



Red Hat Process Automation Manager 7.1

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

Red Hat Process Automation Manager 7.1 Deploying a Red Hat Process Automation Manager authoring environment on Red Hat OpenShift Container Platform

Red Hat Customer Content Services

brms-docs@redhat.com

Legal Notice

Copyright © 2021 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This document describes how to deploy a Red Hat Process Automation Manager 7.1 authoring environment on Red Hat OpenShift Container Platform.

Table of Contents

PREFACE	4
CHAPTER 1. OVERVIEW OF RED HAT PROCESS AUTOMATION MANAGER ON RED HAT OPENSIFT CONTAINER PLATFORM	5
CHAPTER 2. PREPARING TO DEPLOY RED HAT PROCESS AUTOMATION MANAGER IN YOUR OPENSIFT ENVIRONMENT	7
2.1. ENSURING THE AVAILABILITY OF IMAGE STREAMS	7
2.2. CREATING THE SECRETS FOR PROCESS SERVER	7
2.3. CREATING THE SECRETS FOR BUSINESS CENTRAL	8
2.4. CHANGING GLUSTERFS CONFIGURATION	8
CHAPTER 3. AUTHORIZING ENVIRONMENT	10
3.1. DEPLOYING A SINGLE AUTHORIZING ENVIRONMENT	10
3.2. DEPLOYING A HIGH-AVAILABILITY AUTHORIZING ENVIRONMENT	15
3.3. PROVIDING THE GIT HOOKS DIRECTORY	20
3.4. MODIFYING THE TEMPLATE FOR THE SINGLE AUTHORIZING ENVIRONMENT	22
3.5. MODIFYING THE TEMPLATE FOR THE HIGH AVAILABILITY AUTHORIZING ENVIRONMENT	24
3.6. BUILDING A CUSTOM PROCESS SERVER IMAGE FOR AN EXTERNAL DATABASE	25
CHAPTER 4. OPENSIFT TEMPLATE REFERENCE INFORMATION	28
4.1. RHPAM71-AUTHORIZING	28
4.1.1. Parameters	28
4.1.2. Objects	40
4.1.2.1. Services	40
4.1.2.2. Routes	41
4.1.2.3. Deployment Configurations	41
4.1.2.3.1. Triggers	41
4.1.2.3.2. Replicas	41
4.1.2.3.3. Pod Template	42
4.1.2.3.3.1. Service Accounts	42
4.1.2.3.3.2. Image	42
4.1.2.3.3.3. Readiness Probe	42
4.1.2.3.3.4. Liveness Probe	42
4.1.2.3.3.5. Exposed Ports	43
4.1.2.3.3.6. Image Environment Variables	43
4.1.2.3.3.7. Volumes	59
4.1.2.4. External Dependencies	59
4.1.2.4.1. Volume Claims	59
4.1.2.4.2. Secrets	60
4.2. RHPAM71-AUTHORIZING-HA	60
4.2.1. Parameters	60
4.2.2. Objects	76
4.2.2.1. Services	76
4.2.2.2. Routes	77
4.2.2.3. Deployment Configurations	77
4.2.2.3.1. Triggers	77
4.2.2.3.2. Replicas	78
4.2.2.3.3. Pod Template	78
4.2.2.3.3.1. Service Accounts	78
4.2.2.3.3.2. Image	78
4.2.2.3.3.3. Readiness Probe	79

4.2.2.3.3.4. Liveness Probe	79
4.2.2.3.3.5. Exposed Ports	80
4.2.2.3.3.6. Image Environment Variables	80
4.2.2.3.3.7. Volumes	99
4.2.2.4. External Dependencies	99
4.2.2.4.1. Volume Claims	99
4.2.2.4.2. Secrets	100
4.2.2.4.3. Clustering	100
4.3. OPENSIFT USAGE QUICK REFERENCE	101
APPENDIX A. VERSIONING INFORMATION	103

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

- At least four gigabytes of memory must be available in the OpenShift cluster/namespace.
- The OpenShift project for the deployment must be created.
- You must be 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 must be enabled. Alternatively, if dynamic PV provisioning is not enabled, enough persistent volumes must be available. By default, the following sizes are required:
 - The replicated set of Process 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 use the Authoring High Availability template, which scales the Business Central pod:
 - The image streams for Red Hat AMQ version 7.1 or later are available in your OpenShift environment.
 - Your OpenShift environment supports persistent volumes with ReadWriteMany mode. For information about access mode support in OpenShift Online volume plug-ins, see [Access Modes](#).

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, providing as few or as many containers as necessary 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:

- Process Server, also known as *Execution Server* or *KIE 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 Process 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, Process Server can use an H2 database; in this case, the pod cannot be scaled.

You can freely scale up a Process Server pod, providing as many copies as necessary, running on the same host or different hosts. As you scale a pod up or down, all 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 Process 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 Process Server pods as necessary.

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



IMPORTANT

In the current version, high-availability Business Central functionality is a technology preview.

- Business Central Monitoring is a web-based management and monitoring console. It can manage deployment of services to Process 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 Process Servers and other components that interact with them. It is required if you want Business Central or Business Central Monitoring to interact with several different Process Servers. Also, when your environment includes many services

running on different Process Servers, Smart Router provides a single endpoint to all client applications. A client application can make a REST API call requiring any service. Smart Router automatically determines which Process Server must be called for any 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 Process 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 Process 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. For instructions about deploying this environment, see [Deploying a Red Hat Process Automation Manager 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 Process 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 Process 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 Process Server or undeploy any existing ones (you can not 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 Process 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. PREPARING TO DEPLOY RED HAT PROCESS AUTOMATION MANAGER IN YOUR OPENSIFT ENVIRONMENT

Before deploying Red Hat Process Automation Manager in your OpenShift environment, you need to complete several preparatory 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.

2.1. ENSURING THE AVAILABILITY OF IMAGE STREAMS

You must ensure that the image streams that are required for the deployment are available in your OpenShift environment. Some versions of the OpenShift environment include the necessary image streams. You must check if they are available. If they are not available, you must install the **rhpm71-image-streams.yaml** file.

Procedure

1. Run the following commands:

```
$ oc get imagestreamtag -n openshift | grep rhpm71-businesscentral
$ oc get imagestreamtag -n openshift | grep rhpm71-kieserver
```

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

2. If the output of one or both of the commands is empty, download the **rhpm71-image-streams.zip** product deliverable file from the [Software Downloads](#) page. Extract the **rhpm71-image-streams.yaml** file from it. Complete one of the following actions:

- Run the following command:

```
$ oc create -f rhpm71-image-streams.yaml
```

- Using the OpenShift Web UI, select **Add to Project → Import YAML / JSON**, then choose the file or paste its contents.

2.2. CREATING THE SECRETS FOR PROCESS SERVER

OpenShift uses objects called **Secrets** to hold sensitive information, such as passwords or keystores. See the [Secrets chapter](#) in the OpenShift documentation for more information.

You must create an SSL certificate for Process 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 Process Server. In a production environment, generate a valid signed certificate that matches the expected URL of the Process Server. Save the keystore in a file named **keystore.jks**. Record the name of the certificate and the password of the keystore file.
See [Generate a SSL Encryption Key and Certificate](#) for more information on how to create a keystore with self-signed or purchased SSL certificates.

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

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

2.3. CREATING THE SECRETS FOR BUSINESS CENTRAL

If you are planning to deploy Business Central or Business Central Monitoring in your OpenShift environment, you must create an SSL certificate for Business Central and provide it to your OpenShift environment as a secret. Do not use the same certificate and keystore for Business Central and for Process Server.

Procedure

1. Generate an SSL keystore with a private and public key for SSL encryption for Business Central. In a production environment, generate a valid signed certificate that matches the expected URL of the Business Central. Save the keystore in a file named **keystore.jks**. Record the name of the certificate and the password of the keystore file.
See [Generate a SSL Encryption Key and Certificate](#) for more information on how to create a keystore with self-signed or purchased SSL certificates.
2. Use the **oc** command to generate a secret named **businesscentral-app-secret** from the new keystore file:

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

2.4. CHANGING GLUSTERFS CONFIGURATION

Check whether your OpenShift environment uses GlusterFS to provide permanent storage volumes. If it uses GlusterFS, to ensure optimal performance, tune your GlusterFS storage by changing the storage class configuration.

Procedure

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

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

Where **class-name** is the name of the default storage class. For 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. On the line with the **volumeoptions** key, add the following two 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
```

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

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

Where **class-name** is the name of the default storage class. For example:

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

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

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

CHAPTER 3. 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 Process Server for test execution of the processes.

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

A single authoring environment contains two pods. One of the pods runs Business Central, the other runs Process Server. The Process Server includes an embedded in-memory H2 database engine. This type of environment uses the least possible amount of resources. However, because of the in-memory database, restarting the Process Server pod leads to loss of all process information.

An HA authoring environment contains several pods. Both Business Central and Process Server are provided in scalable pods that can run in parallel and share persistent storage. The database is provided by a separate high-availability service. Use a high-availability authoring environment to provide maximum reliability and responsiveness, especially if several users are involved in authoring at the same time.



IMPORTANT

In the current version, the high-availability functionality is a technology preview.

3.1. DEPLOYING A SINGLE AUTHORIZING ENVIRONMENT

To deploy a single authoring environment, use the **rhcam71-authoring.yaml** template file.

You can extract this file from the **rhcam-7.1.0-openshift-templates.zip** product deliverable file. You can download the file from the [Software Downloads](#) page.

If you want to modify the environment defined by the template file, see [Section 3.4, “Modifying the template for the single authoring environment”](#).

Procedure

1. Use one of the following methods to deploy the template:
 - In the OpenShift Web UI, select **Add to Project** → **Import YAML / JSON** and then select or paste the **rhcam71-authoring.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>/rhcam71-authoring.yaml -p
BUSINESS_CENTRAL_HTTPS_SECRET=businesscentral-app-secret -p
KIE_SERVER_HTTPS_SECRET=kieserver-app-secret
```

In this command line:

- Replace **<template-path>** with the path to the downloaded template file.
 - Use as many **-p PARAMETER=value** pairs as needed to set the required parameters. You can view the template file to see descriptions for all parameters.
2. Set the following parameters as necessary:

- **Business Central Server Keystore Secret Name** (**BUSINESS_CENTRAL_HTTPS_SECRET**): The name of the secret for Business Central, as created in [Section 2.3, "Creating the secrets for Business Central"](#) .
 - **KIE Server Keystore Secret Name**(**KIE_SERVER_HTTPS_SECRET**): The name of the secret for Process Server, as created in [Section 2.2, "Creating the secrets for Process Server"](#).
 - **Application Name** (**APPLICATION_NAME**): The name of the OpenShift application. It is used in the default URLs for Business Central and Process Server. OpenShift uses the application name to create a separate set of deployment configurations, services, routes, labels, and artifacts. You can deploy several applications using the same template into the same project, as long as you use different application names. Also, the application name determines the name of the server configuration (server template) on the Business Central that the Process Server is to join.
 - **Business Central Server Certificate Name**(**BUSINESS_CENTRAL_HTTPS_NAME**): The name of the certificate in the keystore that you created in [Section 2.3, "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 2.3, "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 2.2, "Creating the secrets for Process Server"](#) .
 - **KIE Server Keystore Password** (**KIE_SERVER_HTTPS_PASSWORD**): The password for the keystore that you created in [Section 2.2, "Creating the secrets for Process Server"](#) .
 - **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 2.1, "Ensuring the availability of image streams"](#)), the namespace is **openshift**. If you have installed the image streams file, the namespace is the name of the OpenShift project.
You can also set the following user names and passwords:
 - **KIE Admin User** (**KIE_ADMIN_USER**) and **KIE Admin Password** (**KIE_ADMIN_PWD**): The user name and password for the administrative user in Business Central.
 - **KIE Server User** (**KIE_SERVER_USER**) and **KIE Server Password** (**KIE_SERVER_PWD**): The user name and password that a client application must use to connect to the Process Server.
3. If you want to place the built KJAR files into an external Maven repository, set the following parameters:
- **Maven repository URL** (**MAVEN_REPO_URL**): The URL for the Maven repository.
 - **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.
 - **Maven repository ID** (**MAVEN_REPO_ID**): The Maven ID, which must match the **id** setting for the Maven repository.



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. You can use Git hooks to facilitate interaction between the internal Git repository of Business Central and an external Git repository. To configure Git hooks, set the following parameter:
 - **Git hooks directory (GIT_HOOKS_DIR)**: The fully qualified path to a Git hooks directory, for example, **/opt/eap/standalone/data/kie/git/hooks**. You must provide the content of this directory and mount it at the specified path; for instructions, see [Section 3.3, "Providing the Git hooks directory"](#).
5. If you want to use RH-SSO or LDAP authentication, complete the following additional configuration:
 - a. In the RH-SSO or LDAP service, create all user names in the deployment parameters. If you do not set any of the parameters, create users with the default user names. The created users must also be assigned to roles:
 - **KIE_ADMIN_USER**: default user name **adminUser**, roles: **kie-server,rest-all,admin,kiemgmt,Administrators**
 - **KIE_SERVER_CONTROLLER_USER**: default user name **controllerUser**, roles: **kie-server,rest-all,guest**
 - **BUSINESS_CENTRAL_MAVEN_USERNAME** (not needed if you configure the use of an external Maven repository): default user name **mavenUser**. No roles are required.
 - **KIE_SERVER_USER**: default user name **executionUser**, roles **kie-server,rest-all,guest**
 - b. If you want to configure Red Hat Single Sign On (RH-SSO) authentication, an RH-SSO realm that applies to Red Hat Process Automation Manager must exist. Process Server. If the client does not yet exist, the template can create it during deployment. Clients within RH-SSO must also exist for Business Central and for Process Server. If the clients do not yet exist, the template can create them during deployment. For the user roles that you can configure in RH-SSO, see [Roles and users](#).

Use one of the following procedures:

- i. If the clients for Red Hat Process Automation Manager within RH-SSO already exist, set the following parameters in the template:
 - **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.
 - **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 Process Server.
 - **KIE Server RH-SSO Client Secret**(**KIE_SERVER_SSO_SECRET**): The secret string that is set in RH-SSO for the client for Process Server.
 - **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.
- ii. To create the clients for Red Hat Process Automation Manager within RH-SSO, set the following parameters in the template:
- **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.
 - **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.
 - **Business Central Custom http Route Hostname**
(**BUSINESS_CENTRAL_HOSTNAME_HTTP**): The fully qualified host name to use for the HTTP endpoint for Business Central. If you need to create a client in RH-SSO, you can not leave this parameter blank.
 - **Business Central Custom https Route Hostname**
(**BUSINESS_CENTRAL_HOSTNAME_HTTPS**): The fully qualified host name to use for the HTTPS endpoint for Business Central. If you need to create a client in RH-SSO, you can not leave this parameter blank.
 - **KIE Server RH-SSO Client name**(**KIE_SERVER_SSO_CLIENT**): The name of the client to create in RH-SSO for Process Server.
 - **KIE Server RH-SSO Client Secret**(**KIE_SERVER_SSO_SECRET**): The secret string to set in RH-SSO for the client for Process Server.
 - **KIE Server Custom http Route Hostname**(**KIE_SERVER_HOSTNAME_HTTP**): The fully qualified host name to use for the HTTP endpoint for Process Server. If you need to create a client in RH-SSO, you can not leave this parameter blank.
 - **KIE Server Custom https Route Hostname**
(**KIE_SERVER_HOSTNAME_HTTPS**): The fully qualified host name to use for the HTTPS endpoint for Process Server. If you need to create a client in RH-SSO, you can not leave this parameter blank.
 - **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.

- **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.
- c. To configure LDAP, 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](#). Do not configure LDAP authentication and RH-SSO authentication in the same deployment.
6. If you modified the template to use an external database server for the Process Server, as described in [Section 3.4, "Modifying the template for the single authoring environment"](#), 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_HOST**): The JDBC URL for the external database server.
 - **KIE Server External Database Dialect**(**KIE_SERVER_EXTERNALDB_DIALECT**): The Hibernate dialect for the server, depending on the server type:
 - **org.hibernate.dialect.MySQL5Dialect** (used for MySQL and MariaDB)
 - **org.hibernate.dialect.PostgreSQLDialect**
 - **org.hibernate.dialect.SQLServer2012Dialect** (used for MS SQL)
 - **org.hibernate.dialect.DB2Dialect**
 - **org.hibernate.dialect.Oracle12cDialect**
 - **org.hibernate.dialect.SybaseASE15Dialect**
 - **KIE Server External Database Host**(**KIE_SERVER_EXTERNALDB_HOST**): The host name of the external database server.
 - **KIE Server External Database Port**(**KIE_SERVER_EXTERNALDB_PORT**): The port number of the external database server.

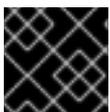
- **KIE Server External Database name(KIE_SERVER_EXTERNALDB_DB)**: The database name to use on the external database server.
7. If you created a custom image for using an external database server other than MySQL or PostgreSQL, as described in [Section 3.6, “Building a custom Process Server image for an external database”](#), set the KIE Server Image Stream Name (**KIE_SERVER_IMAGE_STREAM_NAME**) parameter to the following value:
 - For Microsoft SQL Server, **rhpm71-kieserver-mssql-openshift**
 - For MariaDB, **rhpm71-kieserver-mariadb-openshift**
 - For IBM DB2, **rhpm71-kieserver-db2-openshift**
 - For Oracle Database, **rhpm71-kieserver-oracle-openshift**
 - For Sybase, **rhpm71-kieserver-sybase-openshift**
 8. Complete the creation of the environment, depending on the method that you are using:
 - 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 and run the command line.

3.2. DEPLOYING A HIGH-AVAILABILITY AUTHORIZING ENVIRONMENT

To deploy a high-availability authoring environment, use the **rhpm71-authoring-ha.yaml** template file.

You can download the file from the [Software Downloads](#) page.

If you want to modify the environment defined by the template file, see [Section 3.5, “Modifying the template for the High Availability authoring environment”](#).



IMPORTANT

In the current version, the high-availability functionality is a technology preview.

Procedure

1. Use one of the following methods to deploy the template:
 - In the OpenShift Web UI, select **Add to Project** → **Import YAML / JSON** and then select or paste the **rhpm71-authoring-ha.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>/rhpm71-authoring-ha.yaml -p
BUSINESS_CENTRAL_HTTPS_SECRET=businesscentral-app-secret -p
KIE_SERVER_HTTPS_SECRET=kieserver-app-secret
```

In this command line:

- Replace **<template-path>** with the path to the downloaded template file.
 - Use as many **-p PARAMETER=value** pairs as needed to set the required parameters. You can view the template file to see descriptions for all parameters.
2. Set the following parameters as necessary:
- **Business Central Server Keystore Secret Name** (**BUSINESS_CENTRAL_HTTPS_SECRET**): The name of the secret for Business Central, as created in [Section 2.3, "Creating the secrets for Business Central"](#).
 - **KIE Server Keystore Secret Name**(**KIE_SERVER_HTTPS_SECRET**): The name of the secret for Process Server, as created in [Section 2.2, "Creating the secrets for Process Server"](#).
 - **Application Name** (**APPLICATION_NAME**): The name of the OpenShift application. It is used in the default URLs for Business Central and Process Server. OpenShift uses the application name to create a separate set of deployment configurations, services, routes, labels, and artifacts. You can deploy several applications using the same template into the same project, as long as you use different application names. Also, the application name determines the name of the server configuration (server template) on the Business Central that the Process Server is to join.
 - **Business Central Server Certificate Name**(**BUSINESS_CENTRAL_HTTPS_NAME**): The name of the certificate in the keystore that you created in [Section 2.3, "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 2.3, "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 2.2, "Creating the secrets for Process Server"](#).
 - **KIE Server Keystore Password** (**KIE_SERVER_HTTPS_PASSWORD**): The password for the keystore that you created in [Section 2.2, "Creating the secrets for Process Server"](#).
 - **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 2.1, "Ensuring the availability of image streams"](#)), the namespace is **openshift**. If you have installed the image streams file, the namespace is the name of the OpenShift project.
You can also set the following user names and passwords:
 - **KIE Admin User**(**KIE_ADMIN_USER**) and **KIE Admin Password** (**KIE_ADMIN_PWD**): The user name and password for the administrative user in Business Central.
 - **KIE Server User**(**KIE_SERVER_USER**) and **KIE Server Password** (**KIE_SERVER_PWD**): The user name and password that a client application must use to connect to the Process Server.
3. If you want to place the built KJAR files into an external Maven repository, set the following parameters:
- **Maven repository URL** (**MAVEN_REPO_URL**): The URL for the Maven repository.

- **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.
- **Maven repository ID (MAVEN_REPO_ID)**: The Maven ID, which must match the **id** setting for the Maven repository.



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. You can use Git hooks to facilitate interaction between the internal Git repository of Business Central and an external Git repository. To configure Git hooks, set the following parameter:
 - **Git hooks directory (GIT_HOOKS_DIR)**: The fully qualified path to a Git hooks directory, for example, **/opt/eap/standalone/data/kie/git/hooks**. You must provide the content of this directory and mount it at the specified path; for instructions, see [Section 3.3, "Providing the Git hooks directory"](#).
5. If you want to use RH-SSO or LDAP authentication, complete the following additional configuration:
 - a. In the RH-SSO or LDAP service, create all user names in the deployment parameters. If you do not set any of the parameters, create users with the default user names. The created users must also be assigned to roles:
 - **KIE_ADMIN_USER**: default user name **adminUser**, roles: **kie-server,rest-all,admin,kiemgmt,Administrators**
 - **KIE_SERVER_CONTROLLER_USER**: default user name **controllerUser**, roles: **kie-server,rest-all,guest**
 - **BUSINESS_CENTRAL_MAVEN_USERNAME** (not needed if you configure the use of an external Maven repository): default user name **mavenUser**. No roles are required.
 - **KIE_SERVER_USER**: default user name **executionUser**, roles **kie-server,rest-all,guest**
 - b. If you want to configure Red Hat Single Sign On (RH-SSO) authentication, an RH-SSO realm that applies to Red Hat Process Automation Manager must exist. Process Server. If the client does not yet exist, the template can create it during deployment. Clients within RH-SSO must also exist for Business Central and for Process Server. If the clients do not yet exist, the template can create them during deployment. For the user roles that you can configure in RH-SSO, see [Roles and users](#).

Use one of the following procedures:

- i. If the clients for Red Hat Process Automation Manager within RH-SSO already exist, set the following parameters in the template:
 - **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.
 - **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 Process Server.
 - **KIE Server RH-SSO Client Secret (KIE_SERVER_SSO_SECRET):** The secret string that is set in RH-SSO for the client for Process Server.
 - **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.
- ii. To create the clients for Red Hat Process Automation Manager within RH-SSO, set the following parameters in the template:
- **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.
 - **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.
 - **Business Central Custom http Route Hostname (BUSINESS_CENTRAL_HOSTNAME_HTTP):** The fully qualified host name to use for the HTTP endpoint for Business Central. If you need to create a client in RH-SSO, you can not leave this parameter blank.
 - **Business Central Custom https Route Hostname (BUSINESS_CENTRAL_HOSTNAME_HTTPS):** The fully qualified host name to use for the HTTPS endpoint for Business Central. If you need to create a client in RH-SSO, you can not leave this parameter blank.
 - **KIE Server RH-SSO Client name (KIE_SERVER_SSO_CLIENT):** The name of the client to create in RH-SSO for Process Server.
 - **KIE Server RH-SSO Client Secret (KIE_SERVER_SSO_SECRET):** The secret string to set in RH-SSO for the client for Process Server.
 - **KIE Server Custom http Route Hostname (KIE_SERVER_HOSTNAME_HTTP):** The fully qualified host name to use for the HTTP endpoint for Process Server. If you need to create a client in RH-SSO, you can not leave this parameter blank.
 - **KIE Server Custom https Route Hostname (KIE_SERVER_HOSTNAME_HTTPS):** The fully qualified host name to use for the

HTTPS endpoint for Process Server. If you need to create a client in RH-SSO, you can not leave this parameter blank.

- **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.
 - **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.
- c. To configure LDAP, 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](#). Do not configure LDAP authentication and RH-SSO authentication in the same deployment.
6. If you modified the template to use an external database server for the Process Server, as described in [Section 3.5, "Modifying the template for the High Availability authoring environment"](#), 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_HOST)**: The JDBC URL for the external database server.
 - **KIE Server External Database Dialect (KIE_SERVER_EXTERNALDB_DIALECT)**: The Hibernate dialect for the server, depending on the server type:
 - **org.hibernate.dialect.MySQL5Dialect** (used for MySQL and MariaDB)
 - **org.hibernate.dialect.PostgreSQLDialect**
 - **org.hibernate.dialect.SQLServer2012Dialect** (used for MS SQL)
 - **org.hibernate.dialect.DB2Dialect**
 - **org.hibernate.dialect.Oracle12cDialect**

- **org.hibernate.dialect.SybaseASE15Dialect**
 - **KIE Server External Database Host(KIE_SERVER_EXTERNALDB_HOST)**: The host name of the external database server.
 - **KIE Server External Database Port(KIE_SERVER_EXTERNALDB_PORT)**: The port number of the external database server.
 - **KIE Server External Database name(KIE_SERVER_EXTERNALDB_DB)**: The database name to use on the external database server.
7. If you created a custom image for using an external database server other than MySQL or PostgreSQL, as described in [Section 3.6, “Building a custom Process Server image for an external database”](#), set the KIE Server Image Stream Name (**KIE_SERVER_IMAGE_STREAM_NAME**) parameter to the following value:
- For Microsoft SQL Server, **rhpm71-kieserver-mssql-openshift**
 - For MariaDB, **rhpm71-kieserver-mariadb-openshift**
 - For IBM DB2, **rhpm71-kieserver-db2-openshift**
 - For Oracle Database, **rhpm71-kieserver-oracle-openshift**
 - For Sybase, **rhpm71-kieserver-sybase-openshift**
8. If an AMQ 7.1 image is not available in the **openshift** namespace with default settings, set the following parameters:
- **AMQ ImageStream Namespace (AMQ_IMAGE_STREAM_NAMESPACE)**: Namespace in which the ImageStream for the AMQ image is installed. The default setting is **openshift**.
 - **AMQ ImageStream Name (AMQ_IMAGE_STREAM_NAME)**: The name of the image stream for the AMQ broker. The default setting is **amq-broker71-openshift**.
 - **AMQ ImageStream Tag (AMQ_IMAGE_STREAM_TAG)**: The AMQ image stream tag. The default setting is **1.0**.
9. Complete the creation of the environment, depending on the method that you are using:
- 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 and run the command line.

3.3. 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
```

Where **<myapp>** is the application name that was 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
```

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> ...
```

Where **file_1**, **file_2** and so on are git hook script files. For 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
```

Where **<myapp>** is 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 (where *myapp* is 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 Process 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 Process Server pod is not present, start it:

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

Where **<myapp>** is the application name that was set when configuring the template.

3.4. 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), you need to modify the template before deploying the environment.

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 **rhpam71-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 **rhpam71-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 **rhpam71-kieserver-mysql.yaml** file.)

`rhpan71-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 `rhpan71-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 `rhpan71-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 `rhpan71-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 `rhpan71-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 Process Server image includes drivers for MySQL and PostgreSQL external database servers. If you want to use another database server, you must build a custom Process Server image. For instructions, see [Section 3.6, “Building a custom Process Server image for an external database”](#).

3.5. 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 Process 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 **rhpan71-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 **rhpan71-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 **rhpan71-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 **rhpm71-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 **rhpm71-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 Process Server image includes drivers for MySQL and PostgreSQL external database servers. If you want to use another database server, you must build a custom Process Server image. For instructions, see [Section 3.6, “Building a custom Process Server 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.

3.6. BUILDING A CUSTOM PROCESS SERVER IMAGE FOR AN EXTERNAL DATABASE

If you want to use an external database server for a Process Server and this server is neither MySQL nor PostgreSQL, you must build a custom Process Server image with drivers for this server before deploying your environment.

You can use this build procedure to provide drivers for the following database servers:

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

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

The build procedure creates a custom image that extends the existing Process Server image. It pushes this custom image into a new **ImageStream** in the **openshift** namespace with the same version tag as the original image.

Prerequisites

- You have logged on to your project in the OpenShift environment using the **oc** command as a user with the **cluster-admin** role.
- For IBM DB2, Oracle Database, or Sybase, you have downloaded the JDBC driver from the database server vendor.

Procedure

1. For IBM DB2, Oracle Database, or Sybase, provide the JDBC driver JAR in a local directory or on an HTTP server. Within the local directory or HTTP server, the following paths are expected:
 - For IBM DB2, **<local_path_or_url>/com/ibm/db2/jcc/db2jcc4/10.5/db2jcc4-10.5.jar**
 - For Oracle Database, **<local_path_or_url>/com/oracle/ojdbc7/12.1.0.1/ojdbc7-12.1.0.1.jar**
 - For Sybase, **<local_path_or_url>/com/sybase/jconn4/16.0_PL05/jconn4-16.0_PL05.jar**
Where **<local_path_or_url>** is the path to the local directory or the URL for the HTTP server where the driver is provided.
2. To install the source code for the custom build, download the **rhcam-7.1.0-openshift-templates.zip** product deliverable file from the [Software Downloads](#) page. Unzip the file and, using the command line, change to the **templates/contrib/jdbc** directory of the unzipped file.
3. Change to the following subdirectory:
 - For Microsoft SQL Server, **mssql-driver-image**
 - For MariaDB, **mariadb-driver-image**
 - For IBM DB2, **db2-driver-image**
 - For Oracle Database, **oracle-driver-image**
 - For Sybase, **sybase-driver-image**
4. Run the following command:
 - For Microsoft SQL Server or MariaDB:

```
./build.sh
```

- For IBM DB2, Oracle Database, or Sybase:

```
./build.sh --artifact-repo=<local_path_or_url>
```

Where **<local_path_or_url>** is the path to the local directory or the URL for the HTTP server where the driver is provided. For example:

```
./build.sh --artifact-repo=/home/builder/drivers
./build.sh --artifact-repo=http://nexus.example.com/nexus/content/groups/public
```

If you want to configure your OpenShift docker registry address in the process, add also the **--registry=<registry_name.domain_name:port>** parameter to your build command.

Examples:

```
../build.sh --registry=docker-registry.custom-domain:80
```

```
../build.sh --artifact-repo=/home/builder/drivers --registry=docker-registry.custom-domain:80
```

CHAPTER 4. OPENSIFT TEMPLATE REFERENCE INFORMATION

Red Hat Process Automation Manager provides the following OpenShift templates. To access the templates, download and extract the **rhpmam-7.1.0-openshift-templates.zip** product deliverable file from the [Software Downloads](#) page of the Red Hat customer portal.

- **rhpmam71-authoring.yaml** provides a Business Central and a Process Server connected to the Business Central. The Process 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 4.1, “rhpmam71-authoring”](#).
- **rhpmam71-authoring-ha.yaml** provides a high-availability Business Central, a Process Server connected to the Business Central, and a MySQL instance that the Process Server uses. You can use this environment to author processes, services, and other business assets. The high-availability functionality is in technical preview. For details about this template, see [Section 4.2, “rhpmam71-authoring-ha”](#).

4.1. RHPAM71-AUTHORING

Application template for a non-HA persistent authoring environment, for Red Hat Process Automation Manager 7.1

4.1.1. Parameters

Templates allow you to define parameters which 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. Refer to the [Openshift documentation](#) for more information.

Variable name	Image Environment Variable	Description	Example value	Required
APPLICATION_NAME	–	The name for the application.	myapp	True
KIE_ADMIN_USER	KIE_ADMIN_USER	KIE administrator username	adminUser	False
KIE_ADMIN_PASSWORD	KIE_ADMIN_PASSWORD	KIE administrator password	\${KIE_ADMIN_PASSWORD}	False
KIE_SERVER_CONTROLLER_USER	KIE_SERVER_CONTROLLER_USER	KIE server controller username (Sets the org.kie.server.controller.user system property)	controllerUser	False

Variable name	Image Environment Variable	Description	Example value	Required
KIE_SERVER_CONTROLLER_PWD	KIE_SERVER_CONTROLLER_PWD	KIE server controller password (Sets the org.kie.server.controller.pwd system property)	\${KIE_SERVER_CONTROLLER_PWD}	False
KIE_SERVER_CONTROLLER_TOKEN	KIE_SERVER_CONTROLLER_TOKEN	KIE server controller token for bearer authentication (Sets the org.kie.server.controller.token system property)	\${KIE_SERVER_CONTROLLER_TOKEN}	False
KIE_SERVER_USER	KIE_SERVER_USER	KIE server username (Sets the org.kie.server.user system property)	executionUser	False
KIE_SERVER_PWD	KIE_SERVER_PWD	KIE server password (Sets the org.kie.server.pwd system property)	\${KIE_SERVER_PWD}	False
KIE_SERVER_BYPASS_AUTH_USER	KIE_SERVER_BYPASS_AUTH_USER	KIE server bypass auth user (Sets the org.kie.server.bypass.auth.user system property)	false	False
KIE_SERVER_PERSISTENCE_DS	KIE_SERVER_PERSISTENCE_DS	KIE server persistence datasource (Sets the org.kie.server.persistence.ds system property)	java:/jboss/datasources/rhpam	False
KIE_SERVER_H2_USER	–	KIE server H2 database username	sa	False
KIE_SERVER_H2_PWD	–	KIE server H2 database password	–	False

Variable name	Image Environment Variable	Description	Example value	Required
---------------	----------------------------	-------------	---------------	----------

KIE_MBEANS	KIE_MBEANS	KIE server mbeans enabled/disabled (Sets the kie.mbeans and kie.scanner.mbeans system properties)	enabled	False
DROOLS_SERVER_FILTER_CLASSES	DROOLS_SERVER_FILTER_CLASSES	KIE server class filtering (Sets the org.drools.server.filter.classes system property)	true	False
BUSINESS_CENTRAL_HOSTNAME_HTTP	HOSTNAME_HTTP	Custom hostname for http service route. Leave blank for default hostname, e.g.: <application-name>-rhpamcentr- <project>.<default-domain-suffix>	`\${BUSINESS_CENTRAL_HOSTNAME_HTTP}`	False
BUSINESS_CENTRAL_HOSTNAME_HTTPS	HOSTNAME_HTTP	Custom hostname for https service route. Leave blank for default hostname, e.g.: secure- <application-name>-rhpamcentr- <project>.<default-domain-suffix>	`\${BUSINESS_CENTRAL_HOSTNAME_HTTP}`	False

Variable name	Image Environment Variable	Description	Example value	Required
KIE_SERVER_HOSTNAME_HTTP	HOSTNAME_HTTP	Custom hostname for http service route, if set will also configure the KIE_SERVER_HOST. Leave blank for default hostname, e.g.: <application-name>-kieserver-<project>.<default-domain-suffix>	`\${BUSINESS_CENTRAL_HOSTNAME_HTTP}`	False
KIE_SERVER_HOSTNAME_HTTPS	HOSTNAME_HTTPS	Custom hostname for https service route. Leave blank for default hostname, e.g.: secure-<application-name>-kieserver-<project>.<default-domain-suffix>	`\${BUSINESS_CENTRAL_HOSTNAME_HTTPS}`	False
KIE_SERVER_USE_SECURE_ROUTE_NAME	KIE_SERVER_USE_SECURE_ROUTE_NAME	Use https for the KIE_SERVER_HOST when it is not explicit configured to a custom value.	false	False
BUSINESS_CENTRAL_HTTPS_SECRET	–	The name of the secret containing the keystore file	–	True
BUSINESS_CENTRAL_HTTPS_KEYSTORE	HTTPS_KEYSTORE	The name of the keystore file within the secret	keystore.jks	False
BUSINESS_CENTRAL_HTTPS_NAME	HTTPS_NAME	The name associated with the server certificate	jboss	False
BUSINESS_CENTRAL_HTTPS_PASSWORD	HTTPS_PASSWORD	The password for the keystore and certificate	mykeystorepass	False

Variable name	Image Environment Variable	Description	Example value	Required
KIE_SERVER_HTTPS_SECRET	–	The name of the secret containing the keystore file	–	True
KIE_SERVER_HTTPS_KEYSTORE	HTTPS_KEYSTORE	The name of the keystore file within the secret	keystore.jks	False
KIE_SERVER_HTTPS_NAME	HTTPS_NAME	The name associated with the server certificate	jboss	False
KIE_SERVER_HTTPS_PASSWORD	HTTPS_PASSWORD	The password for the keystore and certificate	mykeystorepass	False
DB_VOLUME_CAPACITY	–	Size of persistent storage for database volume.	1Gi	True
IMAGE_STREAM_NAMESPACE	–	Namespace in which the ImageStreams for Red Hat Middleware images are installed. These ImageStreams are normally installed in the openshift namespace. You should only need to modify this if you've installed the ImageStreams in a different namespace/project.	openshift	True
KIE_SERVER_IMAGE_STREAM_NAME	–	The name of the image stream to use for KIE server. Default is "rhpam71-kieserver-openshift".	rhpam71-kieserver-openshift	True

Variable name	Image Environment Variable	Description	Example value	Required
IMAGE_STREAM_TAG	–	A named pointer to an image in an image stream. Default is "1.1".	1.1	True
MAVEN_REPO_ID	MAVEN_REPO_ID	The id to use for the maven repository, if set. Default is generated randomly.	`\${MAVEN_REPO_ID}`	False
MAVEN_REPO_URL	MAVEN_REPO_URL	Fully qualified URL to a Maven repository or service.	`\${MAVEN_REPO_URL}`	False
MAVEN_REPO_USERNAME	MAVEN_REPO_USERNAME	Username to access the Maven repository, if required.	`\${MAVEN_REPO_USERNAME}`	False
MAVEN_REPO_PASSWORD	MAVEN_REPO_PASSWORD	Password to access the Maven repository, if required.	`\${MAVEN_REPO_PASSWORD}`	False
BUSINESS_CENTRAL_MAVEN_USERNAME	–	Username to access the Maven service hosted by Business Central inside EAP.	mavenUser	True
BUSINESS_CENTRAL_MAVEN_PASSWORD	–	Password to access the Maven service hosted by Business Central inside EAP.	–	True
GIT_HOOKS_DIR	GIT_HOOKS_DIR	The directory to use for git hooks, if required.	`\${GIT_HOOKS_DIR}`	False
BUSINESS_CENTRAL_VOLUME_CAPACITY	–	Size of the persistent storage for Business Central's runtime data.	1Gi	True

Variable name	Image Environment Variable	Description	Example value	Required
BUSINESS_CENTRAL_MEMORY_LIMIT	–	Business Central Container memory limit	2Gi	False
KIE_SERVER_MEMORY_LIMIT	–	KIE server Container memory limit	1Gi	False
SSO_URL	SSO_URL	RH-SSO URL	\${SSO_URL}	False
SSO_REALM	SSO_REALM	RH-SSO Realm name	\${SSO_REALM}	False
BUSINESS_CENTRAL_SSO_CLIENT	SSO_CLIENT	Business Central RH-SSO Client name	\${BUSINESS_CENTRAL_SSO_CLIENT}	False
BUSINESS_CENTRAL_SSO_SECRET	SSO_SECRET	Business Central RH-SSO Client Secret	\${BUSINESS_CENTRAL_SSO_SECRET}	False
KIE_SERVER_SSO_CLIENT	SSO_CLIENT	KIE Server RH-SSO Client name	\${BUSINESS_CENTRAL_SSO_CLIENT}	False
KIE_SERVER_SSO_SECRET	SSO_SECRET	KIE Server RH-SSO Client Secret	\${BUSINESS_CENTRAL_SSO_SECRET}	False
SSO_USERNAME	SSO_USERNAME	RH-SSO Realm Admin Username used to create the Client if it doesn't exist	\${SSO_USERNAME}	False
SSO_PASSWORD	SSO_PASSWORD	RH-SSO Realm Admin Password used to create the Client	\${SSO_PASSWORD}	False
SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	RH-SSO Disable SSL Certificate Validation	false	False

Variable name	Image Environment Variable	Description	Example value	Required
SSO_PRINCIPAL_ATTRIBUTE	SSO_PRINCIPAL_ATTRIBUTE	RH-SSO Principal Attribute to use as username.	preferred_username	False
AUTH_LDAP_URL	AUTH_LDAP_URL	LDAP Endpoint to connect for authentication	\${AUTH_LDAP_URL}	False
AUTH_LDAP_BIND_DN	AUTH_LDAP_BIND_DN	Bind DN used for authentication	\${AUTH_LDAP_BIND_DN}	False
AUTH_LDAP_BIND_CREDENTIAL	AUTH_LDAP_BIND_CREDENTIAL	LDAP Credentials used for authentication	\${AUTH_LDAP_BIND_CREDENTIAL}	False
AUTH_LDAP_JAAS_SECURITY_DOMAIN	AUTH_LDAP_JAAS_SECURITY_DOMAIN	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	\${AUTH_LDAP_JAAS_SECURITY_DOMAIN}	False
AUTH_LDAP_BASE_CTX_DN	AUTH_LDAP_BASE_CTX_DN	LDAP Base DN of the top-level context to begin the user search.	\${AUTH_LDAP_BASE_CTX_DN}	False
AUTH_LDAP_BASE_FILTER	AUTH_LDAP_BASE_FILTER	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}).	\${AUTH_LDAP_BASE_FILTER}	False
AUTH_LDAP_SEARCH_SCOPE	AUTH_LDAP_SEARCH_SCOPE	The search scope to use.	\${AUTH_LDAP_SEARCH_SCOPE}	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_SEARCH_TIME_LIMIT	AUTH_LDAP_SEARCH_TIME_LIMIT	The timeout in milliseconds for user or role searches.	`\${AUTH_LDAP_SEARCH_TIME_LIMIT}`	False
AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE	AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE	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.	`\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}`	False
AUTH_LDAP_PARSE_USERNAME	AUTH_LDAP_PARSE_USERNAME	A flag indicating if the DN is to be parsed for the username. If set to true, the DN is parsed for the username. If set to false the DN is not parsed for the username. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	`\${AUTH_LDAP_PARSE_USERNAME}`	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_USERNAME_BEGIN_STRING	AUTH_LDAP_USERNAME_BEGIN_STRING	Defines the String which is to be removed from the start of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<code>\${AUTH_LDAP_USERNAME_BEGIN_STRING}</code>	False
AUTH_LDAP_USERNAME_END_STRING	AUTH_LDAP_USERNAME_END_STRING	Defines the String which is to be removed from the end of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<code>\${AUTH_LDAP_USERNAME_END_STRING}</code>	False
AUTH_LDAP_ROLE_ATTRIBUTE_ID	AUTH_LDAP_ROLE_ATTRIBUTE_ID	Name of the attribute containing the user roles.	<code>\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}</code>	False
AUTH_LDAP_ROLES_CTX_DN	AUTH_LDAP_ROLES_CTX_DN	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.	<code>\${AUTH_LDAP_ROLES_CTX_DN}</code>	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_ROLE_FILTER	AUTH_LDAP_ROLE_FILTER	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}).	`\${AUTH_LDAP_ROLE_FILTER}`	False
AUTH_LDAP_ROLE_RECURSION	AUTH_LDAP_ROLE_RECURSION	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	`\${AUTH_LDAP_ROLE_RECURSION}`	False
AUTH_LDAP_DEFAULT_ROLE	AUTH_LDAP_DEFAULT_ROLE	A role included for all authenticated users	`\${AUTH_LDAP_DEFAULT_ROLE}`	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID	AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID	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.	`\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}`	False
AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN	AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN	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.	`\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}`	False
AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN	AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN	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 roleNameAttribute ID attribute of the context name. Certain directory schemas, such as Microsoft Active Directory, require this attribute to be set to true.	`\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}`	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK	AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK	If you are not using referrals, this option can be ignored. 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.	<code>\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}</code>	False

4.1.2. 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](#).

4.1.2.1. Services

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

Service	Port	Name	Description
<code>\${APPLICATION_NAME}-rhpamcentr</code>	8080	http	All the Business Central web server's ports.
	8443	https	
	8001	git-ssh	
<code>\${APPLICATION_NAME}-kieserver</code>	8080	http	All the KIE server web server's ports.
	8443	https	

4.1.2.2. 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. Refer to the [OpenShift documentation](#) for more information.

Service	Security	Hostname
\${APPLICATION_NAME}-rhpmcentr-http	none	\${BUSINESS_CENTRAL_HOSTNAME_HTTP}
\${APPLICATION_NAME}-rhpmcentr-https	TLS passthrough	\${BUSINESS_CENTRAL_HOSTNAME_HTTPS}
\${APPLICATION_NAME}-kieserver-http	none	\${KIE_SERVER_HOSTNAME_HTTP}
\${APPLICATION_NAME}-kieserver-https	TLS passthrough	\${KIE_SERVER_HOSTNAME_HTTPS}

4.1.2.3. 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. Refer to the [OpenShift documentation](#) for more information.

4.1.2.3.1. Triggers

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

Deployment	Triggers
\${APPLICATION_NAME}-rhpmcentr	ImageChange
\${APPLICATION_NAME}-kieserver	ImageChange

4.1.2.3.2. 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. Refer to the [container-engine documentation](#) for more information.

Deployment	Replicas
\${APPLICATION_NAME}-rhpmcentr	1

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

4.1.2.3.3. Pod Template

4.1.2.3.3.1. Service Accounts

Service accounts are API objects that exist within each project. They can be created or deleted like any other API object. Refer to 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>

4.1.2.3.3.2. Image

Deployment	Image
<code>\${APPLICATION_NAME}-rhpamcentr</code>	<code>rhpam71-businesscentral-openshift</code>
<code>\${APPLICATION_NAME}-kieserver</code>	<code>\${KIE_SERVER_IMAGE_STREAM_NAME}</code>

4.1.2.3.3.3. Readiness Probe

`${APPLICATION_NAME}-rhpamcentr`

```
/bin/bash -c curl --fail --silent -u '${KIE_ADMIN_USER}:${KIE_ADMIN_PWD}'
http://localhost:8080/kie-wb.jsp
```

`${APPLICATION_NAME}-kieserver`

```
/bin/bash -c curl --fail --silent -u '${KIE_ADMIN_USER}:${KIE_ADMIN_PWD}'
http://localhost:8080/services/rest/server/readychck
```

4.1.2.3.3.4. Liveness Probe

`${APPLICATION_NAME}-rhpamcentr`

```
/bin/bash -c curl --fail --silent -u '${KIE_ADMIN_USER}:${KIE_ADMIN_PWD}'
http://localhost:8080/kie-wb.jsp
```

`${APPLICATION_NAME}-kieserver`

```
/bin/bash -c curl --fail --silent -u '${KIE_ADMIN_USER}:${KIE_ADMIN_PWD}'
http://localhost:8080/services/rest/server/readycheck
```

4.1.2.3.3.5. Exposed Ports

Deployments	Name	Port	Protocol
\${APPLICATION_NAME}-rhpamcentr	jolokia	8778	TCP
	http	8080	TCP
	https	8443	TCP
	git-ssh	8001	TCP
\${APPLICATION_NAME}-kieserver	jolokia	8778	TCP
	http	8080	TCP
	https	8443	TCP

4.1.2.3.3.6. Image Environment Variables

Deployment	Variable name	Description	Example value
\${APPLICATION_NAME}-rhpamcentr	KIE_ADMIN_USER	KIE administrator username	\${KIE_ADMIN_USER}
	KIE_ADMIN_PWD	KIE administrator password	\${KIE_ADMIN_PWD}
	KIE_MBEANS	KIE server mbeans enabled/disabled (Sets the kie.mbeans and kie.scanner.mbeans system properties)	\${KIE_MBEANS}
	KIE_SERVER_CONTROLLER_USER	KIE server controller username (Sets the org.kie.server.controller.user system property)	\${KIE_SERVER_CONTROLLER_USER}
	KIE_SERVER_CONTROLLER_PWD	KIE server controller password (Sets the org.kie.server.controller.pwd system property)	\${KIE_SERVER_CONTROLLER_PWD}

Deployment	Variable name	Description	Example value
	KIE_SERVER_CONTROLLER_TOKEN	KIE server controller token for bearer authentication (Sets the org.kie.server.controller.token system property)	`\${KIE_SERVER_CONTROLLER_TOKEN}`
	KIE_SERVER_USER	KIE server username (Sets the org.kie.server.user system property)	`\${KIE_SERVER_USER}`
	KIE_SERVER_PWD	KIE server password (Sets the org.kie.server.pwd system property)	`\${KIE_SERVER_PWD}`
	MAVEN_REPO_ID	The id to use for the maven repository, if set. Default is generated randomly.	`\${MAVEN_REPO_ID}`
	MAVEN_REPO_URL	Fully qualified URL to a Maven repository or service.	`\${MAVEN_REPO_URL}`
	MAVEN_REPO_USERNAME	Username to access the Maven repository, if required.	`\${MAVEN_REPO_USERNAME}`
	MAVEN_REPO_PASSWORD	Password to access the Maven repository, if required.	`\${MAVEN_REPO_PASSWORD}`
	KIE_MAVEN_USER	–	`\${BUSINESS_CENTRAL_MAVEN_USERNAME}`
	KIE_MAVEN_PWD	–	`\${BUSINESS_CENTRAL_MAVEN_PASSWORD}`
	GIT_HOOKS_DIR	The directory to use for git hooks, if required.	`\${GIT_HOOKS_DIR}`
	HTTPS_KEYSTORE_DIR	–	/etc/businesscentral-secret-volume

Deployment	Variable name	Description	Example value
	HTTPS_KEYSTORE	The name of the keystore file within the secret	\${BUSINESS_CENTRAL_HTTPS_KEYSTORE}
	HTTPS_NAME	The name associated with the server certificate	\${BUSINESS_CENTRAL_HTTPS_NAME}
	HTTPS_PASSWORD	The password for the keystore and certificate	\${BUSINESS_CENTRAL_HTTPS_PASSWORD}
	WORKBENCH_ROUTE_NAME	–	\${APPLICATION_NAME}-rhpamcentr
	SSO_URL	RH-SSO URL	\${SSO_URL}
	SSO_OPENIDCONNECT_DEPLOYMENTS	–	ROOT.war
	SSO_REALM	RH-SSO Realm name	\${SSO_REALM}
	SSO_SECRET	Business Central RH-SSO Client Secret	\${BUSINESS_CENTRAL_SSO_SECRET}
	SSO_CLIENT	Business Central RH-SSO Client name	\${BUSINESS_CENTRAL_SSO_CLIENT}
	SSO_USERNAME	RH-SSO Realm Admin Username used to create the Client if it doesn't exist	\${SSO_USERNAME}
	SSO_PASSWORD	RH-SSO Realm Admin Password used to create the Client	\${SSO_PASSWORD}
	SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	RH-SSO Disable SSL Certificate Validation	\${SSO_DISABLE_SSL_CERTIFICATE_VALIDATION}
	SSO_PRINCIPAL_ATTRIBUTE	RH-SSO Principal Attribute to use as username.	\${SSO_PRINCIPAL_ATTRIBUTE}

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

Deployment	Variable name	Description	Example value
	AUTH_LDAP_SEARCH_SCOPE	The search scope to use.	`\${AUTH_LDAP_SEARCH_SCOPE}`
	AUTH_LDAP_SEARCH_TIME_LIMIT	The timeout in milliseconds for user or role searches.	`\${AUTH_LDAP_SEARCH_TIME_LIMIT}`
	AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE	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.	`\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}`
	AUTH_LDAP_PARSE_USERNAME	A flag indicating if the DN is to be parsed for the username. If set to true, the DN is parsed for the username. If set to false the DN is not parsed for the username. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	`\${AUTH_LDAP_PARSE_USERNAME}`
	AUTH_LDAP_USERNAME_BEGIN_STRING	Defines the String which is to be removed from the start of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	`\${AUTH_LDAP_USERNAME_BEGIN_STRING}`

Deployment	Variable name	Description	Example value
	AUTH_LDAP_USER_NAME_END_STRING	Defines the String which is to be removed from the end of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<code>\${AUTH_LDAP_USER_NAME_END_STRING}</code>
	AUTH_LDAP_ROLE_ATTRIBUTE_ID	Name of the attribute containing the user roles.	<code>\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}</code>
	AUTH_LDAP_ROLE_S_CTX_DN	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.	<code>\${AUTH_LDAP_ROLE_S_CTX_DN}</code>
	AUTH_LDAP_ROLE_FILTER	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> .	<code>\${AUTH_LDAP_ROLE_FILTER}</code>

Deployment	Variable name	Description	Example value
	AUTH_LDAP_ROLE_RECURSION	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	`\${AUTH_LDAP_ROLE_RECURSION}`
	AUTH_LDAP_DEFAULT_ROLE	A role included for all authenticated users	`\${AUTH_LDAP_DEFAULT_ROLE}`
	AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID	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.	`\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}`
	AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN	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.	`\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}`
	AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN	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.	`\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}`

Deployment	Variable name	Description	Example value
	AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK	If you are not using referrals, this option can be ignored. 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.	\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}
\${APPLICATION_NAME}-kieserver	DATASOURCES	–	RHPAM
	RHPAM_DATABASE	–	rhpm7
	RHPAM_JNDI	–	\${KIE_SERVER_PERSISTENCE_DS}
	RHPAM_JTA	–	true
	RHPAM_DRIVER	–	h2
	RHPAM_USERNAME	–	\${KIE_SERVER_H2_USER}
	RHPAM_PASSWORD	–	\${KIE_SERVER_H2_PWD}
	RHPAM_XA_CONNECTION_PROPERTY_URL	–	jdbc:h2:/opt/eap/standalone/data/rhpm
	RHPAM_SERVICE_HOST	–	dummy_ignored
	RHPAM_SERVICE_PORT	–	12345

Deployment	Variable name	Description	Example value
	KIE_SERVER_PERSISTENCE_DIALECT	–	org.hibernate.dialect.H2 Dialect
	DROOLS_SERVER_FILTER_CLASSES	KIE server class filtering (Sets the org.drools.server.filter.classes system property)	\${DROOLS_SERVER_FILTER_CLASSES}
	KIE_ADMIN_USER	KIE administrator username	\${KIE_ADMIN_USER}
	KIE_ADMIN_PWD	KIE administrator password	\${KIE_ADMIN_PWD}
	KIE_MBEANS	KIE server mbeans enabled/disabled (Sets the kie.mbeans and kie.scanner.mbeans system properties)	\${KIE_MBEANS}
	KIE_SERVER_BYPASS_AUTH_USER	KIE server bypass auth user (Sets the org.kie.server.bypass.auth.user system property)	\${KIE_SERVER_BYPASS_AUTH_USER}
	KIE_SERVER_CONTROLLER_USER	KIE server controller username (Sets the org.kie.server.controller.user system property)	\${KIE_SERVER_CONTROLLER_USER}
	KIE_SERVER_CONTROLLER_PWD	KIE server controller password (Sets the org.kie.server.controller.pwd system property)	\${KIE_SERVER_CONTROLLER_PWD}
	KIE_SERVER_CONTROLLER_TOKEN	KIE server controller token for bearer authentication (Sets the org.kie.server.controller.token system property)	\${KIE_SERVER_CONTROLLER_TOKEN}
	KIE_SERVER_CONTROLLER_SERVICE	–	\${APPLICATION_NAME}-rhpamcentr
	KIE_SERVER_CONTROLLER_PROTOCOL	–	ws

Deployment	Variable name	Description	Example value
	KIE_SERVER_ID	–	\${APPLICATION_NAME}-kieserver
	KIE_SERVER_HOST	Custom hostname for http service route, if set will also configure the KIE_SERVER_HOST. Leave blank for default hostname, e.g.: <application-name>-kieserver-<project>. <default-domain-suffix>	\${KIE_SERVER_HOSTNAME_HTTP}
	KIE_SERVER_ROUTE_NAME	–	\${APPLICATION_NAME}-kieserver
	KIE_SERVER_USE_SECURE_ROUTE_NAME	Use https for the KIE_SERVER_HOST when it is not explicit configured to a custom value.	\${KIE_SERVER_USE_SECURE_ROUTE_NAME}
	KIE_SERVER_PERSISTENCE_DS	KIE server persistence datasource (Sets the org.kie.server.persistence.ds system property)	\${KIE_SERVER_PERSISTENCE_DS}
	KIE_SERVER_USER	KIE server username (Sets the org.kie.server.user system property)	\${KIE_SERVER_USER}
	KIE_SERVER_PWD	KIE server password (Sets the org.kie.server.pwd system property)	\${KIE_SERVER_PWD}
	MAVEN_REPOS	–	RHPAMCENTR,EXTERNAL
	RHPAMCENTR_MAVEN_REPO_SERVICE	–	\${APPLICATION_NAME}-rhpamcentr
	RHPAMCENTR_MAVEN_REPO_PATH	–	/maven2/

Deployment	Variable name	Description	Example value
	RHPAMCENTR_MAVEN_REPO_USERNAME	Username to access the Maven repository, if required.	\${BUSINESS_CENTRAL_MAVEN_USERNAME}
	RHPAMCENTR_MAVEN_REPO_PASSWORD	Password to access the Maven repository, if required.	\${BUSINESS_CENTRAL_MAVEN_PASSWORD}
	EXTERNAL_MAVEN_REPO_ID	The id to use for the maven repository, if set. Default is generated randomly.	\${MAVEN_REPO_ID}
	EXTERNAL_MAVEN_REPO_URL	Fully qualified URL to a Maven repository or service.	\${MAVEN_REPO_URL}
	EXTERNAL_MAVEN_REPO_USERNAME	Username to access the Maven repository, if required.	\${MAVEN_REPO_USERNAME}
	EXTERNAL_MAVEN_REPO_PASSWORD	Password to access the Maven repository, if required.	\${MAVEN_REPO_PASSWORD}
	HTTPS_KEYSTORE_DIR	–	/etc/kieserver-secret-volume
	HTTPS_KEYSTORE	The name of the keystore file within the secret	\${KIE_SERVER_HTTPS_KEYSTORE}
	HTTPS_NAME	The name associated with the server certificate	\${KIE_SERVER_HTTPS_NAME}
	HTTPS_PASSWORD	The password for the keystore and certificate	\${KIE_SERVER_HTTPS_PASSWORD}
	SSO_URL	RH-SSO URL	\${SSO_URL}
	SSO_OPENIDCONNECT_DEPLOYMENTS	–	ROOT.war

Deployment	Variable name	Description	Example value
	SSO_REALM	RH-SSO Realm name	\${SSO_REALM}
	SSO_SECRET	Business Central RH-SSO Client Secret	\${KIE_SERVER_SSO_SECRET}
	SSO_CLIENT	Business Central RH-SSO Client name	\${KIE_SERVER_SSO_CLIENT}
	SSO_USERNAME	RH-SSO Realm Admin Username used to create the Client if it doesn't exist	\${SSO_USERNAME}
	SSO_PASSWORD	RH-SSO Realm Admin Password used to create the Client	\${SSO_PASSWORD}
	SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	RH-SSO Disable SSL Certificate Validation	\${SSO_DISABLE_SSL_CERTIFICATE_VALIDATION}
	SSO_PRINCIPAL_ATTRIBUTE	RH-SSO Principal Attribute to use as username.	\${SSO_PRINCIPAL_ATTRIBUTE}
	HOSTNAME_HTTP	Custom hostname for http service route. Leave blank for default hostname, e.g.: <application-name>-rhpamcentr-<project>.<default-domain-suffix>	\${KIE_SERVER_HOSTNAME_HTTP}
	HOSTNAME_HTTPS	Custom hostname for https service route. Leave blank for default hostname, e.g.: secure-<application-name>-rhpamcentr-<project>.<default-domain-suffix>	\${KIE_SERVER_HOSTNAME_HTTPS}
	AUTH_LDAP_URL	LDAP Endpoint to connect for authentication	\${AUTH_LDAP_URL}
	AUTH_LDAP_BIND_DN	Bind DN used for authentication	\${AUTH_LDAP_BIND_DN}

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

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

Deployment	Variable name	Description	Example value
	AUTH_LDAP_ROLE_S_CTX_DN	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.	`\${AUTH_LDAP_ROLE_S_CTX_DN}`
	AUTH_LDAP_ROLE_FILTER	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}).	`\${AUTH_LDAP_ROLE_FILTER}`
	AUTH_LDAP_ROLE_RECURSION	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	`\${AUTH_LDAP_ROLE_RECURSION}`
	AUTH_LDAP_DEFAULT_ROLE	A role included for all authenticated users	`\${AUTH_LDAP_DEFAULT_ROLE}`

Deployment	Variable name	Description	Example value
	AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID	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.	\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}
	AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN	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.	\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}
	AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN	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.	\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}

Deployment	Variable name	Description	Example value
	AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK	If you are not using referrals, this option can be ignored. 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.	\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}

4.1.2.3.3.7. Volumes

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

4.1.2.4. External Dependencies

4.1.2.4.1. 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. Refer to the [Openshift documentation](#) for more information.

Name	Access Mode
\${APPLICATION_NAME}-rhpamcentr-claim	ReadWriteOnce
\${APPLICATION_NAME}-h2-claim	ReadWriteOnce

4.1.2.4.2. Secrets

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

businesscentral-app-secret kieserver-app-secret

4.2. RHPAM71-AUTHORING-HA

Application template for a HA persistent authoring environment, for Red Hat Process Automation Manager 7.1

4.2.1. Parameters

Templates allow you to define parameters which 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. Refer to the [Openshift documentation](#) for more information.

Variable name	Image Environment Variable	Description	Example value	Required
APPLICATION_NAME	–	The name for the application.	myapp	True
KIE_ADMIN_USERNAME	KIE_ADMIN_USERNAME	KIE administrator username	adminUser	False
KIE_ADMIN_PASSWORD	KIE_ADMIN_PASSWORD	KIE administrator password	\${KIE_ADMIN_PASSWORD}	False
KIE_SERVER_CONTROLLER_USERNAME	KIE_SERVER_CONTROLLER_USERNAME	KIE server controller username (Sets the org.kie.server.controller.user system property)	controllerUser	False
KIE_SERVER_CONTROLLER_PASSWORD	KIE_SERVER_CONTROLLER_PASSWORD	KIE server controller password (Sets the org.kie.server.controller.pwd system property)	\${KIE_SERVER_CONTROLLER_PASSWORD}	False

Variable name	Image Environment Variable	Description	Example value	Required
KIE_SERVER_CONTROLLER_TOKEN	KIE_SERVER_CONTROLLER_TOKEN	KIE server controller token for bearer authentication (Sets the org.kie.server.controller.token system property)	\${KIE_SERVER_CONTROLLER_TOKEN}	False
KIE_SERVER_USER	KIE_SERVER_USER	KIE server username (Sets the org.kie.server.user system property)	executionUser	False
KIE_SERVER_PWD	KIE_SERVER_PWD	KIE server password (Sets the org.kie.server.pwd system property)	\${KIE_SERVER_PWD}	False
KIE_SERVER_BYPASS_AUTH_USER	KIE_SERVER_BYPASS_AUTH_USER	KIE server bypass auth user (Sets the org.kie.server.bypass.auth.user system property)	false	False
KIE_SERVER_PERSISTENCE_DS	KIE_SERVER_PERSISTENCE_DS	KIE server persistence datasource (Sets the org.kie.server.persistence.ds system property)	java:/jboss/datasources/rhpam	False
KIE_SERVER_MYSQL_USER	MYSQL_USER	KIE server MySQL database username	rhpam	False
KIE_SERVER_MYSQL_PWD	–	KIE server MySQL database password	–	False
KIE_SERVER_MYSQL_DB	–	KIE server MySQL database name	rhpam7	False

Variable name	Image Environment Variable	Description	Example value	Required
MYSQL_IMAGE_STREAM_NAMESPACE	–	Namespace in which the ImageStream for the MySQL image is installed. The ImageStream is already installed in the openshift namespace. You should only need to modify this if you've installed the ImageStream in a different namespace/project. Default is "openshift".	openshift	False
MYSQL_IMAGE_STREAM_TAG	–	The MySQL image version, which is intended to correspond to the MySQL version. Default is "5.7".	5.7	False
DB_VOLUME_CAPACITY	–	Size of persistent storage for database volume.	1Gi	True
KIE_MBEANS	KIE_MBEANS	KIE server mbeans enabled/disabled (Sets the kie.mbeans and kie.scanner.mbeans system properties)	enabled	False
DROOLS_SERVER_FILTER_CLASSES	DROOLS_SERVER_FILTER_CLASSES	KIE server class filtering (Sets the org.drools.server.filter.classes system property)	true	False

Variable name	Image Environment Variable	Description	Example value	Required
BUSINESS_CENTRAL_HOSTNAME_HTTP	HOSTNAME_HTTP	Custom hostname for http service route. Leave blank for default hostname, e.g.: <application-name>-rhpamcentr-<project>.<default-domain-suffix>	`\${BUSINESS_CENTRAL_HOSTNAME_HTTP}`	False
BUSINESS_CENTRAL_HOSTNAME_HTTPS	HOSTNAME_HTTP	Custom hostname for https service route. Leave blank for default hostname, e.g.: secure-<application-name>-rhpamcentr-<project>.<default-domain-suffix>	`\${BUSINESS_CENTRAL_HOSTNAME_HTTP}`	False
KIE_SERVER_HOSTNAME_HTTP	HOSTNAME_HTTP	Custom hostname for http service route, if set will also configure the KIE_SERVER_HOST. Leave blank for default hostname, e.g.: <application-name>-kieserver-<project>.<default-domain-suffix>	`\${BUSINESS_CENTRAL_HOSTNAME_HTTP}`	False
KIE_SERVER_HOSTNAME_HTTPS	HOSTNAME_HTTP	Custom hostname for https service route. Leave blank for default hostname, e.g.: secure-<application-name>-kieserver-<project>.<default-domain-suffix>	`\${BUSINESS_CENTRAL_HOSTNAME_HTTP}`	False

Variable name	Image Environment Variable	Description	Example value	Required
KIE_SERVER_USE_SECURE_ROUTE_NAME	KIE_SERVER_USE_SECURE_ROUTE_NAME	Use https for the KIE_SERVER_HOST when it is not explicit configured to a custom value.	false	False
BUSINESS_CENTRAL_HTTPS_SECRET	–	The name of the secret containing the keystore file	–	True
BUSINESS_CENTRAL_HTTPS_KEYSTORE	HTTPS_KEYSTORE	The name of the keystore file within the secret	keystore.jks	False
BUSINESS_CENTRAL_HTTPS_NAME	HTTPS_NAME	The name associated with the server certificate	jboss	False
BUSINESS_CENTRAL_HTTPS_PASSWORD	HTTPS_PASSWORD	The password for the keystore and certificate	mykeystorepass	False
KIE_SERVER_HTTPS_SECRET	–	The name of the secret containing the keystore file	–	True
KIE_SERVER_HTTPS_KEYSTORE	HTTPS_KEYSTORE	The name of the keystore file within the secret	keystore.jks	False
KIE_SERVER_HTTPS_NAME	HTTPS_NAME	The name associated with the server certificate	jboss	False
KIE_SERVER_HTTPS_PASSWORD	HTTPS_PASSWORD	The password for the keystore and certificate	mykeystorepass	False
APPFORMER_ELASTIC_RETRIES	APPFORMER_ELASTIC_RETRIES	The number of times that appformer will try to connect to the elasticsearch node before give up.	\${APPFORMER_ELASTIC_RETRIES}	False

Variable name	Image Environment Variable	Description	Example value	Required
APPFORMER_JMS_BROKER_PORT	APPFORMER_JMS_BROKER_PORT	The port to connect in the JMS broker. Defaults to 61616	\${APPFORMER_JMS_BROKER_PORT}	False
APPFORMER_JMS_BROKER_USER	APPFORMER_JMS_BROKER_USER	The username to connect in the JMS broker.	jmsBrokserUser	True
APPFORMER_JMS_BROKER_PASSWORD	APPFORMER_JMS_BROKER_PASSWORD	The password to connect in the JMS broker.	\${APPFORMER_JMS_BROKER_PASSWORD}	True
ES_HOSTNAME_HTTP	HOSTNAME_HTTP	Custom hostname for http service route. Leave blank for default hostname, e.g.: <application-name>-rhpamindex-<project>.<default-domain-suffix>	\${BUSINESS_CENTRAL_HOSTNAME_HTTP}	False
APPFORMER_ELASTIC_CLUSTER_NAME	APPFORMER_ELASTIC_CLUSTER_NAME	Sets the ES cluster.name and configure it on Business Central. Defaults to kie-cluster.	\${APPFORMER_ELASTIC_CLUSTER_NAME}	False
ES_NODE_NAME	ES_NODE_NAME	Sets the ES node.name property. Defaults to HOSTNAME env value.	\${ES_NODE_NAME}	False

Variable name	Image Environment Variable	Description	Example value	Required
ES_TRANSPORT_HOST	ES_TRANSPORT_HOST	Sets the ES transport.host property. This will set the transport address of the main ES cluster node. Used for communication between nodes in the cluster. Defaults to container address.	`\${ES_TRANSPORT_HOST}`	False
APPFORMER_ELASTIC_PORT	APPFORMER_ELASTIC_PORT	Sets the ES http.host property. This will set the http address of the main ES cluster node. Used for communication between nodes in the cluster and the communication with Business Central.	`\${APPFORMER_ELASTIC_PORT}`	False
ES_HTTP_HOST	ES_HTTP_HOST	Sets the ES http.host property. This will set the http address of the main ES cluster node. Used to interact with cluster rest api. Defaults to the container ip address	`\${ES_HTTP_HOST}`	False
ES_HTTP_PORT	ES_HTTP_PORT	Sets the ES http.port property. This will set the http port of the main ES cluster node. Used to interact with cluster rest api.	`\${ES_HTTP_PORT}`	False

Variable name	Image Environment Variable	Description	Example value	Required
ES_JAVA_OPTS	ES_JAVA_OPTS	Appends custom jvm configurations/properties to ES jvm.options configuration file.	\${ES_JAVA_OPTS}	False
AMQ_IMAGE_STREAM_NAMESPACE	–	Namespace in which the ImageStream for the AMQ image is installed. Default is "openshift".	openshift	True
AMQ_IMAGE_STREAM_NAME	–	The name of the image stream to use for the AMQ broker. Default is "amq-broker71-openshift".	amq-broker71-openshift	True
AMQ_IMAGE_STREAM_TAG	–	The AMQ image stream tag. Default is "1.0".	1.0	True
AMQ_ROLE	AMQ_ROLE	User role for standard broker user.	admin	True
AMQ_NAME	AMQ_NAME	The name of the broker	broker	True
AMQ_GLOBAL_MAX_SIZE	AMQ_GLOBAL_MAX_SIZE	Maximum amount of memory which message data may consume (Default: Undefined, half of the system's memory).	100 gb	False
ES_VOLUME_CAPACITY	–	Size of persistent storage for Elasticsearch volume.	1Gi	True

Variable name	Image Environment Variable	Description	Example value	Required
IMAGE_STREAM_NAMESPACE	–	Namespace in which the ImageStreams for Red Hat Middleware images are installed. These ImageStreams are normally installed in the openshift namespace. You should only need to modify this if you've installed the ImageStreams in a different namespace/project.	openshift	True
KIE_SERVER_IMAGE_STREAM_NAME	–	The name of the image stream to use for KIE server. Default is "rhpam71-kieserver-openshift".	rhpam71-kieserver-openshift	True
IMAGE_STREAM_TAG	–	A named pointer to an image in an image stream. Default is "1.1".	1.1	True
MAVEN_REPO_ID	MAVEN_REPO_ID	The id to use for the maven repository, if set. Default is generated randomly.	\${MAVEN_REPO_ID}	False
MAVEN_REPO_URL	MAVEN_REPO_URL	Fully qualified URL to a Maven repository or service.	\${MAVEN_REPO_URL}	False
MAVEN_REPO_USERNAME	MAVEN_REPO_USERNAME	Username to access the Maven repository, if required.	\${MAVEN_REPO_USERNAME}	False

Variable name	Image Environment Variable	Description	Example value	Required
MAVEN_REPO_PASSWORD	MAVEN_REPO_PASSWORD	Password to access the Maven repository, if required.	`\${MAVEN_REPO_PASSWORD}`	False
BUSINESS_CENTRAL_MAVEN_USERNAME	–	Username to access the Maven service hosted by Business Central inside EAP.	mavenUser	True
BUSINESS_CENTRAL_MAVEN_PASSWORD	–	Password to access the Maven service hosted by Business Central inside EAP.	–	True
GIT_HOOKS_DIR	GIT_HOOKS_DIR	The directory to use for git hooks, if required.	`\${GIT_HOOKS_DIR}`	False
TIMER_SERVICE_DATA_STORE_REFRESH_INTERVAL	TIMER_SERVICE_DATA_STORE_REFRESH_INTERVAL	Sets refresh-interval for the EJB timer database data-store service.	60000	True
BUSINESS_CENTRAL_VOLUME_CAPACITY	–	Size of the persistent storage for Business Central's runtime data.	1Gi	True
BUSINESS_CENTRAL_MEMORY_LIMIT	–	Business Central Container memory limit	2Gi	False
KIE_SERVER_MEMORY_LIMIT	–	KIE server Container memory limit	1Gi	False
SSO_URL	SSO_URL	RH-SSO URL	`\${SSO_URL}`	False
SSO_REALM	SSO_REALM	RH-SSO Realm name	`\${SSO_REALM}`	False

Variable name	Image Environment Variable	Description	Example value	Required
BUSINESS_CENTRAL_SSO_CLIENT	SSO_CLIENT	Business Central RH-SSO Client name	\${BUSINESS_CENTRAL_SSO_CLIENT}	False
BUSINESS_CENTRAL_SSO_SECRET	SSO_SECRET	Business Central RH-SSO Client Secret	\${BUSINESS_CENTRAL_SSO_SECRET}	False
KIE_SERVER_SSO_CLIENT	SSO_CLIENT	KIE Server RH-SSO Client name	\${BUSINESS_CENTRAL_SSO_CLIENT}	False
KIE_SERVER_SSO_SECRET	SSO_SECRET	KIE Server RH-SSO Client Secret	\${BUSINESS_CENTRAL_SSO_SECRET}	False
SSO_USERNAME	SSO_USERNAME	RH-SSO Realm Admin Username used to create the Client if it doesn't exist	\${SSO_USERNAME}	False
SSO_PASSWORD	SSO_PASSWORD	RH-SSO Realm Admin Password used to create the Client	\${SSO_PASSWORD}	False
SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	RH-SSO Disable SSL Certificate Validation	false	False
SSO_PRINCIPAL_ATTRIBUTE	SSO_PRINCIPAL_ATTRIBUTE	RH-SSO Principal Attribute to use as username.	preferred_username	False
AUTH_LDAP_URL	AUTH_LDAP_URL	LDAP Endpoint to connect for authentication	\${AUTH_LDAP_URL}	False
AUTH_LDAP_BIND_DN	AUTH_LDAP_BIND_DN	Bind DN used for authentication	\${AUTH_LDAP_BIND_DN}	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_BIND_CREDENTIAL	AUTH_LDAP_BIND_CREDENTIAL	LDAP Credentials used for authentication	`\${AUTH_LDAP_BIND_CREDENTIAL}`	False
AUTH_LDAP_JAAS_SECURITY_DOMAIN	AUTH_LDAP_JAAS_SECURITY_DOMAIN	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	`\${AUTH_LDAP_JAAS_SECURITY_DOMAIN}`	False
AUTH_LDAP_BASE_CTX_DN	AUTH_LDAP_BASE_CTX_DN	LDAP Base DN of the top-level context to begin the user search.	`\${AUTH_LDAP_BASE_CTX_DN}`	False
AUTH_LDAP_BASE_FILTER	AUTH_LDAP_BASE_FILTER	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}).	`\${AUTH_LDAP_BASE_FILTER}`	False
AUTH_LDAP_SEARCH_SCOPE	AUTH_LDAP_SEARCH_SCOPE	The search scope to use.	`\${AUTH_LDAP_SEARCH_SCOPE}`	False
AUTH_LDAP_SEARCH_TIME_LIMIT	AUTH_LDAP_SEARCH_TIME_LIMIT	The timeout in milliseconds for user or role searches.	`\${AUTH_LDAP_SEARCH_TIME_LIMIT}`	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE	AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE	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.	`\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}`	False
AUTH_LDAP_PARSE_USERNAME	AUTH_LDAP_PARSE_USERNAME	A flag indicating if the DN is to be parsed for the username. If set to true, the DN is parsed for the username. If set to false the DN is not parsed for the username. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	`\${AUTH_LDAP_PARSE_USERNAME}`	False
AUTH_LDAP_USERNAME_BEGIN_STRING	AUTH_LDAP_USERNAME_BEGIN_STRING	Defines the String which is to be removed from the start of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	`\${AUTH_LDAP_USERNAME_BEGIN_STRING}`	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_USERNAME_END_STRING	AUTH_LDAP_USERNAME_END_STRING	Defines the String which is to be removed from the end of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<code>\${AUTH_LDAP_USERNAME_END_STRING}</code>	False
AUTH_LDAP_ROLE_ATTRIBUTE_ID	AUTH_LDAP_ROLE_ATTRIBUTE_ID	Name of the attribute containing the user roles.	<code>\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}</code>	False
AUTH_LDAP_ROLES_CTX_DN	AUTH_LDAP_ROLES_CTX_DN	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.	<code>\${AUTH_LDAP_ROLES_CTX_DN}</code>	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_ROLE_FILTER	AUTH_LDAP_ROLE_FILTER	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}).	`\${AUTH_LDAP_ROLE_FILTER}`	False
AUTH_LDAP_ROLE_RECURSION	AUTH_LDAP_ROLE_RECURSION	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	`\${AUTH_LDAP_ROLE_RECURSION}`	False
AUTH_LDAP_DEFAULT_ROLE	AUTH_LDAP_DEFAULT_ROLE	A role included for all authenticated users	`\${AUTH_LDAP_DEFAULT_ROLE}`	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID	AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID	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.	`\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}`	False
AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN	AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN	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.	`\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}`	False
AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN	AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN	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 roleNameAttribute ID attribute of the context name. Certain directory schemas, such as Microsoft Active Directory, require this attribute to be set to true.	`\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}`	False

Variable name	Image Environment Variable	Description	Example value	Required
AUTH_LDAP_REFERRAL_USE_R_ATTRIBUTE_ID_TO_CHECK	AUTH_LDAP_REFERRAL_USE_R_ATTRIBUTE_ID_TO_CHECK	If you are not using referrals, this option can be ignored. 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.	<code>\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}</code>	False

4.2.2. 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](#).

4.2.2.1. Services

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

Service	Port	Name	Description
<code>\${APPLICATION_NAME}-rhpamcentr</code>	8080	http	All the Business Central web server's ports.
	8443	https	
	8001	git-ssh	
<code>\${APPLICATION_NAME}-ping</code>	8888	ping	The JGroups ping port for clustering.
<code>\${APPLICATION_NAME}-kieserver</code>	8080	http	All the KIE server web server's ports.

Service	Port	Name	Description
	8443	https	
\${APPLICATION_NAME}-rhpamindex	9200	rest	All the Business Central Indexing Elasticsearch ports.
	9300	transport	
\${APPLICATION_NAME}-amq-tcp	61616	–	The broker's OpenWire port.
\${APPLICATION_NAME}-mysql	3306	–	The MySQL server's port.

4.2.2.2. 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. Refer to the [Openshift documentation](#) for more information.

Service	Security	Hostname
\${APPLICATION_NAME}-rhpamcentr-http	none	\${BUSINESS_CENTRAL_HOSTNAME_HTTP}
\${APPLICATION_NAME}-rhpamcentr-https	TLS passthrough	\${BUSINESS_CENTRAL_HOSTNAME_HTTPS}
\${APPLICATION_NAME}-kieserver-http	none	\${KIE_SERVER_HOSTNAME_HTTP}
\${APPLICATION_NAME}-kieserver-https	TLS passthrough	\${KIE_SERVER_HOSTNAME_HTTPS}
\${APPLICATION_NAME}-rhpamindex-http	none	\${ES_HOSTNAME_HTTP}

4.2.2.3. 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. Refer to the [Openshift documentation](#) for more information.

4.2.2.3.1. Triggers

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

Deployment	Triggers
\${APPLICATION_NAME}-rhpamcentr	ImageChange
\${APPLICATION_NAME}-kieserver	ImageChange
\${APPLICATION_NAME}-rhpamindex	ImageChange
\${APPLICATION_NAME}-amq	ImageChange
\${APPLICATION_NAME}-mysql	ImageChange

4.2.2.3.2. 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. Refer to the [container-engine documentation](#) for more information.

Deployment	Replicas
\${APPLICATION_NAME}-rhpamcentr	2
\${APPLICATION_NAME}-kieserver	2
\${APPLICATION_NAME}-rhpamindex	1
\${APPLICATION_NAME}-amq	1
\${APPLICATION_NAME}-mysql	1

4.2.2.3.3. Pod Template

4.2.2.3.3.1. Service Accounts

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

Deployment	Service Account
\${APPLICATION_NAME}-rhpamcentr	\${APPLICATION_NAME}-rhpamsvc
\${APPLICATION_NAME}-kieserver	\${APPLICATION_NAME}-rhpamsvc

4.2.2.3.3.2. Image

Deployment	Image
<code>\${APPLICATION_NAME}-rhpamcentr</code>	<code>rhpam71-businesscentral-openshift</code>
<code>\${APPLICATION_NAME}-kieserver</code>	<code>\${KIE_SERVER_IMAGE_STREAM_NAME}</code>
<code>\${APPLICATION_NAME}-rhpamindex</code>	<code>rhpam71-businesscentral-indexing-openshift</code>
<code>\${APPLICATION_NAME}-amq</code>	<code>\${AMQ_IMAGE_STREAM_NAME}</code>
<code>\${APPLICATION_NAME}-mysql</code>	<code>mysql</code>

4.2.2.3.3.3. Readiness Probe

`${APPLICATION_NAME}-rhpamcentr`

```
/bin/bash -c curl --fail --silent -u '${KIE_ADMIN_USER}:${KIE_ADMIN_PWD}'
http://localhost:8080/kie-wb.jsp
```

`${APPLICATION_NAME}-kieserver`

```
/bin/bash -c curl --fail --silent -u '${KIE_ADMIN_USER}:${KIE_ADMIN_PWD}'
http://localhost:8080/services/rest/server/readycheck
```

`${APPLICATION_NAME}-rhpamindex`

```
Http Get on http://localhost:9200/_cluster/health
```

`${APPLICATION_NAME}-amq`

```
/bin/bash -c /opt/amq/bin/readinessProbe.sh
```

`${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'
```

4.2.2.3.3.4. Liveness Probe

`${APPLICATION_NAME}-rhpamcentr`

```
/bin/bash -c curl --fail --silent -u '${KIE_ADMIN_USER}:${KIE_ADMIN_PWD}'
http://localhost:8080/kie-wb.jsp
```

`${APPLICATION_NAME}-kieserver`

```
/bin/bash -c curl --fail --silent -u '${KIE_ADMIN_USER}:${KIE_ADMIN_PWD}'
http://localhost:8080/services/rest/server/readycheck
```

4.2.2.3.3.5. Exposed Ports

Deployments	Name	Port	Protocol
\${APPLICATION_NAME}-rhpamcentr	jolokia	8778	TCP
	http	8080	TCP
	https	8443	TCP
	ping	8888	TCP
\${APPLICATION_NAME}-kieserver	jolokia	8778	TCP
	http	8080	TCP
	https	8443	TCP
\${APPLICATION_NAME}-rhpamindex	es	9300	TCP
	http	9200	TCP
\${APPLICATION_NAME}-amq	jolokia	8161	TCP
	amqp	5672	TCP
	mqtt	1883	TCP
	stomp	61613	TCP
	artemis	61616	TCP
\${APPLICATION_NAME}-mysql	–	3306	TCP

4.2.2.3.3.6. Image Environment Variables

Deployment	Variable name	Description	Example value
\${APPLICATION_NAME}-rhpamcentr	KIE_ADMIN_USER	KIE administrator username	\${KIE_ADMIN_USER}
	KIE_ADMIN_PWD	KIE administrator password	\${KIE_ADMIN_PWD}

Deployment	Variable name	Description	Example value
	KIE_MBEANS	KIE server mbeans enabled/disabled (Sets the kie.mbeans and kie.scanner.mbeans system properties)	\${KIE_MBEANS}
	KIE_SERVER_CONTROLLER_USER	KIE server controller username (Sets the org.kie.server.controller.user system property)	\${KIE_SERVER_CONTROLLER_USER}
	KIE_SERVER_CONTROLLER_PWD	KIE server controller password (Sets the org.kie.server.controller.pwd system property)	\${KIE_SERVER_CONTROLLER_PWD}
	KIE_SERVER_CONTROLLER_TOKEN	KIE server controller token for bearer authentication (Sets the org.kie.server.controller.token system property)	\${KIE_SERVER_CONTROLLER_TOKEN}
	KIE_SERVER_USER	KIE server username (Sets the org.kie.server.user system property)	\${KIE_SERVER_USER}
	KIE_SERVER_PWD	KIE server password (Sets the org.kie.server.pwd system property)	\${KIE_SERVER_PWD}
	MAVEN_REPO_ID	The id to use for the maven repository, if set. Default is generated randomly.	\${MAVEN_REPO_ID}
	MAVEN_REPO_URL	Fully qualified URL to a Maven repository or service.	\${MAVEN_REPO_URL}
	MAVEN_REPO_USERNAME	Username to access the Maven repository, if required.	\${MAVEN_REPO_USERNAME}
	MAVEN_REPO_PASSWORD	Password to access the Maven repository, if required.	\${MAVEN_REPO_PASSWORD}

Deployment	Variable name	Description	Example value
	KIE_MAVEN_USER	–	\${BUSINESS_CENTRAL_MAVEN_USERNAME}
	KIE_MAVEN_PWD	–	\${BUSINESS_CENTRAL_MAVEN_PASSWORD}
	GIT_HOOKS_DIR	The directory to use for git hooks, if required.	\${GIT_HOOKS_DIR}
	HTTPS_KEYSTORE_DIR	–	/etc/businesscentral-secret-volume
	HTTPS_KEYSTORE	The name of the keystore file within the secret	\${BUSINESS_CENTRAL_HTTPS_KEYSTORE}
	HTTPS_NAME	The name associated with the server certificate	\${BUSINESS_CENTRAL_HTTPS_NAME}
	HTTPS_PASSWORD	The password for the keystore and certificate	\${BUSINESS_CENTRAL_HTTPS_PASSWORD}
	WORKBENCH_ROUTE_NAME	–	\${APPLICATION_NAME}-rhpamcentr
	JGROUPS_PING_PROTOCOL	–	openshift.DNS_PING
	OPENSIFT_DNS_PING_SERVICE_NAME	–	\${APPLICATION_NAME}-ping
	OPENSIFT_DNS_PING_SERVICE_PORT	–	8888
	APPFORMER_ELASTIC_PORT	Sets the ES http.host property. This will set the http address of the main ES cluster node. Used for communication between nodes in the cluster and the communication with Business Central.	\${APPFORMER_ELASTIC_PORT}

Deployment	Variable name	Description	Example value
	APPFORMER_ELASTIC_CLUSTER_NAME	Sets the ES cluster.name and configure it on Business Central. Defaults to kie-cluster.	`\${APPFORMER_ELASTIC_CLUSTER_NAME}`
	APPFORMER_ELASTIC_RETRIES	The number of times that appformer will try to connect to the elasticsearch node before give up.	`\${APPFORMER_ELASTIC_RETRIES}`
	APPFORMER_ELASTIC_HOST	–	`\${APPLICATION_NAME}-rhpamindex`
	APPFORMER_JMS_BROKER_ADDRESS	–	`\${APPLICATION_NAME}-amq-tcp`
	APPFORMER_JMS_BROKER_PORT	The port to connect in the JMS broker. Defaults to 61616	`\${APPFORMER_JMS_BROKER_PORT}`
	APPFORMER_JMS_BROKER_USER	The username to connect in the JMS broker.	`\${APPFORMER_JMS_BROKER_USER}`
	APPFORMER_JMS_BROKER_PASSWORD	The password to connect in the JMS broker.	`\${APPFORMER_JMS_BROKER_PASSWORD}`
	SSO_URL	RH-SSO URL	`\${SSO_URL}`
	SSO_OPENIDCONNECT_DEPLOYMENTS	–	ROOT.war
	SSO_REALM	RH-SSO Realm name	`\${SSO_REALM}`
	SSO_SECRET	Business Central RH-SSO Client Secret	`\${BUSINESS_CENTRAL_SSO_SECRET}`
	SSO_CLIENT	Business Central RH-SSO Client name	`\${BUSINESS_CENTRAL_SSO_CLIENT}`
	SSO_USERNAME	RH-SSO Realm Admin Username used to create the Client if it doesn't exist	`\${SSO_USERNAME}`

Deployment	Variable name	Description	Example value
	SSO_PASSWORD	RH-SSO Realm Admin Password used to create the Client	\${SSO_PASSWORD}
	SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	RH-SSO Disable SSL Certificate Validation	\${SSO_DISABLE_SSL_CERTIFICATE_VALIDATION}
	SSO_PRINCIPAL_ATTRIBUTE	RH-SSO Principal Attribute to use as username.	\${SSO_PRINCIPAL_ATTRIBUTE}
	HOSTNAME_HTTP	Custom hostname for http service route. Leave blank for default hostname, e.g.: <application-name>-rhpamcentr-<project>. <default-domain-suffix>	\${BUSINESS_CENTRAL_HOSTNAME_HTTP}
	HOSTNAME_HTTPS	Custom hostname for https service route. Leave blank for default hostname, e.g.: secure-<application-name>-rhpamcentr-<project>. <default-domain-suffix>	\${BUSINESS_CENTRAL_HOSTNAME_HTTPS}
	AUTH_LDAP_URL	LDAP Endpoint to connect for authentication	\${AUTH_LDAP_URL}
	AUTH_LDAP_BIND_DN	Bind DN used for authentication	\${AUTH_LDAP_BIND_DN}
	AUTH_LDAP_BIND_CREDENTIAL	LDAP Credentials used for authentication	\${AUTH_LDAP_BIND_CREDENTIAL}
	AUTH_LDAP_JAAS_SECURITY_DOMAIN	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	\${AUTH_LDAP_JAAS_SECURITY_DOMAIN}
	AUTH_LDAP_BASE_CTX_DN	LDAP Base DN of the top-level context to begin the user search.	\${AUTH_LDAP_BASE_CTX_DN}

Deployment	Variable name	Description	Example value
	AUTH_LDAP_BASE_FILTER	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}).	`\${AUTH_LDAP_BASE_FILTER}`
	AUTH_LDAP_SEARCH_SCOPE	The search scope to use.	`\${AUTH_LDAP_SEARCH_SCOPE}`
	AUTH_LDAP_SEARCH_TIME_LIMIT	The timeout in milliseconds for user or role searches.	`\${AUTH_LDAP_SEARCH_TIME_LIMIT}`
	AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE	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.	`\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}`
	AUTH_LDAP_PARSE_USERNAME	A flag indicating if the DN is to be parsed for the username. If set to true, the DN is parsed for the username. If set to false the DN is not parsed for the username. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	`\${AUTH_LDAP_PARSE_USERNAME}`

Deployment	Variable name	Description	Example value
	AUTH_LDAP_USER_NAME_BEGIN_STRING	Defines the String which is to be removed from the start of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<code>\${AUTH_LDAP_USER_NAME_BEGIN_STRING}</code>
	AUTH_LDAP_USER_NAME_END_STRING	Defines the String which is to be removed from the end of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<code>\${AUTH_LDAP_USER_NAME_END_STRING}</code>
	AUTH_LDAP_ROLE_ATTRIBUTE_ID	Name of the attribute containing the user roles.	<code>\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}</code>
	AUTH_LDAP_ROLE_S_CTX_DN	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.	<code>\${AUTH_LDAP_ROLE_S_CTX_DN}</code>

Deployment	Variable name	Description	Example value
	AUTH_LDAP_ROLE_FILTER	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}).	`\${AUTH_LDAP_ROLE_FILTER}`
	AUTH_LDAP_ROLE_RECURSION	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	`\${AUTH_LDAP_ROLE_RECURSION}`
	AUTH_LDAP_DEFAULT_ROLE	A role included for all authenticated users	`\${AUTH_LDAP_DEFAULT_ROLE}`
	AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID	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.	`\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}`

Deployment	Variable name	Description	Example value
	AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN	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.	\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}
	AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN	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.	\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}
	AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK	If you are not using referrals, this option can be ignored. 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.	\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}
\${APPLICATION_NAME}-kieserver	AUTO_CONFIGURE_EJB_TIMER	–	true

Deployment	Variable name	Description	Example value
	TIMER_SERVICE_DATA_STORE_REFRESH_INTERVAL	Sets refresh-interval for the EJB timer database data-store service.	\${TIMER_SERVICE_DATA_STORE_REFRESH_INTERVAL}
	DATASOURCES	–	RHPAM
	RHPAM_DATABASE	–	\${KIE_SERVER_MYSQL_DB}
	RHPAM_DRIVER	–	mysql
	RHPAM_USERNAME	–	\${KIE_SERVER_MYSQL_USER}
	RHPAM_PASSWORD	–	\${KIE_SERVER_MYSQL_PWD}
	RHPAM_SERVICE_HOST	–	\${APPLICATION_NAME}-mysql
	RHPAM_SERVICE_PORT	–	3306
	KIE_SERVER_PERSISTENCE_DIALECT	–	org.hibernate.dialect.MySQLDialect
	KIE_SERVER_PERSISTENCE_DS	KIE server persistence datasource (Sets the org.kie.server.persistence.ds system property)	\${KIE_SERVER_PERSISTENCE_DS}
	RHPAM_JNDI	–	\${KIE_SERVER_PERSISTENCE_DS}
	RHPAM_JTA	–	true
	DROOLS_SERVER_FILTER_CLASSES	KIE server class filtering (Sets the org.drools.server.filter.classes system property)	\${DROOLS_SERVER_FILTER_CLASSES}
	KIE_ADMIN_PWD	KIE administrator password	\${KIE_ADMIN_PWD}
	KIE_ADMIN_USER	KIE administrator username	\${KIE_ADMIN_USER}

Deployment	Variable name	Description	Example value
	KIE_MBEANS	KIE server mbeans enabled/disabled (Sets the kie.mbeans and kie.scanner.mbeans system properties)	\${KIE_MBEANS}
	KIE_SERVER_BYPASS_AUTH_USER	KIE server bypass auth user (Sets the org.kie.server.bypass.auth.user system property)	\${KIE_SERVER_BYPASS_AUTH_USER}
	KIE_SERVER_CONTROLLER_USER	KIE server controller username (Sets the org.kie.server.controller.user system property)	\${KIE_SERVER_CONTROLLER_USER}
	KIE_SERVER_CONTROLLER_PWD	KIE server controller password (Sets the org.kie.server.controller.pwd system property)	\${KIE_SERVER_CONTROLLER_PWD}
	KIE_SERVER_CONTROLLER_TOKEN	KIE server controller token for bearer authentication (Sets the org.kie.server.controller.token system property)	\${KIE_SERVER_CONTROLLER_TOKEN}
	KIE_SERVER_CONTROLLER_SERVICE	–	\${APPLICATION_NAME}-rhpamcentr
	KIE_SERVER_CONTROLLER_PROTOCOL	–	ws
	KIE_SERVER_ID	–	\${APPLICATION_NAME}-kieserver
	KIE_SERVER_HOST	Custom hostname for http service route, if set will also configure the KIE_SERVER_HOST. Leave blank for default hostname, e.g.: <application-name>-kieserver-<project>. <default-domain-suffix>	\${KIE_SERVER_HOSTNAME_HTTP}
	KIE_SERVER_ROUTE_NAME	–	\${APPLICATION_NAME}-kieserver

Deployment	Variable name	Description	Example value
	KIE_SERVER_USE_SECURE_ROUTE_NAME	Use https for the KIE_SERVER_HOST when it is not explicit configured to a custom value.	\${KIE_SERVER_USE_SECURE_ROUTE_NAME}
	KIE_SERVER_PWD	KIE server password (Sets the org.kie.server.pwd system property)	\${KIE_SERVER_PWD}
	KIE_SERVER_USER	KIE server username (Sets the org.kie.server.user system property)	\${KIE_SERVER_USER}
	MAVEN_REPOS	–	RHPAMCENTR,EXTERNAL
	RHPAMCENTR_MAVEN_REPO_SERVICE	–	\${APPLICATION_NAME}-rhpamcentr
	RHPAMCENTR_MAVEN_REPO_PATH	–	/maven2/
	RHPAMCENTR_MAVEN_REPO_USERNAME	Username to access the Maven repository, if required.	\${BUSINESS_CENTRAL_MAVEN_USERNAME}
	RHPAMCENTR_MAVEN_REPO_PASSWORD	Password to access the Maven repository, if required.	\${BUSINESS_CENTRAL_MAVEN_PASSWORD}
	EXTERNAL_MAVEN_REPO_ID	The id to use for the maven repository, if set. Default is generated randomly.	\${MAVEN_REPO_ID}
	EXTERNAL_MAVEN_REPO_URL	Fully qualified URL to a Maven repository or service.	\${MAVEN_REPO_URL}
	EXTERNAL_MAVEN_REPO_USERNAME	Username to access the Maven repository, if required.	\${MAVEN_REPO_USERNAME}

Deployment	Variable name	Description	Example value
	EXTERNAL_MAVEN_REPO_PASSWORD	Password to access the Maven repository, if required.	`\${MAVEN_REPO_PASSWORD}`
	HTTPS_KEYSTORE_DIR	–	<code>/etc/kieserver-secret-volume</code>
	HTTPS_KEYSTORE	The name of the keystore file within the secret	`\${KIE_SERVER_HTTPS_KEYSTORE}`
	HTTPS_NAME	The name associated with the server certificate	`\${KIE_SERVER_HTTPS_NAME}`
	HTTPS_PASSWORD	The password for the keystore and certificate	`\${KIE_SERVER_HTTPS_PASSWORD}`
	SSO_URL	RH-SSO URL	`\${SSO_URL}`
	SSO_OPENIDCONNECT_DEPLOYMENTS	–	ROOT.war
	SSO_REALM	RH-SSO Realm name	`\${SSO_REALM}`
	SSO_SECRET	Business Central RH-SSO Client Secret	`\${KIE_SERVER_SSO_SECRET}`
	SSO_CLIENT	Business Central RH-SSO Client name	`\${KIE_SERVER_SSO_CLIENT}`
	SSO_USERNAME	RH-SSO Realm Admin Username used to create the Client if it doesn't exist	`\${SSO_USERNAME}`
	SSO_PASSWORD	RH-SSO Realm Admin Password used to create the Client	`\${SSO_PASSWORD}`
	SSO_DISABLE_SSL_CERTIFICATE_VALIDATION	RH-SSO Disable SSL Certificate Validation	`\${SSO_DISABLE_SSL_CERTIFICATE_VALIDATION}`

Deployment	Variable name	Description	Example value
	SSO_PRINCIPAL_tribute	RH-SSO Principal Attribute to use as username.	`\${SSO_PRINCIPAL_tribute}`
	HOSTNAME_HTTP	Custom hostname for http service route. Leave blank for default hostname, e.g.: <application-name>-rhpamcentr-<project>. <default-domain-suffix>	`\${KIE_SERVER_HOSTNAME_HTTP}`
	HOSTNAME_HTTPS	Custom hostname for https service route. Leave blank for default hostname, e.g.: secure-<application-name>-rhpamcentr-<project>. <default-domain-suffix>	`\${KIE_SERVER_HOSTNAME_HTTPS}`
	AUTH_LDAP_URL	LDAP Endpoint to connect for authentication	`\${AUTH_LDAP_URL}`
	AUTH_LDAP_BIND_DN	Bind DN used for authentication	`\${AUTH_LDAP_BIND_DN}`
	AUTH_LDAP_BIND_CREDENTIAL	LDAP Credentials used for authentication	`\${AUTH_LDAP_BIND_CREDENTIAL}`
	AUTH_LDAP_JAAS_SECURITY_DOMAIN	The JMX ObjectName of the JaasSecurityDomain used to decrypt the password.	`\${AUTH_LDAP_JAAS_SECURITY_DOMAIN}`
	AUTH_LDAP_BASE_CTX_DN	LDAP Base DN of the top-level context to begin the user search.	`\${AUTH_LDAP_BASE_CTX_DN}`

Deployment	Variable name	Description	Example value
	AUTH_LDAP_BASE_FILTER	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}).	`\${AUTH_LDAP_BASE_FILTER}`
	AUTH_LDAP_SEARCH_SCOPE	The search scope to use.	`\${AUTH_LDAP_SEARCH_SCOPE}`
	AUTH_LDAP_SEARCH_TIME_LIMIT	The timeout in milliseconds for user or role searches.	`\${AUTH_LDAP_SEARCH_TIME_LIMIT}`
	AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE	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.	`\${AUTH_LDAP_DISTINGUISHED_NAME_ATTRIBUTE}`
	AUTH_LDAP_PARSE_USERNAME	A flag indicating if the DN is to be parsed for the username. If set to true, the DN is parsed for the username. If set to false the DN is not parsed for the username. This option is used together with <code>usernameBeginString</code> and <code>usernameEndString</code> .	`\${AUTH_LDAP_PARSE_USERNAME}`

Deployment	Variable name	Description	Example value
	AUTH_LDAP_USER_NAME_BEGIN_STRING	Defines the String which is to be removed from the start of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<code>\${AUTH_LDAP_USER_NAME_BEGIN_STRING}</code>
	AUTH_LDAP_USER_NAME_END_STRING	Defines the String which is to be removed from the end of the DN to reveal the username. This option is used together with <code>usernameEndString</code> and only taken into account if <code>parseUsername</code> is set to true.	<code>\${AUTH_LDAP_USER_NAME_END_STRING}</code>
	AUTH_LDAP_ROLE_ATTRIBUTE_ID	Name of the attribute containing the user roles.	<code>\${AUTH_LDAP_ROLE_ATTRIBUTE_ID}</code>
	AUTH_LDAP_ROLE_S_CTX_DN	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.	<code>\${AUTH_LDAP_ROLE_S_CTX_DN}</code>

Deployment	Variable name	Description	Example value
	AUTH_LDAP_ROLE_FILTER	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}).	`\${AUTH_LDAP_ROLE_FILTER}`
	AUTH_LDAP_ROLE_RECURSION	The number of levels of recursion the role search will go below a matching context. Disable recursion by setting this to 0.	`\${AUTH_LDAP_ROLE_RECURSION}`
	AUTH_LDAP_DEFAULT_ROLE	A role included for all authenticated users	`\${AUTH_LDAP_DEFAULT_ROLE}`
	AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID	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.	`\${AUTH_LDAP_ROLE_NAME_ATTRIBUTE_ID}`

Deployment	Variable name	Description	Example value
	AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN	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.	\${AUTH_LDAP_PARSE_ROLE_NAME_FROM_DN}
	AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN	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.	\${AUTH_LDAP_ROLE_ATTRIBUTE_IS_DN}
	AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK	If you are not using referrals, this option can be ignored. 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.	\${AUTH_LDAP_REFERRAL_USER_ATTRIBUTE_ID_TO_CHECK}
\${APPLICATION_NAME}-rhpamindex	ES_CLUSTER_NAME	–	\${APPFORMER_ELASTIC_CLUSTER_NAME}

Deployment	Variable name	Description	Example value
	ES_NODE_NAME	Sets the ES node.name property. Defaults to HOSTNAME env value.	\${ES_NODE_NAME}
	ES_TRANSPORT_HOST	Sets the ES transport.host property. This will set the transport address of the main ES cluster node. Used for communication between nodes in the cluster. Defaults to container address.	\${ES_TRANSPORT_HOST}
	ES_TRANSPORT_TCP_PORT	–	\${APPFORMER_ELASTIC_PORT}
	ES_HTTP_PORT	Sets the ES http.port property. This will set the http port of the main ES cluster node. Used to interact with cluster rest api.	\${ES_HTTP_PORT}
	ES_HTTP_HOST	Sets the ES http.host property. This will set the http address of the main ES cluster node. Used to interact with cluster rest api. Defaults to the container ip address	\${ES_HTTP_HOST}
	ES_JAVA_OPTS	Appends custom jvm configurations/properties to ES jvm.options configuration file.	\${ES_JAVA_OPTS}
\${APPLICATION_NAME}-amq	AMQ_USER	–	\${APPFORMER_JMS_BROKER_USER}
	AMQ_PASSWORD	–	\${APPFORMER_JMS_BROKER_PASSWORD}
	AMQ_ROLE	User role for standard broker user.	\${AMQ_ROLE}
	AMQ_NAME	The name of the broker	\${AMQ_NAME}

Deployment	Variable name	Description	Example value
	AMQ_TRANSPORTS	–	openwire
	AMQ_GLOBAL_MAX_SIZE	Maximum amount of memory which message data may consume (Default: Undefined, half of the system's memory).	\${AMQ_GLOBAL_MAX_SIZE}
\${APPLICATION_NAME}-mysql	MYSQL_USER	KIE server MySQL database username	\${KIE_SERVER_MYSQL_USER}
	MYSQL_PASSWORD	–	\${KIE_SERVER_MYSQL_PWD}
	MYSQL_DATABASE	–	\${KIE_SERVER_MYSQL_DB}

4.2.2.3.3.7. Volumes

Deployment	Name	mountPath	Purpose	readOnly
\${APPLICATION_NAME}-rhpamcentr	businesscentral-keystore-volume	/etc/businesscentral-secret-volume	ssl certs	True
\${APPLICATION_NAME}-kieserver	kieserver-keystore-volume	/etc/kieserver-secret-volume	ssl certs	True
\${APPLICATION_NAME}-rhpamindex	\${APPLICATION_NAME}-rhpamindex-pvol	/opt/elasticsearch/data	rhpamindex	false
\${APPLICATION_NAME}-mysql	\${APPLICATION_NAME}-mysql-pvol	/var/lib/mysql/data	mysql	false

4.2.2.4. External Dependencies

4.2.2.4.1. 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. Refer to the [OpenShift documentation](#) for more information.

Name	Access Mode
<code>\${APPLICATION_NAME}-rhpamcentr-claim</code>	ReadWriteMany
<code>\${APPLICATION_NAME}-mysql-claim</code>	ReadWriteOnce
<code>\${APPLICATION_NAME}-rhpamindex-claim</code>	ReadWriteOnce

4.2.2.4.2. Secrets

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

businesscentral-app-secret kieserver-app-secret

4.2.2.4.3. 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 4.1. 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
```

4.3. 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 run 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 run 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 run 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 CLI operation](#).

APPENDIX A. VERSIONING INFORMATION

Documentation last updated on Monday, December 21, 2020.