or manage process definitions in Business Central). In case, the user is not authenticated on the KIE Server, then 401 error is logged in the log file, displaying **Invalid credentials to load data from remote server. Contact your system administrator.** message in Business Central.

This section describes available Red Hat Process Automation Manager user roles.



### NOTE

The **admin**, **analyst**, **developer**, **manager**, **process-admin**, **user**, and **rest-all** roles are reserved for Business Central. The **kie-server** role is reserved for KIE Server. For this reason, the available roles can differ depending on whether Business Central, KIE Server, or both are installed.

- **admin**: Users with the **admin** role are the Business Central administrators. They can manage users and create, clone, and manage the repositories. They have full access to make required changes in the application. Users with the **admin** role have access to all areas within Red Hat Process Automation Manager.
- **analyst**: Users with the **analyst** role have access to all high-level features. They can model and execute their projects. However, these users cannot add contributors to spaces or delete spaces in the **Design → Projects** view. Access to the **Deploy → Execution Servers** view, which is intended for administrators, is not available to users with the **analyst** role. However, the **Deploy** button is available to these users when they access the Library perspective.
- **developer**: Users with the **developer** role have access to almost all features and can manage rules, models, process flows, forms, and dashboards. They can manage the asset repository, they can create, build, and deploy projects, and they can use Red Hat CodeReady Studio to view processes. Only certain administrative functions such as creating and cloning a new repository are hidden from users with the **developer** role.
- **manager**: Users with the **manager** role can view reports. These users are usually interested in statistics about the business processes and their performance, business indicators, and other business-related reporting. A user with this role has access only to process and task reports.
- **process-admin**: Users with the **process-admin** role are business process administrators. They have full access to business processes, business tasks, and execution errors. These users can also view business reports and have access to the Task Inbox list.
- **user**: Users with the **user** role can work on the Task Inbox list, which contains business tasks that are part of currently running processes. Users with this role can view process and task reports and manage processes.

- **rest-all**: Users with the **rest-all** role can access Business Central REST capabilities.
- **kie-server**: Users with the **kie-server** role can access KIE Server (KIE Server) REST capabilities. This role is mandatory for users to have access to **Manage** and **Track** views in Business Central.



## CHAPTER 6. OPENSIFT TEMPLATE REFERENCE INFORMATION

Red Hat Process Automation Manager provides the following OpenShift templates. To access the templates, download and extract the **rhpm-7.8.0-openshift-templates.zip** product deliverable file from the [Software Downloads](#) page of the Red Hat customer portal.

- **rhpm78-authoring.yaml** provides a Business Central and a KIE Server connected to the Business Central. The KIE Server uses an H2 database with persistent storage. You can use this environment to author processes, services, and other business assets. For details about this template, see [Section 6.1, “rhpm78-authoring.yaml template”](#).
- **rhpm78-authoring-ha.yaml** provides a high-availability Business Central, a KIE Server connected to the Business Central, and a MySQL instance that the KIE Server uses. You can use this environment to author processes, services, and other business assets. For details about this template, see [Section 6.2, “rhpm78-authoring-ha.yaml template”](#).

### 6.1. RHPAM78-AUTHORING.YAML TEMPLATE

Application template for a non-HA persistent authoring environment, for Red Hat Process Automation Manager 7.8 - Deprecated

#### 6.1.1. Parameters

Templates allow you to define parameters that take on a value. That value is then substituted wherever the parameter is referenced. References can be defined in any text field in the objects list field. See the [Openshift documentation](#) for more information.

Variable name	Image Environment Variable	Description	Example value	Required
<b>APPLICATION_NAME</b>	–	The name for the application.	myapp	True
<b>CREDENTIALS_SECRET</b>	–	Secret containing the KIE_ADMIN_USER and KIE_ADMIN_PWD values.	rhpm-credentials	True
<b>KIE_SERVER_CONTROLLER_TOKEN</b>	<b>KIE_SERVER_CONTROLLER_TOKEN</b>	KIE server controller token for bearer authentication. (Sets the org.kie.server.controller.token system property)	–	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>KIE_SERVER_BYPASS_AUTH_USER</b>	<b>KIE_SERVER_BYPASS_AUTH_USER</b>	Allows the KIE server to bypass the authenticated user for task-related operations, for example, queries. (Sets the <code>org.kie.server.bypass.auth.user</code> system property)	false	False
<b>KIE_SERVER_PERSISTENCE_DS</b>	<b>RHPAM_JNDI</b>	KIE server persistence datasource. (Sets the <code>org.kie.server.persistence.ds</code> system property)	java:/jboss/datasources/rhpam	False
KIE_SERVER_H2_USER	<b>RHPAM_USERNAME</b>	KIE server H2 database user name.	sa	False
KIE_SERVER_H2_PWD	<b>RHPAM_PASSWORD</b>	KIE server H2 database password.	–	False
<b>KIE_SERVER_MODE</b>	<b>KIE_SERVER_MODE</b>	The KIE Server mode. Valid values are 'DEVELOPMENT' or 'PRODUCTION'. In production mode, you can not deploy SNAPSHOT versions of artifacts on the KIE server and can not change the version of an artifact in an existing container. (Sets the <code>org.kie.server.mode</code> system property)	<b>DEVELOPMENT</b>	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>KIE_MBEANS</b>	<b>KIE_MBEANS</b>	KIE server mbeans enabled/disabled. (Sets the kie.mbeans and kie.scanner.mbeans system properties)	enabled	False
<b>DROOLS_SERVER_FILTER_CLASSES</b>	<b>DROOLS_SERVER_FILTER_CLASSES</b>	KIE server class filtering. (Sets the org.drools.server.filter.classes system property)	true	False
<b>PROMETHEUS_SERVER_EXT_DISABLED</b>	<b>PROMETHEUS_SERVER_EXT_DISABLED</b>	If set to false, the prometheus server extension will be enabled. (Sets the org.kie.prometheus.server.ext.disabled system property)	false	False
<b>BUSINESS_CENTRAL_HOSTNAME_HTTP</b>	<b>HOSTNAME_HTTP</b>	Custom hostname for the http service route for Business Central. Leave blank for default hostname, e.g.: insecure- <application-name>- rhpamcentr- <project>.<default-domain-suffix>	–	False
<b>BUSINESS_CENTRAL_HOSTNAME_HTTPS</b>	<b>HOSTNAME_HTTPS</b>	Custom hostname for the https service route for Business Central. Leave blank for default hostname, e.g.: <application-name>- rhpamcentr- <project>.<default-domain-suffix>	–	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>KIE_SERVER_HOSTNAME_HTTP</b>	<b>HOSTNAME_HTTP</b>	Custom hostname for the http service route for KIE Server. Leave blank for default hostname, e.g.: insecure-<application-name>-kieserver-<project>.<default-domain-suffix>	–	False
<b>KIE_SERVER_HOSTNAME_HTTPS</b>	<b>HOSTNAME_HTTPS</b>	Custom hostname for the https service route for KIE Server. Leave blank for default hostname, e.g.: <application-name>-kieserver-<project>.<default-domain-suffix>	–	False
<b>BUSINESS_CENTRAL_HTTPS_SECRET</b>	–	The name of the secret containing the keystore file for Business Central.	businesscentral-app-secret	True
<b>BUSINESS_CENTRAL_HTTPS_KEYSTORE</b>	<b>HTTPS_KEYSTORE</b>	The name of the keystore file within the secret.	keystore.jks	False
<b>BUSINESS_CENTRAL_HTTPS_NAME</b>	<b>HTTPS_NAME</b>	The name associated with the server certificate.	jboss	False
<b>BUSINESS_CENTRAL_HTTPS_PASSWORD</b>	<b>HTTPS_PASSWORD</b>	The password for the keystore and certificate.	mykeystorepass	False
<b>KIE_SERVER_HTTPS_SECRET</b>	–	The name of the secret containing the keystore file for KIE server.	kieserver-app-secret	True

Variable name	Image Environment Variable	Description	Example value	Required
<b>KIE_SERVER_HTTPS_KEYSTORE</b>	<b>HTTPS_KEYSTORE</b>	The name of the keystore file within the secret.	keystore.jks	False
<b>KIE_SERVER_HTTPS_NAME</b>	<b>HTTPS_NAME</b>	The name associated with the server certificate.	jboss	False
<b>KIE_SERVER_HTTPS_PASSWORD</b>	<b>HTTPS_PASSWORD</b>	The password for the keystore and certificate.	mykeystorepass	False
<b>DB_VOLUME_CAPACITY</b>	–	Size of persistent storage for the database volume.	1Gi	True
<b>KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED</b>	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED</b>	If set to true, turns on KIE server global discovery feature (Sets the org.kie.server.controller.openshift.global.discovery.enabled system property)	false	False
<b>KIE_SERVER_CONTROLLER_OPENSHIFT_PREFER_KIESERVER_SERVICE</b>	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_PREFER_KIESERVER_SERVICE</b>	If OpenShift integration of Business Central is turned on, setting this parameter to true enables connection to KIE Server via an OpenShift internal Service endpoint. (Sets the org.kie.server.controller.openshift.prefer.kieserver.service system property)	true	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>KIE_SERVER_CONTROLLER_TEMPLATE_CACHE_TTL</b>	<b>KIE_SERVER_CONTROLLER_TEMPLATE_CACHE_TTL</b>	KIE ServerTemplate Cache TTL in milliseconds. (Sets the org.kie.server.controller.template.cache.ttl system property)	5000	False
<b>IMAGE_STREAM_NAMESPACE</b>	–	Namespace in which the ImageStreams for Red Hat Process Automation Manager images are installed. These ImageStreams are normally installed in the openshift namespace. You need to modify this parameter only if you installed the ImageStream in a different namespace/project. Default is "openshift".	openshift	True
<b>KIE_SERVER_IMAGE_STREAM_NAME</b>	–	The name of the image stream to use for KIE server. Default is "rhpam-kieserver-rhel8".	rhpam-kieserver-rhel8	True
<b>IMAGE_STREAM_TAG</b>	–	A named pointer to an image in an image stream. Default is "7.8.0".	7.8.0	True

Variable name	Image Environment Variable	Description	Example value	Required
<b>MAVEN_MIRROR_URL</b>	<b>MAVEN_MIRROR_URL</b>	Maven mirror that Business Central and KIE server must use. If you configure a mirror, this mirror must contain all artifacts that are required for building and deploying your services.	–	False
<b>MAVEN_MIRROR_OF</b>	<b>MAVEN_MIRROR_OF</b>	Maven mirror configuration for KIE server.	external:*,!repo-rhpamcentr	False
<b>MAVEN_REPO_ID</b>	<b>MAVEN_REPO_ID</b>	The id to use for the maven repository. If set, it can be excluded from the optionally configured mirror by adding it to MAVEN_MIRROR_OF. For example: external:*,!repo-rhpamcentr,!repo-custom. If MAVEN_MIRROR_URL is set but MAVEN_MIRROR_ID is not set, an id will be generated randomly, but won't be usable in MAVEN_MIRROR_OF.	repo-custom	False
<b>MAVEN_REPO_URL</b>	<b>MAVEN_REPO_URL</b>	Fully qualified URL to a Maven repository or service.	http://nexus.nexus-project.svc.cluster.local:8081/nexus/content/groups/public/	False
<b>MAVEN_REPO_USERNAME</b>	<b>MAVEN_REPO_USERNAME</b>	User name for accessing the Maven repository, if required.	–	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>MAVEN_REPO_PASSWORD</b>	<b>MAVEN_REPO_PASSWORD</b>	Password to access the Maven repository, if required.	–	False
<b>GIT_HOOKS_DIR</b>	<b>GIT_HOOKS_DIR</b>	The directory to use for git hooks, if required.	<b>/opt/kie/data/git/hooks</b>	False
<b>BUSINESS_CENTRAL_VOLUME_CAPACITY</b>	–	Size of the persistent storage for Business Central runtime data.	1Gi	True
<b>BUSINESS_CENTRAL_MEMORY_LIMIT</b>	–	Business Central Container memory limit.	2Gi	False
<b>KIE_SERVER_MEMORY_LIMIT</b>	–	KIE server Container memory limit.	1Gi	False
<b>SSO_URL</b>	<b>SSO_URL</b>	RH-SSO URL.	https://rh-sso.example.com/auth	False
<b>SSO_REALM</b>	<b>SSO_REALM</b>	RH-SSO Realm name.	–	False
<b>BUSINESS_CENTRAL_SSO_CLIENT</b>	<b>SSO_CLIENT</b>	Business Central RH-SSO Client name.	–	False
<b>BUSINESS_CENTRAL_SSO_SECRET</b>	<b>SSO_SECRET</b>	Business Central RH-SSO Client Secret.	252793ed-7118-4ca8-8dab-5622fa97d892	False
<b>KIE_SERVER_SSO_CLIENT</b>	<b>SSO_CLIENT</b>	KIE Server RH-SSO Client name.	–	False
<b>KIE_SERVER_SSO_SECRET</b>	<b>SSO_SECRET</b>	KIE Server RH-SSO Client Secret.	252793ed-7118-4ca8-8dab-5622fa97d892	False



Variable name	Image Environment Variable	Description	Example value	Required
<b>SSO_USERNAME</b>	<b>SSO_USERNAME</b>	RH-SSO Realm admin user name for creating the Client if it doesn't exist.	–	False
<b>SSO_PASSWORD</b>	<b>SSO_PASSWORD</b>	RH-SSO Realm Admin Password used to create the Client.	–	False
<b>SSO_DISABLE_SSL_CERTIFICATE_VALIDATION</b>	<b>SSO_DISABLE_SSL_CERTIFICATE_VALIDATION</b>	RH-SSO Disable SSL Certificate Validation.	false	False
<b>SSO_PRINCIPAL_ATTRIBUTE</b>	<b>SSO_PRINCIPAL_ATTRIBUTE</b>	RH-SSO Principal Attribute to use as user name.	preferred_username	False
<b>AUTH_LDAP_URL</b>	<b>AUTH_LDAP_URL</b>	LDAP Endpoint to connect for authentication	ldap://myldap.example.com	False
<b>AUTH_LDAP_BIND_DN</b>	<b>AUTH_LDAP_BIND_DN</b>	Bind DN used for authentication.	uid=admin,ou=users,ou=example,ou=com	False
<b>AUTH_LDAP_BIND_CREDENTIAL</b>	<b>AUTH_LDAP_BIND_CREDENTIAL</b>	LDAP Credentials used for authentication.	Password	False
<b>AUTH_LDAP_JAAS_SECURITY_DOMAIN</b>	<b>AUTH_LDAP_JAAS_SECURITY_DOMAIN</b>	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	–	False
<b>AUTH_LDAP_BASE_CTX_DN</b>	<b>AUTH_LDAP_BASE_CTX_DN</b>	LDAP Base DN of the top-level context to begin the user search.	ou=users,ou=example,ou=com	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_B ASE_FILTER</b>	<b>AUTH_LDAP_B ASE_FILTER</b>	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}).	(uid={0})	False
<b>AUTH_LDAP_S EARCH_SCOPE</b>	<b>AUTH_LDAP_S EARCH_SCOPE</b>	The search scope to use.	<b>SUBTREE_SCO PE</b>	False
<b>AUTH_LDAP_S EARCH_TIME_L IMIT</b>	<b>AUTH_LDAP_S EARCH_TIME_L IMIT</b>	The timeout in milliseconds for user or role searches.	10000	False
<b>AUTH_LDAP_DI STINGUISHED_ NAME_ATTRIB UTE</b>	<b>AUTH_LDAP_DI STINGUISHED_ NAME_ATTRIB UTE</b>	The name of the attribute in the user entry that contains the DN of the user. This may be necessary if the DN of the user itself contains special characters, backslash for example, that prevent correct user mapping. If the attribute does not exist, the entry's DN is used.	distinguishedName	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_PARSE_USERNAME</b>	<b>AUTH_LDAP_PARSE_USERNAME</b>	A flag indicating if the DN is to be parsed for the user name. If set to true, the DN is parsed for the user name. If set to false the DN is not parsed for the user name. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	true	False
<b>AUTH_LDAP_USERNAME_BEGIN_STRING</b>	<b>AUTH_LDAP_USERNAME_BEGIN_STRING</b>	Defines the String which is to be removed from the start of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	–	False
<b>AUTH_LDAP_USERNAME_END_STRING</b>	<b>AUTH_LDAP_USERNAME_END_STRING</b>	Defines the String which is to be removed from the end of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	–	False
<b>AUTH_LDAP_ROLE_ATTRIBUTE_ID</b>	<b>AUTH_LDAP_ROLE_ATTRIBUTE_ID</b>	Name of the attribute containing the user roles.	memberOf	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_ROLES_CTX_DN</b>	<b>AUTH_LDAP_ROLES_CTX_DN</b>	The fixed DN of the context to search for user roles. This is not the DN where the actual roles are, but the DN where the objects containing the user roles are. For example, in a Microsoft Active Directory server, this is the DN where the user account is.	ou=groups,ou=example,ou=com	False
<b>AUTH_LDAP_ROLE_FILTER</b>	<b>AUTH_LDAP_ROLE_FILTER</b>	A search filter used to locate the roles associated with the authenticated user. The input username or userDN obtained from the login module callback is substituted into the filter anywhere a {0} expression is used. The authenticated userDN is substituted into the filter anywhere a {1} is used. An example search filter that matches on the input username is (member={0}). An alternative that matches on the authenticated userDN is (member={1}).	(memberOf={1})	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_ROLE_RECURSION</b>	<b>AUTH_LDAP_ROLE_RECURSION</b>	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	1	False
<b>AUTH_LDAP_DEFAULT_ROLE</b>	<b>AUTH_LDAP_DEFAULT_ROLE</b>	A role included for all authenticated users	user	False
<b>AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID</b>	<b>AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID</b>	Name of the attribute within the roleCtxDN context which contains the role name. If the roleAttributelsDN property is set to true, this property is used to find the role object's name attribute.	name	False
<b>AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN</b>	<b>AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN</b>	A flag indicating if the DN returned by a query contains the roleNameAttribute ID. If set to true, the DN is checked for the roleNameAttribute ID. If set to false, the DN is not checked for the roleNameAttribute ID. This flag can improve the performance of LDAP queries.	false	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN</b>	<b>AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN</b>	Whether or not the roleAttributeID contains the fully-qualified DN of a role object. If false, the role name is taken from the value of the roleNameAttributeId attribute of the context name. Certain directory schemas, such as Microsoft Active Directory, require this attribute to be set to true.	false	False
<b>AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK</b>	<b>AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK</b>	If you are not using referrals, you can ignore this option. When using referrals, this option denotes the attribute name which contains users defined for a certain role, for example member, if the role object is inside the referral. Users are checked against the content of this attribute name. If this option is not set, the check will always fail, so role objects cannot be stored in a referral tree.	–	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_ROLE_MAPPER_ROLES_PROPERTIES</b>	<b>AUTH_ROLE_MAPPER_ROLES_PROPERTIES</b>	When present, the RoleMapping Login Module will be configured to use the provided file. This parameter defines the fully-qualified file path and name of a properties file or resource which maps roles to replacement roles. The format is original_role=role1,role2,role3	–	False
<b>AUTH_ROLE_MAPPER_REPLACE_ROLE</b>	<b>AUTH_ROLE_MAPPER_REPLACE_ROLE</b>	Whether to add to the current roles, or replace the current roles with the mapped ones. Replaces if set to true.	–	False

## 6.1.2. Objects

The CLI supports various object types. A list of these object types as well as their abbreviations can be found in the [Openshift documentation](#).

### 6.1.2.1. Services

A service is an abstraction which defines a logical set of pods and a policy by which to access them. See the [container-engine documentation](#) for more information.

Service	Port	Name	Description
<b>\${APPLICATION_NAME}-rhpamcentr</b>	8080	http	All the Business Central web server's ports.
	8443	https	
<b>\${APPLICATION_NAME}-kieserver</b>	8080	http	All the KIE server web server's ports.
	8443	https	

### 6.1.2.2. Routes

A route is a way to expose a service by giving it an externally reachable hostname such as **www.example.com**. A defined route and the endpoints identified by its service can be consumed by a router to provide named connectivity from external clients to your applications. Each route consists of a route name, service selector, and (optionally) security configuration. See the [OpenShift documentation](#) for more information.

Service	Security	Hostname
insecure- \${APPLICATION_NAME}- rhpamcentr-http	none	<b>\${BUSINESS_CENTRAL_HOSTNAME_HTTP}</b>
<b>\${APPLICATION_NAME}- rhpamcentr-https</b>	TLS passthrough	<b>\${BUSINESS_CENTRAL_HOSTNAME_HTTPS}</b>
insecure- \${APPLICATION_NAME}- kieserver-http	none	<b>\${KIE_SERVER_HOSTNAME_HTTP}</b>
<b>\${APPLICATION_NAME}- kieserver-https</b>	TLS passthrough	<b>\${KIE_SERVER_HOSTNAME_HTTPS}</b>

### 6.1.2.3. Deployment Configurations

A deployment in OpenShift is a replication controller based on a user-defined template called a deployment configuration. Deployments are created manually or in response to triggered events. See the [OpenShift documentation](#) for more information.

#### 6.1.2.3.1. Triggers

A trigger drives the creation of new deployments in response to events, both inside and outside OpenShift. See the [OpenShift documentation](#) for more information.

Deployment	Triggers
<b>\${APPLICATION_NAME}-rhpamcentr</b>	ImageChange
<b>\${APPLICATION_NAME}-kieserver</b>	ImageChange

#### 6.1.2.3.2. Replicas

A replication controller ensures that a specified number of pod "replicas" are running at any one time. If there are too many, the replication controller kills some pods. If there are too few, it starts more. See the [container-engine documentation](#) for more information.



Deployment	Replicas
<b><code>\${APPLICATION_NAME}-rhpamcentr</code></b>	1
<b><code>\${APPLICATION_NAME}-kieserver</code></b>	1

### 6.1.2.3.3. Pod Template

#### 6.1.2.3.3.1. Service Accounts

Service accounts are API objects that exist within each project. They can be created or deleted like any other API object. See the [OpenShift documentation](#) for more information.

Deployment	Service Account
<b><code>\${APPLICATION_NAME}-rhpamcentr</code></b>	<b><code>\${APPLICATION_NAME}-rhpamsvc</code></b>
<b><code>\${APPLICATION_NAME}-kieserver</code></b>	<b><code>\${APPLICATION_NAME}-rhpamsvc</code></b>

#### 6.1.2.3.3.2. Image

Deployment	Image
<b><code>\${APPLICATION_NAME}-rhpamcentr</code></b>	<code>rhpam-businesscentral-rhel8</code>
<b><code>\${APPLICATION_NAME}-kieserver</code></b>	<b><code>\${KIE_SERVER_IMAGE_STREAM_NAME}</code></b>

#### 6.1.2.3.3.3. Readiness Probe

**`${APPLICATION_NAME}-rhpamcentr`**

Http Get on `http://localhost:8080/rest/ready`

**`${APPLICATION_NAME}-kieserver`**

Http Get on `http://localhost:8080/services/rest/server/readycheck`

#### 6.1.2.3.3.4. Liveness Probe

**`${APPLICATION_NAME}-rhpamcentr`**

Http Get on `http://localhost:8080/rest/healthy`

**`${APPLICATION_NAME}-kieserver`**

Http Get on `http://localhost:8080/services/rest/server/healthcheck`

## 6.1.2.3.3.5. Exposed Ports

Deployments	Name	Port	Protocol
<b>\${APPLICATION_NAME}-rhpamcentr</b>	jolokia	8778	<b>TCP</b>
	http	8080	<b>TCP</b>
	https	8443	<b>TCP</b>
<b>\${APPLICATION_NAME}-kieserver</b>	jolokia	8778	<b>TCP</b>
	http	8080	<b>TCP</b>
	https	8443	<b>TCP</b>

## 6.1.2.3.3.6. Image Environment Variables

Deployment	Variable name	Description	Example value
<b>\${APPLICATION_NAME}-rhpamcentr</b>	<b>APPLICATION_USE_RS_PROPERTIES</b>	–	<b>/opt/kie/data/configuration/application-users.properties</b>
	<b>APPLICATION_ROLES_PROPERTIES</b>	–	<b>/opt/kie/data/configuration/application-roles.properties</b>
	<b>KIE_ADMIN_USER</b>	Admin user name	Set according to the credentials secret
	<b>KIE_ADMIN_PWD</b>	Admin user password	Set according to the credentials secret
	<b>KIE_MBEANS</b>	KIE server mbeans enabled/disabled. (Sets the kie.mbeans and kie.scanner.mbeans system properties)	<b>\${KIE_MBEANS}</b>
	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_ENABLED</b>	–	false

Deployment	Variable name	Description	Example value
	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED</b>	If set to true, turns on KIE server global discovery feature (Sets the org.kie.server.controller.openshift.global.discovery.enabled system property)	<b>`\${KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED}`</b>
	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_PREFER_KIESERVER_SERVICE</b>	If OpenShift integration of Business Central is turned on, setting this parameter to true enables connection to KIE Server via an OpenShift internal Service endpoint. (Sets the org.kie.server.controller.openshift.prefer.kieserver.service system property)	<b>`\${KIE_SERVER_CONTROLLER_OPENSHIFT_PREFER_KIESERVER_SERVICE}`</b>
	<b>KIE_SERVER_CONTROLLER_TEMPLATE_CACHE_TTL</b>	KIE ServerTemplate Cache TTL in milliseconds. (Sets the org.kie.server.controller.template.cache.ttl system property)	<b>`\${KIE_SERVER_CONTROLLER_TEMPLATE_CACHE_TTL}`</b>
	<b>KIE_SERVER_CONTROLLER_TOKEN</b>	KIE server controller token for bearer authentication. (Sets the org.kie.server.controller.token system property)	<b>`\${KIE_SERVER_CONTROLLER_TOKEN}`</b>
	<b>MAVEN_MIRROR_URL</b>	Maven mirror that Business Central and KIE server must use. If you configure a mirror, this mirror must contain all artifacts that are required for building and deploying your services.	<b>`\${MAVEN_MIRROR_URL}`</b>

Deployment	Variable name	Description	Example value
	<b>MAVEN_REPO_ID</b>	The id to use for the maven repository. If set, it can be excluded from the optionally configured mirror by adding it to MAVEN_MIRROR_OF. For example: external:*,!repo-rhpamcentr,!repo-custom. If MAVEN_MIRROR_URL is set but MAVEN_MIRROR_ID is not set, an id will be generated randomly, but won't be usable in MAVEN_MIRROR_OF.	<b>\${MAVEN_REPO_ID}</b>
	<b>MAVEN_REPO_URL</b>	Fully qualified URL to a Maven repository or service.	<b>\${MAVEN_REPO_URL}</b>
	<b>MAVEN_REPO_USERNAME</b>	User name for accessing the Maven repository, if required.	<b>\${MAVEN_REPO_USERNAME}</b>
	<b>MAVEN_REPO_PASSWORD</b>	Password to access the Maven repository, if required.	<b>\${MAVEN_REPO_PASSWORD}</b>
	<b>GIT_HOOKS_DIR</b>	The directory to use for git hooks, if required.	<b>\${GIT_HOOKS_DIR}</b>
	<b>HTTPS_KEYSTORE_DIR</b>	–	<b>/etc/businesscentral-secret-volume</b>
	<b>HTTPS_KEYSTORE</b>	The name of the keystore file within the secret.	<b>\${BUSINESS_CENTRAL_HTTPS_KEYSTORE}</b>
	<b>HTTPS_NAME</b>	The name associated with the server certificate.	<b>\${BUSINESS_CENTRAL_HTTPS_NAME}</b>
	<b>HTTPS_PASSWORD</b>	The password for the keystore and certificate.	<b>\${BUSINESS_CENTRAL_HTTPS_PASSWORD}</b>

Deployment	Variable name	Description	Example value
	<b>WORKBENCH_ROUTE_NAME</b>	–	<b>\${APPLICATION_NAME}-rhpamcentr</b>
	<b>SSO_URL</b>	RH-SSO URL.	<b>\${SSO_URL}</b>
	<b>SSO_OPENIDCONNECT_DEPLOYMENTS</b>	–	ROOT.war
	<b>SSO_REALM</b>	RH-SSO Realm name.	<b>\${SSO_REALM}</b>
	<b>SSO_SECRET</b>	Business Central RH-SSO Client Secret.	<b>\${BUSINESS_CENTRAL_SSO_SECRET}</b>
	<b>SSO_CLIENT</b>	Business Central RH-SSO Client name.	<b>\${BUSINESS_CENTRAL_SSO_CLIENT}</b>
	<b>SSO_USERNAME</b>	RH-SSO Realm admin user name for creating the Client if it doesn't exist.	<b>\${SSO_USERNAME}</b>
	<b>SSO_PASSWORD</b>	RH-SSO Realm Admin Password used to create the Client.	<b>\${SSO_PASSWORD}</b>
	<b>SSO_DISABLE_SSL_CERTIFICATE_VALIDATION</b>	RH-SSO Disable SSL Certificate Validation.	<b>\${SSO_DISABLE_SSL_CERTIFICATE_VALIDATION}</b>
	<b>SSO_PRINCIPAL_ATTRIBUTE</b>	RH-SSO Principal Attribute to use as user name.	<b>\${SSO_PRINCIPAL_ATTRIBUTE}</b>
	<b>HOSTNAME_HTTP</b>	Custom hostname for the http service route for Business Central. Leave blank for default hostname, e.g.: insecure-<application-name>-rhpamcentr-<project>.<default-domain-suffix>	<b>\${BUSINESS_CENTRAL_HOSTNAME_HTTP}</b>

Deployment	Variable name	Description	Example value
	<b>HOSTNAME_HTTPS</b>	Custom hostname for the https service route for Business Central. Leave blank for default hostname, e.g.: <application-name>-rhpamcentr-<project>. <default-domain-suffix>	<b>\${BUSINESS_CENTRAL_HOSTNAME_HTTPS}</b>
	<b>AUTH_LDAP_URL</b>	LDAP Endpoint to connect for authentication	<b>\${AUTH_LDAP_URL}</b>
	<b>AUTH_LDAP_BIND_DN</b>	Bind DN used for authentication.	<b>\${AUTH_LDAP_BIND_DN}</b>
	<b>AUTH_LDAP_BIND_CREDENTIAL</b>	LDAP Credentials used for authentication.	<b>\${AUTH_LDAP_BIND_CREDENTIAL}</b>
	<b>AUTH_LDAP_JAAS_SECURITY_DOMAIN</b>	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	<b>\${AUTH_LDAP_JAAS_SECURITY_DOMAIN}</b>
	<b>AUTH_LDAP_BASE_CTX_DN</b>	LDAP Base DN of the top-level context to begin the user search.	<b>\${AUTH_LDAP_BASE_CTX_DN}</b>
	<b>AUTH_LDAP_BASE_FILTER</b>	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}).	<b>\${AUTH_LDAP_BASE_FILTER}</b>
	<b>AUTH_LDAP_SEARCH_SCOPE</b>	The search scope to use.	<b>\${AUTH_LDAP_SEARCH_SCOPE}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_SEARCH_TIME_LIMIT</b>	The timeout in milliseconds for user or role searches.	<b>`\${AUTH_LDAP_SEARCH_TIME_LIMIT}`</b>
	<b>AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE</b>	The name of the attribute in the user entry that contains the DN of the user. This may be necessary if the DN of the user itself contains special characters, backslash for example, that prevent correct user mapping. If the attribute does not exist, the entry's DN is used.	<b>`\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}`</b>
	<b>AUTH_LDAP_PARSE_USERNAME</b>	A flag indicating if the DN is to be parsed for the user name. If set to true, the DN is parsed for the user name. If set to false the DN is not parsed for the user name. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	<b>`\${AUTH_LDAP_PARSE_USERNAME}`</b>
	<b>AUTH_LDAP_USERNAME_BEGIN_STRING</b>	Defines the String which is to be removed from the start of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<b>`\${AUTH_LDAP_USERNAME_BEGIN_STRING}`</b>
	<b>AUTH_LDAP_USERNAME_END_STRING</b>	Defines the String which is to be removed from the end of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<b>`\${AUTH_LDAP_USERNAME_END_STRING}`</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_ROLE_ATTRIBUTE_ID</b>	Name of the attribute containing the user roles.	<b>`\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}`</b>
	<b>AUTH_LDAP_ROLE_S_CTX_DN</b>	The fixed DN of the context to search for user roles. This is not the DN where the actual roles are, but the DN where the objects containing the user roles are. For example, in a Microsoft Active Directory server, this is the DN where the user account is.	<b>`\${AUTH_LDAP_ROLE_S_CTX_DN}`</b>
	<b>AUTH_LDAP_ROLE_FILTER</b>	A search filter used to locate the roles associated with the authenticated user. The input username or userDN obtained from the login module callback is substituted into the filter anywhere a <code>{0}</code> expression is used. The authenticated userDN is substituted into the filter anywhere a <code>{1}</code> is used. An example search filter that matches on the input username is <code>(member={0})</code> . An alternative that matches on the authenticated userDN is <code>(member={1})</code> .	<b>`\${AUTH_LDAP_ROLE_FILTER}`</b>
	<b>AUTH_LDAP_ROLE_RECURSION</b>	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	<b>`\${AUTH_LDAP_ROLE_RECURSION}`</b>
	<b>AUTH_LDAP_DEFAULT_ROLE</b>	A role included for all authenticated users	<b>`\${AUTH_LDAP_DEFAULT_ROLE}`</b>



Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID</b>	Name of the attribute within the roleCtxDN context which contains the role name. If the roleAttributesDN property is set to true, this property is used to find the role object's name attribute.	<b>\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}</b>
	<b>AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN</b>	A flag indicating if the DN returned by a query contains the roleNameAttributeID. If set to true, the DN is checked for the roleNameAttributeID. If set to false, the DN is not checked for the roleNameAttributeID. This flag can improve the performance of LDAP queries.	<b>\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}</b>
	<b>AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN</b>	Whether or not the roleAttributeID contains the fully-qualified DN of a role object. If false, the role name is taken from the value of the roleNameAttributeID attribute of the context name. Certain directory schemas, such as Microsoft Active Directory, require this attribute to be set to true.	<b>\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK</b>	If you are not using referrals, you can ignore this option. When using referrals, this option denotes the attribute name which contains users defined for a certain role, for example member, if the role object is inside the referral. Users are checked against the content of this attribute name. If this option is not set, the check will always fail, so role objects cannot be stored in a referral tree.	<b>`\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}`</b>
	<b>AUTH_ROLE_MAPPER_ROLES_PROPERTIES</b>	When present, the RoleMapping Login Module will be configured to use the provided file. This parameter defines the fully-qualified file path and name of a properties file or resource which maps roles to replacement roles. The format is original_role=role1,role2,role3	<b>`\${AUTH_ROLE_MAPPER_ROLES_PROPERTIES}`</b>
	<b>AUTH_ROLE_MAPPER_REPLACE_ROLE</b>	Whether to add to the current roles, or replace the current roles with the mapped ones. Replaces if set to true.	<b>`\${AUTH_ROLE_MAPPER_REPLACE_ROLE}`</b>
<b>`\${APPLICATION_NAME}`-kieserver</b>	<b>WORKBENCH_SERVICE_NAME</b>	–	<b>`\${APPLICATION_NAME}`-rhpamcentr</b>
	<b>DATASOURCES</b>	–	<b>RHPAM</b>
	<b>RHPAM_DATABASE</b>	–	rhpam7

Deployment	Variable name	Description	Example value
	<b>RHPAM_JNDI</b>	KIE server persistence datasource. (Sets the org.kie.server.persistence.ds system property)	<b>\${KIE_SERVER_PERSISTENCE_DS}</b>
	<b>RHPAM_JTA</b>	–	true
	<b>RHPAM_DRIVER</b>	–	h2
	<b>RHPAM_USERNAME</b>	KIE server H2 database user name.	<b>\${KIE_SERVER_H2_USER}</b>
	<b>RHPAM_PASSWORD</b>	KIE server H2 database password.	<b>\${KIE_SERVER_H2_PWD}</b>
	<b>RHPAM_NONXA</b>	–	false
	<b>RHPAM_XA_CONNECTION_PROPERTY_URL</b>	–	jdbc:h2:/opt/kie/data/h2/rhpam;AUTO_SERVER=TRUE
	<b>KIE_SERVER_PERSISTENCE_DIALECT</b>	–	org.hibernate.dialect.H2Dialect
	<b>KIE_ADMIN_USER</b>	Admin user name	Set according to the credentials secret
	<b>KIE_ADMIN_PWD</b>	Admin user password	Set according to the credentials secret
	<b>KIE_SERVER_MODE</b>	The KIE Server mode. Valid values are 'DEVELOPMENT' or 'PRODUCTION'. In production mode, you can not deploy SNAPSHOT versions of artifacts on the KIE server and can not change the version of an artifact in an existing container. (Sets the org.kie.server.mode system property)	<b>\${KIE_SERVER_MODE}</b>

Deployment	Variable name	Description	Example value
	<b>KIE_MBEANS</b>	KIE server mbeans enabled/disabled. (Sets the kie.mbeans and kie.scanner.mbeans system properties)	<b>\${KIE_MBEANS}</b>
	<b>DROOLS_SERVER_FILTER_CLASSES</b>	KIE server class filtering. (Sets the org.drools.server.filter.classes system property)	<b>\${DROOLS_SERVER_FILTER_CLASSES}</b>
	<b>PROMETHEUS_SERVER_EXT_DISABLED</b>	If set to false, the prometheus server extension will be enabled. (Sets the org.kie.prometheus.server.ext.disabled system property)	<b>\${PROMETHEUS_SERVER_EXT_DISABLED}</b>
	<b>KIE_SERVER_BYPASS_AUTH_USER</b>	Allows the KIE server to bypass the authenticated user for task-related operations, for example, queries. (Sets the org.kie.server.bypass.auth.user system property)	<b>\${KIE_SERVER_BYPASS_AUTH_USER}</b>
	<b>KIE_SERVER_CONTROLLER_SERVICE</b>	–	<b>\${APPLICATION_NAME}-rhpamcentr</b>
	<b>KIE_SERVER_CONTROLLER_PROTOCOL</b>	–	ws
	<b>KIE_SERVER_ID</b>	–	–
	<b>KIE_SERVER_ROUTE_NAME</b>	–	insecure-\${APPLICATION_NAME}-kieserver
	<b>KIE_SERVER_PERSISTENCE_DS</b>	KIE server persistence datasource. (Sets the org.kie.server.persistence.ds system property)	<b>\${KIE_SERVER_PERSISTENCE_DS}</b>

Deployment	Variable name	Description	Example value
	<b>KIE_SERVER_STARTUP_STRATEGY</b>	–	ControllerBasedStartupStrategy
	<b>MAVEN_MIRROR_URL</b>	Maven mirror that Business Central and KIE server must use. If you configure a mirror, this mirror must contain all artifacts that are required for building and deploying your services.	<b>\${MAVEN_MIRROR_URL}</b>
	<b>MAVEN_MIRROR_OFF</b>	Maven mirror configuration for KIE server.	<b>\${MAVEN_MIRROR_OFF}</b>
	<b>MAVEN_REPOS</b>	–	RHPAMCENTR,EXTERNAL
	<b>RHPAMCENTR_MAVEN_REPO_ID</b>	–	repo-rhpamcentr
	<b>RHPAMCENTR_MAVEN_REPO_SERVICE</b>	–	<b>\${APPLICATION_NAME}-rhpamcentr</b>
	<b>RHPAMCENTR_MAVEN_REPO_PATH</b>	–	<b>/maven2/</b>
	<b>RHPAMCENTR_MAVEN_REPO_USERNAME</b>	–	Set according to the credentials secret
	<b>RHPAMCENTR_MAVEN_REPO_PASSWORD</b>	–	Set according to the credentials secret

Deployment	Variable name	Description	Example value
	<b>EXTERNAL_MAVEN_REPO_ID</b>	The id to use for the maven repository. If set, it can be excluded from the optionally configured mirror by adding it to MAVEN_MIRROR_OF. For example: external:*,!repo-rhpamcentr,!repo-custom. If MAVEN_MIRROR_URL is set but MAVEN_MIRROR_ID is not set, an id will be generated randomly, but won't be usable in MAVEN_MIRROR_OF.	<b>`\${MAVEN_REPO_ID}`</b>
	<b>EXTERNAL_MAVEN_REPO_URL</b>	Fully qualified URL to a Maven repository or service.	<b>`\${MAVEN_REPO_URL}`</b>
	<b>EXTERNAL_MAVEN_REPO_USERNAME</b>	User name for accessing the Maven repository, if required.	<b>`\${MAVEN_REPO_USERNAME}`</b>
	<b>EXTERNAL_MAVEN_REPO_PASSWORD</b>	Password to access the Maven repository, if required.	<b>`\${MAVEN_REPO_PASSWORD}`</b>
	<b>HTTPS_KEYSTORE_DIR</b>	–	<b>/etc/kieserver-secret-volume</b>
	<b>HTTPS_KEYSTORE</b>	The name of the keystore file within the secret.	<b>`\${KIE_SERVER_HTTPS_KEYSTORE}`</b>
	<b>HTTPS_NAME</b>	The name associated with the server certificate.	<b>`\${KIE_SERVER_HTTPS_NAME}`</b>
	<b>HTTPS_PASSWORD</b>	The password for the keystore and certificate.	<b>`\${KIE_SERVER_HTTPS_PASSWORD}`</b>
	<b>SSO_URL</b>	RH-SSO URL.	<b>`\${SSO_URL}`</b>

Deployment	Variable name	Description	Example value
	<b>SSO_OPENIDCONNECT_DEPLOYMENTS</b>	–	ROOT.war
	<b>SSO_REALM</b>	RH-SSO Realm name.	<b>\${SSO_REALM}</b>
	<b>SSO_SECRET</b>	KIE Server RH-SSO Client Secret.	<b>\${KIE_SERVER_SSO_SECRET}</b>
	<b>SSO_CLIENT</b>	KIE Server RH-SSO Client name.	<b>\${KIE_SERVER_SSO_CLIENT}</b>
	<b>SSO_USERNAME</b>	RH-SSO Realm admin user name for creating the Client if it doesn't exist.	<b>\${SSO_USERNAME}</b>
	<b>SSO_PASSWORD</b>	RH-SSO Realm Admin Password used to create the Client.	<b>\${SSO_PASSWORD}</b>
	<b>SSO_DISABLE_SSL_CERTIFICATE_VALIDATION</b>	RH-SSO Disable SSL Certificate Validation.	<b>\${SSO_DISABLE_SSL_CERTIFICATE_VALIDATION}</b>
	<b>SSO_PRINCIPAL_ATTRIBUTE</b>	RH-SSO Principal Attribute to use as user name.	<b>\${SSO_PRINCIPAL_ATTRIBUTE}</b>
	<b>HOSTNAME_HTTP</b>	Custom hostname for the http service route for KIE Server. Leave blank for default hostname, e.g.: insecure-<application-name>-kieserver-<project>.<default-domain-suffix>	<b>\${KIE_SERVER_HOSTNAME_HTTP}</b>
	<b>HOSTNAME_HTTPS</b>	Custom hostname for the https service route for KIE Server. Leave blank for default hostname, e.g.: <application-name>-kieserver-<project>.<default-domain-suffix>	<b>\${KIE_SERVER_HOSTNAME_HTTPS}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_URL</b>	LDAP Endpoint to connect for authentication	<b>\${AUTH_LDAP_URL}</b>
	<b>AUTH_LDAP_BIND_DN</b>	Bind DN used for authentication.	<b>\${AUTH_LDAP_BIND_DN}</b>
	<b>AUTH_LDAP_BIND_CREDENTIAL</b>	LDAP Credentials used for authentication.	<b>\${AUTH_LDAP_BIND_CREDENTIAL}</b>
	<b>AUTH_LDAP_JAAS_SECURITY_DOMAIN</b>	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	<b>\${AUTH_LDAP_JAAS_SECURITY_DOMAIN}</b>
	<b>AUTH_LDAP_BASE_CTX_DN</b>	LDAP Base DN of the top-level context to begin the user search.	<b>\${AUTH_LDAP_BASE_CTX_DN}</b>
	<b>AUTH_LDAP_BASE_FILTER</b>	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}).	<b>\${AUTH_LDAP_BASE_FILTER}</b>
	<b>AUTH_LDAP_SEARCH_SCOPE</b>	The search scope to use.	<b>\${AUTH_LDAP_SEARCH_SCOPE}</b>
	<b>AUTH_LDAP_SEARCH_TIME_LIMIT</b>	The timeout in milliseconds for user or role searches.	<b>\${AUTH_LDAP_SEARCH_TIME_LIMIT}</b>



Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE</b>	The name of the attribute in the user entry that contains the DN of the user. This may be necessary if the DN of the user itself contains special characters, backslash for example, that prevent correct user mapping. If the attribute does not exist, the entry's DN is used.	<b>`\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}`</b>
	<b>AUTH_LDAP_PARSE_USERNAME</b>	A flag indicating if the DN is to be parsed for the user name. If set to true, the DN is parsed for the user name. If set to false the DN is not parsed for the user name. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	<b>`\${AUTH_LDAP_PARSE_USERNAME}`</b>
	<b>AUTH_LDAP_USERNAME_BEGIN_STRING</b>	Defines the String which is to be removed from the start of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<b>`\${AUTH_LDAP_USERNAME_BEGIN_STRING}`</b>
	<b>AUTH_LDAP_USERNAME_END_STRING</b>	Defines the String which is to be removed from the end of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<b>`\${AUTH_LDAP_USERNAME_END_STRING}`</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_ROLE_ATTRIBUTE_ID</b>	Name of the attribute containing the user roles.	<b>\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}</b>
	<b>AUTH_LDAP_ROLE_S_CTX_DN</b>	The fixed DN of the context to search for user roles. This is not the DN where the actual roles are, but the DN where the objects containing the user roles are. For example, in a Microsoft Active Directory server, this is the DN where the user account is.	<b>\${AUTH_LDAP_ROLE_S_CTX_DN}</b>
	<b>AUTH_LDAP_ROLE_FILTER</b>	A search filter used to locate the roles associated with the authenticated user. The input username or userDN obtained from the login module callback is substituted into the filter anywhere a {0} expression is used. The authenticated userDN is substituted into the filter anywhere a {1} is used. An example search filter that matches on the input username is (member={0}). An alternative that matches on the authenticated userDN is (member={1}).	<b>\${AUTH_LDAP_ROLE_FILTER}</b>
	<b>AUTH_LDAP_ROLE_RECURSION</b>	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	<b>\${AUTH_LDAP_ROLE_RECURSION}</b>
	<b>AUTH_LDAP_DEFAULT_ROLE</b>	A role included for all authenticated users	<b>\${AUTH_LDAP_DEFAULT_ROLE}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID</b>	Name of the attribute within the roleCtxDN context which contains the role name. If the roleAttributelsDN property is set to true, this property is used to find the role object's name attribute.	<b>\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}</b>
	<b>AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN</b>	A flag indicating if the DN returned by a query contains the roleNameAttributeID. If set to true, the DN is checked for the roleNameAttributeID. If set to false, the DN is not checked for the roleNameAttributeID. This flag can improve the performance of LDAP queries.	<b>\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}</b>
	<b>AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN</b>	Whether or not the roleAttributeID contains the fully-qualified DN of a role object. If false, the role name is taken from the value of the roleNameAttributeID attribute of the context name. Certain directory schemas, such as Microsoft Active Directory, require this attribute to be set to true.	<b>\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK</b>	If you are not using referrals, you can ignore this option. When using referrals, this option denotes the attribute name which contains users defined for a certain role, for example member, if the role object is inside the referral. Users are checked against the content of this attribute name. If this option is not set, the check will always fail, so role objects cannot be stored in a referral tree.	<b><code>\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}</code></b>
	<b>AUTH_ROLE_MAPPER_ROLES_PROPERTIES</b>	When present, the RoleMapping Login Module will be configured to use the provided file. This parameter defines the fully-qualified file path and name of a properties file or resource which maps roles to replacement roles. The format is original_role=role1,role2,role3	<b><code>\${AUTH_ROLE_MAPPER_ROLES_PROPERTIES}</code></b>
	<b>AUTH_ROLE_MAPPER_REPLACE_ROLE</b>	Whether to add to the current roles, or replace the current roles with the mapped ones. Replaces if set to true.	<b><code>\${AUTH_ROLE_MAPPER_REPLACE_ROLE}</code></b>

#### 6.1.2.3.3.7. Volumes

Deployment	Name	mountPath	Purpose	readOnly
<b><code>\${APPLICATION_NAME}-rhcamcentr</code></b>	businesscentral-keystore-volume	<b><code>/etc/businesscentral-secret-volume</code></b>	ssl certs	True

Deployment	Name	mountPath	Purpose	readOnly
<b>\${APPLICATION_NAME}-kieserver</b>	kieserver-keystore-volume	<b>/etc/kieserver-secret-volume</b>	ssl certs	True

#### 6.1.2.4. External Dependencies

##### 6.1.2.4.1. Volume Claims

A **PersistentVolume** object is a storage resource in an OpenShift cluster. Storage is provisioned by an administrator by creating **PersistentVolume** objects from sources such as GCE Persistent Disks, AWS Elastic Block Stores (EBS), and NFS mounts. See the [OpenShift documentation](#) for more information.

Name	Access Mode
<b>\${APPLICATION_NAME}-rhpamcentr-claim</b>	ReadWriteOnce
<b>\${APPLICATION_NAME}-kie-claim</b>	ReadWriteOnce

##### 6.1.2.4.2. Secrets

This template requires the following secrets to be installed for the application to run.

businesscentral-app-secret kieserver-app-secret

## 6.2. RHPAM78-AUTHORING-HA.YAML TEMPLATE

Application template for a HA persistent authoring environment, for Red Hat Process Automation Manager 7.8 - Deprecated

### 6.2.1. Parameters

Templates allow you to define parameters that take on a value. That value is then substituted wherever the parameter is referenced. References can be defined in any text field in the objects list field. See the [OpenShift documentation](#) for more information.

Variable name	Image Environment Variable	Description	Example value	Required
<b>APPLICATION_NAME</b>	–	The name for the application.	myapp	True

Variable name	Image Environment Variable	Description	Example value	Required
<b>CREDENTIALS_SECRET</b>	–	Secret containing the KIE_ADMIN_USER and KIE_ADMIN_PWD values.	rhpm-credentials	True
<b>KIE_SERVER_CONTROLLER_TOKEN</b>	<b>KIE_SERVER_CONTROLLER_TOKEN</b>	KIE server controller token for bearer authentication. (Sets the org.kie.server.controller.token system property)	–	False
<b>KIE_SERVER_BYPASS_AUTH_USER</b>	<b>KIE_SERVER_BYPASS_AUTH_USER</b>	Allows the KIE server to bypass the authenticated user for task-related operations, for example, queries. (Sets the org.kie.server.bypass.auth.user system property)	false	False
<b>KIE_SERVER_PERSISTENCE_DS</b>	<b>KIE_SERVER_PERSISTENCE_DS</b>	KIE server persistence datasource. (Sets the org.kie.server.persistence.ds system property)	java:/jboss/datasources/rhpm	False
<b>MYSQL_USER</b>	<b>RHPAM_USERNAME</b>	MySQL database user name.	rhpm	False
<b>MYSQL_PWD</b>	<b>RHPAM_PASSWORD</b>	MySQL database password.	–	False
<b>MYSQL_DB</b>	<b>RHPAM_DATABASE</b>	MySQL database name.	rhpm7	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>MYSQL_DB_VOLUME_CAPACITY</b>	–	Size of persistent storage for the KIE server database volume.	1Gi	True
<b>MYSQL_IMAGE_STREAM_NAMESPACE</b>	–	Namespace in which the ImageStream for the MySQL image is installed. The ImageStream is already installed in the openshift namespace. You need to modify this parameter only if you installed the ImageStream in a different namespace/project. Default is "openshift".	openshift	False
<b>MYSQL_IMAGE_STREAM_TAG</b>	–	The MySQL image version, which is intended to correspond to the MySQL version. Default is "8.0".	8.0	False
<b>KIE_SERVER_MYSQL_DIALECT</b>	<b>KIE_SERVER_PERSISTENCE_DIALECT</b>	KIE server MySQL Hibernate dialect.	org.hibernate.dialect.MySQL8Dialect	True

Variable name	Image Environment Variable	Description	Example value	Required
<b>KIE_SERVER_MODE</b>	<b>KIE_SERVER_MODE</b>	The KIE Server mode. Valid values are 'DEVELOPMENT' or 'PRODUCTION'. In production mode, you can not deploy SNAPSHOT versions of artifacts on the KIE server and can not change the version of an artifact in an existing container. (Sets the org.kie.server.mode system property).	<b>DEVELOPMENT</b>	False
<b>KIE_MBEANS</b>	<b>KIE_MBEANS</b>	KIE server mbeans enabled/disabled. (Sets the kie.mbeans and kie.scanner.mbeans system properties)	enabled	False
<b>DROOLS_SERVER_FILTER_CLASSES</b>	<b>DROOLS_SERVER_FILTER_CLASSES</b>	KIE server class filtering. (Sets the org.drools.server.filter.classes system property)	true	False
<b>PROMETHEUS_SERVER_EXT_DISABLED</b>	<b>PROMETHEUS_SERVER_EXT_DISABLED</b>	If set to false, the prometheus server extension will be enabled. (Sets the org.kie.prometheus.server.ext.disabled system property)	false	False



Variable name	Image Environment Variable	Description	Example value	Required
<b>BUSINESS_CENTRAL_HOSTNAME_HTTP</b>	<b>HOSTNAME_HTTP</b>	Custom hostname for http service route for Business Central. Leave blank for default hostname, e.g.: insecure- <application-name>- rhpamcentr- <project>.<default-domain-suffix>	–	False
<b>BUSINESS_CENTRAL_HOSTNAME_HTTPS</b>	<b>HOSTNAME_HTTPS</b>	Custom hostname for https service route for Business Central. Leave blank for default hostname, e.g.: <application-name>- rhpamcentr- <project>.<default-domain-suffix>	–	False
<b>KIE_SERVER_HOSTNAME_HTTP</b>	<b>HOSTNAME_HTTP</b>	Custom hostname for http service route for KIE Server. Leave blank for default hostname, e.g.: insecure- <application-name>-kieserver- <project>.<default-domain-suffix>	–	False
<b>KIE_SERVER_HOSTNAME_HTTPS</b>	<b>HOSTNAME_HTTPS</b>	Custom hostname for https service route for KIE Server. Leave blank for default hostname, e.g.: <application-name>-kieserver- <project>.<default-domain-suffix>	–	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>BUSINESS_CENTRAL_HTTPS_SECRET</b>	–	The name of the secret containing the keystore file for Business Central.	businesscentral-app-secret	True
<b>BUSINESS_CENTRAL_HTTPS_KEYSTORE</b>	<b>HTTPS_KEYSTORE</b>	The name of the keystore file within the secret for Business Central.	keystore.jks	False
<b>BUSINESS_CENTRAL_HTTPS_NAME</b>	<b>HTTPS_NAME</b>	The name associated with the server certificate for Business Central.	jboss	False
<b>BUSINESS_CENTRAL_HTTPS_PASSWORD</b>	<b>HTTPS_PASSWORD</b>	The password for the keystore and certificate for Business Central.	mykeystorepass	False
<b>KIE_SERVER_HTTPS_SECRET</b>	–	The name of the secret containing the keystore file for KIE Server.	kieserver-app-secret	True
<b>KIE_SERVER_HTTPS_KEYSTORE</b>	<b>HTTPS_KEYSTORE</b>	The name of the keystore file within the secret for KIE Server.	keystore.jks	False
<b>KIE_SERVER_HTTPS_NAME</b>	<b>HTTPS_NAME</b>	The name associated with the server certificate for KIE Server.	jboss	False
<b>KIE_SERVER_HTTPS_PASSWORD</b>	<b>HTTPS_PASSWORD</b>	The password for the keystore and certificate for KIE Server.	mykeystorepass	False
<b>APPFORMER_JMS_BROKER_USER</b>	<b>APPFORMER_JMS_BROKER_USER</b>	The user name for connecting to the JMS broker.	jmsBrokerUser	True

Variable name	Image Environment Variable	Description	Example value	Required
<b>APPFORMER_JMS_BROKER_PASSWORD</b>	<b>APPFORMER_JMS_BROKER_PASSWORD</b>	The password to connect to the JMS broker.	–	True
<b>DATAGRID_IMAGE</b>	–	DataGrid image.	registry.redhat.io/jboss-datagrid-7/datagrid73-openshift:1.5	True
<b>DATAGRID_CPU_LIMIT</b>	–	DataGrid Container CPU limit.	1000m	True
<b>DATAGRID_MEMORY_LIMIT</b>	–	DataGrid Container memory limit.	2Gi	True
<b>DATAGRID_VOLUME_CAPACITY</b>	–	Size of the persistent storage for DataGrid's runtime data.	1Gi	True
<b>AMQ_BROKER_IMAGE</b>	–	AMQ Broker Image.	registry.redhat.io/amq7/amq-broker:7.6	True
<b>AMQ_ROLE</b>	–	User role for standard broker user.	admin	True
<b>AMQ_NAME</b>	–	The name of the broker.	broker	True
<b>AMQ_GLOBAL_MAX_SIZE</b>	–	Specifies the maximum amount of memory that message data can consume. If no value is specified, half of the system's memory is allocated.	10 gb	False
<b>AMQ_VOLUME_CAPACITY</b>	–	Size of persistent storage for AMQ broker volume.	1Gi	True

Variable name	Image Environment Variable	Description	Example value	Required
<b>AMQ_REPLICAS</b>	–	Number of broker replicas for a cluster.	2	True
<b>KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED</b>	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED</b>	If set to true, turns on KIE server global discovery feature (Sets the org.kie.server.controller.openshift.global.discovery.enabled system property)	false	False
<b>KIE_SERVER_CONTROLLER_OPENSHIFT_PREFER_KIESERVER_SERVICE</b>	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_PREFER_KIESERVER_SERVICE</b>	Enables connection to KIE Server via OpenShift internal Service endpoint. (Sets the org.kie.server.controller.openshift.prefer.kieserver.service system property)	true	False
<b>KIE_SERVER_CONTROLLER_TEMPLATE_CACHE_TTL</b>	<b>KIE_SERVER_CONTROLLER_TEMPLATE_CACHE_TTL</b>	KIE ServerTemplate Cache TTL in milliseconds. (Sets the org.kie.server.controller.template.cache.ttl system property)	5000	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>IMAGE_STREAM_NAMESPACE</b>	–	Namespace in which the ImageStreams for Red Hat Process Automation Manager images are installed. These ImageStreams are normally installed in the openshift namespace. You need to modify this parameter only if you installed the ImageStreams in a different namespace/project.	openshift	True
<b>BUSINESS_CENTRAL_IMAGE_STREAM_NAME</b>	–	The name of the image stream to use for Business Central. Default is "rhpm-businesscentral-rhel8".	rhpm-businesscentral-rhel8	True
<b>KIE_SERVER_IMAGE_STREAM_NAME</b>	–	The name of the image stream to use for KIE server. Default is "rhpm-kieserver-rhel8".	rhpm-kieserver-rhel8	True
<b>IMAGE_STREAM_TAG</b>	–	A named pointer to an image in an image stream. Default is "7.8.0".	7.8.0	True

Variable name	Image Environment Variable	Description	Example value	Required
<b>MAVEN_MIRROR_URL</b>	<b>MAVEN_MIRROR_URL</b>	Maven mirror that Business Central and KIE server must use. If you configure a mirror, this mirror must contain all artifacts that are required for building and deploying your services.	–	False
<b>MAVEN_MIRROR_OF</b>	<b>MAVEN_MIRROR_OF</b>	Maven mirror configuration for KIE server.	external:*,!repo-rhpamcentr	False
<b>MAVEN_REPOSITORY_ID</b>	<b>MAVEN_REPOSITORY_ID</b>	The id to use for the maven repository. If set, it can be excluded from the optionally configured mirror by adding it to MAVEN_MIRROR_OF. For example: external:*,!repo-rhpamcentr,!repo-custom. If MAVEN_MIRROR_URL is set but MAVEN_MIRROR_ID is not set, an id will be generated randomly, but won't be usable in MAVEN_MIRROR_OF.	repo-custom	False
<b>MAVEN_REPOSITORY_URL</b>	<b>MAVEN_REPOSITORY_URL</b>	Fully qualified URL to a Maven repository or service.	http://nexus.nexus-project.svc.cluster.local:8081/nexus/content/groups/public/	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>MAVEN_REPO_USERNAME</b>	<b>MAVEN_REPO_USERNAME</b>	User name for accessing the Maven repository, if required.	–	False
<b>MAVEN_REPO_PASSWORD</b>	<b>MAVEN_REPO_PASSWORD</b>	Password to access the Maven repository, if required.	–	False
<b>GIT_HOOKS_DIR</b>	<b>GIT_HOOKS_DIR</b>	The directory to use for git hooks, if required.	<b>/opt/kie/data/git/hooks</b>	False
<b>TIMER_SERVICE_DATA_STORE_REFRESH_INTERVAL</b>	<b>TIMER_SERVICE_DATA_STORE_REFRESH_INTERVAL</b>	Sets refresh-interval for the EJB timer database data-store service.	60000	True
<b>BUSINESS_CENTRAL_VOLUME_CAPACITY</b>	–	Size of the persistent storage for Business Central runtime data.	1Gi	True
<b>BUSINESS_CENTRAL_MEMORY_LIMIT</b>	–	Business Central Container memory limit.	8Gi	True
<b>BUSINESS_CENTRAL_JAVA_MAX_MEMORY_RATIO</b>	<b>JAVA_MAX_MEMORY_RATIO</b>	Business Central Container JVM max memory ratio. <b>-Xmx</b> is set to a ratio of the memory available on the container. The default is 80, which means the upper boundary is 80% of the available memory. To skip adding the <b>-Xmx</b> option, set this value to 0.	80	True

Variable name	Image Environment Variable	Description	Example value	Required
<b>BUSINESS_CENTRAL_CPU_LIMIT</b>	–	Business Central Container CPU limit.	2000m	True
<b>KIE_SERVER_MEMORY_LIMIT</b>	–	KIE server Container memory limit.	1Gi	True
<b>KIE_SERVER_CPU_LIMIT</b>	–	KIE server Container CPU limit.	1000m	True
<b>BUSINESS_CENTRAL_CONTAINER_REPLICAS</b>	–	Business Central Container Replicas, defines how many Business Central containers will be started.	2	True
<b>KIE_SERVER_CONTAINER_REPLICAS</b>	–	KIE Server Container Replicas, defines how many KIE Server containers will be started.	2	True
<b>SSO_URL</b>	<b>SSO_URL</b>	RH-SSO URL.	https://rh-sso.example.com/auth	False
<b>SSO_REALM</b>	<b>SSO_REALM</b>	RH-SSO Realm name.	–	False
<b>BUSINESS_CENTRAL_SSO_CLIENT</b>	<b>SSO_CLIENT</b>	Business Central RH-SSO Client name.	–	False
<b>BUSINESS_CENTRAL_SSO_SECRET</b>	<b>SSO_SECRET</b>	Business Central RH-SSO Client Secret.	252793ed-7118-4ca8-8dab-5622fa97d892	False
<b>KIE_SERVER_SSO_CLIENT</b>	<b>SSO_CLIENT</b>	KIE Server RH-SSO Client name.	–	False



Variable name	Image Environment Variable	Description	Example value	Required
<b>KIE_SERVER_SSO_SECRET</b>	<b>SSO_SECRET</b>	KIE Server RH-SSO Client Secret.	252793ed-7118-4ca8-8dab-5622fa97d892	False
<b>SSO_USERNAME</b>	<b>SSO_USERNAME</b>	RH-SSO Realm admin user name for creating the Client if it doesn't exist.	–	False
<b>SSO_PASSWORD</b>	<b>SSO_PASSWORD</b>	RH-SSO Realm Admin Password used to create the Client.	–	False
<b>SSO_DISABLE_SSL_CERTIFICATE_VALIDATION</b>	<b>SSO_DISABLE_SSL_CERTIFICATE_VALIDATION</b>	RH-SSO Disable SSL Certificate Validation.	false	False
<b>SSO_PRINCIPAL_ATTRIBUTE</b>	<b>SSO_PRINCIPAL_ATTRIBUTE</b>	RH-SSO Principal Attribute to use as user name.	preferred_username	False
<b>AUTH_LDAP_URL</b>	<b>AUTH_LDAP_URL</b>	LDAP Endpoint to connect for authentication.	ldap://myldap.example.com	False
<b>AUTH_LDAP_BIND_DN</b>	<b>AUTH_LDAP_BIND_DN</b>	Bind DN used for authentication.	uid=admin,ou=users,ou=example,ou=com	False
<b>AUTH_LDAP_BIND_CREDENTIAL</b>	<b>AUTH_LDAP_BIND_CREDENTIAL</b>	LDAP Credentials used for authentication.	Password	False
<b>AUTH_LDAP_JAAS_SECURITY_DOMAIN</b>	<b>AUTH_LDAP_JAAS_SECURITY_DOMAIN</b>	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	–	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_BASE_CTX_DN</b>	<b>AUTH_LDAP_BASE_CTX_DN</b>	LDAP Base DN of the top-level context to begin the user search.	ou=users,ou=example,ou=com	False
<b>AUTH_LDAP_BASE_FILTER</b>	<b>AUTH_LDAP_BASE_FILTER</b>	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}).	(uid={0})	False
<b>AUTH_LDAP_SEARCH_SCOPE</b>	<b>AUTH_LDAP_SEARCH_SCOPE</b>	The search scope to use.	<b>SUBTREE_SCOPE</b>	False
<b>AUTH_LDAP_SEARCH_TIME_LIMIT</b>	<b>AUTH_LDAP_SEARCH_TIME_LIMIT</b>	The timeout in milliseconds for user or role searches.	10000	False
<b>AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE</b>	<b>AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE</b>	The name of the attribute in the user entry that contains the DN of the user. This may be necessary if the DN of the user itself contains special characters, backslash for example, that prevent correct user mapping. If the attribute does not exist, the entry's DN is used.	distinguishedName	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_PARSE_USERNAME</b>	<b>AUTH_LDAP_PARSE_USERNAME</b>	A flag indicating if the DN is to be parsed for the user name. If set to true, the DN is parsed for the user name. If set to false the DN is not parsed for the user name. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	true	False
<b>AUTH_LDAP_USERNAME_BEGIN_STRING</b>	<b>AUTH_LDAP_USERNAME_BEGIN_STRING</b>	Defines the String which is to be removed from the start of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	–	False
<b>AUTH_LDAP_USERNAME_END_STRING</b>	<b>AUTH_LDAP_USERNAME_END_STRING</b>	Defines the String which is to be removed from the end of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	–	False
<b>AUTH_LDAP_ROLE_ATTRIBUTE_ID</b>	<b>AUTH_LDAP_ROLE_ATTRIBUTE_ID</b>	Name of the attribute containing the user roles.	memberOf	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_ROLES_CTX_DN</b>	<b>AUTH_LDAP_ROLES_CTX_DN</b>	The fixed DN of the context to search for user roles. This is not the DN where the actual roles are, but the DN where the objects containing the user roles are. For example, in a Microsoft Active Directory server, this is the DN where the user account is.	ou=groups,ou=example,ou=com	False
<b>AUTH_LDAP_ROLE_FILTER</b>	<b>AUTH_LDAP_ROLE_FILTER</b>	A search filter used to locate the roles associated with the authenticated user. The input username or userDN obtained from the login module callback is substituted into the filter anywhere a {0} expression is used. The authenticated userDN is substituted into the filter anywhere a {1} is used. An example search filter that matches on the input username is (member={0}). An alternative that matches on the authenticated userDN is (member={1}).	(memberOf={1})	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_ROLE_RECURSION</b>	<b>AUTH_LDAP_ROLE_RECURSION</b>	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	1	False
<b>AUTH_LDAP_DEFAULT_ROLE</b>	<b>AUTH_LDAP_DEFAULT_ROLE</b>	A role included for all authenticated users	user	False
<b>AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID</b>	<b>AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID</b>	Name of the attribute within the roleCtxDN context which contains the role name. If the roleAttributelsDN property is set to true, this property is used to find the role object's name attribute.	name	False
<b>AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN</b>	<b>AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN</b>	A flag indicating if the DN returned by a query contains the roleNameAttribute ID. If set to true, the DN is checked for the roleNameAttribute ID. If set to false, the DN is not checked for the roleNameAttribute ID. This flag can improve the performance of LDAP queries.	false	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN</b>	<b>AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN</b>	Whether or not the roleAttributeID contains the fully-qualified DN of a role object. If false, the role name is taken from the value of the roleNameAttributeId attribute of the context name. Certain directory schemas, such as Microsoft Active Directory, require this attribute to be set to true.	false	False
<b>AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK</b>	<b>AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK</b>	If you are not using referrals, you can ignore this option. When using referrals, this option denotes the attribute name which contains users defined for a certain role, for example member, if the role object is inside the referral. Users are checked against the content of this attribute name. If this option is not set, the check will always fail, so role objects cannot be stored in a referral tree.	–	False

Variable name	Image Environment Variable	Description	Example value	Required
<b>AUTH_ROLE_MAPPER_ROLES_PROPERTIES</b>	<b>AUTH_ROLE_MAPPER_ROLES_PROPERTIES</b>	When present, the RoleMapping Login Module will be configured to use the provided file. This parameter defines the fully-qualified file path and name of a properties file or resource which maps roles to replacement roles. The format of every entry in the file is <code>original_role=role1,role2,role3</code>	–	False
<b>AUTH_ROLE_MAPPER_REPLACE_ROLE</b>	<b>AUTH_ROLE_MAPPER_REPLACE_ROLE</b>	Whether to add to the current roles, or replace the current roles with the mapped ones. Replaces if set to true.	–	False

## 6.2.2. Objects

The CLI supports various object types. A list of these object types as well as their abbreviations can be found in the [Openshift documentation](#).

### 6.2.2.1. Services

A service is an abstraction which defines a logical set of pods and a policy by which to access them. See the [container-engine documentation](#) for more information.

Service	Port	Name	Description
<b>`\${APPLICATION_NAME}-rhpamcentr</b>	8080	http	All the Business Central web server's ports.
	8443	https	

Service	Port	Name	Description
<b>\${APPLICATION_NAME}-rhpamcentr-ping</b>	8888	ping	The JGroups ping port for rhpamcentr clustering.
<b>\${APPLICATION_NAME}-datagrid-ping</b>	8888	ping	The JGroups ping port for clustering.
<b>\${APPLICATION_NAME}-datagrid</b>	11222	hotrod	Provides a service for accessing the application over Hot Rod protocol.
<b>\${APPLICATION_NAME}-kieserver</b>	8080	http	All the KIE server web server's ports.
	8443	https	
<b>\${APPLICATION_NAME}-amq-tcp</b>	61616	–	The broker's OpenWire port.
<b>ping</b>	8888	–	The JGroups ping port for amq clustering.
<b>\${APPLICATION_NAME}-mysql</b>	3306	–	The MySQL server's port.

### 6.2.2.2. Routes

A route is a way to expose a service by giving it an externally reachable hostname such as **www.example.com**. A defined route and the endpoints identified by its service can be consumed by a router to provide named connectivity from external clients to your applications. Each route consists of a route name, service selector, and (optionally) security configuration. See the [Openshift documentation](#) for more information.

Service	Security	Hostname
insecure- <b>\${APPLICATION_NAME}-rhpamcentr-http</b>	none	<b>\${BUSINESS_CENTRAL_HOSTNAME_HTTP}</b>
<b>\${APPLICATION_NAME}-rhpamcentr-https</b>	TLS passthrough	<b>\${BUSINESS_CENTRAL_HOSTNAME_HTTPS}</b>
insecure- <b>\${APPLICATION_NAME}-kieserver-http</b>	none	<b>\${KIE_SERVER_HOSTNAME_HTTP}</b>



Service	Security	Hostname
<b><code>\${APPLICATION_NAME}-kieserver-https</code></b>	TLS passthrough	<b><code>\${KIE_SERVER_HOSTNAME_HTTPS}</code></b>

### 6.2.2.3. Deployment Configurations

A deployment in OpenShift is a replication controller based on a user-defined template called a deployment configuration. Deployments are created manually or in response to triggered events. See the [OpenShift documentation](#) for more information.

#### 6.2.2.3.1. Triggers

A trigger drives the creation of new deployments in response to events, both inside and outside OpenShift. See the [OpenShift documentation](#) for more information.

Deployment	Triggers
<b><code>\${APPLICATION_NAME}-rhpamcentr</code></b>	ImageChange
<b><code>\${APPLICATION_NAME}-kieserver</code></b>	ImageChange
<b><code>\${APPLICATION_NAME}-mysql</code></b>	ImageChange

#### 6.2.2.3.2. Replicas

A replication controller ensures that a specified number of pod "replicas" are running at any one time. If there are too many, the replication controller kills some pods. If there are too few, it starts more. See the [container-engine documentation](#) for more information.

Deployment	Replicas
<b><code>\${APPLICATION_NAME}-rhpamcentr</code></b>	2
<b><code>\${APPLICATION_NAME}-kieserver</code></b>	2
<b><code>\${APPLICATION_NAME}-mysql</code></b>	1

#### 6.2.2.3.3. Pod Template

##### 6.2.2.3.3.1. Service Accounts

Service accounts are API objects that exist within each project. They can be created or deleted like any other API object. See the [OpenShift documentation](#) for more information.

Deployment	Service Account
<code>\${APPLICATION_NAME}-rhpamcentr</code>	<code>\${APPLICATION_NAME}-rhpamsvc</code>
<code>\${APPLICATION_NAME}-kieserver</code>	<code>\${APPLICATION_NAME}-rhpamsvc</code>

#### 6.2.2.3.3.2. Image

Deployment	Image
<code>\${APPLICATION_NAME}-rhpamcentr</code>	<code>\${BUSINESS_CENTRAL_IMAGE_STREAM_NAME}</code>
<code>\${APPLICATION_NAME}-kieserver</code>	<code>\${KIE_SERVER_IMAGE_STREAM_NAME}</code>
<code>\${APPLICATION_NAME}-mysql</code>	mysql

#### 6.2.2.3.3.3. Readiness Probe

`${APPLICATION_NAME}-rhpamcentr`

Http Get on `http://localhost:8080/rest/ready`

`${APPLICATION_NAME}-kieserver`

Http Get on `http://localhost:8080/services/rest/server/readycheck`

`${APPLICATION_NAME}-mysql`

```
/bin/sh -i -c MYSQL_PWD="${MYSQL_PASSWORD}" mysql -h 127.0.0.1 -u $MYSQL_USER -D
$MYSQL_DATABASE -e 'SELECT 1'
```

#### 6.2.2.3.3.4. Liveness Probe

`${APPLICATION_NAME}-rhpamcentr`

Http Get on `http://localhost:8080/rest/healthy`

`${APPLICATION_NAME}-kieserver`

Http Get on `http://localhost:8080/services/rest/server/healthcheck`

`${APPLICATION_NAME}-mysql`

tcpSocket on port 3306

## 6.2.2.3.3.5. Exposed Ports

Deployments	Name	Port	Protocol
<b>\${APPLICATION_NAME}-rhpamcentr</b>	jolokia	8778	<b>TCP</b>
	http	8080	<b>TCP</b>
	https	8443	<b>TCP</b>
	ping	8888	<b>TCP</b>
<b>\${APPLICATION_NAME}-kieserver</b>	jolokia	8778	<b>TCP</b>
	http	8080	<b>TCP</b>
	https	8443	<b>TCP</b>
<b>\${APPLICATION_NAME}-mysql</b>	–	3306	<b>TCP</b>

## 6.2.2.3.3.6. Image Environment Variables

Deployment	Variable name	Description	Example value
<b>\${APPLICATION_NAME}-rhpamcentr</b>	<b>APPLICATION_USE_RS_PROPERTIES</b>	–	<b>/opt/kie/data/configuration/application-users.properties</b>
	<b>APPLICATION_ROLES_PROPERTIES</b>	–	<b>/opt/kie/data/configuration/application-roles.properties</b>
	<b>KIE_ADMIN_USER</b>	Admin user name	Set according to the credentials secret
	<b>KIE_ADMIN_PWD</b>	Admin user password	Set according to the credentials secret
	<b>KIE_MBEANS</b>	KIE server mbeans enabled/disabled. (Sets the kie.mbeans and kie.scanner.mbeans system properties)	<b>\${KIE_MBEANS}</b>
	<b>KIE_SERVER_CONTROLLER_OPENSIFT_ENABLED</b>	–	true

Deployment	Variable name	Description	Example value
	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED</b>	If set to true, turns on KIE server global discovery feature (Sets the <code>org.kie.server.controller.openshift.global.discovery.enabled</code> system property)	<b><code>\${KIE_SERVER_CONTROLLER_OPENSHIFT_GLOBAL_DISCOVERY_ENABLED}</code></b>
	<b>KIE_SERVER_CONTROLLER_OPENSHIFT_PREFER_KIESERVER_SERVICE</b>	Enables connection to KIE Server via OpenShift internal Service endpoint. (Sets the <code>org.kie.server.controller.openshift.prefer.kieserver.service</code> system property)	<b><code>\${KIE_SERVER_CONTROLLER_OPENSHIFT_PREFER_KIESERVER_SERVICE}</code></b>
	<b>KIE_SERVER_CONTROLLER_TEMPLATE_CACHE_TTL</b>	KIE ServerTemplate Cache TTL in milliseconds. (Sets the <code>org.kie.server.controller.template.cache.ttl</code> system property)	<b><code>\${KIE_SERVER_CONTROLLER_TEMPLATE_CACHE_TTL}</code></b>
	<b>KIE_SERVER_CONTROLLER_TOKEN</b>	KIE server controller token for bearer authentication. (Sets the <code>org.kie.server.controller.token</code> system property)	<b><code>\${KIE_SERVER_CONTROLLER_TOKEN}</code></b>
	<b>WORKBENCH_ROUTE_NAME</b>	–	<b><code>\${APPLICATION_NAME}-rhpamcentr</code></b>
	<b>MAVEN_MIRROR_URL</b>	Maven mirror that Business Central and KIE server must use. If you configure a mirror, this mirror must contain all artifacts that are required for building and deploying your services.	<b><code>\${MAVEN_MIRROR_URL}</code></b>

Deployment	Variable name	Description	Example value
	<b>MAVEN_REPO_ID</b>	The id to use for the maven repository. If set, it can be excluded from the optionally configured mirror by adding it to MAVEN_MIRROR_OF. For example: external:*,!repo-rhpamcentr,!repo-custom. If MAVEN_MIRROR_URL is set but MAVEN_MIRROR_ID is not set, an id will be generated randomly, but won't be usable in MAVEN_MIRROR_OF.	<b>\${MAVEN_REPO_ID}</b>
	<b>MAVEN_REPO_URL</b>	Fully qualified URL to a Maven repository or service.	<b>\${MAVEN_REPO_URL}</b>
	<b>MAVEN_REPO_USERNAME</b>	User name for accessing the Maven repository, if required.	<b>\${MAVEN_REPO_USERNAME}</b>
	<b>MAVEN_REPO_PASSWORD</b>	Password to access the Maven repository, if required.	<b>\${MAVEN_REPO_PASSWORD}</b>
	<b>GIT_HOOKS_DIR</b>	The directory to use for git hooks, if required.	<b>\${GIT_HOOKS_DIR}</b>
	<b>HTTPS_KEYSTORE_DIR</b>	–	<b>/etc/businesscentral-secret-volume</b>
	<b>HTTPS_KEYSTORE</b>	The name of the keystore file within the secret for Business Central.	<b>\${BUSINESS_CENTRAL_HTTPS_KEYSTORE}</b>
	<b>HTTPS_NAME</b>	The name associated with the server certificate for Business Central.	<b>\${BUSINESS_CENTRAL_HTTPS_NAME}</b>

Deployment	Variable name	Description	Example value
	<b>HTTPS_PASSWORD</b>	The password for the keystore and certificate for Business Central.	<b>\${BUSINESS_CENTRAL_HTTPS_PASSWORD}</b>
	<b>JGROUPS_PING_PROTOCOL</b>	–	openshift.DNS_PING
	<b>OPENSIFT_DNS_PING_SERVICE_NAME</b>	–	<b>\${APPLICATION_NAME}-rhpamcentr-ping</b>
	<b>OPENSIFT_DNS_PING_SERVICE_PORT</b>	–	8888
	<b>APPFORMER_INFISPAN_SERVICE_NAME</b>	–	<b>\${APPLICATION_NAME}-datagrid</b>
	<b>APPFORMER_INFISPAN_PORT</b>	–	11222
	<b>APPFORMER_JMS_BROKER_ADDRESS</b>	–	<b>\${APPLICATION_NAME}-amq-tcp</b>
	<b>APPFORMER_JMS_BROKER_PORT</b>	–	61616
	<b>APPFORMER_JMS_BROKER_USER</b>	The user name for connecting to the JMS broker.	<b>\${APPFORMER_JMS_BROKER_USER}</b>
	<b>APPFORMER_JMS_BROKER_PASSWORD</b>	The password to connect to the JMS broker.	<b>\${APPFORMER_JMS_BROKER_PASSWORD}</b>
	<b>JAVA_MAX_MEMORY_RATIO</b>	Business Central Container JVM max memory ratio. <b>-Xmx</b> is set to a ratio of the memory available on the container. The default is 80, which means the upper boundary is 80% of the available memory. To skip adding the <b>-Xmx</b> option, set this value to 0.	<b>\${BUSINESS_CENTRAL_JAVA_MAX_MEMORY_RATIO}</b>

Deployment	Variable name	Description	Example value
	<b>SSO_URL</b>	RH-SSO URL.	<b>\${SSO_URL}</b>
	<b>SSO_OPENIDCONNECT_DEPLOYMENTS</b>	–	ROOT.war
	<b>SSO_REALM</b>	RH-SSO Realm name.	<b>\${SSO_REALM}</b>
	<b>SSO_SECRET</b>	Business Central RH-SSO Client Secret.	<b>\${BUSINESS_CENTRAL_SSO_SECRET}</b>
	<b>SSO_CLIENT</b>	Business Central RH-SSO Client name.	<b>\${BUSINESS_CENTRAL_SSO_CLIENT}</b>
	<b>SSO_USERNAME</b>	RH-SSO Realm admin user name for creating the Client if it doesn't exist.	<b>\${SSO_USERNAME}</b>
	<b>SSO_PASSWORD</b>	RH-SSO Realm Admin Password used to create the Client.	<b>\${SSO_PASSWORD}</b>
	<b>SSO_DISABLE_SSL_CERTIFICATE_VALIDATION</b>	RH-SSO Disable SSL Certificate Validation.	<b>\${SSO_DISABLE_SSL_CERTIFICATE_VALIDATION}</b>
	<b>SSO_PRINCIPAL_ATTRIBUTE</b>	RH-SSO Principal Attribute to use as user name.	<b>\${SSO_PRINCIPAL_ATTRIBUTE}</b>
	<b>HOSTNAME_HTTP</b>	Custom hostname for http service route for Business Central. Leave blank for default hostname, e.g.: insecure-<application-name>-rhpamcentr-<project>.<default-domain-suffix>	<b>\${BUSINESS_CENTRAL_HOSTNAME_HTTP}</b>

Deployment	Variable name	Description	Example value
	<b>HOSTNAME_HTTPS</b>	Custom hostname for https service route for Business Central. Leave blank for default hostname, e.g.: <application-name>-rhpamcentr-<project>.<default-domain-suffix>	<b>`\${BUSINESS_CENTRAL_HOSTNAME_HTTPS}`</b>
	<b>AUTH_LDAP_URL</b>	LDAP Endpoint to connect for authentication.	<b>`\${AUTH_LDAP_URL}`</b>
	<b>AUTH_LDAP_BIND_DN</b>	Bind DN used for authentication.	<b>`\${AUTH_LDAP_BIND_DN}`</b>
	<b>AUTH_LDAP_BIND_CREDENTIAL</b>	LDAP Credentials used for authentication.	<b>`\${AUTH_LDAP_BIND_CREDENTIAL}`</b>
	<b>AUTH_LDAP_JAAS_SECURITY_DOMAIN</b>	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	<b>`\${AUTH_LDAP_JAAS_SECURITY_DOMAIN}`</b>
	<b>AUTH_LDAP_BASE_CTX_DN</b>	LDAP Base DN of the top-level context to begin the user search.	<b>`\${AUTH_LDAP_BASE_CTX_DN}`</b>
	<b>AUTH_LDAP_BASE_FILTER</b>	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}).	<b>`\${AUTH_LDAP_BASE_FILTER}`</b>
	<b>AUTH_LDAP_SEARCH_SCOPE</b>	The search scope to use.	<b>`\${AUTH_LDAP_SEARCH_SCOPE}`</b>
	<b>AUTH_LDAP_SEARCH_TIME_LIMIT</b>	The timeout in milliseconds for user or role searches.	<b>`\${AUTH_LDAP_SEARCH_TIME_LIMIT}`</b>



Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE</b>	The name of the attribute in the user entry that contains the DN of the user. This may be necessary if the DN of the user itself contains special characters, backslash for example, that prevent correct user mapping. If the attribute does not exist, the entry's DN is used.	<b>\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}</b>
	<b>AUTH_LDAP_PARSE_USERNAME</b>	A flag indicating if the DN is to be parsed for the user name. If set to true, the DN is parsed for the user name. If set to false the DN is not parsed for the user name. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	<b>\${AUTH_LDAP_PARSE_USERNAME}</b>
	<b>AUTH_LDAP_USERNAME_BEGIN_STRING</b>	Defines the String which is to be removed from the start of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<b>\${AUTH_LDAP_USERNAME_BEGIN_STRING}</b>
	<b>AUTH_LDAP_USERNAME_END_STRING</b>	Defines the String which is to be removed from the end of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<b>\${AUTH_LDAP_USERNAME_END_STRING}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_ROLE_ATTRIBUTE_ID</b>	Name of the attribute containing the user roles.	<b>`\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}`</b>
	<b>AUTH_LDAP_ROLE_S_CTX_DN</b>	The fixed DN of the context to search for user roles. This is not the DN where the actual roles are, but the DN where the objects containing the user roles are. For example, in a Microsoft Active Directory server, this is the DN where the user account is.	<b>`\${AUTH_LDAP_ROLE_S_CTX_DN}`</b>
	<b>AUTH_LDAP_ROLE_FILTER</b>	A search filter used to locate the roles associated with the authenticated user. The input username or userDN obtained from the login module callback is substituted into the filter anywhere a <code>{0}</code> expression is used. The authenticated userDN is substituted into the filter anywhere a <code>{1}</code> is used. An example search filter that matches on the input username is <code>(member={0})</code> . An alternative that matches on the authenticated userDN is <code>(member={1})</code> .	<b>`\${AUTH_LDAP_ROLE_FILTER}`</b>
	<b>AUTH_LDAP_ROLE_RECURSION</b>	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	<b>`\${AUTH_LDAP_ROLE_RECURSION}`</b>
	<b>AUTH_LDAP_DEFAULT_ROLE</b>	A role included for all authenticated users	<b>`\${AUTH_LDAP_DEFAULT_ROLE}`</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID</b>	Name of the attribute within the roleCtxDN context which contains the role name. If the roleAttributesDN property is set to true, this property is used to find the role object's name attribute.	<b>\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}</b>
	<b>AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN</b>	A flag indicating if the DN returned by a query contains the roleNameAttributeID. If set to true, the DN is checked for the roleNameAttributeID. If set to false, the DN is not checked for the roleNameAttributeID. This flag can improve the performance of LDAP queries.	<b>\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}</b>
	<b>AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN</b>	Whether or not the roleAttributeID contains the fully-qualified DN of a role object. If false, the role name is taken from the value of the roleNameAttributeID attribute of the context name. Certain directory schemas, such as Microsoft Active Directory, require this attribute to be set to true.	<b>\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK</b>	If you are not using referrals, you can ignore this option. When using referrals, this option denotes the attribute name which contains users defined for a certain role, for example member, if the role object is inside the referral. Users are checked against the content of this attribute name. If this option is not set, the check will always fail, so role objects cannot be stored in a referral tree.	<b>\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}</b>
	<b>AUTH_ROLE_MAPPER_ROLES_PROPERTIES</b>	When present, the RoleMapping Login Module will be configured to use the provided file. This parameter defines the fully-qualified file path and name of a properties file or resource which maps roles to replacement roles. The format of every entry in the file is original_role=role1,role2,role3	<b>\${AUTH_ROLE_MAPPER_ROLES_PROPERTIES}</b>
	<b>AUTH_ROLE_MAPPER_REPLACE_ROLE</b>	Whether to add to the current roles, or replace the current roles with the mapped ones. Replaces if set to true.	<b>\${AUTH_ROLE_MAPPER_REPLACE_ROLE}</b>
<b>\${APPLICATION_NAME}-kieserver</b>	<b>WORKBENCH_SERVICE_NAME</b>	–	<b>\${APPLICATION_NAME}-rhpamcentr</b>
	<b>TIMER_SERVICE_DATA_STORE_REFRESH_INTERVAL</b>	Sets refresh-interval for the EJB timer database data-store service.	<b>\${TIMER_SERVICE_DATA_STORE_REFRESH_INTERVAL}</b>
	<b>DATASOURCES</b>	–	<b>RHPAM</b>

Deployment	Variable name	Description	Example value
	<b>RHPAM_DATABASE</b>	MySQL database name.	<b>`\${MYSQL_DB}`</b>
	<b>RHPAM_DRIVER</b>	–	mariadb
	<b>RHPAM_USERNAME</b>	MySQL database user name.	<b>`\${MYSQL_USER}`</b>
	<b>RHPAM_PASSWORD</b>	MySQL database password.	<b>`\${MYSQL_PWD}`</b>
	<b>RHPAM_SERVICE_HOST</b>	–	<b>`\${APPLICATION_NAME}-mysql`</b>
	<b>RHPAM_SERVICE_PORT</b>	–	3306
	<b>KIE_SERVER_PERSISTENCE_DIALECT</b>	KIE server MySQL Hibernate dialect.	<b>`\${KIE_SERVER_MYSQL_DIALECT}`</b>
	<b>KIE_SERVER_PERSISTENCE_DS</b>	KIE server persistence datasource. (Sets the org.kie.server.persistence.ds system property)	<b>`\${KIE_SERVER_PERSISTENCE_DS}`</b>
	<b>RHPAM_JNDI</b>	KIE server persistence datasource. (Sets the org.kie.server.persistence.ds system property)	<b>`\${KIE_SERVER_PERSISTENCE_DS}`</b>
	<b>RHPAM_JTA</b>	–	true
	<b>KIE_ADMIN_USER</b>	Admin user name	Set according to the credentials secret
	<b>KIE_ADMIN_PWD</b>	Admin user password	Set according to the credentials secret
	<b>KIE_MBEANS</b>	KIE server mbeans enabled/disabled. (Sets the kie.mbeans and kie.scanner.mbeans system properties)	<b>`\${KIE_MBEANS}`</b>

Deployment	Variable name	Description	Example value
	<b>KIE_SERVER_MODE</b>	The KIE Server mode. Valid values are 'DEVELOPMENT' or 'PRODUCTION'. In production mode, you can not deploy SNAPSHOT versions of artifacts on the KIE server and can not change the version of an artifact in an existing container. (Sets the org.kie.server.mode system property).	<b>`\${KIE_SERVER_MODE}`</b>
	<b>DROOLS_SERVER_FILTER_CLASSES</b>	KIE server class filtering. (Sets the org.drools.server.filter.classes system property)	<b>`\${DROOLS_SERVER_FILTER_CLASSES}`</b>
	<b>PROMETHEUS_SERVER_EXT_DISABLED</b>	If set to false, the prometheus server extension will be enabled. (Sets the org.kie.prometheus.server.ext.disabled system property)	<b>`\${PROMETHEUS_SERVER_EXT_DISABLED}`</b>
	<b>KIE_SERVER_BYPASS_AUTH_USER</b>	Allows the KIE server to bypass the authenticated user for task-related operations, for example, queries. (Sets the org.kie.server.bypass.auth.user system property)	<b>`\${KIE_SERVER_BYPASS_AUTH_USER}`</b>
	<b>KIE_SERVER_CONTROLLER_SERVICE</b>	–	<b>`\${APPLICATION_NAME}-rhpamcentr</b>
	<b>KIE_SERVER_CONTROLLER_PROTOCOL</b>	–	ws
	<b>KIE_SERVER_ID</b>	–	–
	<b>KIE_SERVER_ROUTE_NAME</b>	–	insecure- <b>`\${APPLICATION_NAME}`</b> -kieserver

Deployment	Variable name	Description	Example value
	<b>KIE_SERVER_STARTUP_STRATEGY</b>	–	OpenShiftStartupStrategy
	<b>MAVEN_MIRROR_URL</b>	Maven mirror that Business Central and KIE server must use. If you configure a mirror, this mirror must contain all artifacts that are required for building and deploying your services.	<b>\${MAVEN_MIRROR_URL}</b>
	<b>MAVEN_MIRROR_OFF</b>	Maven mirror configuration for KIE server.	<b>\${MAVEN_MIRROR_OFF}</b>
	<b>MAVEN_REPOS</b>	–	RHPAMCENTR,EXTERNAL
	<b>RHPAMCENTR_MAVEN_REPO_ID</b>	–	repo-rhpamcentr
	<b>RHPAMCENTR_MAVEN_REPO_SERVICE</b>	–	<b>\${APPLICATION_NAME}-rhpamcentr</b>
	<b>RHPAMCENTR_MAVEN_REPO_PATH</b>	–	<b>/maven2/</b>
	<b>RHPAMCENTR_MAVEN_REPO_USERNAME</b>	–	Set according to the credentials secret
	<b>RHPAMCENTR_MAVEN_REPO_PASSWORD</b>	–	Set according to the credentials secret

Deployment	Variable name	Description	Example value
	<b>EXTERNAL_MAVEN_REPO_ID</b>	The id to use for the maven repository. If set, it can be excluded from the optionally configured mirror by adding it to MAVEN_MIRROR_OF. For example: external:*,!repo-rhpamcentr,!repo-custom. If MAVEN_MIRROR_URL is set but MAVEN_MIRROR_ID is not set, an id will be generated randomly, but won't be usable in MAVEN_MIRROR_OF.	<b>\${MAVEN_REPO_ID}</b>
	<b>EXTERNAL_MAVEN_REPO_URL</b>	Fully qualified URL to a Maven repository or service.	<b>\${MAVEN_REPO_URL}</b>
	<b>EXTERNAL_MAVEN_REPO_USERNAME</b>	User name for accessing the Maven repository, if required.	<b>\${MAVEN_REPO_USERNAME}</b>
	<b>EXTERNAL_MAVEN_REPO_PASSWORD</b>	Password to access the Maven repository, if required.	<b>\${MAVEN_REPO_PASSWORD}</b>
	<b>HTTPS_KEYSTORE_DIR</b>	–	<b>/etc/kieserver-secret-volume</b>
	<b>HTTPS_KEYSTORE</b>	The name of the keystore file within the secret for KIE Server.	<b>\${KIE_SERVER_HTTPS_KEYSTORE}</b>
	<b>HTTPS_NAME</b>	The name associated with the server certificate for KIE Server.	<b>\${KIE_SERVER_HTTPS_NAME}</b>
	<b>HTTPS_PASSWORD</b>	The password for the keystore and certificate for KIE Server.	<b>\${KIE_SERVER_HTTPS_PASSWORD}</b>
	<b>SSO_URL</b>	RH-SSO URL.	<b>\${SSO_URL}</b>



Deployment	Variable name	Description	Example value
	<b>SSO_OPENIDCONNECT_DEPLOYMENTS</b>	–	ROOT.war
	<b>SSO_REALM</b>	RH-SSO Realm name.	<b>\${SSO_REALM}</b>
	<b>SSO_SECRET</b>	KIE Server RH-SSO Client Secret.	<b>\${KIE_SERVER_SSO_SECRET}</b>
	<b>SSO_CLIENT</b>	KIE Server RH-SSO Client name.	<b>\${KIE_SERVER_SSO_CLIENT}</b>
	<b>SSO_USERNAME</b>	RH-SSO Realm admin user name for creating the Client if it doesn't exist.	<b>\${SSO_USERNAME}</b>
	<b>SSO_PASSWORD</b>	RH-SSO Realm Admin Password used to create the Client.	<b>\${SSO_PASSWORD}</b>
	<b>SSO_DISABLE_SSL_CERTIFICATE_VALIDATION</b>	RH-SSO Disable SSL Certificate Validation.	<b>\${SSO_DISABLE_SSL_CERTIFICATE_VALIDATION}</b>
	<b>SSO_PRINCIPAL_ATTRIBUTE</b>	RH-SSO Principal Attribute to use as user name.	<b>\${SSO_PRINCIPAL_ATTRIBUTE}</b>
	<b>HOSTNAME_HTTP</b>	Custom hostname for http service route for KIE Server. Leave blank for default hostname, e.g.: insecure-<application-name>-kieserver-<project>.<default-domain-suffix>	<b>\${KIE_SERVER_HOSTNAME_HTTP}</b>
	<b>HOSTNAME_HTTPS</b>	Custom hostname for https service route for KIE Server. Leave blank for default hostname, e.g.: <application-name>-kieserver-<project>.<default-domain-suffix>	<b>\${KIE_SERVER_HOSTNAME_HTTPS}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_URL</b>	LDAP Endpoint to connect for authentication.	<b>\${AUTH_LDAP_URL}</b>
	<b>AUTH_LDAP_BIND_DN</b>	Bind DN used for authentication.	<b>\${AUTH_LDAP_BIND_DN}</b>
	<b>AUTH_LDAP_BIND_CREDENTIAL</b>	LDAP Credentials used for authentication.	<b>\${AUTH_LDAP_BIND_CREDENTIAL}</b>
	<b>AUTH_LDAP_JAAS_SECURITY_DOMAIN</b>	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	<b>\${AUTH_LDAP_JAAS_SECURITY_DOMAIN}</b>
	<b>AUTH_LDAP_BASE_CTX_DN</b>	LDAP Base DN of the top-level context to begin the user search.	<b>\${AUTH_LDAP_BASE_CTX_DN}</b>
	<b>AUTH_LDAP_BASE_FILTER</b>	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}).	<b>\${AUTH_LDAP_BASE_FILTER}</b>
	<b>AUTH_LDAP_SEARCH_SCOPE</b>	The search scope to use.	<b>\${AUTH_LDAP_SEARCH_SCOPE}</b>
	<b>AUTH_LDAP_SEARCH_TIME_LIMIT</b>	The timeout in milliseconds for user or role searches.	<b>\${AUTH_LDAP_SEARCH_TIME_LIMIT}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE</b>	The name of the attribute in the user entry that contains the DN of the user. This may be necessary if the DN of the user itself contains special characters, backslash for example, that prevent correct user mapping. If the attribute does not exist, the entry's DN is used.	<b>`\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}`</b>
	<b>AUTH_LDAP_PARSE_USERNAME</b>	A flag indicating if the DN is to be parsed for the user name. If set to true, the DN is parsed for the user name. If set to false the DN is not parsed for the user name. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	<b>`\${AUTH_LDAP_PARSE_USERNAME}`</b>
	<b>AUTH_LDAP_USERNAME_BEGIN_STRING</b>	Defines the String which is to be removed from the start of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<b>`\${AUTH_LDAP_USERNAME_BEGIN_STRING}`</b>
	<b>AUTH_LDAP_USERNAME_END_STRING</b>	Defines the String which is to be removed from the end of the DN to reveal the user name. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<b>`\${AUTH_LDAP_USERNAME_END_STRING}`</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_ROLE_ATTRIBUTE_ID</b>	Name of the attribute containing the user roles.	<b>`\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}`</b>
	<b>AUTH_LDAP_ROLE_S_CTX_DN</b>	The fixed DN of the context to search for user roles. This is not the DN where the actual roles are, but the DN where the objects containing the user roles are. For example, in a Microsoft Active Directory server, this is the DN where the user account is.	<b>`\${AUTH_LDAP_ROLE_S_CTX_DN}`</b>
	<b>AUTH_LDAP_ROLE_FILTER</b>	A search filter used to locate the roles associated with the authenticated user. The input username or userDN obtained from the login module callback is substituted into the filter anywhere a <code>{0}</code> expression is used. The authenticated userDN is substituted into the filter anywhere a <code>{1}</code> is used. An example search filter that matches on the input username is <code>(member={0})</code> . An alternative that matches on the authenticated userDN is <code>(member={1})</code> .	<b>`\${AUTH_LDAP_ROLE_FILTER}`</b>
	<b>AUTH_LDAP_ROLE_RECURSION</b>	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	<b>`\${AUTH_LDAP_ROLE_RECURSION}`</b>
	<b>AUTH_LDAP_DEFAULT_ROLE</b>	A role included for all authenticated users	<b>`\${AUTH_LDAP_DEFAULT_ROLE}`</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID</b>	Name of the attribute within the roleCtxDN context which contains the role name. If the roleAttributesDN property is set to true, this property is used to find the role object's name attribute.	<b>\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}</b>
	<b>AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN</b>	A flag indicating if the DN returned by a query contains the roleNameAttributeID. If set to true, the DN is checked for the roleNameAttributeID. If set to false, the DN is not checked for the roleNameAttributeID. This flag can improve the performance of LDAP queries.	<b>\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}</b>
	<b>AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN</b>	Whether or not the roleAttributeID contains the fully-qualified DN of a role object. If false, the role name is taken from the value of the roleNameAttributeID attribute of the context name. Certain directory schemas, such as Microsoft Active Directory, require this attribute to be set to true.	<b>\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}</b>

Deployment	Variable name	Description	Example value
	<b>AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK</b>	If you are not using referrals, you can ignore this option. When using referrals, this option denotes the attribute name which contains users defined for a certain role, for example member, if the role object is inside the referral. Users are checked against the content of this attribute name. If this option is not set, the check will always fail, so role objects cannot be stored in a referral tree.	<b>`\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}`</b>
	<b>AUTH_ROLE_MAPPER_ROLES_PROPERTIES</b>	When present, the RoleMapping Login Module will be configured to use the provided file. This parameter defines the fully-qualified file path and name of a properties file or resource which maps roles to replacement roles. The format of every entry in the file is original_role=role1,role2,role3	<b>`\${AUTH_ROLE_MAPPER_ROLES_PROPERTIES}`</b>
	<b>AUTH_ROLE_MAPPER_REPLACE_ROLE</b>	Whether to add to the current roles, or replace the current roles with the mapped ones. Replaces if set to true.	<b>`\${AUTH_ROLE_MAPPER_REPLACE_ROLE}`</b>
<b>`\${APPLICATION_NAME}`-mysql</b>	<b>MYSQL_USER</b>	MySQL database user name.	<b>`\${MYSQL_USER}`</b>
	<b>MYSQL_PASSWORD</b>	MySQL database password.	<b>`\${MYSQL_PWD}`</b>
	<b>MYSQL_DATABASE</b>	MySQL database name.	<b>`\${MYSQL_DB}`</b>

Deployment	Variable name	Description	Example value
	<b>MYSQL_DEFAULT_AUTHENTICATION_PLUGIN</b>	–	mysql_native_password

#### 6.2.2.3.3.7. Volumes

Deployment	Name	mountPath	Purpose	readOnly
<b>\${APPLICATION_NAME}-rhpmcentr</b>	businesscentral-keystore-volume	<b>/etc/businesscentral-secret-volume</b>	ssl certs	True
<b>\${APPLICATION_NAME}-kieserver</b>	kieserver-keystore-volume	<b>/etc/kieserver-secret-volume</b>	ssl certs	True
<b>\${APPLICATION_NAME}-mysql</b>	<b>\${APPLICATION_NAME}-mysql-pvol</b>	<b>/var/lib/mysql/data</b>	mysql	false

#### 6.2.2.4. External Dependencies

##### 6.2.2.4.1. Volume Claims

A **PersistentVolume** object is a storage resource in an OpenShift cluster. Storage is provisioned by an administrator by creating **PersistentVolume** objects from sources such as GCE Persistent Disks, AWS Elastic Block Stores (EBS), and NFS mounts. See the [OpenShift documentation](#) for more information.

Name	Access Mode
<b>\${APPLICATION_NAME}-rhpmcentr-claim</b>	ReadWriteMany
<b>\${APPLICATION_NAME}-mysql-claim</b>	ReadWriteOnce

##### 6.2.2.4.2. Secrets

This template requires the following secrets to be installed for the application to run.

businesscentral-app-secret kieserver-app-secret

##### 6.2.2.4.3. Clustering

Clustering in OpenShift EAP is achieved through one of two discovery mechanisms: Kubernetes or DNS. This is done by configuring the JGroups protocol stack in standalone-openshift.xml with either the **<openshift.KUBE\_PING/>** or **<openshift.DNS\_PING/>** elements. The templates are configured to use **DNS\_PING**, however ``KUBE_PING`` is the default used by the image.

The discovery mechanism used is specified by the **JGROUPS\_PING\_PROTOCOL** environment variable which can be set to either **openshift.DNS\_PING** or **openshift.KUBE\_PING**. **openshift.KUBE\_PING** is the default used by the image if no value is specified for **JGROUPS\_PING\_PROTOCOL**.

For **DNS\_PING** to work, the following steps must be taken:

1. The **OPENSIFT\_DNS\_PING\_SERVICE\_NAME** environment variable must be set to the name of the ping service for the cluster (see table above). If not set, the server will act as if it is a single-node cluster (a "cluster of one").
2. The **OPENSIFT\_DNS\_PING\_SERVICE\_PORT** environment variables should be set to the port number on which the ping service is exposed (see table above). The **DNS\_PING** protocol will attempt to discern the port from the SRV records, if it can, otherwise it will default to 8888.
3. A ping service which exposes the ping port must be defined. This service should be "headless" (ClusterIP=None) and must have the following:
  - a. The port must be named for port discovery to work.
  - b. It must be annotated with **service.alpha.kubernetes.io/tolerate-unready-endpoints** set to **"true"**. Omitting this annotation will result in each node forming their own "cluster of one" during startup, then merging their cluster into the other nodes' clusters after startup (as the other nodes are not detected until after they have started).

### Example ping service for use with **DNS\_PING**

```
kind: Service
apiVersion: v1
spec:
  clusterIP: None
  ports:
  - name: ping
    port: 8888
  selector:
    deploymentConfig: eap-app
metadata:
  name: eap-app-ping
  annotations:
    service.alpha.kubernetes.io/tolerate-unready-endpoints: "true"
    description: "The JGroups ping port for clustering."
```

For **KUBE\_PING** to work, the following steps must be taken:

1. The **OPENSIFT\_KUBE\_PING\_NAMESPACE** environment variable must be set (see table above). If not set, the server will act as if it is a single-node cluster (a "cluster of one").
2. The **OPENSIFT\_KUBE\_PING\_LABELS** environment variables should be set (see table above). If not set, pods outside of your application (albeit in your namespace) will try to join.
3. Authorization must be granted to the service account the pod is running under to be allowed to access Kubernetes' REST api. This is done on the command line.

### Example 6.1. Policy commands

Using the default service account in the myproject namespace:



```
oc policy add-role-to-user view system:serviceaccount:myproject:default -n myproject
```

Using the eap-service-account in the myproject namespace:

```
oc policy add-role-to-user view system:serviceaccount:myproject:eap-service-account -n myproject
```

## 6.3. OPENSIFT USAGE QUICK REFERENCE

To deploy, monitor, manage, and undeploy Red Hat Process Automation Manager templates on Red Hat OpenShift Container Platform, you can use the OpenShift Web console or the **oc** command.

For instructions about using the Web console, see [Create and build an image using the Web console](#).

For detailed instructions about using the **oc** command, see [CLI Reference](#). The following commands are likely to be required:

- To create a project, use the following command:

```
$ oc new-project <project-name>
```

For more information, see [Creating a project using the CLI](#).

- To deploy a template (create an application from a template), use the following command:

```
$ oc new-app -f <template-name> -p <parameter>=<value> -p <parameter>=<value> ...
```

For more information, see [Creating an application using the CLI](#).

- To view a list of the active pods in the project, use the following command:

```
$ oc get pods
```

- To view the current status of a pod, including information whether or not the pod deployment has completed and it is now in a running state, use the following command:

```
$ oc describe pod <pod-name>
```

You can also use the **oc describe** command to view the current status of other objects. For more information, see [Application modification operations](#).

- To view the logs for a pod, use the following command:

```
$ oc logs <pod-name>
```

- To view deployment logs, look up a **DeploymentConfig** name in the template reference and enter the following command:

```
$ oc logs -f dc/<deployment-config-name>
```

For more information, see [Viewing deployment logs](#).

- To view build logs, look up a **BuildConfig** name in the template reference and enter the command:

```
$ oc logs -f bc/<build-config-name>
```

For more information, see [Accessing build logs](#).

- To scale a pod in the application, look up a **DeploymentConfig** name in the template reference and enter the command:

```
$ oc scale dc/<deployment-config-name> --replicas=<number>
```

For more information, see [Manual scaling](#).

- To undeploy the application, you can delete the project by using the command:

```
$ oc delete project <project-name>
```

Alternatively, you can use the **oc delete** command to remove any part of the application, such as a pod or replication controller. For details, see [Application modification operations](#).

## APPENDIX A. VERSIONING INFORMATION

Documentation last updated on Friday, June 25, 2021.