



Red Hat Process Automation Manager 7.2

Installing and configuring Red Hat Process Automation Manager in a Red Hat JBoss EAP clustered environment

Red Hat Process Automation Manager 7.2 Installing and configuring Red Hat Process Automation Manager in a Red Hat JBoss EAP clustered environment

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

Legal Notice

Copyright © 2020 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 create a Red Hat Process Automation Manager 7.2 clustered environment on Red Hat JBoss Enterprise Application Platform 7.2.

Table of Contents

PREFACE	3
CHAPTER 1. RED HAT PROCESS AUTOMATION MANAGER CLUSTERS	4
CHAPTER 2. RED HAT PROCESS AUTOMATION MANAGER CLUSTERS IN A DEVELOPMENT (AUTHORING) ENVIRONMENT	5
2.1. INSTALLING AND CONFIGURING ELASTICSEARCH	5
2.2. DOWNLOADING AND CONFIGURING AMQ BROKER	6
2.3. DOWNLOADING AND EXTRACTING RED HAT JBOSS EAP 7.2 AND RED HAT PROCESS AUTOMATION MANAGER	6
2.4. CONFIGURING AND RUNNING BUSINESS CENTRAL IN A CLUSTER	8
2.5. VERIFYING THE RED HAT PROCESS AUTOMATION MANAGER CLUSTER	9
CHAPTER 3. PROCESS SERVER CLUSTERS IN A RUNTIME ENVIRONMENT	11
3.1. DOWNLOADING AND EXTRACTING RED HAT JBOSS EAP 7.2 AND PROCESS SERVER	11
3.2. CONFIGURING AND RUNNING A RED HAT JBOSS EAP 7.2 CLUSTER FOR PROCESS SERVER	12
3.3. CLUSTERING WITH THE HEADLESS PROCESS AUTOMATION MANAGER CONTROLLER	15
3.4. CLUSTERING SMART ROUTER FOR A RUNTIME ENVIRONMENT	16
CHAPTER 4. ADDITIONAL RESOURCES	19
APPENDIX A. VERSIONING INFORMATION	20

PREFACE

As a system engineer, you can create a Red Hat Process Automation Manager clustered environment to provide high availability and load balancing for your development and runtime environments.

Prerequisite

You have reviewed the information in [Planning a Red Hat Process Automation Manager installation](#) .

CHAPTER 1. RED HAT PROCESS AUTOMATION MANAGER CLUSTERS

By clustering two or more computers, you have the benefits of high availability, enhanced collaboration, and load balancing. High availability decreases the chance of a loss of data when a single computer fails. When a computer fails, another computer fills the gap by providing a copy of the data that was on the failed computer. When the failed computer comes online again, it resumes its place in the cluster. Load balancing shares the computing load across the nodes of the cluster. Doing this improves the overall performance.

There are several ways that you can cluster Red Hat Process Automation Manager components. This document describes how to cluster the following scenarios:

- [Chapter 2, *Red Hat Process Automation Manager clusters in a development \(authoring\) environment*](#)
- [Chapter 3, *Process Server clusters in a runtime environment*](#)

CHAPTER 2. RED HAT PROCESS AUTOMATION MANAGER CLUSTERS IN A DEVELOPMENT (AUTHORING) ENVIRONMENT



NOTE

Configuration of Business Central for high availability is currently technology preview.

Developers use Red Hat Process Automation Manager to author rules and processes that assist users with decision making. The primary benefit of clustering Red Hat Process Automation Manager development environments is high availability. If clustering is configured and Developer X is working on Node1 and that node fails, Developer X's work is preserved and visible on any other node of the cluster.

Most development environments consist of Business Central for creating rules and processes, and at least one Process Server to test those rules and processes.

To create a Red Hat Process Automation Manager clustered development environment, you download and extract Red Hat JBoss EAP 7.2 and Red Hat Process Automation Manager 7.2. Then, you configure Red Hat JBoss EAP 7.2 with Elasticsearch and AMQ Broker, a Java messaging server (JMS) broker. Finally, you start the cluster and install Business Central on each cluster node.

Elasticsearch is a highly scalable open source full-text search and analytics engine. It enables you to store, search, and analyze high volumes of data quickly and in near-real time. In a Red Hat Process Automation Manager clustered environment, it enables you to perform complex and efficient searches across cluster nodes.

A JMS broker is a software component that receives messages, stores them locally, and forwards the messages to a recipient. AMQ Broker enables your applications to communicate with any messaging provider. It specifies how components such as message-driven beans, Enterprise JavaBeans, and servlets can send or receive messages.

2.1. INSTALLING AND CONFIGURING ELASTICSEARCH

To use Elasticsearch for more efficient searching across cluster nodes, you must install and configure Elasticsearch for the Red Hat Process Automation Manager clustered environment. Use the following instructions to configure a simplified, non-high availability environment.

To configure Elasticsearch for a high availability Red Hat Process Automation Manager clustered environment, see the [Elasticsearch documentation](#).

Procedure

1. Download and unzip the [Elasticsearch](#) installation file.
2. Replace the contents of the `elasticsearch/config/elasticsearch.yml` file with the following content, where `<ELASTICSEARCH_NODE_IP>` is the IP address of the system where you unzipped the Elasticsearch file:

```
cluster.name: kie-cluster
transport.host: <ELASTICSEARCH_NODE_IP>
http.host: <ELASTICSEARCH_NODE_IP>
transport.tcp.port: 9300
discovery.zen.minimum_master_nodes: 1
```

3. Run Elasticsearch:

- On Linux or UNIX-based systems, enter:

```
ELASTICSEARCH_HOME/bin/elasticsearch
```

On Windows, enter:

```
ELASTICSEARCH_HOME\bin\elasticsearch.bat
```

2.2. DOWNLOADING AND CONFIGURING AMQ BROKER

AMQ Broker enables your applications to communicate with any messaging provider. It specifies how components such as message-driven beans, Enterprise JavaBeans, and servlets can send or receive messages. Use the following instructions to configure a simplified, non-high availability environment.

To configure AMQ Broker for a high availability Red Hat Process Automation Manager clustered environment, see [Using AMQ Broker](#).

Procedure

1. Navigate to the [Software Downloads](#) page in the Red Hat Customer Portal (login required), and select the product and version from the drop-down options:

- **Product: AMQ Broker**
- **Version: 7.2.0**

2. Click **Download** next to **Red Hat AMQ Broker 7.2.0(amq-broker-7.2.0-bin.zip)**.

3. Extract the **amq-broker-7.2.0-bin.zip** file.

4. Change directory to **amq-broker-7.2.0-bin/amq-broker-7.2.0/bin**.

5. Enter the following command and replace the following placeholders to create the broker and broker user:

- **<HOST>** is the IP address or host name of the server where you installed AMQ Broker.
- **<AMQ_USER>** and **<AMQ_PASSWORD>** is a user name and password combination of your choice.
- **<BROKER_NAME>** is a name for the broker that you are creating.

```
./artemis create --host <HOST> --user <AMQ_USER> --password <AMQ_PASSWORD>
--require-login <BROKER_NAME>
```

6. To run AMQ Broker, enter the following command in the **amq-broker-7.2.0-bin/amq-broker-7.2.0/bin** directory:

```
amq-broker-7.2.0/bin/<BROKER_NAME>/bin/artemis run
```

2.3. DOWNLOADING AND EXTRACTING RED HAT JBOSS EAP 7.2 AND RED HAT PROCESS AUTOMATION MANAGER

Download and install Red Hat JBoss EAP 7.2 and Red Hat Process Automation Manager 7.2 on each node of the cluster.

Procedure

1. Download Red Hat JBoss EAP 7.2 on each node of the cluster:
 - a. Navigate to the [Software Downloads](#) page in the Red Hat Customer Portal (login required), and select the product and version from the drop-down options:
 - **Product: Enterprise Application Platform**
 - **Version: 7.2**
 - b. Click **Download** next to **Red Hat JBoss Enterprise Application Platform 7.2.0 (JBEAP-7.2.0/jboss-eap-7.2.0.zip)**.
2. Extract the **jboss-eap-7.2.0.zip** file. In the following steps, **EAP_HOME** is the **jboss-eap-7.2/jboss-eap-7.2** directory.
3. Download and apply the latest Red Hat JBoss EAP patch, if available.
4. Download Red Hat Process Automation Manager on each node of the cluster:
 - a. Navigate to the [Software Downloads](#) page in the Red Hat Customer Portal, and select the product and version from the drop-down options:
 - **Product: Process Automation Manager**
 - **Version: 7.2**
 - b. Download **Red Hat Process Automation Manager 7.2.0 Business Central Deployable for Red Hat JBoss EAP 7 (rhpam-7.2.0-business-central-eap7-deployable.zip)**.
5. Extract the **rhpam-7.2.0-business-central-eap7-deployable.zip** file to a temporary directory. In the following commands this directory is called **TEMP_DIR**.
6. Open the **TEMP_DIR/rhpam-7.2.0-business-central-eap7-deployable/jboss-eap-7.2/standalone/deployments/business-central.war/WEB-INF/web.xml/** file, uncomment the **<distributed/>** tag, and save the **web.xml** file.
7. Copy the contents of **TEMP_DIR/rhpam-7.2.0-business-central-eap7-deployable/jboss-eap-7.2** to **EAP_HOME**.
8. Download and apply the latest Red Hat Process Automation Manager patch, if available.
9. Navigate to the **EAP_HOME/bin** directory.
10. Create a user with the **admin** role that you will use to log in to Business Central. In the following command, replace **<username>** and **<password>** with the user name and password of your choice.

```
$ ./add-user.sh -a --user <USERNAME> --password <PASSWORD> --role admin,rest-all
```

**NOTE**

Make sure that the specified user name is not the same as an existing user, role, or group. For example, do not create a user with the user name **admin**.

The password must have at least eight characters and must contain at least one number and one non-alphanumeric character, but not & (ampersand).

11. Create a user with the **kie-server** role that you will use to log in to Process Server.

```
$ ./add-user.sh -a --user <USERNAME> --password <PASSWORD> --role kie-server
```

12. Make a note of your user names and passwords.

2.4. CONFIGURING AND RUNNING BUSINESS CENTRAL IN A CLUSTER

After you install Red Hat JBoss EAP and Business Central you can use Elasticsearch and the AMQ Broker to configure the cluster.

**NOTE**

These steps describe a basic cluster configuration. For more complex configurations, see the [Red Hat JBoss EAP 7.2 Configuration Guide](#).

Prerequisites

- Elasticsearch is installed as described in [Section 2.1, “Installing and configuring Elasticsearch”](#).
- AMQ Broker is installed and configured, as described in [Section 2.2, “Downloading and configuring AMQ Broker”](#).
- Red Hat JBoss EAP and Red Hat Process Automation Manager are installed on each node of the cluster as described in [Section 2.3, “Downloading and extracting Red Hat JBoss EAP 7.2 and Red Hat Process Automation Manager”](#).
- An NFS server with a mounted partition is available.

Procedure

1. Open the ***EAP_HOME/standalone/configuration/standalone-full-ha.xml*** file in a text editor.
2. Edit or add the following properties to the **<system-properties>** element and replace the following placeholders:
 - **<AMQ_USER>** and **<AMQ_PASSWORD>** are the credentials that you defined when creating the AMQ Broker.
 - **<AMQ_BROKER_IP_ADDRESS>** is the IP address of the AMQ Broker.
 - **<ELASTICSEARCH_NODE_IP>** is the IP address where Elasticsearch is installed.
 - **<SHARED_NETWORK_FOLDER>** is the network folder that all Business Centrals in the cluster use to exchange data. This folder must be the same on all nodes of the cluster.

```
<system-properties>
```

```

<property name="appformer-jms-connection-mode" value="REMOTE"/>
<property name="appformer-jms-username" value="<AMQ_USER> "/>
<property name="appformer-jms-password" value="<AMQ_USER_PASSWORD>"/>
<property name="appformer-jms-url" value="tcp://<AMQ_BROKER_IP_ADDRESS>:61616?
ha=true&amp;retryInterval=1000&amp;retryIntervalMultiplier=1.0&amp;reconnectAttempts=-
1"/>
<property name="org.appformer.ext.metadata.elastic.port" value="9300"/>
<property name="org.appformer.ext.metadata.elastic.host"
  value="<ELASTICSEARCH_NODE_IP>"/>
<property name="org.appformer.ext.metadata.elastic.cluster"
  value="kie-cluster"/>
<property name="org.appformer.ext.metadata.index" value="elastic"/>
<property name="org.appformer.ext.metadata.elastic.retries" value="10"/>
<property name="org.uberfire.nio.git.dir" value="<SHARED_NETWORK_FOLDER>"/>
<property name="es.set.netty.runtime.available.processors" value="false"/>
</system-properties>

```

3. Save the **standalone-full-ha.xml** file.
4. To start the cluster, navigate to **EAP_HOME/bin** and enter one of the following commands:
 - On Linux or UNIX-based systems:

```
$ ./standalone.sh -c standalone-full-ha.xml
```

- On Windows:

```
standalone.bat -c standalone-full-ha.xml
```

2.5. VERIFYING THE RED HAT PROCESS AUTOMATION MANAGER CLUSTER

After configuring the cluster for Red Hat Process Automation Manager, create an asset to verify that the installation is working.

Procedure

1. In a web browser, enter **<node-IP-address>:8080/business-central**. Replace **<node-IP-address>** with the IP address of a particular node.
2. Enter the **admin** user credentials that you created during installation. The Business Central home page appears.
3. Select **Menu → Design → Projects**.
4. Click **Try Samples → Mortgages → OK**. The **Assets** window appears.
5. Click **Create New Asset → Data Object**.
6. Enter **MyDataObject** in the **Data Object** field and click **OK**.
7. Click **Spaces → MySpace → Mortgages** and confirm that **MyDataObject** is in the list of assets.
8. Enter the following URL in a web browser, where **<node_IP_address>** is the address of a different node of the cluster:

http://<node_IP_address>:8080/business-central

9. Enter the same credentials that you used to log in to Business Central on the first node, where you created the **MyDataObject** asset.
10. Select **Menu** → **Design** → **Projects**.
11. Select the **Mortgages** project.
12. Verify that **MyDataObject** is in the asset list.
13. Delete the **Mortgages** project.

CHAPTER 3. PROCESS SERVER CLUSTERS IN A RUNTIME ENVIRONMENT

In a runtime environment, Process Server runs services that contain rules and processes that support business decisions. The primary benefit of clustering a Process Server runtime environment is load balancing. If activity on one node of the cluster increases, that activity can be shared among the other nodes of the cluster to improve performance.

To create a Process Server clustered runtime environment, you download and extract Red Hat JBoss EAP 7.2 and Process Server. Then, you configure Red Hat JBoss EAP 7.2 for a cluster, start the cluster, and install Process Server on each cluster node.

If desired, you can then cluster the headless Process Automation Manager controller and Smart Router.

3.1. DOWNLOADING AND EXTRACTING RED HAT JBOSS EAP 7.2 AND PROCESS SERVER

Complete the steps in this section to download and install Red Hat JBoss EAP 7.2 and Process Server for installation in a clustered environment.

Procedure

1. Download Red Hat JBoss EAP 7.2 on each node of the cluster:
 - a. Navigate to the [Software Downloads](#) page in the Red Hat Customer Portal (login required) and select the product and version from the drop-down options:
 - **Product: Red Hat JBoss EAP**
 - **Version: 7.2**
 - b. Click **Download** next to **Red Hat JBoss Enterprise Application Platform 7.2.0 (jboss-eap-7.2.0.zip)**.
2. Extract the **jboss-eap-7.2.0.zip** file. The **jboss-eap-7.2/jboss-eap-7.2** directory is referred to as **EAP_HOME**.
3. Download and apply the latest Red Hat JBoss EAP patch, if available.
4. Download Process Server:
 - a. Navigate to the [Software Downloads](#) page in the Red Hat Customer Portal and select the product and version from the drop-down options:
 - **Product: Process Automation Manager**
 - **Version: 7.2**
 - b. Download **Red Hat Process Automation Manager 7.2.0 Process Server for All Supported EE7 Containers (rhpam-7.2.0-kie-server-ee7.zip)**.
5. Extract the **rhpam-7.2.0-kie-server-ee7.zip** archive to a temporary directory. In the following examples this directory is called **TEMP_DIR**.
6. Copy the **TEMP_DIR/rhpam-7.2.0-kie-server-ee7/rhpam-7.2.0-kie-server-ee7/kie-server.war** directory to **EAP_HOME/standalone/deployments/**.



WARNING

Ensure the names of the Red Hat Process Automation Manager deployments you are copying do not conflict with your existing deployments in the Red Hat JBoss EAP instance.

7. Copy the contents of the **TEMP_DIR/rhpam-7.2.0-kie-server-ee7/rhpam-7.2.0-kie-server-ee7/SecurityPolicy/** to **EAP_HOME/bin**. When asked to overwrite files, click **Replace**.
8. In the **EAP_HOME/standalone/deployments/** directory, create an empty file named **kie-server.war.dodeploy**. This file ensures that Process Server is automatically deployed when the server starts.
9. Download and apply the latest Red Hat Process Automation Manager patch, if available.
10. Navigate to the **EAP_HOME/bin** directory.
11. Create a user with the **kie-server** role that you will use to log in to Process Server.

```
$ ./add-user.sh -a --user <USERNAME> --password <PASSWORD> --role kie-server
```

12. Make a note of your user names and passwords.

3.2. CONFIGURING AND RUNNING A RED HAT JBOSS EAP 7.2 CLUSTER FOR PROCESS SERVER

Configure the Red Hat JBoss EAP cluster for Process Server, and then start the cluster.

Procedure

1. Install the JDBC driver on all Red Hat JBoss EAP instances that are part of this cluster. For more information, see the "[JDBC Drivers](#)" section of the *Red Hat JBoss EAP 7.2 Configuration Guide*.
2. Open the **EAP_HOME/standalone/configuration/standalone-full.xml** file in a text editor.
3. Edit the **data-stores** property and the **timer-service thread-pool-name** above it:
 - The **datasource-jndi-name** is the JNDI name of the database specified at the beginning of this procedure.
 - You can enter any name for the value of the **partition** property. However, a node will only see timers from other nodes that have the same partition name. Grouping nodes into partitions by assigning partition names enables you to break a large cluster up into several smaller clusters. Doing this improves performance. For example, instead of having a cluster of 100 nodes, where all 100 nodes are trying to execute and refresh the same timers, you can create 20 clusters of 5 nodes by giving every group of 5 a different partition name.
 - Replace the **default-data-store** attribute value with **ejb_timer_ds**.

- Set the value of **refresh-interval** in milliseconds to specify how often the EJB timer connects to the database to synchronize and load tasks to be processed.

```
<timer-service thread-pool-name="default" default-data-store="ejb_timer_ds">
<data-stores>
  <database-data-store name="ejb_timer_ds" datasource-jndi-
name="java:jboss/datasources/ejb_timer" database="postgresql"
partition="ejb_timer_part" refresh-interval="30000"/>
</data-stores>
</timer-service>
```

The following table lists the supported databases and the corresponding **database** attribute value:

Table 3.1. Supported databases

Database	Attribute value
Hyper SQL (for demonstration purposes only, not supported)	hsq
PostgreSQL	postgresql
Oracle	oracle
IBM DB2	db2
Microsoft SQL Server	mssql
MySQL and MariaDB	mysql

4. Add the Process Server and EJB timer data sources to the **standalone-full.xml** file. In these examples, **<DATABASE>** is the name of the database, **<SERVER_NAME>** is the host name of the JNDI database, and **<USER_NAME>** and **<USER_PWD>** are the credentials for that database.

- Add the data source to allow Process Server to connect to the database, for example:

```
<xa-datasource jndi-name="java:/jboss/datasources/rhpam" pool-name="rhpam-RHPAM"
use-java-context="true" enabled="true">
  <xa-datasource-property name="DatabaseName"><DATABASE></xa-datasource-
property>
  <xa-datasource-property name="PortNumber">5432</xa-datasource-property>
  <xa-datasource-property name="ServerName"><SERVER_NAME></xa-datasource-
property>
  <driver>postgresql</driver>
  <security>
    <user-name><USER_NAME></user-name>
    <password><USER_PWD></password>
  </security>
</xa-datasource>
```

- Add the data source to enable the EJB timer, for example:

```

<xa-datasource jndi-name="java:jboss/datasources/ejb_timer" pool-name="ejb_timer"
use-java-context="true" enabled="true">
  <xa-datasource-property name="DatabaseName"><DATABASE></xa-datasource-
property>
  <xa-datasource-property name="PortNumber">5432</xa-datasource-property>
  <xa-datasource-property name="ServerName"><SERVER_NAME></xa-datasource-
property>
  <driver>postgresql</driver>
  <transaction-isolation>TRANSACTION_READ_COMMITTED</transaction-isolation>
  <security>
    <user-name><USER_NAME></user-name>
    <password><USER_PWD></password>
  </security>
</xa-datasource>

```



WARNING

You must use two different databases for Process Server runtime data and EJB timer data.

5. Add the following properties to the **<system-properties>** element and replace the following placeholders:

- **<JNDI_NAME>** is the JNDI name of your data source. For Red Hat Process Automation Manager, this is **java:/jboss/datasources/rhpam**.
- **<DIALECT>** is the hibernate dialect for your database. The following dialects are supported:
 - DB2: **org.hibernate.dialect.DB2Dialect**
 - MSSQL: **org.hibernate.dialect.SQLServer2012Dialect**
 - MySQL: **org.hibernate.dialect.MySQL5InnoDBDialect**
 - MariaDB: **org.hibernate.dialect.MySQL5InnoDBDialect**
 - Oracle: **org.hibernate.dialect.Oracle10gDialect**
 - PostgreSQL: **org.hibernate.dialect.PostgreSQL82Dialect**
 - PostgreSQL plus: **org.hibernate.dialect.PostgresPlusDialect**

```

<system-properties>
  <property name="org.kie.server.persistence.ds" value="<JNDI_NAME>"/>
  <property name="org.kie.server.persistence.dialect" value="<DIALECT>"/>
  <property name="org.jbpm.ejb.timer.tx" value="true"/>
</system-properties>

```

6. Save the **standalone-full.xml** file.

7. To start the cluster, navigate to **EAP_HOME/bin** and enter one of the following commands:

- On Linux or UNIX-based systems:

```
┆ $ ./standalone.sh -c standalone-full.xml
```

- On Windows:

```
┆ standalone.bat -c standalone-full.xml
```

3.3. CLUSTERING WITH THE HEADLESS PROCESS AUTOMATION MANAGER CONTROLLER

The Process Automation Manager controller is integrated with Business Central. However, if you do not install Business Central, you can install the headless Process Automation Manager controller and use the REST API or the Process Server Java Client API to interact with it.

Prerequisites

- A backed-up Red Hat JBoss EAP installation version 7.2 or later is available. The base directory of the Red Hat JBoss EAP installation is referred to as **EAP_HOME**.
- Sufficient user permissions to complete the installation are granted.
- An NFS server with a mounted partition is available.

Procedure

1. Navigate to the [Software Downloads](#) page in the Red Hat Customer Portal (login required), and select the product and version from the drop-down options:
 - **Product: Process Automation Manager**
 - **Version: 7.2**
2. Download **Red Hat Process Automation Manager 7.2.0 Add Ons**(the **rhpmam-7.2.0-add-ons.zip** file).
3. Unzip the **rhpmam-7.2.0-add-ons.zip** file. The **rhpmam-7.2-controller-ee7.zip** file is in the unzipped directory.
4. Extract the **rhpmam-7.2-controller-ee7** archive to a temporary directory. In the following examples this directory is called **TEMP_DIR**.
5. Copy the **TEMP_DIR/rhpmam-7.2-controller-ee7/controller.war** directory to **EAP_HOME/standalone/deployments/**.



WARNING

Ensure that the names of the headless Process Automation Manager controller deployments you are copying do not conflict with your existing deployments in the Red Hat JBoss EAP instance.

6. Copy the contents of the **`TEMP_DIR/rhpam-7.2-controller-ee7/SecurityPolicy/`** directory to **`EAP_HOME/bin`**. When asked to overwrite files, select **Yes**.
7. In the **`EAP_HOME/standalone/deployments/`** directory, create an empty file named **`controller.war.dodeploy`**. This file ensures that the headless Process Automation Manager controller is automatically deployed when the server starts.
8. Open the **`EAP_HOME/standalone/configuration/standalone.xml`** file in a text editor.
9. Add the following properties to the **`<system-properties>`** element and replace **`<NFS_STORAGE>`** with the absolute path to the NFS storage where the template configuration is stored:

```
<system-properties>
  <property name="org.kie.server.controller.templatefile.watcher.enabled" value="true"/>
  <property name="org.kie.server.controller.templatefile" value="<NFS_STORAGE>"/>
</system-properties>
```

If the value of the **`org.kie.server.controller.templatefile.watcher.enabled`** property is set to true, a separate thread is started to watch for modifications of the template file. The default interval for these checks is 30000 milliseconds and can be further controlled by the **`org.kie.server.controller.templatefile.watcher.interval`** system property. If the value of this property is set to false, changes to the template file are detected only when the server restarts.

10. To start the headless Process Automation Manager controller, navigate to **`EAP_HOME/bin`** and enter the following command:

- On Linux or UNIX-based systems:

```
$ ./standalone.sh
```

- On Windows:

```
standalone.bat
```

3.4. CLUSTERING SMART ROUTER FOR A RUNTIME ENVIRONMENT

Smart Router enables you to aggregate multiple independent Process Server instances as though they are a single server. It performs the role of an intelligent load balancer because it can both route requests to individual Process Servers and aggregate data from different Process Servers. Through aliases, Smart Router is a proxy to Process Servers. Smart Router performs the following tasks:

- Collects information from various server instances in a single client request

- Finds the right server for a specific request
- Aggregates responses from different servers
- Provides efficient load-balancing
- Manages changing environments, for example adding and removing server instances
- Manages registration with the Process Automation Manager controller

This section describes how to install Smart Router and configure it for a Red Hat Process Automation Manager runtime environment.



NOTE

Load balancing requests for Smart router cluster must be managed externally, using standard load balancing tools.

Prerequisites

- Process Server is installed on each node of a Red Hat JBoss EAP 7.2 cluster.

Procedure

1. Navigate to the [Software Downloads](#) page in the Red Hat Customer Portal (login required), and select the product and version from the drop-down options:
 - **Product: Process Automation Manager**
 - **Version: 7.2**
2. Download **Red Hat Process Automation Manager 7.2.0 Add-Ons**
3. Extract the downloaded **rhpam-7.2.0-add-ons.zip** file to a temporary directory. The **rhpam-7.2-smart-router.jar** file is in the extracted **rhpam-7.2.0-add-ons** directory.
4. Copy the **rhpam-7.0-smart-router.jar** file to location where you will run the file.
5. From the temporary directory, enter the following command to start SmartRouter:

```
java
-Dorg.kie.server.router.host=<ROUTER_HOST>
-Dorg.kie.server.router.port=<ROUTER_PORT>
-Dorg.kie.server.controller=<CONTROLLER_URL>
-Dorg.kie.server.controller.user=<CONTROLLER_USER>
-Dorg.kie.server.controller.pwd=<CONTROLLER_PWD>
-Dorg.kie.server.router.config.watcher.enabled=true
-Dorg.kie.server.router.repo=<NFS_STORAGE>
-jar rhpam-7.0-smart-router.jar
```

The properties in the preceding command have the following default values:

```
org.kie.server.router.host=localhost
org.kie.server.router.port=9000
org.kie.server.controller= N/A
org.kie.server.controller.user=kieserver
```

```
org.kie.server.controller.pwd=kieserver1!  
org.kie.server.router.repo= <CURRENT_WORKING_DIR>  
org.kie.server.router.config.watcher.enabled=false
```



NOTE

The router is capable of providing an aggregate sort, however the data returned when used through the management console is in raw format. Therefore, sorting is in whatever way individual servers return.

Paging is supported in its standard format.

6. To use Smart Router from the client side, use the Smart Router URL instead of the Process Server URL, for example:

```
KieServicesConfiguration config =  
KieServicesFactory.newRestConfiguration("http://smartrouter.example.com:9000",  
"USERNAME", "PASSWORD");
```

In this example, **smartrouter.example.com** is the Smart Router URL, and **USERNAME** and **PASSWORD** are the log in credentials for the Smart Router configuration.

CHAPTER 4. ADDITIONAL RESOURCES

- *Installing and configuring Red Hat Process Automation Manager on Red Hat JBoss EAP*
- *Planning a Red Hat Process Automation Manager installation*
- *Deploying a Red Hat Process Automation Manager immutable server environment on Red Hat OpenShift Container Platform*
- *Deploying a Red Hat Process Automation Manager authoring environment on Red Hat OpenShift Container Platform*
- *Deploying a Red Hat Process Automation Manager managed server environment on Red Hat OpenShift Container Platform*

APPENDIX A. VERSIONING INFORMATION

Documentation last updated on Friday, May 22, 2020.