Red Hat JBoss BPM Suite 6.3 Administration and Configuration Guide

The Administration and Configuration Guide for Red Hat JBoss BPM Suite

Red Content Services   Gemma Sheldon   Klara Kufova
Marek Czernek   Tomas Radej   Vidya Iyengar

Red Hat JBoss BPM Suite 6.3 Administration and Configuration Guide

The Administration and Configuration Guide for Red Hat JBoss BPM Suite

Red Content Services

Gemma Sheldon
gsheldon@redhat.com

Klara Kufova
kkufova@redhat.com

Marek Czernek
mczernek@redhat.com

Tomas Radej
tradej@redhat.com

Vidya Iyengar
viyengar@redhat.com

Legal Notice

Abstract

A guide for administrators and advanced users dealing with Red Hat JBoss BPM Suite setup, configuration, and advanced usage.

Table of Contents

PART I. INTRODUCTION

CHAPTER 1. BUSINESS PROCESS MODEL AND NOTATION

Business Process Model and Notation (BPMN) is a standard notation for business process modeling. It aspires to link the gap between business analysts and programmers by providing a workflow language that can be clearly understood by both.

1.1. COMPONENTS

Red Hat JBoss BPM Suite integrates multiple components to support business processes throughout their entire life cycle and to provide process management features and tools for business analysts, developers, and business users. The product can be deployed on various JEE-compliant servers; the recommended option is Red Hat JBoss Enterprise Application Platform 6.

Red Hat JBoss BPM Suite consists of the following main components:

- *Execution Engine*: Provides the runtime environment for Processes and Business Rules. It encompasses a workflow library that can be embedded into a user web application. Runtime manager is the root object and contains the following components:

  - *Runtime Engine*: Implements the core behavior of the computer language and it is provided by the runtime manager.

    - *Process Engine*: The environment for business process model execution.

    - *Task Service*: Handles human task lifecycles.

  - *Rule Engine*: Can be used with the process engine or on its own.

    - *Rules Evaluation*: Executes business rules on the provided set of facts.

    - *Complex Event Processing*: Applies business rules on incoming stream of events.

- *Business Central*: A web-based application that accommodates tooling for asset creation, management, and monitoring by providing an integrated web environment.

  - *Asset Repository*: The central sharing location (Knowledge Store) for business assets, processes, rules, forms, etc. Users access this repository through the Project Explorer view of Business Central via **Authoring → Project Authoring**. By default, the product initializes a local GIT repository as its Asset Repository. However, other repositories may be added or removed as necessary.

  - *Artifact Repository*: A Maven based repository for storage of project jar artifacts.

  - *Execution Server*: Provides an execution environment for business process instances and tasks.

  - *Business Activity Monitor*: Provides customizable view on business performance.

> **Note**
>
> Red Hat JBoss BRMS comes with its own Business Central application that is a subset of the Business Central application in Red Hat JBoss BPM Suite.

1.2. PROJECT

A project is a container for asset packages (business processes, rules, work definitions, decision tables, fact models, data models, and DSLs) that lives in the Knowledge Repository. It is this container that defines the properties of the KIE Base and KIE Session that are applied to its content. In the GUI, you can edit these entities in the Project Editor.

As a project is a Maven project, it contains the Project Object Model file (`pom.xml`) with information on how to build the output artifact. It also contains the Module Descriptor file, `kmodule.xml`, that contains the KIE Base and KIE Session configuration for the assets in the project.

1.3. CREATING A PROJECT

It is possible to create a project either in the *Project Authoring* perspective of Business Central or using the REST API calls.

Creating a Project in Business Central

Important

Note that only users with the `admin` role in Business Central can create projects.

Procedure: Using Business Central to Create a Project

1. In Business Central, go to **Authoring → Project Authoring**.

2. In the **Project Explorer**, select the organizational unit and the repository in which you want to create the project.

3. On the perspective menu, click **New Item → Project**.

   The **New Project** dialog window opens.

4. Define the *Project General Settings* and *Group artifact version* details of the new project. These parameters are stored in the `pom.xml` Maven configuration file.

   See the detailed description of the parameters:

   » **Project Name**: name of the project (for example `MortgageProject`).

   » **Project Description**: description of the project, which may be useful for the project documentation purposes.

   » **Group ID**: group ID of the project (for example `org.mycompany.commons`).

   » **Artifact ID**: artifact ID unique in the group (for example `myframework`). Avoid using a space or any other special character that might lead to an invalid name.

   » **Version**: version of the project (for example `2.1.1`).

5. Click **Finish**.

   The project screen view is updated with the new project details as defined in the `pom.xml` file. You can switch between project descriptor files and edit their content by clicking the **Project Settings: Project General Settings** button at the top of the project screen view.

Creating a Project Using the REST API

Important

Note that only users with the `rest-all` or `rest-project` role can create projects.

To create a project in the repository, issue the **POST** REST API call. Details of the project are defined by the corresponding JSON entity.

Input parameter of the call is an **Entity** instance. The call returns a **CreateProjectRequest** instance.

> **Example 1.1. Creating a Project Using the Curl Utility**
>
> Example JSON entity containing details of a project to be created:
>
> ```
> {
>   "name"        : "MortgageProject",
>   "description" :  null,
>   "groupId"     : "org.mycompany.commons",
>   "version"     : "2.1.1"
> }
> ```
>
> Execute the following command:
>
> ```
> curl -X POST 'localhost:8080/business-
> central/rest/repositories/REPOSITORY_NAME/projects/' -u
> USERNAME:PASSWORD -H 'Accept: application/json' -H 'Content-Type:
> application/json' -d
> '{"name":"MortgageProject","description":null,"groupId":"org.myco
> mpany.commons","version":"2.1.1"}'
> ```

For further information, refer to chapter *Knowledge Store REST API*, section *Repository Calls* of the *Red Hat JBoss BPM Suite Development Guide*.

1.4. ADDING DEPENDENCIES

To add dependencies to your project, do the following:

1. Open the Project Editor for the given project:

   a. In the **Project Explorer** view of the **Project Authoring** perspective, open the project directory.

   b. Click **Open Project Editor** to open the project view.

2. In the **Project Screen** view, select in the **Project Settings** drop-down box the **Dependencies** item.

3. On the updated **Project Screen**, click the **Add** button to add a maven dependency or click the **Add from repository** button to add a dependency from the Knowledge Store (Artifact repository):

   a. When adding a maven dependency, a user has to define the **Group ID**, **Artifact ID** and the **Version ID** in the **Dependency** dialogue window.

b. When adding a dependency from the Knowledge Store, select the dependency in the displayed dialog box: the dependency will be added to the dependency table.

4. To apply the various changes, the dependencies must be saved.

Additionally, you can use the **Package white list** when working with dependencies. When you add a repository, you can click the gear icon and select **Add all** or **Add none**, which results in including all or none of the packages from the added dependency.

> **Warning**
>
> If working with modified artifacts, do not re-upload modified non-snapshot artifacts as Maven will not know these artifacts have been updated, and it will not work if it is deployed in this manner.

PART II. CONFIGURATION

CHAPTER 2. BUSINESS CENTRAL CONFIGURATION

As Business Central is a web application, any configuration settings are loaded from **DEPLOY_DIRECTORY/business-central.war/WEB-INF/web.xml** and the referenced files, and if deployed on Red Hat JBoss EAP 6, also in **jboss-web.xml** and **jboss-deployment-structure.xml**.

Note that the entire application can be run in different profiles (refer to the *Red Hat JBoss BPM Suite Installation Guide*).

2.1. ACCESS CONTROL

The access control mechanism includes authorization and authentication. In the unified environment of Red Hat JBoss BPM Suite, users are able to update the default user roles located within **JBOSS_HOME/standalone/deployments/business-central.war/WEB-INF/classes/userinfo.properties**.

To grant a user access to JBoss BPM Suite, the user needs to have the respective role assigned:

➤ **admin**: Administrates JBoss BPM Suite system and has full access rights to make any changes necessary including the ability to add and remove users from the system.

➤ **developer**: Implements code required for processes to work and has access to everything except administration tasks.

➤ **analyst**: Creates and designs processes and forms, instantiates the processes and deploys artifacts. This role is the similar to a developer, without access to asset repository and deployments.

➤ **user**: Claims, performs, and invokes other actions (such as, escalation, rejection, etc.) on the assigned Tasks and has no access to authoring functions.

➤ **manager**: Monitors the system and its statistics and only has access to the dashboard.

➤ **business user**: Takes action on business tasks that are required for processes to continue forward. Works primarily with the task list.

If using Red Hat JBoss EAP, to create a user with particular roles, run the **JBOSS_HOME/add-user.sh** script and create an Application User in the **ApplicationRealm** with the respectives roles.

Workbench Configuration

Within Red Hat JBoss BPM Suite, users may set up roles using LDAP to modify existing roles. Users may modify the roles in the workbench configuration to ensure the unique LDAP based roles conform to enterprise standards by editing the deployments directory located at **JBOSS_HOME/standalone/deployments/business-central.war/WEB-INF/classes/workbench-policy.propeties**.

If authenticating user via LDAP over Git, administrators must set system property **org.uberfire.domain** to the name of login module it should use to authenticate users via the Git service. This must be set in the **standalone.xml** file in EAP.

Note

You can further customize Business Central with parameters *no_build* or *no_search*. The parameters disable the build and search functionality. Include one or both parameters in the Business Central URL, for example `http://SERVER:PORT/business-central/kie-wb.html?no_build&no_search`.

Authentication in Human Tasks

Every Task that needs to be executed is assigned to one or multiple roles or groups, so that any user with the given role or the given group assigned can claim the Task instance and execute it. Tasks can also be assigned to one or multiple users directly. JBoss BPM Suite uses the `UserGroupCallback` interface to assign tasks to user.

Warning

A group for a Human Task must not be named after an existing user of the system. Doing so causes intermittent issues.

2.2. BUSINESS CENTRAL PROFILE CONFIGURATION

Red Hat JBoss BPM Suite 6 (or better) server is capable of starting the Business Central application in three different modes:

» *Full profile*: Default profile that is active without additional configuration required (UI and remote services e.g. REST).

» *Execution server profile*: Disables completely UI components of the application and allows only remote access e.g. via REST interface.

» *UI server profile*: Disables remote services e.g REST and allows only UI access to the application.

To change the profile use the following configuration steps.

Procedure: Configuring Business Central Profiles

1. Select the desired `web.xml` inside `BPMS_HOME/standalone/deployments/business-central.war/WEB-INF/`. The following files are provided:

   » `web.xml` (default) for full profile

   » `web-exec-server.xml` for execution server profile

   » `web-ui-server.xml` for UI server profile

2. To activate a profile other than the default full profile, the web-<PROFILE>.xml file must be renamed to `web.xml`. The following steps demonstrate one way to enable the execution server profile:

a. Backup the **web.xml** file from the full profile

```
$ mv web.xml web-full.xml
```

b. Rename the **web-exec-server.xml** file:

```
$ mv web-exec-server.xml web.xml
```

3. Start application server with additional system property to instruct the profile manager to activate given profile.

   » **Dorg.kie.active.profile=full** - To activate full profile or skip the property completely

   » **Dorg.kie.active.profile=exec-server** - To activate execution server profile

   » **Dorg.kie.active.profile=ui-server** - To activate UI server profile

2.3. BRANDING THE BUSINESS CENTRAL APPLICATION

The Business Central web application enables you to customize its look and feel by allowing you to override some of its default styles. The ability to customize the Business Central branding allows you to get a consistent appearance across all your applications thereby improving the user experience. It also helps in cases when multiple teams are using the application. Each team can develop their own customized user interface. The customizable elements are built using cascading style sheets (CSS), images, and HTML files, providing an easy and flexible approach to customize without having to recompile the code.

You can modify the following elements in the Business Central application to make it inline with your company's brand:

Login screen

You can customize the following attributes of the Business Central login screen:

» The background image

» The company logo

» The application logo

Application header

You can customize the following attributes of the Business Central application header:

» The Business Central header containing the title and banner logo

Help pop-up windows

You can customize the following attributes of the splash help pop-up windows:

» The splash help images

» The label text

2.3.1. Customizing Business Central Login Page

Procedure: Changing the Business Central Login Page Background Image

1. Start the EAP server and open http://*localhost*:8080/business-central in a web browser.

2. Copy the new background image to the **_EAP_HOME_/standalone/deployments/business-central.war/images** directory in your Red Hat JBoss BPM Suite installation.

3. Navigate to **_EAP_HOME_/standalone/deployments/business-central.war/styles** directory and open the **login-screen.css** file in a text editor.

4. In the **login-screen.css** file, provide the location of your new background image in the following **background-image** attribute:

   ```
   background-image: url("../images/login-screen-
   background.jpg");
   ```

   The **background-image** attribute points to the default **login-screen-background.jpg** image.

   In addition to the background image, you can modify other attributes such as image size, position, and background color in the **login-screen.css** file.

Refresh the Business Central login page to view your changes.

Procedure: Changing the Business Central Login Page Company Logo and Project Logo

1. Start the EAP server and open http://*localhost*:8080/business-central in a web browser.

2. Navigate to the **_EAP_HOME_/standalone/deployments/business-central.war/images** directory in your Red Hat JBoss BPM Suite installation.

3. Replace the default image **login-screen-logo.png** with a new one. This is the company logo that appears on the top right hand corner of the login page.

4. Replace the default image **RH_JBoss_BPMS_Logo.png** with a new one.

   This is the project logo that appears on the center left hand side of the login page.

Refresh the Business Central login page to view your changes.

2.3.2. Customizing Business Central Application Header

Procedure: Changing the Business Central Application Header (Banner)

1. Start the EAP server and open http://*localhost*:8080/business-central in a web browser.

2. Log in to the Business Central application with your user credentials.

3. Copy your new application header image to the **EAP_HOME/standalone/deployments/business-central.war/banner** directory in your Red Hat JBoss BPM Suite installation.

4. Open **EAP_HOME/standalone/deployments/business-central.war/banner/banner.html** file in a text editor.

5. In the **banner.html** file, edit the following **<img>** tag to provide the name of your new header image:

```
<img src="banner/logo.png"/>
```

The default image is **logo.png**.

Refresh the Business Central Home page to view your changes.

2.3.3. Customizing Business Central Splash Help Windows

The **EAP_HOME/standalone/deployments/business-central.war/plugins** directory contains the splash pages and the corresponding html files. Each splash page holds the name of the html file, which contains information about the image(s) and the text to be displayed. For example, the **authoring_perspective.splash.js** splash page points to the **authoring_perspective.splash.html** file. The **authoring_perspective.splash.html** contains the names and location of all the image files that appear on the Authoring Perspective splash help and also their captions. You can customize the images and the corresponding captions of the existing splash help pop-up windows.

Procedure: Changing the Business Central Splash Help Pop-Up Images and Captions

1. Start the EAP server and open http://*localhost*:8080/business-central in a web browser.

2. Log in to the Business Central application with your user credentials.

3. Copy your new splash help image(s) to the **EAP_HOME/standalone/deployments/business-central.war/images** directory in your Red Hat JBoss BPM Suite installation.

4. Open the corresponding html file from **EAP_HOME/standalone/deployments/business-central.war/plugins** directory in a text editor.

5. Edit the html file to point to your new splash help image. For example, to change the first image that appears in the Authoring Perspective splash help, edit the following **<img>** tag in the **authoring_perspective.splash.html** file to add your new image:

```
<img src="images/authoring_perspective1.png" alt="">
```

The default image is **authoring_perspective1.png**, which appears on the first page of the Authoring Perspective splash help.

6. To change the image caption that appears on the splash help, edit the **<h4>** and **<p>** tag contents below the **<img>** tag:

```
<h4>Authoring</h4>
<p>Modularized and customizable workbench</p>
```

Refresh the Business Central Home page and access the splash help pop-up windows to view your changes.

2.4. DEPLOYMENT DESCRIPTORS

Processes and rules within Red Hat JBoss BPM Suite 6 onwards are stored in Apache Maven based packaging, and are known as knowledge archives or KJAR. The rules, processes, assets, etc. are part of a jar file built and managed by Maven. A file kept inside the **META-INF** directory of the KJAR called **kmodule.xml** can be used to define the knowledge bases and sessions. This **kmodule.xml** file, by default, is empty.

Whenever a runtime component such as Business Central is about to process the KJAR, it looks up **kmodule.xml** to build the runtime representation.

Deployment Descriptors, a new feature introduced in the 6.1 branch of Red Hat JBoss BPM Suite, allows you fine grained control over your deployment and supplements the **kmodule.xml** file. The presence of these descriptors is optional and your deployment will proceed successfully without them. The properties that you can set using these descriptors are purely technical in nature and include meta values like persistence, auditing and runtime strategy.

These descriptors allow you to configure the execution server on multiple levels (server level default, different deployment descriptor per KJAR and so on). This allows you to make simple customizations to the execution server's out-of-the-box configuration (possibly per KJAR).

You define these descriptors in a file called **kie-deployment-descriptor.xml** and place this file next to your **kmodule.xml** file in the **META-INF** folder. You can change this default location (and the filename) by specifying it as a system parameter:

```
-Dorg.kie.deployment.desc.location=file:/path/to/file/company-
deployment-descriptor.xml
```

2.4.1. Deployment Descriptor Configuration

Deployment descriptors allow the user to configure the execution server on multiple levels:

- *Server level*: The main level and the one that applies to all KJARs deployed on the server.

- *Kjar level*: This allows you to configure descriptors on a per KJAR basis.

- *Deploy time level*: Descriptors that apply while a KJAR is being deployed.

The granular configuration items specified by the deployment descriptors take precedence over the server level ones, except in case of configuration items that are collection based, which are merged. The hierarchy works like this: *deploy time configuration > KJAR configuration > server configuration*.

Note

The deploy time configuration applies to deployments done via the REST API.

For example, if the persistence mode (one of the items you can configure) defined at the server level is **NONE** but the same mode is specified as **JPA** at the KJAR level, the actual mode will be **JPA** for that KJAR. If nothing is specified for the persistence mode in the deployment descriptor for that KJAR (or if there is no deployment descriptor), it will fall back to the server level configuration, which in this case is **NONE** (or to **JPA** if there is no server level deployment descriptor).

Can You Override this Hierarchal Merge Mode Behavior?

Yes. In the default way, if there are deployment descriptors present at multiple levels, the configuration properties are merged with the granular ones overriding the coarse values, and with missing configuration items at the granular level being supplied with those values from the higher levels. The end result is a merged Deployment Descriptor configuration. This default merge mode is called the **MERGE_COLLECTIONS** mode. However, you can change it (see Section 2.4.2, "Managing Deployment Descriptors") if it does not suit your environment to one of the following modes:

» **KEEP_ALL**: In this mode, all higher level values override all lower level values (server level values replace KJAR level values)

» **OVERRIDE_ALL**: In this mode, all lower level values override all higher level values (KJAR values replace server level values)

» **OVERRIDE_EMPTY**: In this mode, all *non empty* configuration items from lower levels replace those at higher levels, including items that are represented as collections.

» **MERGE_COLLECTIONS (DEFAULT)**: In this mode, all non empty configuration items from lower level replace those from higher levels (like in **OVERRIDE_EMPTY**), but collection properties are merged (combined).

Deployment Descriptors from dependent KJARs are placed lower than the actual KJAR being deployed, but they still have higher hierarchy than the server level.

Do I Need to Provide a Full Deployment Descriptor for All Kjars?

No, and this is where the beauty of the merge between different files can help you. Providing partial Deployment Descriptors is possible and recommended. For example, if you want to only override the audit mode in a KJAR, then you just need to provide that and the rest of the values will be merged from server level or higher level KJARs.

It is worth noting that when using **OVERRIDE_ALL** merge mode, all configuration items should be specified since the relevant KJAR will always use them and will not merge with any other deployment descriptor in the hierarchy.

What Can You Configure?

High level technical configuration details can be configured via deployment descriptors. The following table lists these along with the permissible and default values for each.

Table 2.1. Deployment Descriptors

Configuration	XML Entry	Permissible Values	Default Value
Persistence unit name for runtime data	persistence-unit	Any valid persistence package name	org.jbpm.domain
Persistence unit name for audit data	audit-persistence-unit	Any valid persistence package name	org.jbpm.domain
Persistence mode	persistence-mode	JPA, NONE	JPA
Audit mode	audit-mode	JPA, JMS or NONE	JPA
Runtime Strategy	runtime-strategy	SINGLETON, PER_REQUEST or PER_PROCESS_INSTANCE	SINGLETON
List of Event Listeners to be registered	event-listeners	Valid listener class names as **ObjectModel**	No default value
List of Task Event Listeners to be registered	task-event-listeners	Valid listener class names as **ObjectModel**	No default value
List of Work Item Handlers to be registered	work-item-handlers	Valid Work Item Handler classes given as **NamedObjectHandler**	No default value
List of Globals to be registered	globals	Valid Global variables given as **NamedObjectModel**	No default value

Configuration	XML Entry	Permissible Values	Default Value
Marshalling strategies to be registered (for pluggable variable persistence)	marshalling-strategies	Valid **ObjectModel** classes	No default value
Required Roles to be granted access to the resources of the KJAR	required-roles	String role names	No default value
Additional Environment Entries for Knowledge Session	environment-entries	Valid **NamedObjectModel**	No default value
Additional configuration options of Knowledge Session	configurations	Valid **NamedObjectModel**	No default value

How Do You Provide Values For Collections-Based Configuration Items?

In the table of valid configuration items earlier, you would have noticed that the valid values for the collection based items are either **ObjectModel** or **NamedObjectModel**. Both are similar and provide a definition of the object to be built or created at runtime, with the exception that the **NamedObjectModel** object details name the object to be looked. Both these types are defined using an identifier, optional parameters and resolver (to resolve the object).

Identifier

Defines all the information about the object, such as fully qualified class name, Spring bean id or an MVEL expression.

Parameters

Optional parameters that should be used while creating instances of objects from this model.

Resolver

Identifier of the resolver that will be used to create object instances from the model, that is reflection, mvel, or Spring.

As an example, if you have built a custom marshaling strategy and want your deployments

to use that strategy instead of the default, you will need to provide that strategy as an **ObjectModel**, with the identifier being **com.mycompany.MyStrategy**, resolver being reflection (the easiest and the default) and any parameters that are required for your strategy to work. Reflection will then be used to create an instance of this strategy using the fully qualified class name that you have provided as the identifier.

```
<marshalling-strategy>
 <resolver>reflection</resolver>
 <identifier>com.myCompany.MyStrategy</identifier>
 <parameters>
    <parameter xsi:type="xs:string"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
       param
    </parameter>
  </parameters>
</marshalling-strategy>
```

In the case that reflection based on resolver is not enough (as demonstrated in the previous example), you can use a resolver based on MVEL expression as the identifier of the object model. While evaluating expressions, you can substitute out-of-the-box parameters. For example:

```
<marshalling-strategy>
   <resolver>mvel</resolver>
   <identifier>new com.myCompany.CustomStrategy(runtimeManager)
</identifier>
</marshalling-strategy>
```

The Spring based resolver allows you to look up a bean by its identifier from a Spring application context. Whenever JBoss BPM Suite is used with Spring, this resolver helps in deploying KJARs into the runtime. As an example (note that the identifier in this case is a named bean in the Spring context):

```
<marshalling-strategy>
 <resolver>spring</resolver>
 <identifier>customStrategy</identifier>
</marshalling-strategy>
```

2.4.2. Managing Deployment Descriptors

Deployment Descriptors can be edited via the Business Central in one of two ways. Either graphically (by clicking on **Authoring → Project Authoring → Deployment Descriptor** or by clicking on **Authoring → Administration** menu and then clicking through to the **META-INF** folder in the File Explorer. Click on the **kie-deployment-descriptor.xml** file to edit it manually.

Every time a project is created, a stock **kie-deployment-descriptor.xml** file is generated with default values as described earlier.

Overriding Hierarchical Merge Mode Behavior

To change the default mode of **MERGE_COLLECTIONS** to one of **KEEP_ALL**, **OVERRIDE_ALL**, or **OVERRIDE_EMPTY**, you can use the following methods, depending on the requirement.

- Set the system property **org.kie.dd.mergemode** to one of these values. This merge mode will become default for all KJARs deployed in the system, unless you override it at a KJAR level via the next method.

- When deploying a new deployment unit via Business Central (**Deploy → Deployments**) you can select what merge mode should be used for that particular KJAR.

- When deploying via the REST API, you can add **mergemode** query parameter to the command URL to one of these modes to set the merge mode for that deployment.

Restricting Access to the Runtime Engine

One of the configuration items discussed earlier, **required-roles**, can be edited via the Deployment Descriptors. This property restricts access to the runtime engine on a per KJAR or per server level by ensuring that access to certain processes is only granted to users that belong to groups defined by this property.

The security role can be used to restrict access to process definitions or restrict access at runtime.

The default behavior is to add required roles to this property based on repository restrictions. You can of course, edit these properties manually if required, as described above by providing roles that match actual roles defined in the security realm.

2.5. MANAGING DEPLOYMENT OVERRIDE POLICY

If a user tries to deploy an artifact with a GAV (Group-Id, Artifact-Id and Version) that already exists in the system, the deployment will fail and an error message will be displayed in the **Messages** panel.

This feature prevents the user from overwriting an existing deployment by mistake.

By default this feature is enabled, that is, by default the system will prevent the user from overwriting an existing installation with the same GAV.

However, there may be cases when the user *may* want to overwrite existing deployments with the same GAV. Although you cannot enable overwriting on a per-deployment basis, you can set this up for the system as a whole by using the system setting **org.kie.override.deploy.enabled**. This setting, is **false** by default. Change it to **true** to enable overwriting of deployments with the same GAV by providing it at startup time of your server (**-Dorg.kie.override.deploy.enabled=true**).

2.6. EXTENDING BUSINESS CENTRAL

Starting with version 6.1 of Red Hat JBoss BPM Suite, Business Central can be configured to add new screens, menus, editors, splashscreens and perspectives by the Administrator. These elements can extend functionality of Business Central and can be accessed through the **Extensions → Plugin Management**.

You can now define your own Javascript and HTML based plugins to extend Business Central and add them without having to worry about copying files in the underlying filesystem. Let us add a new screen in the system to show you the basics of this functionality.

2.6.1. Plugin Management

You access the **Plugin Management** screen by clicking on **Extensions → Plugin Management**. This brings up the *Plugin Explorer* screen that lists all the existing plugins under their respective categories:

- *Perspective Plugin*

- *Screen Plugin*

- *Editor Plugin*

- *Splashscreen Plugin*

- and *Dynamic Menu*

Open any of these, and you will see the existing plugins in each category, including the uneditable system generated ones.

Let us create a new plugin that echoes "Hello World" when users visit the screen for that plugin. In general, the steps to creating a new plugin are:

1. Create a new screen

2. Create a new perspective (and add the new screen to it)

3. Create a new menu (and add the new perspective to it)

4. Apps (optional)

Adding a New Screen

Click the [⌄] button and select **New Screen**. You will be prompted to enter the name of this new screen. Enter "HelloWorldJS" and press the **OK** button. The Screen plugin editor will open, divided into 4 sections: **Template**, **CSS**, **JavaScript** and **Media**.

> **Note**
>
> All manually created elements go into their respective categories in case you want to edit them later. In this case, to open the Screen plugin editor again if you close it, open the **Screen Plugin** category and scroll past the system generated screens to your manually created plugin and click on it to open the Screen plugin editor again.

Template is where your HTML goes, CSS is for styling, JavaScript is for your functions and Media is for uploading and managing images.

Since we are making a simple Hello World plugin, enter the following code in the Template section: `<div>My Hello World Screen</div>`. This can be any HTML code, and you can use the supplied **Angular** and **Knockout** frameworks. For the purposes of this example, we are not using any of those frameworks, but you can choose to by selecting them from the drop down in the Template section.

Enter your JavaScript code in the JavaScript section. Some common methods and properties are defined for you, including **main**, **on_close** and **on_open**. For this demo, select the **on_open** and enter the following: `function () { alert('Hello World'); }`

Click the **Save** button to finish creating the screen. After you save the screen, refresh business central so that the Screen Plugin is listed in the Screen Component of Perspective plugin.

Adding New Perspective

Once a screen has been created, you need to create a perspective on which this screen will reside. Perspectives can also be created similar to the way a screen is created by clicking on the New button and then selecting **New Perspective**. You can now provide a name for this perspective, say **HelloWorldPerspective**. This will open the Perspective plugin editor, similar to the Screen plugin editor.

The Perspective Editor is like a drag and drop grid builder for screens and HTML components. Remove any existing grids and then drag a 6×6 grid on the right hand side to the left hand side.

Next, open the **Components** category and drag a Screen Component on the right hand side to the left hand side (in any grid). This will open the **Edit Component** dialog box that allows you to select the screen created in the previous step (**HelloWorldJS**). Click the **OK** button and then click **Save** to save this perspective. To tag your perspective, enter **Home** in the tag name field and click **Tags**. Click **OK** and save the changes.

You can open this perspective again from the Perspective plugins listed on the left hand side.

Adding New Menu

The final step in creating our plugin is to add a dynamic menu from where the new screen/perspective can be called up. To do so, go to **Extensions → Plugin Management** and then click on the *New* button to select *New Dynamic Menu*. Give this dynamic menu a name (**HelloWorldMenu**) and then click the **OK** button. The dynamic menu editor opens up.

Enter the perspective name (**HelloWorldPerspective**) as the **Activity Id** and the name for the drop down menu (**HelloWorldMenuDropDown**). Click **OK** and then **Save**.

This new menu will be added to your workbench the next time you refresh Business Central. Refresh it now to see **HelloWorldMenu** added to your top level menu. Click on it to reveal **HelloWorldMenuDropDown**, which when clicked will open your perspective/screen with the message **Hello World**.

You have created your first Plugin!

Working with Apps (Optional)

If you create multiple plugins, you can use the Apps directory feature to organize your own components and plugins, instead of having to rely on just the top menu entries.

When you save a new perspective, you can add labels (tags) for them and these labels (tags) are used to associate a perspective with an App directory. You can open the App directories by clicking on **Extensions → Apps**.

The Apps directory provides an alternate way to open your perspective. When you created your **HelloWorldPerspective**, you entered the tag **Home**. The Apps directory by default contains a single directory called **Home** with which you associated your perspective. This is where you will find it when you open the Apps directory. You can click on it to run the perspective now.

You can create multiple directories and associate perspectives with those directories depending on functional and vertical business requirements. For example, you could create an HR directory and then associate all HR related perspectives with that directory to better manage Apps.

You can create a new directory by clicking the ⊞ button.

2.6.2. The JavaScript (JS) API for Extensions

The extensibility of Business Central is achieved by an underlying JavaScript (JS) API which is automatically loaded if it is placed in the **plugins** folder of the Business Central webapp (typically: *INSTALL_DIR***/business-central.war/plugins/**), or it can be loaded via regular JavaScript calls.

This API is divided into multiple sets depending on the functionality it performs.

Register Perspective API

Allows for the dynamic creation of perspectives. The example below creates a panel using the **registerPerspective** method:

```
$registerPerspective({
    id: "Home",
    is_default: true,
    panel_type:
"org.uberfire.client.workbench.panels.impl.MultiListWorkbenchP
anelPresenter",
    view: {
        parts: [
            {
                place: "welcome",
                min_height: 100,
                parameters: {}
            }
        ],
        panels: [
            {
                width: 250,
                min_width: 200,
                position: "west",
                panel_type:
"org.uberfire.client.workbench.panels.impl.MultiListWorkbenchP
anelPresenter",
                parts: [
                    {
                        place: "YouTubeVideos",
                        parameters: {}
                    }
                ]
            },
            {
                position: "east",
                panel_type:
"org.uberfire.client.workbench.panels.impl.MultiListWorkbenchP
anelPresenter",
```

```
                        parts: [
                            {
                                place: "TodoListScreen",
                                parameters: {}
                            }
                        ]
                },
                {
                        height: 400,
                        position: "south",
                        panel_type:
"org.uberfire.client.workbench.panels.impl.MultiTabWorkbenchPa
nelPresenter",
                        parts: [
                            {
                                place: "YouTubeScreen",
                                parameters: {}
                            }
                        ]
                }
            ]
        }
});
```

Editor API

Allows you to dynamically create editors and associate them with a file type. The
example below creates a sample editor and associates it with **filename** file type.

```
$registerEditor({
    "id": "sample editor",
    "type": "editor",
    "templateUrl": "editor.html",
    "resourceType":
"org.uberfire.client.workbench.type.AnyResourceType",
    "on_concurrent_update":function(){
        alert('on_concurrent_update callback')

$vfs_readAllString(document.getElementById('filename').innerHT
ML, function(a) {
            document.getElementById('editor').value= a;
        });
    },
    "on_startup": function (uri) {
        $vfs_readAllString(uri, function(a) {
            alert('sample on_startup callback')
        });
    },
    "on_open":function(uri){
        $vfs_readAllString(uri, function(a) {
            document.getElementById('editor').value=a;
        });
        document.getElementById('filename').innerHTML = uri;
    }
});
```

In addition to **on_startup** and **on_open** methods seen in the previous example, the API exposes the following callback events for managing the editor's lifecycle:

- **on_concurrent_update;**

- **on_concurrent_delete;**

- **on_concurrent_rename;**

- **on_concurrent_copy;**

- **on_rename;**

- **on_delete;**

- **on_copy;**

- **on_update;**

- **on_open;**

- **on_close;**

- **on_focus;**

- **on_lost_focus;**

- **on_may_close;**

- **on_startup;**

- **on_shutdown;**

You can display this editor via an HTML template:

```
<div id="sampleEditor">
    <p>Sample JS editor (generated by editor-sample.js)</p>
    <textarea id="editor"></textarea>

    <p>Current file:</p><span id="filename"></span>
    <button id="save" type="button"
onclick="$vfs_write(document.getElementById('filename').innerH
TML, document.getElementById('editor').value,  function(a)
{});">Save</button>
    <br>

    <p>This button change the file content, and uberfire send
a callback to the editor:</p>
    <button id="reset" type="button"
onclick="$vfs_write(document.getElementById('filename').innerH
TML, 'Something else',  function(a) {});">Reset File</button>
</div>
```

PlaceManager API

The methods of this API allow you to request that the Business Central display a particular component associated with a target:
$goToPlace("componentIdentifier");

Register plugin API

The methods of this API allow you to create dynamic plugins (that will be transformed in Business Central screens) via the JS API.

```
$registerPlugin( {
    id: "my_angular_js",
    type: "angularjs",
    templateUrl: "angular.sample.html",
    title: function () {
        return "angular " + Math.floor(Math.random() * 10);
    },
    on_close: function () {
        alert("this is a pure JS alert!");
    }
});
```

The plugin references the **angular.sample.html** template:

```
<div ng-controller="TodoCtrl">
    <span>{{remaining()}} of {{todos.length}} remaining</span>
    [ <a href="" ng-click="archive()">archive</a> ]
    <ul class="unstyled">
        <li ng-repeat="todo in todos">
            <input type="checkbox" ng-model="todo.done">
            <span class="done-
{{todo.done}}">{{todo.text}}</span>
        </li>
    </ul>
    <form ng-submit="addTodo()">
        <input type="text" ng-model="todoText" size="30"
placeholder="add new todo here">
        <input class="btn-primary" type="submit" value="add">
    </form>
    <form ng-submit="goto()">
        <input type="text" ng-model="placeText" size="30"
placeholder="place to go">
        <input class="btn-primary" type="submit" value="goTo">
    </form>
</div>
```

A plugin can be hooked to Business Central events via a series of JavaScript callbacks:

- **on_concurrent_update;**

- **on_concurrent_delete;**

- **on_concurrent_rename;**

- **on_concurrent_copy;**

- **on_rename;**

- **on_delete;**

- **on_copy;**

- **on_update;**

- **on_open;**

- **on_close;**

- **on_focus;**

- **on_lost_focus;**

- **on_may_close;**

- **on_startup;**

- **on_shutdown;**

Register splash screens API

use the methods in this API to create splash screens.

```
$registerSplashScreen({
    id: "home.splash",
    templateUrl: "home.splash.html",
    body_height: 325,
    title: function () {
        return "Cool Home Splash " + Math.floor(Math.random() *
10);
    },
    display_next_time: true,
    interception_points: ["Home"]
});
```

Virtual File System (VFS) API

with this API, you can read and write a file saved in the file system using an asynchronous call.

```
$vfs_readAllString(uri,  function(a) {
  //callback logic
});

$vfs_write(uri,content,  function(a) {
  //callback logic
})
```

2.7. CONFIGURING TABLE COLUMNS

Business Central allows you to configure views that contain lists of items in the form of tables. You can resize columns, move columns, add or remove the default list of columns and sort the columns. This functionality is provided for all views that contain tables.

Once you make changes to the columns of a table view, these changes are persisted for the current logged in user.

Adding and Removing Columns

Tables that allow columns to be configured have the [≡] button in the top right corner. Clicking on this button opens up the list of columns that can added or removed to the current table with a check box next to each column:

Resizing Columns

To resize columns, place your cursor between the edges of the column header and move in the direction that you want:

Moving Columns

To re-order and drag and drop a column in a different position, hover your mouse over the rightmost area of the column header:

You can now grab the column and move it:

Drop it over the column header that you want to move it to.

Sorting Columns

To sort columns, click on the desired column's header. To reverse-sort, click on the header again.

CHAPTER 3. REPOSITORY HOOKS

In Business Central, it is possible to trigger a chosen action every time a particular event happens. For this purpose, you can configure the repository to use scripts called hooks.

3.1. CONFIGURING GIT HOOKS

Business Central can automatically push changes to a remote repository using the Git hooks. Git hooks support has been introduced with the release of Red Hat JBoss BPM Suite 6.2.0.

> **Note**
>
> Please note that currently only the **post-commit** hook is supported. **Post-commit** hooks are triggered after finishing the entire commit process.

The following procedure shows how to configure the **post-commit** hook to automatically push your changes to the remote repository.

1. In Business Central, go to **Authoring** → **Administration**.

2. Below the main menu, click **Repositories** → **Clone repository**.

3. In the displayed **Clone repository** dialog box, fill in the repository information:

   ▸ Repository Name

   ▸ Organizational Unit

   ▸ Git URL: For example
   https://github.com/USERNAME/REPOSITORY_NAME.git

   > **Important**
   >
   > It is important to use the HTTPS or Git protocol instead of a SCP-style SSH URL. Business Central does not support the basic SSH URL and fails with *Invalid URL format*.

   Figure 3.1. An invalid SCP-style SSH URL.

Clone Repository   ✕

Repository Information

Repository Name *   NewRepository

Organizational Unit *   example   ⌄

Git URL *   git@github.com:user/ExampleRepository/repo.git
                           Invalid URL format

User Name   user

Password   ••••••••••••

                                                      Cancel   ☁ Clone

- » User Name: your Git user name

- » Password: your Git password

4. Go to the created repository:

   ```
   ~]$ cd JBOSS_HOME/bin/.niogit/REPOSITORY_NAME.git
   ```

5. Change the remote URL:

   ```
   ~]$ git remote set-url origin
   git@github.com:USERNAME/REPOSITORY_NAME.git
   ```

   Make sure that you can access the remote repository through command line using
   SSH. For example, the private SSH key for the repository should exist under the
   ~/.ssh/ directory.

6. Verify that the remote repository was successfully added:

   ```
   ~]$ git remote -v
   ```

   The command should list the following:

   ```
   origin   git@github.com:USERNAME/REPOSITORY_NAME.git (fetch)
   origin   git@github.com:USERNAME/REPOSITORY_NAME.git (push)
   ```

7. Create a file named **post-commit** with the permissions set to **rwxr--r--** under
   JBOSS_HOME/bin/.niogit/REPOSITORY_NAME.git/hooks with the
   following content:

```
#!/bin/sh
git push origin master
```

8. Make sure that the configuration was successful by creating a new guided rule in
   Business Central: Go to **Authoring** → **Project Authoring** → **New Item** →
   Guided Rule. Fill in the required information in the displayed **Create new Guided
   Rule** window and click **Ok**.

   All of the changes should be pushed automatically.

For further information about remote Git repositories, refer to this article: How to configure the
BxMS 6 server to use a remote Git repository for storing assets?

Furthermore, it is also possible to specify the system property
org.uberfire.nio.git.hooks. Its value determines a directory with default hook files,
which will be copied to the newly created Git repositories. See the example of a
standalone.xml file with this setting below:

```
<system-properties>
  <property name="org.uberfire.nio.git.hooks" value="/opt/jboss-
as/git-hooks">
  </property>
  ...
</system-properties>
```

CHAPTER 4. COMMAND LINE CONFIGURATION

The `kie-config-cli` tool is a command line configuration tool that provides capabilities to manage the system repository from the command line and can be used in an online or offline mode.

Online mode (default and recommended)

On startup, the tool connects to a Git repository using a Git server provided by **kie-wb**. All changes are made locally and published to upstream only after explicitly executing the push-changes command. Use the exit command to publish local changes. To discard local changes on exit, use the discard command.

Offline mode (installer style)

Creates and manipulates the system repository directly on the server (there is no discard option).

The tool is available on the Red Hat Customer Portal. To download the `kie-config-cli` tool, do the following:

1. Go to the Red Hat Customer Portal and log in.

2. Click **DOWNLOADS** at the top of the page.

3. In the **Product Downloads** page that opens, click **Red Hat JBoss BPM Suite**.

   a. From the **Version** drop-down menu, select **6.3.0**.

   b. In the displayed table, navigate to the **Supplementary Tools** row and then click **Download**.

Extract the zip package for supplementary tools you downloaded from the Red Hat Customer Portal. It contains the directory `kie-config-cli-6.MINOR_VERSION-redhat-x-dist` with file `kie-config-cli.sh`.

4.1. STARTING THE KIE-CONFIG-CLI TOOL IN ONLINE MODE

1. To start the `kie-config-cli` tool in online mode, navigate to the `kie-config-cli-6.MINOR_VERSION-redhat-x-dist` directory where you installed the tool and then execute the following command.

2. In a Unix environment run:

```
./kie-config-cli.sh
```

In a Windows environment run:

```
./kie-config-cli.bat
```

By default, the tool starts in online mode and asks for user credentials and a Git URL to connect to (the default value is `git://localhost/system`). To connect to a remote server, replace the host and port with appropriate values.

Example: `git://kie-wb-host:9148/system`

4.2. STARTING THE KIE-CONFIG-CLI TOOL IN OFFLINE MODE

To operate in offline mode, append the offline parameter to the command as below.

1. Navigate to the `kie-config-cli-6.MINOR_VERSION-redhat-x-dist` directory where you installed the tool.

2. In a Unix environment, run:

```
./kie-config-cli.sh offline
```

In a Windows environment, run:

```
./kie-config-cli.bat offline
```

Executing this command changes the tool's behaviour and displays a request to specify the folder where the system repository (`.niogit`) is located. If .niogit does not yet exist, the folder value can be left empty and a brand new setup is created.

4.3. COMMANDS AVAILABLE FOR THE KIE-CONFIG-CLI TOOL

The following commands are available for managing the Git repository using the `kie-config-cli` tool:

- `add-deployment`: Adds a new deployment unit

- `add-repo-org-unit`: Adds a repository to the organizational unit

- `add-role-org-unit`: Adds role(s) to an organizational unit

- `add-role-project`: Adds role(s) to a project

- `add-role-repo`: Adds role(s) to a repository

- `create-org-unit`: Creates new organizational unit

- `create-repo`: Creates a new git repository

- `discard`: Does not publish local changes, cleans up temporary directories and closes the tool

- `exit`: Publishes work, cleans up temporary directories and closes the tool

- `fetch-changes`: Fetches changes from upstream repository

- `help`: Prints available commands with descriptions

- `list-deployment`: Lists available deployments

- `list-org-units`: Lists available organizational units

- `list-repo`: Lists available repositories

- `push-changes`: Pushes changes to upstream repository (in online mode only)

- **remove-deployment**: Removes existing deployment

- **remove-org-unit**: Removes existing organizational unit

- **remove-repo**: Removes an existing repository from config only

- **remove-repo-org-unit**: Removes a repository from the organizational unit

- **remove-role-org-unit**: Removes role(s) from an organizational unit

- **remove-role-project**: Removes role(s) from a project

- **remove-role-repo**: Removes role(s) from a repository

CHAPTER 5. MIGRATION

Migrating your projects from Red Hat JBoss BPM Suite 5 to Red Hat JBoss BPM Suite 6 requires careful planning and step by step evaluation of the various issues. You can plan for migration either manually, or by using automatic processes. Most real world migration will require a combination of these two processes.

Because Red Hat JBoss BPM Suite 6 uses Git for storing assets, artifacts and code repositories including processes and rules, you should start by creating an empty project in Red Hat JBoss BPM Suite 6 as the basis for your migration with dummy files as placeholders for the various assets and artifacts. Running a Git clone of this empty project into your favorite IDE will initiate the migration process.

Based on the placeholder files in your cloned project, you can start adding assets at the correct locations. The Red Hat JBoss BPM Suite 6 system is smart enough to pick these changes and apply them correctly. Ensure that when you are importing old rule files that they are imported with the right package name structure.

Since Maven is used for building projects, the projects assets like the rules, processes and models are accessible as a simple JAR file.

This section lists the generally accepted step by step ways to migrate your project. These are just guidelines though, and actual migration may vary a lot from this.

In general, you should:

1. Migrate the data first: These are your business assets.

2. Next, migrate your runtime processes.

3. Finally, convert old API calls to new ones one by one.

Let us look at these steps in more detail in the next few sections:

5.1. DATA MIGRATION

To migrate data from Red Hat JBoss BPM Suite 5, do the following:

1. Download the migration tool by logging in at the Red Hat Customer Portal and then navigating to Red Hat JBoss BPM Suite Software Downloads section. Click on **Red Hat JBoss BPM Suite Migration Tool** to download the zip archive.

2. Unzip the downloaded zip archive in a directory of your choice and navigate to this directory in a command prompt. This directory contains four folders:

   » `bin` - contains the launch scripts.

   » `jcr-exporter-libs` - contains the libs specific to the `export-from-JCR` part of the migration.

   » `vfs-importer-libs` - contains the libs specific to the `import-into-Git` part of the migration.

   » `conf` - contains global migration tool configuration.

3. For production databases, copy the JDBC driver for the database that is used by the JCR repository into the **jcr-exporter-libs** directory of the migration tool.

4. Execute the following command:

```
./bin/runMigration.sh -i <source-path> -o <destination-path> -r <repository-name>
```

Where:

- ➤ **<source-path>** is a path to a source JCR repository.

- ➤ **<desintation-path>** is a path to a destination Git VFS. This folder must not exist already.

- ➤ **<repository-name>** an arbitrary name for the new repository.

The repository is migrated at the specified destination.

Besides the **-i** command, you can also use **-h** to print out a help message and **-f** which forces an overwrite of the output directory, thus eliminating the need for manual deletion of this directory.

Importing the Repository in Business Central

The repository can be imported in business central by cloning it. In the Administration perspective, click on the **Repositories → Clone Repository** menu to start the process.

> **Note**
>
> Assets can also be migrated manually as they are all just text files. The BPMN2 specification and the DRL syntax did not change between the different versions.

Importing the Repository in JBDS

To import the repository in JBoss Developer Studio, do the following

1. Start JBoss Developer Studio.

2. Start the Red Hat JBoss BPM Suite server (if not already running) by selecting the server from the server tab and click the start icon.

3. Select **File → Import...** and navigate to the Git folder. Open the Git folder to select **Projects from Git** and click next.

4. Select the repository source as **Existing local repository** and click next.

5. Select the repository that is to be configured from the list of available repositories.

6. Import the project as a general project in the next window and click next. Name this project and click Finish.

5.2. RUNTIME MIGRATION

To run Red Hat JBoss BPM Suite 5 processes in Red Hat JBoss BPM Suite 6, do the following:

1. Set the system property **jbpm.v5.id.strategy** to true in the JBoss BPM Suite **standalone.xml** file:

   ```
   <property name="jbpm.v5.id.strategy" value="true"/>
   ```

2. Load the KieSession as shown here:

   ```
   KieSession ksession =
   JPAKnowledgeService.loadStatefulKnowledgeSession(sessionID,
   kbase, sessionConf, env);
   ```

3. Continue the normal execution of the process using KieSession methods:

   ```
   ksession.signalEvent("SomeEvent", null);
   ```

5.3. API AND BACKWARDS COMPATIBILITY

Migrating to Version 6.1

In version 6.1, 5.*X* APIs are no longer officially supported.

Red Hat JBoss BPM Suite no longer provides backward compatibility with the rule, event, and process application programming interface (API) from Red Hat JBoss BRMS 5. The content of the **knowledge-api** JAR file is no longer supported in version 6.1 and is replaced by APIs contained in the **kie-api** JAR file that were introduced in Red Hat JBoss BPM Suite 6.0.

If you used the legacy 5.x API (located in **knowledge-api.jar**), please migrate (rewrite) the API calls to the new KIE API. Please be aware that several other APIs have changed between Red Hat JBoss BRMS 5.x and Red Hat JBoss BPM Suite 6.x, namely the task service API and the REST API.

Migrating to Version 6.0

The Red Hat JBoss BPM Suite 6 system provides backward compatibility with the rule, event and process interactions from JBoss BRMS 5. You should eventually migrate (rewrite) these interactions to the all new revamped core API because this backward compatibility is likely to be deprecated.

If you cannot migrate your code to use the new API, then you can use the API provided by the purpose built **knowledge-api** JAR for backwards compatible code. This API is the public interface for working with JBoss BPM Suite and JBoss BRMS and is backwards compatible.

If you are instead using the REST API in Red Hat JBoss BPM Suite 5, note that this has changed as well and there is no mechanism in it for backwards compatibility.

5.4. MIGRATING TASK SERVICE

Red Hat JBoss BPM Suite 6 provides support for a locally running task server only. This means that you do not need to setup any messaging service in your project. This differs from Red Hat JBoss BPM Suite 5 because it provided a task server that was bridged from the core engine by using, most commonly, the messaging system provided by HornetQ.

To help you bridge the gap until you can migrate this in your current architecture, there is a helper or utility method, **LocalHTWorkItemHandler**.

Since the TaskService API is part of the public API you will now need to refactor your imports because of package changes and refactor your methods due to API changes themselves.

CHAPTER 6. DATA MANAGEMENT

6.1. DATA BACKUPS

When applying a backup mechanism to Red Hat JBoss BPM Suite, make sure you back up the following resources:

- Any customized deployment descriptors (such as, `web.xml`, `jboss-web.xml`, `jboss.xml`)

- Any customized properties files

> **Note**
>
> Consider backing up the entire `business-central.war` and `dashbuilder.war` files.

6.2. SETUP INDEXES

Setup Foreign Key Indexes

Some databases, for instance Oracle and PostgreSQL, do not automatically create an index for each foreign key. This can result in deadlocks occurring. To avoid this situation, it is necessary to create an index on all foreign keys, especially in the Oracle database.

Setup Indexes for Process and Task Dashboard

Process and Task Dashboard in 6.1 has been refactored in order to cope with high volume of task and process instances. In order to get good response times while querying the database the following JBoss BPM Suite tables need to be indexed: `processinstancelog` and `bamtasksummary`.

Note that *ALL* the columns in these two tables need to be indexed and not just the primary and foreign keys.

6.3. SETTING UP THE DATABASE

Dashbuilder application requires an existing database, previously created before running the application. To create a database you can use any database client tool and run the following commands:

PostgreSQL

The following sql sentence is used to create a PostgreSQL database:

```
CREATE DATABASE dashbuilder
  WITH ENCODING='UTF8'
       OWNER=dashbuilder
       CONNECTION LIMIT=-1
```

> **Note**
>
> The database encoding must be UTF8

DB2

DB2 database can be created using the following SQL sentence:

```
CREATE DATABASE dashb PAGESIZE 16384
```

> **Note**
>
> The default pagesize for DB2 systems is 4k which is not enough for the
> dashbuilder table columns size. The pagesize should be forced to 16384
> as shown in the above sentence.

Once the database is created, the application server datasource must be configured. You
must edit the JBoss EAP configuration file and configure the datasource subsystem as the
following examples:

```
<datasource jndi-name="java:jboss/datasources/jbpmDS" enabled="true"
use-java-context="true" pool-name="postgresqlDS">
  <connection-url>jdbc:postgresql://localhost/test</connection-url>
  <driver>postgresql</driver>
  <pool></pool>
  <security>
    <user-name>sa</user-name>
    <password>sa</password>
  </security>
</datasource>
<drivers>
  <driver name="postgresql" module="org.postgresql">
    <xa-datasource-class>org.postgresql.xa.PGXADataSource</xa-
datasource-class>
  </driver>
</drivers>
```

6.4. EDITING THE DATABASE

Dashbuilder requires the JBoss BPM Suite to have history log's database tables. It is
mandatory to deploy the Human Task console (or a superset, i.e: kie-wb) first. Otherwise, the
Dashboard will not be initialized correctly and it will not be possible to display its key
performance indicators.

By default, the application is configured to use a datasource with the following JNDI name:

```
java:jboss/datasources/ExampleDS
```

This is specified in JBoss EAP configuration file, such as **standalone.xml**.

> **Note**
>
> This datasource is intended for development/demo purposes; it is present by default in any JBoss installation.

If you want to deploy on a database different from H2 like Oracle, MySQL, Postgres or MS SQL Server, please perform the following steps:

Procedure: Changing Database

1. Install the database driver on JBoss (refer to JBoss driver documentation).

2. Create an empty database and a JBoss data source which connects to the database driver.

3. Modify the file **dashbuilder.war/WEB-INF/jboss-web.xml**:

   ```
   <jboss-web>
     <context-root>/dashbuilder</context-root>
     <resource-ref>
       <res-ref-name>jdbc/dashbuilder</res-ref-name>
       <res-type>javax.sql.DataSource</res-type>
       <jndi-name>java:jboss/datasources/myDataSource</jndi-name>
     </resource-ref>
     ...
   ```

4. Replace the jndi-name parameter value by the JNDI path of the JBoss data source you created.

5. Modify the file **dashbuilder.war/WEB-INF/jboss-deployment-structure.xml**

6. Add the following snippet of configuration inside the **deployment** tag, where **jdbcDriverModuleName** is the name of the JBoss JDBC driver module:

   ```
   <dependencies>
     <module name="jdbcDriverModuleName" />
   </dependencies>
   ```

6.5. DDL SCRIPTS

DDL scripts for database tables for both JBoss BRMS and BPM Suite are available for download via the Customer Portal. These scripts allow you to study the tables and use them to create the tables and indexes manually or in databases that are not directly supported.

To download these scripts, login to the Customer Portal and click on Red Hat JBoss BPM Suite. Select the version of the product for your requirements and then click on *Download* in the row *Red Hat JBoss BPM Suite 6.x.x Supplementary Tools* to download the supplementary tools.

Unzip the file on your machine. The DDL scripts are located in the **ddl-scripts** folder. Database scripts are provided for DB2, H2, MySQL5, Oracle, PostgreSQL, and SQLServer.

The complete Entity Relationship diagram can be viewed in this Red Hat Solution.

CHAPTER 7. ASSET REPOSITORY

Business Rules, Process definition files and other assets and resources created in Business Central are stored in Asset repository, which is otherwise known as the Knowledge Store.

Knowledge Store is a centralized repository for your business knowledge. It connects with the Git repository that allows you to store different kinds of knowledge assets and artifacts at a single location. Business Central provides a web front-end that allows you to view and update the stored content. You can access it using the *Project Explorer* from the unified environment of Red Hat JBoss BPM Suite.

All business assets are stored in repositories. These repositories are then saved in directories called organizational units. By default, the Artifact repository does not contain any organizational unit. Therefore, to be able to create your own business assets, you need to create an organizational unit and a repository first.

7.1. CREATING AN ORGANIZATIONAL UNIT

It is possible to create an organizational unit either in the **Administration** perspective of Business Central, using the `kie-config-cli` tool or the REST API calls.

Creating an Organizational Unit in Business Central

Important

Note that only users with the `admin` role in Business Central can create organizational units.

Procedure: Using Business Central to Create an Organizational Unit

1. In Business Central, go to **Authoring → Administration**.

2. On the perspective menu, click **Organizational Units → Manage Organizational Units**.

3. In the **Organization Unit Manager** view, click **Add**.

   The **Add New Organizational Unit** dialog window opens.

   Figure 7.1. Add New Organizational Unit_Dialog Window

Add New Organizational Unit ✕

Organizational Unit Information

Name * [Organizational Unit name...]

Default Group ID * ⓘ [Default Group Id...]

Owner [Organizational Unit owner...]

[**+Ok**] [Cancel]

4. Enter the two mandatory parameters (name and default group ID) and click **Ok**.

Creating an Organizational Unit Using the kie-config-cli Tool

Organizational units can be created using the **kie-config-cli** tool as well. To do so, run the **create-org-unit** command. The tool then guides you through the entire process of creating an organizational unit by asking for other required parameters. Type **help** for a list of all commands.

For more information about the **kie-config-cli** tool, see Chapter 4, *Command Line Configuration*.

Creating an Organizational Unit Using the REST API

Important

Note that only users with the **rest-all** role can create organizational units.

To create an organizational unit in Knowledge Store, issue the **POST** REST API call. Details of the organizational unit are defined by the JSON entity.

Input parameter of the call is a **OrganizationalUnit** instance. Call returns a **CreateOrganizationalUnitRequest** instance.

Example 7.1. Creating an Organizational Unit Using the Curl Utility

Example JSON entity containing details of an organizational unit to be created:

```
{
  "name"        : "helloWorldUnit",
  "owner"       : "tester",
  "description" :  null,
```

```
    "repositories" :   []
}
```

Execute the following command:

```
curl -X POST 'localhost:8080/business-
central/rest/organizationalunits/' -u USERNAME:PASSWORD -H
'Accept: application/json' -H 'Content-Type: application/json' -d
'{"name":"helloWorldUnit","owner":"tester","description":null,"re
positories":[]}'
```

For further information, refer to chapter *Knowledge Store REST API*, section *Organizational Unit Calls* of *Red Hat JBoss BPM Suite Development Guide*.

7.2. CREATING A REPOSITORY

There are three ways to create a repository: through the **Administration** perspective of Business Central, the **kie-config-cli** tool, or using the REST API calls.

Creating a Repository in Business Central

Important

Note that only users with the **admin** role in Business Central can create repositories.

Procedure: Using Business Central to Create a Repository

1. In Business Central, go to **Authoring** → **Administration**.

2. On the perspective menu, click **Repositories** → **New repository**.

    The **New Repository** pop-up window is displayed.

    Figure 7.2. New Repository Dialog Window

3. Specify the two mandatory parameters:

    » **Repository name**

> **Note**
>
> Make sure that the repository name is a valid file name. Avoid using a space or any special character that might lead to an invalid name.

> **Organizational unit**: Specifies the location of the newly created repository.

4. Click **Finish**.

You can view the newly created repository either in the **File Explorer** or the **Project Explorer**.

Creating Repository Using kie-config-cli Tool

To create a new Git repository using the **kie-config-cli** tool, run the **create-repo** command. The tool then guides you through the entire process of creating a repository by asking for other required parameters. Type **help** for a list of all commands.

For more information about the **kie-config-cli** tool, see Chapter 4, *Command Line Configuration*.

Creating Repository Using REST API

> **Important**
>
> Note that only users with the **rest-all** role can create repositories.

To create a repository in the Knowledge Store, issue the **POST** REST API call. Details of the repository are defined by the JSON entity. Make sure you established an authenticated HTTP session before executing this call.

Input parameter of the call is a **RepositoryRequest** instance. Call returns a **CreateOrCloneRepositoryRequest** instance.

Example 7.2. Creating Repository Using Curl Utility

Example JSON entity containing details of a repository to be created:

```
{
  "name"                 : "newRepository",
  "description"          :  null,
  "gitURL"               :  null,
  "requestType"          : "new",
  "organizationalUnitName" : "helloWorldUnit"
}
```

Execute the following command:

```
curl -X POST 'localhost:8080/business-central/rest/repositories/'
-u USERNAME:PASSWORD -H 'Accept: application/json' -H 'Content-
Type: application/json' -d
'{"name":"newRepository","description":null,"requestType":"new","
```

```
gitURL":null,"organizationalUnitName":"helloWorldUnit"}'
```

For further information, refer to chapter *Knowledge Store REST API*, section *Repository Calls* of *Red Hat JBoss BPM Suite Development Guide*.

7.3. CLONING A REPOSITORY

It is possible to clone a repository either in Business Central or using the REST API calls. The **kie-config-cli** tool cannot be used to clone arbitrary repositories. Run **git clone**, or use one of the following options instead:

Cloning a Repository in Business Central

> **Important**
>
> Note that only users with the **admin** role in Business Central can clone repositories.

Procedure: Using Business Central to Clone a Repository

1. In Business Central, go to **Authoring → Administration**.

2. On the perspective menu, choose **Repositories → Clone repository**.

   The **Clone Repository** pop-up window is displayed.

   Figure 7.3. Clone Repository Dialog Window

Clone Repository   ✕

Repository Information

Repository Name *   [repository name...]

Organizational Unit *   [--- Select ---   ⌄]

Git URL *   [git url...]

User Name   [user name...]

Password   [password...]

[Cancel]   [☁ Clone]

3. In the **Clone Repository** dialog window, enter the repository details:

   a. Enter the **Repository Name** to be used as the repository identifier in the
      Asset repository and select the **Organizational Unit** it should be added to.

   b. Enter the URL of the Git repository:

      » For a local repository, use
        `file:///`*PATH_TO_REPOSITORY/REPOSITORY_NAME*.

      » For a remote or preexisting repository, use
        `https://github.com/`*USERNAME/REPOSITORY_NAME*`.git` or
        `git://`*HOST_NAME/REPOSITORY_NAME*.

        Important

        It is important to use the HTTPS or Git protocol instead of a SCP-
        style SSH URL. Business Central does not support the basic
        SSH URL and fails with `Invalid URL format`.

        Note

        The file protocol is only supported for READ operations. WRITE
        operations are *not* supported.

c. If applicable, enter the **User Name** and **Password** of your Git account to be used for authentication.

4. Click **Clone**.

   A confirmation prompt with the notification that the repository was created successfully is displayed.

5. Click **Ok**.

   The repository is now being indexed. Some workbench features may be unavailable until the indexing has completed.

You can view the cloned repository either in the **File Explorer** or the **Project Explorer**.

> **Note**
>
> If you are deploying Business Central on WebLogic server, set the following Java system property in the **setDomainEnv.sh** file (for Linux) or **setDomainEnv.cmd** file (for Windows):
>
> ```
> JAVA_OPTIONS="%JAVA_OPTIONS% -DUseSunHttpHandler=true"
> ```
>
> This enables the WebLogic server to use the HTTP handlers.

Cloning a Repository Using the REST API

To clone a repository, issue the **POST** REST API call. This call creates or clones (according to the value of the **requestType** parameter) the repository defined by the JSON entity.

The input parameter of the call is a **RepositoryRequest** instance. The Call returns a **CreateOrCloneRepositoryRequest** instance.

> **Important**
>
> Note that, only users with the **rest-all** role can clone repositories.

Example 7.3. Cloning a Repository Using the Curl Utility

Example JSON entity containing details of a repository to be cloned:

```
{
  "name"                  : "clonedRepository",
  "description"           :  null,
  "requestType"           : "clone",
  "gitURL"                :
"git://localhost:9418/newRepository",
  "organizationalUnitName" : "helloWorldUnit"
}
```

Execute the following command:

```
curl -X POST 'localhost:8080/business-central/rest/repositories/'
-u USERNAME:PASSWORD -H 'Accept: application/json' -H 'Content-
Type: application/json' -d
'{"name":"clonedRepository","description":null,"requestType":"clo
ne","gitURL":"git://localhost:9418/newRepository","organizational
UnitName":"helloWorldUnit"}'
```

For further information, refer to chapter *Knowledge Store REST API*, section *Repository Calls* of
Red Hat JBoss BPM Suite Development Guide.

7.4. REMOVING A REPOSITORY

Repositories can be removed using any of the following procedures.

Removing a Repository in Business Central

The simplest way to remove a repository is using the **RepositoryEditor** in Business
Central.

Procedure: Using Business Central to Remove a Repository

1. In Business Central, go to **Authoring → Administration**.

2. Select **Repositories** from the tree menu on the left.

3. In the **RepositoryEditor** on the right side of the page, locate the repository you
   want to delete from the list of available repositories.

4. From the drop-down menu, select **master → Delete**.

   The following message will appear:

   ```
   Are you sure you want to remove Repository "REPOSITORY_NAME"?
   Some editors may become inoperable if their content is
   inaccessible.
   ```

5. Press **OK** to delete the repository.

Removing a Repository Using the kie-config-cli Tool

Repositories can be removed using the **kie-config-cli** tool as well. To do so, run the
remove-repo command.

For further information about the **kie-config-cli** tool, see Chapter 4, *Command Line
Configuration*.

Removing a Repository Using the REST API

> **Important**
>
> Note that only users with the **rest-all** role can remove repositories.

To remove a repository from the Knowledge Store, issue the **DELETE** REST API call. Make sure you established an authenticated HTTP session before executing this call.

The call returns a **RemoveRepositoryRequest** instance.

> **Example 7.4. Removing a Repository Using the Curl Utility**
>
> Execute the following command:
>
> ```
> curl -X DELETE 'localhost:8080/business-
> central/rest/repositories/REPOSITORY_NAME' -u USERNAME:PASSWORD -
> H 'Accept: application/json' -H 'Content-Type: application/json'
> ```

For further information, refer to chapter *Knowledge Store REST API*, section *Repository Calls* of *Red Hat JBoss BPM Suite Development Guide*.

7.5. MANAGING ASSETS

Technical Preview

The content in this section is classified as Technical Preview for the 6.1 release of Red Hat JBoss BPM Suite. It is provided as is and no support is provided.

To activate and use the featured described here, you will need to login to Business Central with a user that has been given the special role of **kiemgmt**.

To make management of projects easier, Red Hat JBoss BPM Suite now provides a way to manage multiple projects based on standards. This allows you to create repository structures using industry standard best practices for maintenance, versioning and distribution of your projects.

To start with, repositories can now be managed or unmanaged.

Managed and Unmanaged Repositories

Unmanaged Repositories are the repository structures that you are used to. They can contain multiple unrelated projects.

Managed Repositories, on the other hand, provide version control at the project level and project branches for managing the release cycle. Further, Managed Repositories can be restricted to just a single project or encompass multiple projects. When you create a Managed Repository, the asset management configuration process is automatically launched in order to create the repository branches. Corresponding project structure is created as well.

Procedure: Creating an Unmanaged Repository

1. In Business Central, go to **Authoring → Administration**.

2. Click **Repositories → New Repository**.

The **New Repository** window is displayed.

3. Enter the repository name and select an organizational unit the repository belongs to.

4. Click **Finish**.

Procedure: Creating a Managed Repository

1. In Business Central, go to **Authoring → Administration**.

2. Click **Repositories → New Repository**.

   The **New Repository** window is displayed.

3. Enter the repository name and select an organizational unit the repository belongs to.

4. Select the **Managed Repository** check box and click **Next** to enter additional details of the Managed Repository.

5. Choose either the **Single-project Repository** or the **Multi-project Repository** radio button.

   If the project you are creating is simple and self-contained, select the **Single-project Repository** radio button. Note that you will not be able to add more projects to this repository later.

   For more complex projects, where there is likely to be a parent project that encompasses smaller ones, select the **Multi-project Repository** radio button.

6. Enter the details of the managed project along with the GAV (Group, Artifact, Version) details.

   Note that all projects created in a **Multi-project Repository** will be managed together, with their version numbers being incremented together as well. Details of the parent project will be inherited by all future projects that you create in this Managed Repository.

7. Click **Finish** .

Managed Branches

With Managed Repositories comes the added advantage of Managed Branches. As in Git, you can choose to work on different branches of your project (for example: master, dev and release). This process of branching can also be automated for you, by selecting the checkbox while creating a new Managed Repository (for both single and multi-projects).

You can switch between branches by selecting the desired branch while working in the Project Explorer.

Repository Structure

If you do not select automatic branch management while creating a repository, you can create branches manually afterwards. For Managed Repositories, you can do so by using the **Configure** button. This button, along with **Promote** and **Release** buttons, is provided in the **Repository Structure** view. You can access this view, by clicking on **Repository → Repository Structure** in the Project Explorer perspective menu.

Clicking on the **Configure** button allows you to create branches or edit automatically created ones.

You can promote assets from the master branch to other branches using the **Promote** button. Similarly, you can Release branches and deploy them on the server using the **Release** button.

Both these functions are controlled internally by the use of pre-defined processes that are deployed on your instance. For example, when you click on **Promote** button after having done work on your development branch, a Promote Changes process is started in the background. A user, with the role of `kiemgmt` will have a user task appear in this task list to review the assets being promoted. This user can claim this task, and decide to promote all, some or none of the assets. The underlying process will cherry-pick the commits selected by the user to a release branch. This user can also request another review of these assets and this process can be repeated multiple times till all the assets are ready for release. The flow for this process is shown below:
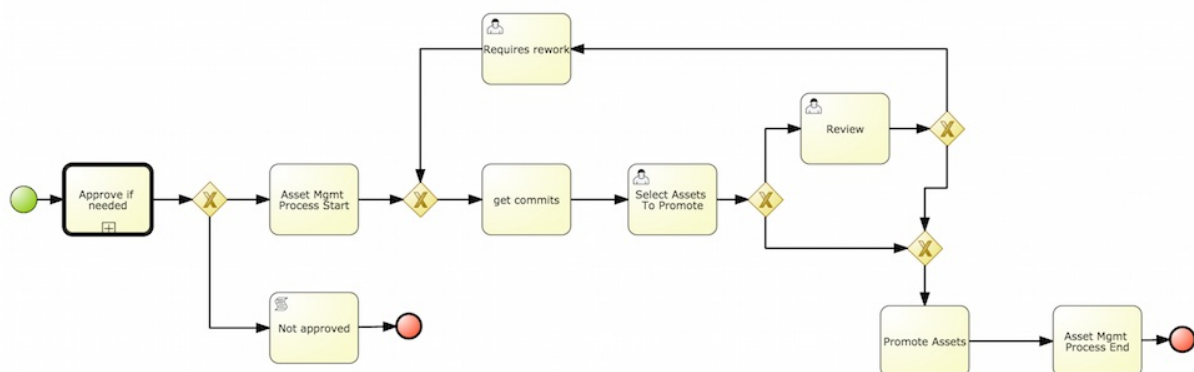
Similarly, when you click on the **Release** button, a release process flow is initiated. This process flow builds the project and updates all the Maven artifacts to the next version, and deploys the project to the runtime, if runtime deployment details are supplied.

> **Warning**
>
> Project branches to be released must start with the keyword `release`.

Release Configuration ×

Repository	TestManagedRepository
Source Branch	master
Release Version *	1.0.0
	The current repository version is: 1.0.0-SNAPSHOT
Deploy To Runtime *	☑
User Name *	admin
Password *	Password
Server URL *	http://localhost:8080/business-central/

＋Ok Cancel

7.6. MAVEN REPOSITORY

Maven is a software project management tool which uses a project object model (POM) file to manage:

» Builds

» Documentation

» Reporting

» Dependencies

» Releases

» SCMs

» Distribution

A Maven repository is used to hold or store the build artifacts and project dependencies and is generally of two types:

Local

> Refers to a local repository where all the project dependencies are stored and is located with the current installation in the default folder as "m2". It is a cache of the remote downloads, and also contains the temporary build artifacts which have not yet been released.

Remote

> Refers to any other type of repository that can be accessed by a variety of protocols such as `file://` or `http://`. These repositories can be at a remote location set up by a third-party for downloading of artifacts or an internal repository set up on a file or HTTP server, used to share private artifacts between the development teams for managing internal releases.

7.7. CONFIGURING DEPLOYMENT TO A REMOTE NEXUS REPOSITORY

Nexus is a repository manager frequently used in organizations to centralize storage and management of software development artifacts. It is possible to configure your project so that artifacts produced by every build are automatically deployed to a repository on a remote Nexus server.

To configure your project to deploy artifacts to a remote Nexus repository, add a `distributionManagement` element to your project's `pom.xml` file as demonstrated in the code example below.

```xml
<distributionManagement>
  <repository>
    <id>deployment</id>
    <name>Internal Releases</name>

<url>http://your_nexus_host:8081/nexus/content/repositories/releases</url>
  </repository>
  <snapshotRepository>
    <id>deployment</id>
    <name>Internal Releases</name>

<url>http://your_nexus_host:8081/nexus/content/repositories/snapshots/</url>
  </snapshotRepository>
</distributionManagement>
```

Replace the URLs in the example with real URLs of your Nexus repositories. The repository specified in the `snapshotRepository` element is used when the `-SNAPSHOT` qualifier is appended to the project's current version number. In other cases the repository specified in the `repository` element is used.

If your Nexus server requires authentication, you will also need to modify your projects Maven settings to add your credentials in the `settings-security.xml` file, using a master password. By default, this file is in the `$M2_HOME/conf` folder, unless you have changed its location by modifying the `kie.maven.settings.custom` system property. It

is also possible to create new user-specific **settings.xml** and **settings-security.xml** files in the **~/.m2** folder. In that case, these files will override the original ones.

See the following example of the **settings-security.xml**:

```
<servers>
  <server>
    <id>deployment</id>
    <username>admin</username>
    <password>{COQLCE6DU6GtcS5P=}</password>
  </server>
</servers>
```

> **Important**
>
> Note that keeping your server authentication credentials (for example the passwords) as a plain text in the **settings.xml** file is *not* recommended. All the information should be hashed with a master password in the **settings-security.xml** file.
>
> For further information about password encryption and creating a master password, see the Apache Maven documentation, article Password Encryption.

With this configuration in place, clicking the **Build & Deploy** button in Business Central executes a Maven build and deploys the built artifacts both to the local repository and to one of the Nexus repositories specified in the **pom.xml** file.

7.8. SYSTEM CONFIGURATION

In JBoss EAP, to change a Business Central property, such as the configuration for SSH, do the following:

Procedure: Changing System Properties

1. Edit the file *JBOSS_HOME*/**domain/configuration/host.xml**

2. Locate the XML elements server that belong to the **main-server-group** and add the system property. For example:

   ```
   <system-properties>
     <property name="org.uberfire.nio.git.dir" value="..." boot-time="false"/>
     ...
   </system-properties>
   ```

Here is a list of all the available system properties:

Table 7.1. Available System Properties

Property	Usage	Default
`org.uberfire.nio.git.dir`	Location of the directory .niogit.	Working directory
`org.uberfire.nio.git.daemon.enabled`	Enables/disables Git daemon.	**true**
`org.uberfire.nio.git.daemon.host`	If Git daemon enabled, uses this property as the localhost identifier.	**localhost**
`org.uberfire.nio.git.daemon.port`	If Git daemon is enabled, uses this property as the port number.	**9418**
`org.uberfire.nio.git.ssh.enabled`	Enables/Disables SSH daemon.	**true**
`org.uberfire.nio.git.ssh.host`	If SSH daemon is enabled, uses this property as the localhost identifier.	**localhost**
`org.uberfire.nio.git.ssh.port`	If SSH daemon is enabled, uses this property as the port number.	**8001**
`org.uberfire.nio.git.ssh.cert.dir`	Location of the `.security` directory where local certificates will be stored.	Working directory
`org.uberfire.metadata.index.dir`	Location of the `.index` folder for Lucene.	Working directory
`org.uberfire.cluster.id`	Name of the Helix cluster, for example: kie-cluster	

Property	Usage	Default
`org.uberfire.cluster.zk`	Connection string to Zookeeper. This is of the form `host1:port1,host2:port2,host3:port3`. For example: `localhost:2188`.	
`org.uberfire.cluster.local.id`	Unique id of the Helix cluster node. Note that ':' is replaced with '_'. For example: node1_12345.	
`org.uberfire.cluster.vfs.lock`	Name of the resource defined on the Helix cluster, for example: kie-vfs	
`org.uberfire.cluster.autostart`	Delays VFS clustering until the application is fully initialized to avoid conflicts when all cluster members create local clones.	false
`org.uberfire.sys.repo.monitor.disabled`	Disable configuration monitor (do not disable unless you know what you're doing).	false
`org.uberfire.secure.key`	Secret password used by password encryption.	`org.uberfire.admin`
`org.uberfire.secure.alg`	Crypto algorithm used by password encryption.	`PBEWithMD5AndDES`
`org.guvnor.m2repo.dir`	Place where Maven repository folder will be stored.	`working-directory/repositories/kie`

Property	Usage	Default
`org.kie.example.rep ositories`	Folder from where demo repositories will be cloned. The demo repositories need to have been obtained and placed in this folder. This system property takes precedence over org.kie.demo and org.kie.example properties.	Not used.
`org.kie.demo`	Enables external clone of a demo application from GitHub. This system property takes precedence over org.kie.example.	true
`org.kie.example`	Enables example structure composed by Repository, Organization Unit and Project.	false
`org.kie.build.disab le-project-explorer`	Disables automatic project building when navigating the Project Explorer.	false

CHAPTER 8. PROCESS EXPORT AND IMPORT

8.1. CREATING A PROCESS DEFINITION

Make sure you have logged in to JBoss BPM Suite or you are in JBoss Developer Studio with the repository connected.

To create a Process, do the following:

1. Open the Project Authoring perspective (**Authoring → Project Authoring**).

2. In Project Explorer (**Project Authoring → Project Explorer**), navigate to the project where you want to create the Process definition (in the **Project** view, select the respective repository and project in the drop-down lists; in the **Repository** view, navigate to **REPOSITORY/PROJECT/src/main/resources/** directory).

   > **Creating packages**
   >
   > It is recommended to create your resources, including your Process definitions, in a package of a Project to allow importing of resources and their referencing. To create a package, do the following:
   >
   > » In the **Repository** view of the Project Explorer, navigate to the **REPOSITORY/PROJECT/src/main/resources/** directory.
   >
   > » Go to **New Item → Package**.
   >
   > » In the **New resource** dialog, define the package name and check the location of the package in the repository.

3. From the perspective menu, go to **New Item → Business Process**.

4. In the **New Processes** dialog box, enter the Process name and click **OK**. Wait until the Process Editor with the Process diagram appears.

8.2. IMPORTING A PROCESS DEFINITION

To import an existing BPMN2 or JSON definition, do the following:

1. In the **Project Explorer**, select a Project and the respective package to which you want to import the Process definition.

2. Create a new Business Process to work in by going to **New Item → Business Process**.

3. In the Process Designer toolbar, click the **Import** icon in the editor toolbar and pick the format of the imported process definition. Note that you have to choose to overwrite the existing process definition in order to import.

4. From the **Import** window, locate the Process file and click **Import**.

   Figure 8.1. Import Window

Whenever a process definition is imported, the existing imported definition is overwritten. Make sure you are not overwriting a process definition you have edited so as not to lose any changes.

A process can also be imported to the git repository in the filesystem by cloning the repository, adding the process files, and pushing the changes back to git. In addition to alternative import methods, you can copy and paste a process or just open a file in the import dialog.

When importing processes, the Process Designer provides visual support for Process elements and therefore requires information on element positions on the canvas. If the information is not provided in the imported Process, you need to add it manually.

8.3. IMPORTING JPDL 3.2 TO BPMN2

To migrate and import a jPDL definition to BPMN2, in the Process Designer, click on the import button then scroll down and select **Migrate jPDL 3.2 to BPMN2**.

Figure 8.2. Migrate jPDL 3.2 to BPMN2

In the **Migrate to BPMN2** dialog box, select the process definition file and the name of the **gpd** file. Confirm by clicking the **Migrate** button.

Figure 8.3. Migrate to BPMN2 dialog box

> **Important**
>
> The migration tool for jPDL 3.2 to BPMN2 is a technical preview feature, and therefore not currently supported in Red Hat JBoss BPM Suite.

8.4. EXPORTING A PROCESS

Procedure: Exporting a business process

To export a business process, do the following:

1. Open the **Project Authoring** perspective: on the main menu, click **Authoring → Project Authoring**.

2. Select the business process which is to be exported, to view it in the Process Designer.

3. Click on the button of the process designer toolbar and select **View Process Sources** from the drop-down options.

   The **Process Sources** window is displayed.

4. Click on the **Download BPMN2** button and save the business process at the desired location.

PART III. INTEGRATION

CHAPTER 9. DEPLOYING RED HAT JBOSS BPM SUITE TO AMAZON EC2 WEB SERVICE

Amazon Elastic Compute Cloud (hereinafter referred to as Amazon EC2) is a web service providing a computing environment in the Amazon Web Services (AWS) cloud. Deploying Red Hat JBoss BPM Suite to Amazon EC2 has been supported since version 6.2.

Amazon EC2 can be used to create a virtual machine and its computing environments called instances. Templates for this instances are then known as Amazon Machine Images or AMIs. The service allows you to modify the instance to use a required configuration of CPU, memory, storage, and other properties. The cost of the service is calculated based on the used compute capacity.

For more information about Amazon EC2 and all the provided features, see the official AWS Documentation.

9.1. GETTING STARTED WITH AMAZON EC2

Logging into AWS

1. Go to the Amazon Web Services main page, click **Sign in to the Console** and enter your user credentials.

2. On the AWS home page that opens, choose **EC2 Virtual Servers in the Cloud** in the left hand corner.

   You are redirected to the **EC2 Dashboard**.

   Figure 9.1. Amazon EC2 Dashboard
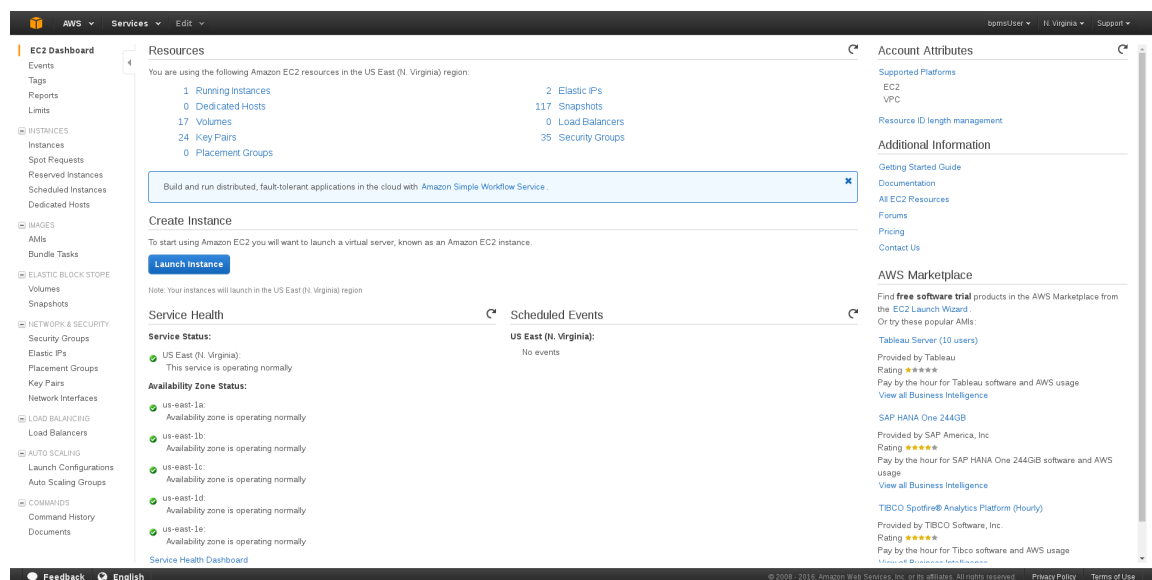
EC2 Dashboard provides an overview of your AWS account as well as the information about the overall service health and scheduled events. The **Resources** section of the dashboard contains links to all running instances, created key pairs, security groups, and more. To launch a virtual server, click the **Launch Instance** quick link in the **Create Instance** section. The main menu on the left side of the page contains links necessary for configuring instances and images, setting up networks and controlling security issues.

Important

Note that login information in Amazon EC2 is encrypted. Before creating or starting an instance, you need to import your public key to AWS in order to be able to connect to the instance.

Importing Public Key

1. In the main menu on the left, click **Key Pairs** under the **NETWORK & SECURITY** drop-down menu.

2. Click **Import Key Pair** at the top of the page.

   The **Import Key Pair** dialog window opens.

3. Define the name of the key pair and either upload the public key from file, or copy and paste the content of your public key into the **Public key contents** text field.

4. Click **Import**.

   The imported key pair appears in the displayed list of key pairs.

Alternatively, you can create your own key pair using Amazon EC2. For more information, see the Amazon EC2 Key Pairs chapter in the *AWS Documentation*.

9.2. CREATING AND STARTING VIRTUAL MACHINE INSTANCES

To deploy Red Hat JBoss BPM Suite to Amazon EC2, you must have a valid membership in the JBoss Cloud Access program. The program provides a number of supported AMIs created by Red Hat, while each of them has Red Hat Enterprise Linux 6 and Red Hat JBoss EAP 6 (together with other required software) pre-installed.

For further information about supported Amazon EC2 instance types, see the Deploy JBoss EAP 6 on Amazon EC2 chapter of *Red Hat JBoss Enterprise Application Platform 6.4 Administration and Configuration Guide*.

To create a virtual machine instance, do the following:

Creating Virtual Machine Instance

1. In the main menu, click **AMIs** under the **IMAGES** drop-down menu.

2. Choose one of the supported AMIs.

   For example **RHEL-6.7_HVM_GA-JBEAP-6.4.7-20160421-x86_64-1-Access2-GP2 (ami-50bba23a)**.

3. Click **Launch.**

   The **Choose an Instance Type** window opens. An *instance type* is a custom configuration of CPU, memory, storage, and networking capacity of your virtual machine instance.

4. Choose one of the listed instance types.

Instance type depends on the needs of your application. KIE Server and Business Central are usually two separated instances in AWS and each of them requires different settings. However, at least 2 vCPUs (the number of virtual CPUs) and 7.5 GiB of memory is recommended.

It is possible to launch the instance directly from here by clicking **Review and Launch**. If you want to change additional settings, see the following steps.

5. Click **Next: Configure Instance Details**.

   In this step, you can choose a number of launched instances, network (either EC2-Classic or VPC, the Amazon Virtual Private Cloud), or enable detailed monitoring. If you want to keep the default settings, continue to the next tab.

6. Click **Next: Add Storage**.

   This screen allows you to add more Amazon Elastic Block Store (Amazon EBS) volumes and modify the settings of the root volume. The minimum recommended size of the root volume is `30 GiB`.

   > **Important**
   >
   > Make sure that the set volume size is greater than the size of the used snapshot.

7. Click **Next: Tag Instance** and specify the tags for the instance that can help you search and filter the instances later. For example, the key **Name** can indicate a logical name of the instance in the instance list.

8. Click **Next: Configure Security Group**.

   A *security group* is a virtual firewall containing a number of firewall rules that control the traffic for your instance. Each instance can have one or more security groups assigned and these groups can be modified at any time.

   You need to create a new security group in order to be able to deploy Red Hat JBoss BPM Suite:

   Creating New Security Group

   1. If not already selected, select the **Create a new security group** check box.

   2. Fill in the **Security group name** and **Description**.

   3. Keep the default SSH type TCP protocol rule.

   4. Define three new firewall rules by clicking **Add Rule**:

      ≫ Default SSH access to Business Central:

         ▪ **Type**: Custom TCP Rule

         ▪ **Protocol**: TCP

         ▪ **Port Range**: 8001

      ≫ Default HTTP port of Red Hat JBoss EAP:

- **Type**: Custom TCP Rule

- **Protocol**: TCP

- **Port Range**: 8080

  - Default Git port:

    - **Type**: Custom TCP Rule

    - **Protocol**: TCP

    - **Port Range**: 9418

Type ⓘ	Protocol ⓘ	Port Range ⓘ	Source ⓘ
SSH	TCP	22	0.0.0.0/0
Custom TCP Rule	TCP	8001	0.0.0.0/0
Custom TCP Rule	TCP	8080	0.0.0.0/0
Custom TCP Rule	TCP	9418	0.0.0.0/0

> **Warning**
>
> All the configured ports must be allowed in the **PORTS_ALLOWED** property of user data. For more information, see Configuring User Data.

9. Click **Review and Launch**.

10. Review your instance details and if you want to launch the instance, click **Launch**.

A list of all created instances can be found under the **Instances** option in the main menu on the left. You can filter, search and launch the instances as well as view the instance description and edit the settings.

Figure 9.2. List of Instances Searched by Name

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS	Public IP	Key Name	Monitoring	Launch Time	Security Groups
bpms-kku	i-f158e66b	m3.large	us-east-1d	stopped	None				kkufova	disabled	May 24, 2016 at 11:30:43 AM...	bpms-sec-group
bpms-eap...	i-08519092	m3.large	us-east-1d	stopped	None				rsynek-key	disabled	May 10, 2016 at 10:15:11 AM...	bpms-sec-group
bpms-doc	i-d683c951	m1.large	us-east-1a	stopped	None				rsynek-key	disabled	May 10, 2016 at 9:25:41 AM...	bpms-sec-group
bpm-busin...	i-3bf0118d	m1.large	us-east-1a	stopped	None				rsynek-key	disabled	April 8, 2016 at 12:38:54 PM...	bpms-6.2-security
bpmsuite	i-395940bd	m1.large	us-east-1a	stopped	None				rsynek-key	disabled	April 5, 2016 at 8:37:37 AM ...	bpms-sec-group
bpms-6.2...	i-0c9b7aba	m1.large	us-east-1a	stopped	None				rsynek-key	disabled	November 25, 2015 at 4:20:4...	bpms-6.2-security
bpms6.2	i-d2410779	m1.large	us-east-1e	stopped	None				rsynek-key	disabled	November 25, 2015 at 8:40:0...	launch-wizard-13

In case you created a custom security group and defined new firewall rules with specific ports, you have to allow these ports in user data for the new virtual machine instance.

Configuring User Data

1. In the main menu, go to **INSTANCES → Instances**.

2. In the list of instances, find and select the instance you want to configure.

3. Click **Actions → Instance Settings → View/Change User Data**.
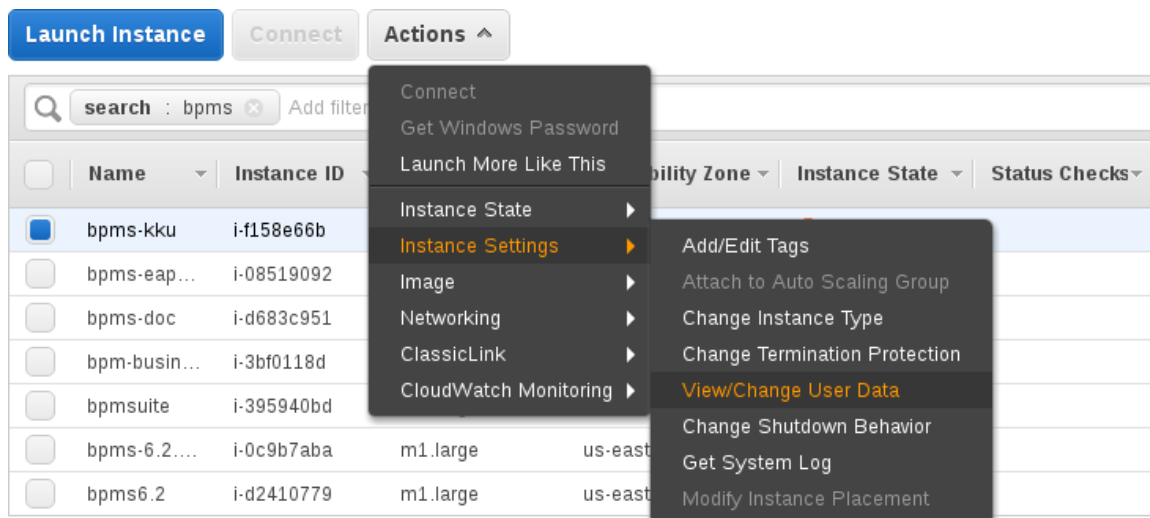
4. In the **View/Change User Data** dialog window that opens, define the **PORTS_ALLOWED** property. For example:

```
PORTS_ALLOWED="8001 8080 9418"
```


5. Click **Save**.

To start the created instance, do the following:

Starting Virtual Machine Instance

1. In the main menu, go to **INSTANCES → Instances**.

2. Select the instance from the list and click **Actions → Instance State → Start**.
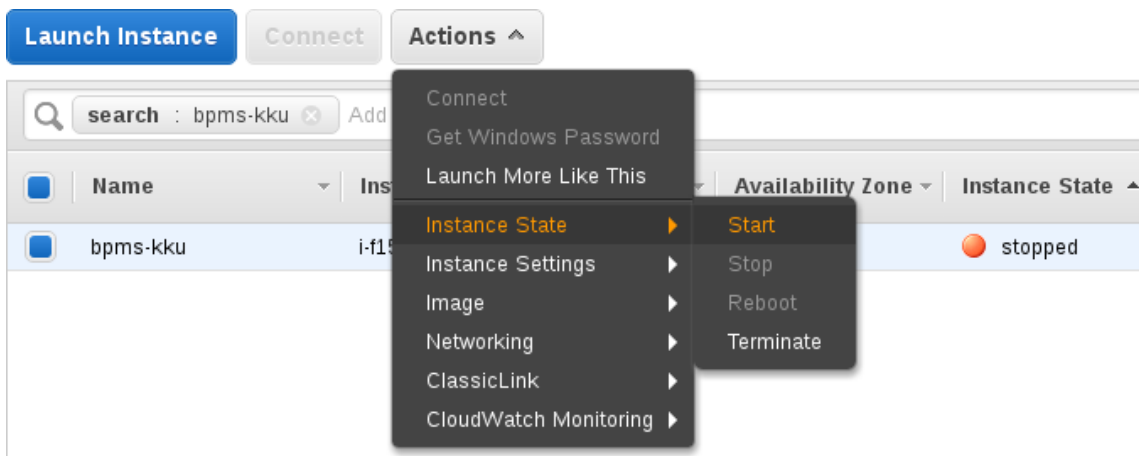
3. In the **Start Instances** dialog window that opens, click **Yes, Start** to confirm that you want to start the instance.

4. Wait until the result of a status check changes to **2/2 checks passed**.

> **Note**
>
> Do not forget to stop your instance after you no longer need it.

9.3. INSTALLING RED HAT JBOSS BPM SUITE ON VIRTUAL MACHINE INSTANCE

First, follow the steps below to connect to your instance:

Connecting to Instance

1. Start the instance.

2. Click **Connect** to view details required for accessing your instance.

   The **Connect To Your Instance** pop-up window opens.

3. Copy the public DNS address.

4. Execute the following command:

   ```
   ~]$ ssh ec2-user@PUBLIC_DNS_ADDRESS
   ```

5. Enter **yes** to confirm that you want to connect to your instance.

Installing Red Hat JBoss BPM Suite on VM Instance

1. Connect to the instance on which you want to install Red Hat JBoss BPM Suite.

2. Update your instance in order to get the latest version of Red Hat JBoss EAP 6.*x* by executing the following command:

```
~]$ sudo yum -y update
```

3. Switch to root:

```
~]$ sudo su
```

4. Copy the following configuration files to **/etc/jbossas/standalone/** and **/etc/jbossas/domain/**:

   » **standalone.xml**

   » **standalone-full.xml**

   » **standalone-full-ha.xml**

   » **standalone-ha.xml**

   » **standalone-osgi.xml**

   Follow the instructions:

   a. Download Red Hat JBoss BPM Suite 6.3.0 Deployable for EAP 6 from Customer Portal.

   b. Change directory to the location of the downloaded zip file.

   c. Copy the zip file into the **/tmp/** directory of your virtual machine instance:

   ```
   ~]$ scp jboss-bpmsuite-6.3.0.GA-deployable-eap6.x.zip
   ec2-user@PUBLIC_DNS_ADDRESS:/tmp
   ```

   d. Go to the **/tmp/** directory of your VM instance and unzip the file:

   ```
   ~]# unzip jboss-bpmsuite-6.3.0.GA-deployable-eap6.x.zip
   ```

   e. Change directory to **jboss-eap-6.4/standalone/configuration/** and copy all the configuration files listed above into **/etc/jbossas/standalone/** and **/etc/jbossas/domain/**, for example:

   ```
   ~]# cp standalone.xml /etc/jbossas/standalone/
   ```

   Enter **y** to overwrite the file.

   f. Make sure that the files were overwritten: **standalone.xml** should contain the following:

   ```
   <system-properties>
     <property name="org.kie.example" value="true"/>
     <property name="org.jbpm.designer.perspective"
   value="full"/>
     <property name="designerdataobjects" value="false"/>
   </system-properties>
   ```

5. Copy the required applications (for example Business Central) to **/usr/share/jbossas/standalone/deployments/**:

a. Change directory to **/tmp/jboss-eap-6.4/standalone/deployments/**.

b. Copy Business Central:

```
~]# cp -r business-central.war
/usr/share/jbossas/standalone/deployments/
```

```
~]# cp -r business-central.war.dodeploy
/usr/share/jbossas/standalone/deployments/
```

6. Set the proper Red Hat JBoss EAP binding address.

Each virtual machine has private and public IP addresses. The public IP address can be accessed outside the AWS, while the private IP address is accessible only inside the virtual network of your Amazon EC2 account.

> **Important**
>
> Every time you start the instance, the addresses change. Refer to the Red Hat JBoss Enterprise Application Platform 6.4 Administration and Configuration Guide for more information.

One way to set the address is by doing the following:

a. Run **ifconfig** and copy the **inet addr** address (for example **10.233.159.148**).

b. Open **/etc/jbossas/standalone/standalone.xml** and set the proper binding address:

```
<interfaces>
  ...
  <interface name="public">
    <inet-address value="10.233.159.148"/>
  </interface>
  ...
</interfaces>
```

7. Enable Git and SSH access to Business Central: add the following system properties (with *your* private IP address) into **/etc/jbossas/standalone/standalone.xml**.

```
<property name="org.uberfire.nio.git.daemon.host"
value="10.233.159.148"/>
<property name="org.uberfire.nio.git.ssh.host"
value="10.233.159.148"/>
```

Once the setup has been done, you can create a new AMI from your instance.

Creating new AMI from Configured Instance

1. In the main menu, go to **INSTANCES → Instances**.

2. Locate your configured instance and click **Actions** → **Image** → **Create Image**.
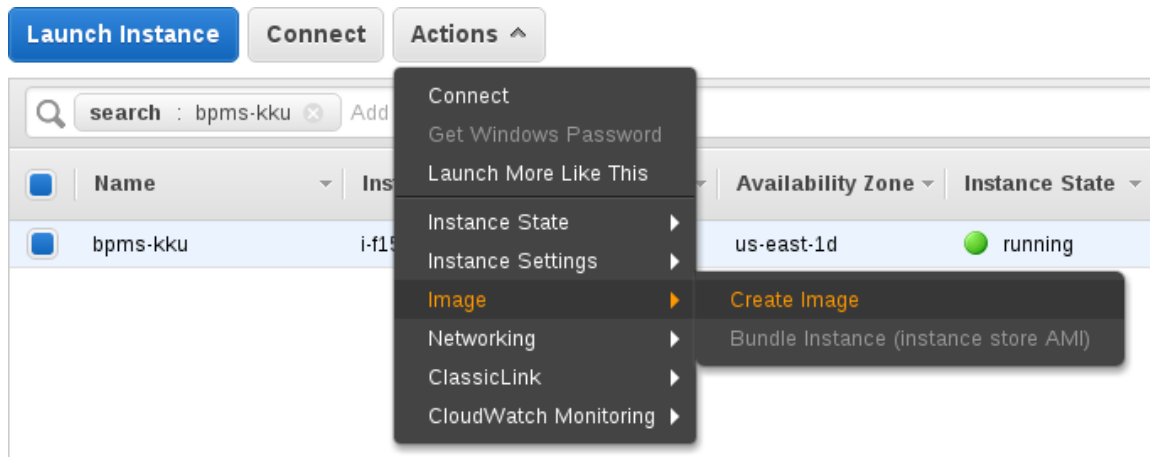
3. In the **Create Image** dialog window, specify the image details and click **Create Image**.

9.4. RUNNING RED HAT JBOSS BPM SUITE ON VIRTUAL MACHINE INSTANCE

Running Red Hat JBoss BPM Suite on VM Instance

1. Start Red Hat JBoss EAP by executing the command:

```
~]# /etc/init.d/jbossas start
```

Alternatively, you can use the command with option **stop** (or **restart**) to stop (or restart) the Red Hat JBoss EAP.

The following message appears:

```
Starting jbossas:                                            [
OK  ]
```

If the deployment finished successfully, the **business-central.war.dodeploy** file in **/usr/share/jbossas/standalone/deployments/** changes to **business-central.war.deployed**.

2. To log in to Business Central, navigate to **http://*PUBLIC_DNS_ADDRESS*:8080/business-central** in a web browser.

CHAPTER 10. DEPLOYING RED HAT JBOSS BPM SUITE ARTIFACTS TO SOA REPOSITORY ARTIFACT MODEL AND PROTOCOL (S-RAMP) REPOSITORY

While Red Hat JBoss BPM Suite and S-RAMP are two independent products, it is possible to move artifacts between them. You can move artifacts from JBoss BPM Suite to S-RAMP using Maven or via a user interface.

This section provides information about these two processes.

10.1. DEPLOYING RED HAT JBOSS BPM SUITE ARTIFACTS TO SOA REPOSITORY ARTIFACT MODEL AND PROTOCOL (S-RAMP) USING MAVEN

Before you can deploy Red Hat JBoss BPM Suite artifacts to S-RAMP using Maven, you will need to enable the S-RAMP Maven Wagon. The Maven Wagon is a key feature that supports the S-RAMP Atom based REST API protocol. By enabling the S-RAMP Maven Wagon, users will be able to access artifacts from the S-RAMP repository as dependencies in a Maven project.

Enable the S-RAMP Maven Wagon by making an edit in the **pom.xml** file as shown below:

```
<build>
  <extensions>
    <extension>
      <groupId>org.overlord.sramp</groupId>
      <artifactId>s-ramp-wagon</artifactId>
      <version>${s-ramp-wagon.version}</version>
    </extension>
  </extensions>
</build>
```

Once the S-RAMP Maven Wagon is enabled, you can deploy the JBoss BPM Suite artifacts to that S-RAMP repository. To do this, follow the steps below:

1. Clone the git repository where you have saved the BPM Suite project by running this command:

   ```
   git clone http://localhost:8001/REPOSITORY_NAME
   ```

2. On the command line, move into the folder that contains the project.

3. Follow the instructions in *Red Hat JBoss Fuse Service Works 6 Development Guide, Volume 3: Governance*, section Deploying to S-RAMP. Use the URL from the example below:

   ```
   <distributionManagement>
     <repository>
       <id>local-sramp-repo</id>
       <name>S-RAMP Releases Repository</name>
       <url>sramp://S-RAMP_SERVER_URL/s-ramp-server/</url>
     </repository>
   ```

```
      <snapshotRepository>
        <id>local-sramp-repo-snapshots</id>
        <name>S-RAMP Snapshots Repository</name>
        <url>sramp://S-RAMP_SERVER_URL/s-ramp-server/</url>
      </snapshotRepository>
    </distributionManagement>
```

With these settings, Maven deployments are sent directly to the S-RAMP repository using the S-RAMP API. Note that artifacts are added to the S-RAMP repository with an artifact type based on the Maven type of the project. You can override this behavior by adding a query parameter to the repository URL in the **pom.xml** file. For example:

```
<distributionManagement>
  <repository>
    <id>local-sramp-repo</id>
    <name>S-RAMP Releases Repository</name>
    <url>sramp://S-RAMP_SERVER_URL/s-ramp-server/?
artifactType=KieJarArchive</url>
  </repository>
</distributionManagement>
```

The above example causes the Maven artifact to be uploaded with an S-RAMP artifact type of KieJarArchive.

4. Amend the maven plug-in in file **pom.xml** and add a dependency to it as follows in case the project does not contain decision tables:

```
<plugins>
    <plugin>
      <groupId>org.kie</groupId>
      <artifactId>kie-maven-plugin</artifactId>
      <version>6.4.0.Final-redhat-3</version>
      <extensions>true</extensions>
      <dependencies>
          <dependency>
              <groupId>org.jbpm</groupId>
              <artifactId>jbpm-bpmn2</artifactId>
              <version>6.4.0.Final-redhat-3</version>
          </dependency>
      </dependencies>
     </plugin>
    </plugins>
```

If the project contains decision tables, use this dependency for the kie-maven-plugin instead:

```
<plugins>
    <plugin>
      <groupId>org.kie</groupId>
      <artifactId>kie-maven-plugin</artifactId>
      <version>6.4.0.Final-redhat-3</version>
      <extensions>true</extensions>
      <dependencies>
          <dependency>
              <groupId>org.drools</groupId>
```

```
                <artifactId>drools-decisiontables</artifactId>
                <version>6.4.0.Final-redhat-3</version>
            </dependency>
        </dependencies>
      </plugin>
    </plugins>
```

5. Run a clean Maven deployment using the following command:

```
mvn -s sramp-settings.xml deploy
```

Note

For the Maven deployment to the S-RAMP repository, it is necessary to have credentials set in the `settings.xml` file. For further details on the credentials, refer to *Red Hat JBoss Fuse Service Works* (FSW) documentation on Authentication.

10.2. DEPLOYING RED HAT JBOSS BPM SUITE ARTIFACTS TO SOA REPOSITORY ARTIFACT MODEL AND PROTOCOL (S-RAMP) USING GRAPHICAL USER INTERFACE (GUI)

To deploy Red Hat JBoss BPM Suite artifacts to an S-RAMP repository using the user interface, do the following:

1. In a web browser, navigate to http://localhost:8080/s-ramp-ui/. If the user interface has been configured to run from a domain name, substitute `localhost` for the domain name. For example http://www.example.com:8080/s-ramp-ui/.

2. Click on **Artifacts**.

3. In the **Manage Artifacts** section, select **Import**.

4. Locate the kie archive you want to deploy. In the dialog that opens, fill out **KieJarArchive** as the type, and select **Import**.

5. The deployment then creates these entries in the S-RAMP repository:

   KieJarArchive, from which it derives:

   » **KieXmlDocument** (if the archive contains **kmodule.xml**)

   » **BpmnDocument** (if the archive contains **bpmn** definitions)

   » **DroolsDocument** (if the archive contains **drl** definitions)

CHAPTER 11. INTEGRATING RED HAT JBOSS BPM SUITE WITH RED HAT JBOSS FUSE

Red Hat JBoss Fuse integration allows users of Red Hat JBoss Fuse to complement their integration solution with additional features provided by Red Hat JBoss BPM Suite and Red Hat JBoss BRMS.

Red Hat JBoss BPM Suite integration is provided by two **features.xml** files:

- **drools-karaf-features-VERSION-features.xml**

  This file provides core Red Hat JBoss BPM Suite and Red Hat JBoss BRMS features, which defines the OSGi features that can be deployed into Red Hat JBoss Fuse. This file is a part of the Red Hat JBoss BPM Suite and Red Hat JBoss BRMS product. OSGi users can install features from this file in order to install Red Hat JBoss BRMS engine or Red Hat JBoss BPM Suite engine into Red Hat JBoss Fuse and use it in their applications.

- **karaf-features-VERSION-features.xml**

  This file provides additional features used for integrating Red Hat JBoss BPM Suite and Red Hat JBoss BRMS with Apache Camel, primarily in Red Hat JBoss Fuse. This file is part of the Integration Pack and it defines OSGi features that enable integration with Apache Camel and SwitchYard. In addition to the **karaf-features** XML file, the Integration Pack also contains a **features.xml** file for quick starts.

For further information about integration of Red Hat JBoss BPM Suite with Red Hat JBoss Fuse, see the Install Integration Pack chapter of the *Red Hat JBoss Fuse Integration Guide*.

11.1. CORE RED HAT JBOSS BPM SUITE AND RED HAT JBOSS BRMS FEATURES

Core Red Hat JBoss BPM Suite and Red Hat JBoss BRMS features are provided by the **drools-karaf-features-VERSION-features.xml** file present in your product Maven repository or the **jboss-brms-bpmsuite-VERSION-redhat-VERSION-fuse-features.zip** file. It provides the following features:

- **drools-common**

- **drools-module**

- **drools-templates**

- **drools-decisiontable**

- **drools-jpa**

- **kie**

- **kie-ci**

- **kie-spring**

- **kie-aries-blueprint**

- **jbpm-commons**

- **jbpm-human-task**

- **jbpm**

- **droolsjbpm-hibernate**

- **h2**

The following table provides example of use cases for some of the features listed above.

Table 11.1. Features and Use Case Examples

Feature	Use Case
drools-module	Use the Red Hat JBoss BRMS engine for rules evaluation, without requiring persistence, processes, or decision tables.
drools-jpa	Use the Red Hat JBoss BRMS engine for rules evaluation with persistence and transactions, but without requiring processes or decision tables. The **drools-jpa** feature already includes **drools-module**, however you may also need to install the **droolsjbpm-hibernate** feature, or ensure there is a compatible hibernate bundle installed.
drools-decisiontable	Use the Red Hat JBoss BRMS engine with decision tables.
jbpm	Use the Red Hat JBoss BPM Suite (or Red Hat JBoss BRMS engine with processes). The **jbpm** feature already includes **drools-module** and **drools-jpa**. You may also need to install the **droolsjbpm-hibernate** feature, or ensure that there is a compatible hibernate bundle installed.
jbpm and **jbpm-human-task**	Use the Red Hat JBoss BPM Suite (or Red Hat JBoss BRMS engine with processes) with Human Task.
core engine JARs and **kie-ci**	Use Red Hat JBoss BRMS or Red Hat JBoss BPM Suite with **KieScanner** (KIE-CI) to download kJARs from a Maven repository.

Feature	Use Case
`kie-spring`	Use KIE-Spring integration. See the section called "kie-spring Feature Further Information" for more information.
`kie-spring` and `kie-aries-blueprint`	Use KIE-Aries-Blueprint integration.

kie-spring Feature Further Information

- Use `org.drools.osgi.spring.OsgiKModuleBeanFactoryPostProcessor` instead of `org.kie.spring.KModuleBeanFactoryPostProcessor` to postprocess KIE elements in an OSGi environment.

- Do not install the `drools-module` feature before the `kie-spring` feature. Otherwise, the `drools-compiler` bundle does not detect packages exported by `kie-spring`.

  Run *osgi:refresh drools-compiler_bundle_ID* if you have installed the features in the incorrect order to force `drools-compiler` to rebuild its Import-Package metadata.

11.2. ADDITIONAL FEATURES FOR SWITCHYARD AND APACHE CAMEL INTEGRATION

The following additional features for integration with SwitchYard and Apache Camel on Red Hat JBoss Fuse are provided by the integration pack:

- `fuse-bxms-switchyard-common-knowledge`

- `fuse-bxms-switchyard-rules`

- `fuse-bxms-switchyard-bpm`

- `kie-camel`

- `jbpm-workitems-camel`

The integration pack features are defined in the `karaf-features-VERSION-features.xml` file. This file (and supporting repositories) is located in http://repository.jboss.org/nexus/content/repositories/public, which is already configured for use on Red Hat JBoss Fuse 6.2 out of the box in *INSTALLATION_DIRECTORY*`/etc/org.ops4j.pax.url.mvn.cfg`.

The file can also be downloaded from either the Red Hat JBoss Fuse 6.2 or Red Hat JBoss BPM Suite product page in the Red Hat Customer Portal.

11.3. INSTALLING AND UPDATING CORE INTEGRATION FEATURES

> **Note**
>
> This section refers to features in the **drools-karaf-features-*VERSION*-features.xml** file. For additional integration features, see Section 11.4, "Installing Additional Integration Features".

If you have already installed an older version of the core Red Hat JBoss BPM Suite and Red Hat JBoss BRMS features (for example, **drools-karaf-features-6.2.0.Final-redhat-6-features.xml**), you need to remove them and all associated files before installing the most recent **features.xml** file.

Procedure: Removing Existing drools-karaf-features Installation

1. Start the Red Hat JBoss Fuse console using:

   ```
   $ ./INSTALLATION_DIRECTORY/bin/fuse
   ```

2. Unistall old features or applications that used the previous **features.xml** file. For example:

   ```
   JBossFuse:karaf@root> features:uninstall drools-module
   JBossFuse:karaf@root> features:uninstall jbpm
   JBossFuse:karaf@root> features:uninstall kie-ci
   ```

3. Search for references of bundles using **drools**, **kie**, or **jbpm**, and remove them:

   ```
   karaf@root> list -t 0 -s | grep drools
   karaf@root> list -t 0 -s | grep kie
   karaf@root> list -t 0 -s | grep jbpm
   ```

   To remove the bundles:

   ```
   karaf@root> osgi:uninstall BUNDLE_ID
   ```

4. Remove the old **drools-karaf-features** URL:

   ```
   karaf@root> features:removeurl mvn:org.drools/drools-karaf-
   features/6.2.0.Final-redhat-VERSION/xml/features
   ```

5. Restart Red Hat JBoss Fuse.

To install the **drools-karaf-features**:

Procedure: Installing Core Red Hat JBoss BPM Suite and Red Hat JBoss BRMS Features

1. Configure required repositories:

   » Edit the ***INSTALLATION_DIRECTORY*/etc/org.ops4j.pax.url.mvn.cfg** file in your Red Hat JBoss Fuse installation and add the following entry to the **org.ops4j.pax.url.mvn.repositories** variable (note that entries are separated by **,** \):

- **https://maven.repository.redhat.com/ga/**

2. Start Red Hat JBoss Fuse:

```
$ ./INSTALLATION_DIRECTORY/bin/fuse
```

3. Add a reference to the core features file by running the following console command:

```
JBossFuse:karaf@root> features:addurl mvn:org.drools/drools-
karaf-features/VERSION/xml/features
```

For example:

```
features:addurl mvn:org.drools/drools-karaf-
features/6.4.0.Final-redhat-3/xml/features
```

4. You can now install the features provided by this file by running, for example, the following console command:

```
JBossFuse:karaf@root> features:install jbpm
```

11.4. INSTALLING ADDITIONAL INTEGRATION FEATURES

Use the following procedure for additional integration with SwitchYard and Apache Camel.

Procedure: SwitchYard and Apache Camel Integration

1. Download the **fuse-integration** package that is aligned with your version of Red Hat JBoss Fuse.

   Note

   For instance, if you want to use the **6.2.0.redhat-117** version of Red Hat JBoss Fuse, you need to install the **fuse-6.2.0.redhat-117** Red Hat JBoss Fuse integration features.

2. Add the remote Maven repository that contains the fuse dependencies to your **karaf** instance:

   - Edit *FUSE_HOME*/etc/org.ops4j.pax.url.mvn.cfg

3. Update the Drools features URL:

```
JBossFuse:karaf@root> features:addurl
mvn:org.switchyard.karaf/mvn:org.switchyard.karaf/switchyard/S
WITCHYARD_VERSION/xml/core-features
JBossFuse:karaf@root> features:addurl mvn:org.drools/drools-
karaf-features/DROOLS_VERSION/xml/features
JBossFuse:karaf@root> features:addurl
mvn:org.jboss.integration.fuse/karaf-features/1.0.0.redhat-
VERSION/xml/features
```

4. You can now install the features provided for SwitchYard and Apache Camel integration by running, for example, the following console command:

```
JBossFuse:karaf@root> features:install fuse-bxms-switchyard-
rules
JBossFuse:karaf@root> features:install kie-camel
JBossFuse:karaf@root> features:install jbpm-workitems-camel
```

11.5. CONFIGURING DEPENDENCIES

When you configure KIE, Red Hat JBoss BRMS, or Red Hat Jboss BPM Suite in your application, you can follow one of the following approaches to build your OSGi application bundles:

➤ Bundle required dependencies into your application bundle. In this approach, you declare all required artifacts as runtime dependencies in your **pom.xml**. Hence, you need not import the packages that provide these artifacts that you have already added as dependencies.

➤ Import the required dependencies into the application bundle. This is a preferred approach for building OSGi bundles as it adheres to the principles of OSGi framework. In this approach, you declare only the API JARs (such as **org.kie:kie-api**) as dependencies in your application bundle. You will need to install the required Red Hat JBoss BRMS and Red Hat JBoss BPM Suite bundles and then import them in your application.

> **Warning**
>
> The **MVELUserGroupCallback** class fails to initialize in an OSGi environment.
> Do *not* use or include **MVELUserGroupCallback** as it is not designed for
> production purposes.

11.6. INSTALLING RED HAT JBOSS FUSE INTEGRATION QUICK START APPLICATIONS

The following features for Red Hat JBoss Fuse integration quick start applications are provided by **org/jboss/integration/fuse/quickstarts/karaf-features/1.0.0.redhat-*VERSION*/karaf-features-1.0.0.redhat-*VERSION*-features.xml**:

➤ **fuse-bxms-switchyard-quickstart-bpm-service**

➤ **fuse-bxms-switchyard-quickstart-rules-camel-cbr**

➤ **fuse-bxms-switchyard-quickstart-rules-interview**

➤ **fuse-bxms-switchyard-quickstart-rules-interview-container**

➤ **fuse-bxms-switchyard-quickstart-rules-interview-dtable**

➤ **fuse-bxms-switchyard-demo-library**

➤ **fuse-bxms-switchyard-demo-helpdesk**

- **`fuse-bxms-camel-blueprint-drools-decision-table`**

- **`fuse-bxms-camel-spring-drools-decision-table`**

- **`fuse-bxms-jbpm-workitems-camel-quickstart`**

- **`fuse-bxms-spring-jbpm-osgi-example`**

This file (and supporting repositories) is located in
http://repository.jboss.org/nexus/content/repositories/public, which is already configured for
use on Red Hat JBoss Fuse 6.2 out of the box in
***INSTALLATION_DIRECTORY*`/etc/org.ops4j.pax.url.mvn.cfg`**.

Procedure: Installing Quick Start Applications

1. Add a reference to the features file by running the following console command:

   ```
   JBossFuse:karaf@root> features:addurl
   mvn:org.jboss.integration.fuse.quickstarts/karaf-
   features/1.0.0.redhat-VERSION/xml/features
   ```

2. You can now install the quick start applications provided by this features file by
   running, for example, the following console command:

   ```
   JBossFuse:karaf@root> features:install fuse-bxms-switchyard-
   quickstart-bpm-service
   ```

Procedure: Downloading and Installing Quick Start ZIP Files

1. Download the quick start application ZIP file.

2. Unpack the contents of the quick starts directory into your existing
   ***INSTALLATION_DIRECTORY*`/quickstarts`** directory.

3. Unpack the contents of the system directory into your existing
   ***INSTALLATION_DIRECTORY*`/system`** directory.

11.6.1. Testing Your First Quick Start Application

Procedure: Testing Quick Start Application

1. Start Red Hat JBoss Fuse:

   ```
   $ ./INSTALLATION_DIRECTORY/bin/fuse
   ```

2. Install and start the **`switchyard-bpm-service`** by running the following console
   command:

   ```
   JBossFuse:karaf@root> features:install fuse-bxms-switchyard-
   quickstart-bpm-service
   ```

> **Note**
>
> Any dependent features specified by the application's features file will be installed automatically.

3. Submit a web service request to invoke the SOAP gateway.

   a. Open a terminal window and navigate to the associated quick start directory that was unpacked from the quick start application ZIP file (in this case, **switchyard-bpm-service**).

   b. Run the following command:

   ```
   $ mvn clean install
   ```

   > **Note**
   >
   > You will need the following repositories configured in your **settings.xml** file:
   >
   > » **http://maven.repository.redhat.com/ga/**
   >
   > » **http://repository.jboss.org/nexus/content/repositories/public/**

   c. Run the following command:

   ```
   $ mvn exec:java -Pkaraf
   ```

4. You will receive the following response:

```
SOAP Reply:
<soap:Envelope
xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <SOAP-ENV:Header xmlns:SOAP-
ENV="http://schemas.xmlsoap.org/soap/envelope/"/>
  <soap:Body>
    <ns2:submitOrderResponse xmlns:ns2="urn:switchyard-
quickstart:bpm-service:1.0">
      <orderId>test1</orderId>
      <accepted>true</accepted>
      <status>Thanks for your order, it has been
shipped!</status>
    </ns2:submitOrderResponse>
  </soap:Body>
</soap:Envelope>
```

CHAPTER 12. INTEGRATING RED HAT JBOSS BPM SUITE WITH RED HAT SINGLE SIGN-ON

Red Hat Single Sign-On (RH-SSO) is a Single Sign-On solution that you can use for securing your browser applications and your REST web services. This chapter describes how you can integrate RH-SSO with Red Hat JBoss BPM Suite and leverage its features.

Integrating with RH-SSO brings an integrated SSO and IDM (Identity Management) environment for Red Hat JBoss BPM Suite. The session management feature of RH-SSO allows you to use different Red Hat JBoss BPM Suite environments on the web by authenticating only once.

For more information on RH-SSO, see the RH-SSO documentation.

Red Hat Single Sign On Integration Points

You can integrate RH-SSO with Intelligent Process Servers using the following integration points:

> **Business Central authentication through an RH-SSO server**

Authenticating Red Hat JBoss BPM Suite Business Central through RH-SSO involves securing both the Business Central web client and remote services through RH-SSO. This integration enables you to connect to Business Central using either web interface or a remote service consumer through RH-SSO.

> **Intelligent Process Server authentication through an RH-SSO server**

Authenticating Red Hat JBoss BPM Suite Intelligent Process Server through RH-SSO involves securing the remote services provided by the Intelligent Process Server as it does not provide a web interface for server authentication. This enables any remote Red Hat JBoss BPM Suite service consumer (user or a service) to authenticate through RH-SSO.

> **Third-party client authentication through an RH-SSO server**

Authenticating a third-party client through an RH-SSO server involves third-party clients to authenticate themselves using RH-SSO to consume the remote service endpoints provided by Business Central and Intelligent Process Server.

The following sections describe how to achieve RH-SSO integration through these integration points:

12.1. BUSINESS CENTRAL AUTHENTICATION THROUGH RH-SSO

To authenticate Business Central through RH-SSO:

1. Set up and run an RH-SSO server with a realm client for Business Central.

2. Install and set up the RH-SSO client adapter for EAP.

3. Secure Business Central Remote Service using RH-SSO.

4. Create an RH-SSO client and configure the RH-SSO client adapter for Dashbuilder (BAM).

12.1.1. Setting Up RH-SSO with Realm Client for Business Central

Security realms are used to restrict access for the different applications resources. It is advised to create a new realm whether your RH-SSO instance is private or shared amongst other products. You can keep the master realm as a place for super administrators to create and manage the realms in your system. If you are integrating with an RH-SSO instance that is shared with other product installations to achieve Single Sign-On with those applications, all those applications must use the same realm.

Here is how you can install an RH-SSO server and create a security realm for Business Central:

Procedure: Setting Up RH-SSO with Realm Client

1. Install and configure a basic RH-SSO standalone server. To do this, follow the instructions in the RH-SSO Installation Guide.

   > **Note**
   >
   > If you want to run both RH-SSO and Red Hat JBoss BPM Suite servers on the same machine, ensure that you avoid port conflicts. To do so, do one of the following:
   >
   > ❧ Update the **RHSSO_HOME/standalone/configuration/standalone.xml** file and set a port offset to 100. For example:
   >
   > ```
   > <socket-binding-group name="standard-sockets"
   > default-interface="public" port-
   > offset="${jboss.socket.binding.port-offset:100}">
   > ```
   >
   > ❧ Use an environment variable to run the server:
   >
   > ```
   > bin/standalone.sh -Djboss.socket.binding.port-
   > offset=100
   > ```

2. Start the RH-SSO server using the following command from **RHSSO_HOME/bin**:

   ```
   ./standalone.sh
   ```

   Once the RH-SSO server starts, open http://*localhost*:8180/auth/admin in a web browser and log in using your admin credentials that you created while installing RH-SSO. When you login for the first time, you can set up the initial user using the new user registration form.

3. On the RH-SSO admin console, click **Realm Settings** tab.

4. On the **Realm Settings** page, click **Add Realm**.

   The **Add realm** page opens.

5. On the **Add realm** page, provide a name for the realm and click **Create**.

6. Click **Client** tab from the main admin console menu and click **Create**.

The **Add Client** page opens.

7. On the **Add Client** page, provide the required information to create a new client for your realm. For example:

   » **Client ID**: kie

   » **Client protocol**: openid-connect

   » **Root URL**: http://*localhost*:8080/business-central

8. Click **Save** to save your changes.

   At this point, the RH-SSO server is configured with a realm with a client for Red Hat JBoss BPM Suite applications (Business Central, in this example) and running and listening for HTTP connections at `localhost:8180`. This realm provides different users, roles, and sessions for the Red Hat JBoss BPM Suite applications.

12.1.2. Setting Up RH-SSO Client Adapter for EAP

To set up the RH-SSO client adapter for EAP:

1. Install the RH-SSO adapter for EAP.

2. Configure the Red Hat JBoss BPM Suite application and the RH-SSO client adapter.

Procedure: Installing the RH-SSO Adapter for EAP

1. Install EAP 6.4.*X*. See Chapter 2. Installation Instructions from the *Red Hat JBoss Enterprise Application Platform Installation Guide*.

2. Install Red Hat JBoss BPM Suite in the freshly installed JBoss EAP. This step is important because if you configure RH-SSO adapter by making changes in `standalone.xml`, and then unzip Red Hat JBoss BPM Suite, you may overwrite and lose the RH-SSO adapter configuration.

3. Download the EAP adapter from the Red Hat Customer Portal.

4. Unzip and install the adapter. For installation instructions, see Installing Adapters for Red Hat Single Sign-On in the *RH-SSO Installation Guide*.

Procedure: Configuring the RH-SSO Adapter

1. Navigate to ***EAP_HOME*`/standalone/configuration`** in your EAP installation and edit `standalone.xml` to add the RH-SSO subsystem configuration. For example:

```
<subsystem xmlns="urn:jboss:domain:keycloak:1.1">
 <secure-deployment name="business-central.war">
   <realm>demo</realm>
   <realm-public-
key>MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCrVrCuTtArbgaZzL1hvh0x
tL5mc7o0NqPVnYXkLvgcwiC3BjLGw1tGEGoJaXDuSaRllobm53JBhjx33UNv+5z/U
MG4kytBWxheNVKnL6GgqlNabMaFfPLPCF8kAgKnsi79NMo+n6KnSY8YeUmec/p2v
jO2NjsSAVcWEQMVhJ31LwIDAQAB</realm-public-key>
   <auth-server-url>http://localhost:8180/auth</auth-server-url>
```

```
    <ssl-required>external</ssl-required>
    <enable-basic-auth>true</enable-basic-auth>
    <resource>kie</resource>
    <credential name="secret">759514d0-dbb1-46ba-b7e7-
ff76e63c6891</credential>
    <principal-attribute>preferred_username</principal-attribute>
  </secure-deployment>
</subsystem>
```

Here,

- **secure-deployment name**: Name of your application's WAR file.

- **realm**: Name of the realm that you created for the applications to use.

- **realm-public-key**: The public key of the realm you created. You can find it in the *Keys* tab in the *Realm settings* page of the realm you created in the RH-SSO admin console. If you do not provide a value for realm-public-key, the server retrieves it automatically.

- **auth-server-url**: The URL for the RH-SSO authentication server.

- **enable-basic-auth**: The setting to enable basic authentication mechanism, so that the clients can use both token-based and basic authentication approaches to perform the requests.

- **resource**: The name for the client that you created.

- **credential name**: The secret key for the client you created. You can find it in the *Credentials* tab on the *Clients* page of the RH-SSO admin console.

- **principal-attribute**: The login name of the user. If you do not provide this value, your User Id is displayed in the application instead of your user name.

> **Note**
>
> The RH-SSO server converts the user names to lowercase. Therefore, after integration with RH-SSO, your user name will appear in lowercase in Business Central. If you have user names in upper-case letters hard coded in business processes, the application may not be able to identify the upper-case user.

2. Add the following sub element under the **<extensions>** section of **standalone.xml**:

```
<extension module="org.keycloak.keycloak-adapter-subsystem"/>
```

This secures the application using the urn:jboss:domain:keycloak subsystem in **standalone.xml** as opposed to securing it through **web.xml** file inside each WAR.

3. Navigate to *EAP_HOME***/bin/** and start the EAP server using the following command:

```
./standalone.sh
```

You can now login to your Red Hat JBoss BPM Suite application (in this example, Business Central) once the server is running using the RH-SSO admin user credentials.

> **Note**
>
> You can also configure RH-SSO adapter for EAP by updating your applications WAR file to use the RH-SSO security subsystem. However, the recommended approach is configuring the adapter through the RH-SSO subsystem. This means that you are updating EAP configuration instead of applying the configuration on each WAR file.

12.1.3. Adding a New User

To add new users and assign them a role to access Business Central:

1. Login to the RH-SSO Admin console and open the realm you wish to add a user to.

2. Click **Users** under **Manage** section in the left menu bar.

   An empty user list page called **Users** opens.

3. Click the **Add User** button on the empty user list to start creating your new user.

   An **Add user** page opens.

4. Provide user information on the **Add user** page and click **Save**.

5. Set new password under the **Credentials** tab.

6. Assign the new user one of the roles that allow access to Business Central (For example, `admin` or `analyst` role for Red Hat JBoss BRMS and `admin` or `developer` role for Red Hat JBoss BPM Suite).

   Define the roles as realm roles in the **Roles** page under the **Configure** section.

7. Click **Role Mappings** tab on the **Users** page to assign roles.

12.1.4. Securing Business Central Remote Service using RH-SSO

Business Central provides different remote service endpoints that can be consumed by third-party clients using remote API. To authenticate those services through RH-SSO, you must disable a security filter called **BasicAuthSecurityFilter**. To do this, follow these steps:

Procedure: Disabling BasicAuthSecurityFilter

1. Open your application deployment descriptor file (`WEB-INF/web.xml`) and apply the following changes to it:

   ≫ Remove the following lines to remove the servlet filter and its mapping for class `org.uberfire.ext.security.server.BasicAuthSecurityFilter`:

   ```
   <filter>
     <filter-name>HTTP Basic Auth Filter</filter-name>
     <filter-
   ```

```
class>org.uberfire.ext.security.server.BasicAuthSecurityFilter
</filter-class>
  <init-param>
    <param-name>realmName</param-name>
    <param-value>KIE Workbench Realm</param-value>
  </init-param>
</filter>

<filter-mapping>
  <filter-name>HTTP Basic Auth Filter</filter-name>
  <url-pattern>/rest/*</url-pattern>
  <url-pattern>/maven2/*</url-pattern>
  <url-pattern>/ws/*</url-pattern>
</filter-mapping>
```

Add the following lines to add the **security-constraint** for the url-patterns that you have removed from the filter mapping:

```
<security-constraint>
  <web-resource-collection>
    <web-resource-name>remote-services</web-resource-name>
    <url-pattern>/rest/*</url-pattern>
    <url-pattern>/maven2/*</url-pattern>
    <url-pattern>/ws/*</url-pattern>
  </web-resource-collection>
  <auth-constraint>
    <role-name>rest-all</role-name>
  </auth-constraint>
</security-constraint>
```

2. Save your changes.

12.1.5. Creating and Configuring RH-SSO Client Adapter for BAM

Once you have created a realm client for Business Central and set up the RH-SSO client adapter for EAP, you can repeat the same steps to integrate BAM with RH-SSO. After successful integration of RH-SSO with BAM, you can navigate within Business Central and BAM without having to re-login.

Procedure: Integrating RH-SSO with BAM

1. On the RH-SSO admin console, open the security realm that you created.

2. Click **Client** tab from the main admin console menu and click **Create**.

   The **Add Client** page opens.

3. On the **Add Client** page, provide the required information to create a new BAM client for your realm. For example:

   - **Client ID**: dashbuilder

   - **Root URL**: http://*localhost*:8080/dashbuilder

   - **Client protocol**: openid-connect

4. Configure the RH-SSO client adapter for BAM. To do so, navigate to
   EAP_HOME/**standalone/configuration** in your EAP installation and edit
   standalone.xml to add the RH-SSO subsystem configuration. For example:

```
<secure-deployment name="dashbuilder.war">
  <realm>demo</realm>
  <realm-public-
key>MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCrVrCuTtArbgaZzL1hvh0x
tL5mc7o0NqPVnYXkLvgcwiC3BjLGw1tGEGoJaXDuSaRllobm53JBhjx33UNv+5z/U
MG4kytBWxheNVKnL6GgqlNabMaFfPLPCF8kAgKnsi79NMo+n6KnSY8YeUmec/p2v
jO2NjsSAVcWEQMVhJ31LwIDAQAB</realm-public-key>
  <auth-server-url>http://localhost:8180/auth</auth-server-url>
  <ssl-required>external</ssl-required>
  <resource>dashbuilder</resource>
  <enable-basic-auth>true</enable-basic-auth>
  <credential name="secret">e92ec68d-6177-4239-be05-
28ef2f3460ff</credential>
  <principal-attribute>preferred_username</principal-attribute>
</secure-deployment>
```

Here,

- **secure-deployment name**: The name of the BAM war file being deployed.

- **realm-public-key**: The public key of the realm you created.

- **resource**: The name of the client that you created for BAM (dashbuilder).

- **enable-basic-auth**: The setting to enable basic authentication mechanism, so that the clients can use both token-based and basic authentication approaches to perform the requests.

- **credential name**: The secret key for the client you created for BAM. You can find it in the *Credentials* tab on the *Clients* page of the RH-SSO admin console.

- **principal-attribute**: The login name of the user. If you do not provide this value, your User Id is displayed in the application instead of your user name.

5. Restart the EAP server and open http://localhost:8080/dashbuilder in a web browser to access BAM. Login to BAM using your RH-SSO admin user credentials.

   You can now access both the applications (Business Central and BAM) without having to re-login.

12.2. INTELLIGENT PROCESS SERVER AUTHENTICATION THROUGH RH-SSO

The Red Hat JBoss BPM Suite Intelligent Process Server provides a REST API for third-party clients. You can integrate the Intelligent Process Server with RH-SSO to delegate the third-party clients identity management to the RH-SSO server.

Once you have created a realm client for Business Central and set up the RH-SSO client adapter for EAP, you can repeat the same steps to integrate the Intelligent Process Server with RH-SSO.

12.2.1. Creating Client for Intelligent Process Server on RH-SSO

To create a new client on your already created realm on RH-SSO admin console:

Procedure: Creating a Client for the Intelligent Process Server

1. On the RH-SSO admin console, open the security realm that you created.

2. Click **Client** tab from the main menu and click **Create**.

   The **Add Client** page opens.

3. On the **Add Client** page, provide the required information to create a new client for your realm. For example:

   » **Client ID**: kie-execution-server

   » **Root URL**: http://*localhost*:8080/kie-server

   » **Client protocol**: openid-connect

4. Navigate to the **Credentials** tab and copy the secret key and and paste it on the kie-execution-server client configuration screen.

5. Click **Save** to save your changes.

   Once you create a new client, its access value is **public** by default. Change it to **confidential**.

12.2.2. Installing and Setting Up Intelligent Process Server with Client Adapter

To consume the Intelligent Process Server remote service endpoints, you must first create and assign the **kie-server** role in the RH-SSO admin console.

Procedure: Setting Up the Intelligent Process Server

1. Navigate to *EAP_HOME*/**standalone/configuration** in your EAP installation and edit **standalone.xml** to add the RH-SSO subsystem configuration. For example:

```
<subsystem xmlns="urn:jboss:domain:keycloak:1.1">
  <secure-deployment name="kie-server.war">
    <realm>demo</realm>
    <realm-public-
key>MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCrVrCuTtArbgaZzL1hvh0x
tL5mc7o0NqPVnYXkLvgcwiC3BjLGw1tGEGoJaXDuSaRllobm53JBhjx33UNv+5z/U
MG4kytBWxheNVKnL6GgqlNabMaFfPLPCF8kAgKnsi79NMo+n6KnSY8YeUmec/p2v
jO2NjsSAVcWEQMVhJ31LwIDAQAB</realm-public-key>
    <auth-server-url>http://localhost:8180/auth</auth-server-
url>
    <ssl-required>external</ssl-required>
    <resource>kie-execution-server</resource>
    <enable-basic-auth>true</enable-basic-auth>
    <credential name="secret">03c2b267-7f64-4647-8566-
572be673f5fa</credential>
    <principal-attribute>preferred_username</principal-
attribute>
  </secure-deployment>
```

```
</subsystem>

<system-properties>
  <property name="org.kie.server.sync.deploy" value="false"/>
</system-properties>
```

Here,

- **secure-deployment name**: Name of your application WAR file.

- **realm**: Name of the realm that you created for the applications to use.

- **realm-public-key**: The public key of the realm you created. You can

- find it in the *Keys* tab in the *Realm settings* page of the realm you created in the RH-SSO admin console. If you do not provide a value for this public key, the server retrieves it automatically.

- **auth-server-url**: The URL for the RH-SSO authentication server.

- **resource**: The name for the server client that you created.

- **enable-basic-auth**: The setting to enable basic authentication mechanism, so that the clients can use both token-based and basic authentication approaches to perform the requests.

- **credential name**: The secret key of the server client you created. You can find it in the *Credentials* tab on the *Clients* page of the RH-SSO admin console.

- **principal-attribute**: The login name of the user. If you do not provide this value, your User Id is displayed in the application instead of your user name.

2. Save your configuration changes in **standalone.xml**.

3. Use the following command to restart the EAP server and run the Intelligent Process Server.

```
EXEC_SERVER_HOME/bin/standalone.sh -Dorg.kie.server.id=<ID> -
Dorg.kie.server.user=<USER> -Dorg.kie.server.pwd=<PWD> -
Dorg.kie.server.location=<LOCATION_URL> -
Dorg.kie.server.controller=<CONTROLLER_URL> -
Dorg.kie.server.controller.user=<CONTROLLER_USER> -
Dorg.kie.server.controller.pwd=<CONTOLLER_PASSWORD>
```

Here is an example:

```
EXEC_SERVER_HOME/bin/standalone.sh -
Dorg.kie.server.id=kieserver1 -Dorg.kie.server.user=kieserver -
Dorg.kie.server.pwd=password -
Dorg.kie.server.location=http://localhost:8080/kie-
server/services/rest/server -
Dorg.kie.server.controller=http://localhost:8080/business-
central/rest/controller -
Dorg.kie.server.controller.user=kiecontroller -
Dorg.kie.server.controller.pwd=password
```

4. Once the Intelligent Process Server is running, you can check the server status using the following command:

```
curl http://kieserver:password@localhost:8080/kie-
server/services/rest/server/
```

> **Note**
>
> Token-based authorization is also supported for communication between Business Central and the Intelligent Process Server. Additionally, you can use the complete token as system property (instead of username and password) for your applications. However, you must ensure that the token does not expire for the period of interaction between the applications, as it is not automatically refreshed.

12.3. THIRD-PARTY CLIENT AUTHENTICATION THROUGH RH-SSO

To use the different remote services provided by the Red Hat JBoss BPM Suite or by an Intelligent Process Server, your client must be authenticated on the RH-SSO server and have a valid token to perform the requests. To use the remote services, the authenticated user must have assigned the following roles:

» **rest-all**: For using the Business Central remote services

> **Note**
>
> The rest-all role must have user specified in **org.kie.server.controller.user** property.

» **kie-server**: For using the Intelligent Process Server remote services

Use the RH-SSO Administrator Console to create these roles and assign them to the users that will consume the remote services.

To achieve third-party client authentication through RH-SSO, you can choose between one of these options:

» Basic authentication (if the application's client supports it)

» Token-based authentication

12.3.1. Basic Authentication

If you have enabled the basic authentication in the RH-SSO client adapter configuration for both Business Central and Intelligent Process Server, you can avoid the token grant/refresh calls and call the services as shown in the examples below:

» For web based remote repositories endpoint:

```
curl http://admin:password@localhost:8080/business-
central/rest/repositories
```

» For the Intelligent Process Server:

```
curl http://admin:password@localhost:8080/kie-
server/services/rest/server/
```

12.3.2. Token-Based Authentication

If you want to opt for a more secure option of authentication, you can consume the remote services from both Business Central and Intelligent Process Server using a granted token provided by a new RH-SSO client.

Procedure: Obtaining and Using Token for Authorizing Remote Calls

1. Click **Client** tab from the main admin console menu and click **Create** to create a new client.

   The **Add Client** page opens.

2. On the **Add Client** page, provide the required information to create a new client for your realm. For example:

   » **Client ID**: kie-remote

   » **Client protocol**: openid-connect

3. Click **Save** to save your changes.

   Once you create a new client, its access value **public** by default. Change it to **confidential**.

4. Obtain a token from **Realm Settings**:

   a. On the RH-SSO admin console, click **Realm Settings** tab.

   b. Click **Tokens** tab.

   c. Change the value for **Access Token Lifespan** to **15** minutes.

      This gives you enough time to obtain a token and invoke the service before it expires.

   d. Click **Save** to save your changes.

5. Once a public client for your remote clients is created, you can now obtain the token by making an HTTP request to the RH-SSO server's token endpoint using:

   ```
   RESULT=`curl --data "grant_type=password&client_id=kie-
   remote&username=admin&password=password"
   http://localhost:8180/auth/realms/demo/protocol/openid-
   connect/token`
   ```

6. To view the token obtained from the RH-SSO server, use the following command:

   ```
   TOKEN=`echo RESULT | sed 's/.*access_token":"//g' | sed
   's/".*//g'`
   ```

You can now use this token to authorize the remote calls. For example, if you want to check the internal Red Hat JBoss BPM Suite repositories, use the token as shown below:

```
curl -H "Authorization: bearer TOKEN"
http://localhost:8080/business-central/rest/repositories
```

CHAPTER 13. INTEGRATION WITH SPRING

13.1. CONFIGURING RED HAT JBOSS BPM SUITE WITH SPRING

The **jboss-bpms-engine.zip** file contains the Spring module, which is called **kie-spring-*VERSION*-redhat-*MINOR_VERSION*.jar**.

You can configure the Spring modules:

As a Self Managed Process Engine

> If you require a single runtime manager instance, use the **RuntimeManager** API. The **RuntimeManager** API synchronizes the process engine and task service internally.

As a Shared Task Service

> If you require multiple runtime manager instances, use the jBPM services.

13.1.1. Integrating Spring with Runtime Manager API

To integrate Spring with the Runtime Manager API, include the following factory beans:

» **org.kie.spring.factorybeans.RuntimeEnvironmentFactoryBean**

» **org.kie.spring.factorybeans.RuntimeManagerFactoryBean**

» **org.kie.spring.factorybeans.TaskServiceFactoryBean**

> **Note**
>
> **TaskServiceFactoryBean** is required only for shared task service.

RuntimeEnvironmentFactoryBean

RuntimeEnvironmentFactoryBean produces **RuntimeEnvironment** instances consumed by **RuntimeManager**.

You can create the following types of **RuntimeEnvironment** instances:

» **DEFAULT**: The default type.

» **EMPTY**: An empty environment.

» **DEFAULT_IN_MEMORY**: Same as **DEFAULT** with no persistence of the runtime engine.

» **DEFAULT_KJAR**: Same as **DEFAULT** with knowledge assets taken from kJAR and identified by **releaseID** or GAV (Group, Artifact, Version).

» **DEFAULT_KJAR_CL**: Built from class path that consists of a **kmodule.xml** descriptor.

Knowledge information is required for all the **RuntimeEnvironment** types. Provide one or more of the following:

» **knowledgeBase**

» **assets**

- **releaseId**

- **groupId**, **artifactId**, **version**

For the **DEFAULT**, **DEFAULT_KJAR**, **DEFAULT_KJAR_CL** types, configure persistence using **entity manager factory** or **transaction manager**.

RuntimeManagerFactoryBean

RuntimeManagerFactoryBean creates **RuntimeManager** instances based on **runtimeEnvironment**. You can create the following **runtimeEnvironment** instances:

- **SINGLETON** (default)

- **PER_REQUEST**

- **PER_PROCESS_INSTANCE**

Every **RuntimeManager** instance must have a unique ID. You can dispose of any **RuntimeManager** instance created by **RuntimeManagerFactoryBean** by calling the **close()** method.

TaskServiceFactoryBean

TaskServiceFactoryBean creates **TaskService** instance based on the given properties. Creates a single instance only.

Properties required:

- **entity manager factory**

- **transaction manager**

When using the **TaskServiceFactoryBean**, provide the Spring transaction manager. When using a shared entity manager from Spring, you can also provide **EntityManager** instance instead of entity manager factory.

Optional properties:

- **userGroupCallback**: **MVELUserGroupCallbackImpl** by default.

- **userInfo**: **DefaultUserInfo** by default.

- **listener**: List of **TaskLifeCycleEventListener** instances.

Sample RuntimeManager Configuration with Spring

To create a single runtime manager Spring configuration:

1. Configure the entity manager factory and the transaction manager, for example:

   ```
   <bean id="jbpmEMF"
   class="org.springframework.orm.jpa.LocalContainerEntityManagerFa
   ctoryBean">
     <property name="persistenceUnitName"
   value="org.jbpm.persistence.spring.jta"/>
   </bean>

   <bean id="btmConfig" factory-method="getConfiguration"
   ```

```xml
class="bitronix.tm.TransactionManagerServices"></bean>

<bean id="BitronixTransactionManager" factory-
method="getTransactionManager"
        class="bitronix.tm.TransactionManagerServices" depends-
on="btmConfig" destroy-method="shutdown" />

<bean id="jbpmTxManager"
class="org.springframework.transaction.jta.JtaTransactionManager"
>
  <property name="transactionManager"
ref="BitronixTransactionManager" />
  <property name="userTransaction"
ref="BitronixTransactionManager" />
</bean>
```

This configuration provides:

> JTA transaction manager, backed by Bitronix for unit tests or servlet containers.

> The **org.jbpm.persistence.spring.jta** entity manager factory for persistence unit.

2. Configure resources you use, for example a business process:

```xml
<bean id="process" factory-method="newclass pathResource"
class="org.kie.internal.io.ResourceFactory">
  <constructor-arg>
    <value>jbpm/processes/sample.bpmn</value>
  </constructor-arg>
</bean>
```

The **sample.bpmn** process is included from the class path.

3. Configure **RuntimeEnvironment** using your entity manager, transaction manager, and resources:

```xml
<bean id="runtimeEnvironment"
class="org.kie.spring.factorybeans.RuntimeEnvironmentFactoryBean
">
  <property name="type" value="DEFAULT"/>
  <property name="entityManagerFactory" ref="jbpmEMF"/>
  <property name="transactionManager" ref="jbpmTxManager"/>
  <property name="assets">
    <map>
      <entry key-ref="process"><util:constant static-
field="org.kie.api.io.ResourceType.BPMN2"/></entry>
    </map>
  </property>
</bean>
```

4. Create **RuntimeManager**:

```xml
<bean id="runtimeManager"
class="org.kie.spring.factorybeans.RuntimeManagerFactoryBean"
destroy-method="close">
```

```xml
      <property name="identifier" value="spring-rm"/>
      <property name="runtimeEnvironment" ref="runtimeEnvironment"/>
    </bean>
```

An example of complete configuration:

```xml
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:util="http://www.springframework.org/schema/util"
       xsi:schemaLocation="http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans-3.0.xsd
                           http://www.springframework.org/schema/util
http://www.springframework.org/schema/util/spring-util-3.0.xsd">

  <import resource="classpath:jbpm/configuration-template/assets.xml"
/>

  <bean id="jbpmEMF"
class="org.springframework.orm.jpa.LocalContainerEntityManagerFactoryB
ean">
    <property name="persistenceUnitName"
value="org.jbpm.persistence.spring.jta"/>
    <property name="persistenceXmlLocation"
value="classpath:jbpm/persistence-jta.xml"/>
  </bean>

  <bean id="btmConfig" factory-method="getConfiguration"
class="bitronix.tm.TransactionManagerServices"/>

  <bean id="BitronixTransactionManager" factory-
method="getTransactionManager"
       class="bitronix.tm.TransactionManagerServices" depends-
on="btmConfig" destroy-method="shutdown" />

  <bean id="jbpmTxManager"
class="org.springframework.transaction.jta.JtaTransactionManager">
    <property name="transactionManager"
ref="BitronixTransactionManager" />
    <property name="userTransaction" ref="BitronixTransactionManager"
/>
  </bean>

  <bean id="runtimeEnvironment"
class="org.kie.spring.factorybeans.RuntimeEnvironmentFactoryBean">
    <property name="type" value="DEFAULT"/>
    <property name="entityManagerFactory" ref="jbpmEMF"/>
    <property name="transactionManager" ref="jbpmTxManager"/>
    <property name="assets" ref="assets"/>
  </bean>


  <bean id="logService"
class="org.jbpm.process.audit.JPAAuditLogService" depends-
on="runtimeEnvironment">
```

```xml
      <constructor-arg value="#{runtimeEnvironment.environment}" />
      <constructor-arg value="STANDALONE_JTA" />
   </bean>
</beans>
```

13.1.2. Spring and jBPM Services

If you require multiple runtime managers, you can use jBPM services directly in your application. Due to the dynamic nature of jBPM services, processes and other assets can be added and removed without restarting your application.

To configure jBPM services, implement the **IdentityProvider** interface:

```java
import java.util.ArrayList;
import java.util.Collections;
import java.util.List;

import org.kie.internal.identity.IdentityProvider;
import org.springframework.security.core.Authentication;
import org.springframework.security.core.GrantedAuthority;
import org.springframework.security.core.context.SecurityContextHolder;

public class SpringSecurityIdentityProvider implements IdentityProvider {

 public String getName() {

   Authentication auth =
SecurityContextHolder.getContext().getAuthentication();
   if (auth != null && auth.isAuthenticated()) {
    return auth.getName();
   }
   return "system";
 }

 public List<String> getRoles() {
   Authentication auth =
SecurityContextHolder.getContext().getAuthentication();
   if (auth != null && auth.isAuthenticated()) {
    List<String> roles = new ArrayList<String>();

    for (GrantedAuthority ga : auth.getAuthorities()) {
     roles.add(ga.getAuthority());
    }

    return roles;
   }

   return Collections.emptyList();
 }

 public boolean hasRole(String role) {
```

```
  return false;
 }

}
```

To configure jBPM services in a Spring application:

1. Configure the transaction manager:

```
<context:annotation-config />
<tx:annotation-driven />
<tx:jta-transaction-manager />

<bean id="transactionManager"
class="org.springframework.transaction.jta.JtaTransactionManager"
/>
```

2. Configure JPA and persistence:

```
<bean id="entityManagerFactory"
class="org.springframework.orm.jpa.LocalContainerEntityManagerFa
ctoryBean" depends-on="transactionManager">
    <property name="persistenceXmlLocation"
value="classpath:/META-INF/jbpm-persistence.xml" />
</bean>
```

3. Configure security providers:

```
<util:properties id="roleProperties"
location="classpath:/roles.properties" />

<bean id="userGroupCallback"
class="org.jbpm.services.task.identity.JBossUserGroupCallbackImpl
">
  <constructor-arg name="userGroups"
ref="roleProperties"></constructor-arg>
</bean>

<bean id="identityProvider"
class="org.jbpm.spring.SpringSecurityIdentityProvider"/>
```

4. Configure the runtime manager factory:

```
<bean id="runtimeManagerFactory"
class="org.kie.spring.manager.SpringRuntimeManagerFactoryImpl">
  <property name="transactionManager" ref="transactionManager"/>
  <property name="userGroupCallback" ref="userGroupCallback"/>
</bean>

<bean id="transactionCmdService"
class="org.jbpm.shared.services.impl.TransactionalCommandService
">
  <constructor-arg name="emf"
ref="entityManagerFactory"></constructor-arg>
</bean>
```

```xml
<bean id="taskService"
class="org.kie.spring.factorybeans.TaskServiceFactoryBean"
destroy-method="close">
  <property name="entityManagerFactory"
ref="entityManagerFactory"/>
  <property name="transactionManager" ref="transactionManager"/>
  <property name="userGroupCallback" ref="userGroupCallback"/>
  <property name="listeners">
    <list>
      <bean
class="org.jbpm.services.task.audit.JPATaskLifeCycleEventListener
">
        <constructor-arg value="true"/>
      </bean>
    </list>
  </property>
</bean>
```

The runtime manager factory is Spring context aware and can interact with Spring containers.

5. Configure jBPM services as Spring beans:

```xml
<!-- definition service -->
<bean id="definitionService"
class="org.jbpm.kie.services.impl.bpmn2.BPMN2DataServiceImpl"/>

<!-- runtime data service -->
<bean id="runtimeDataService"
class="org.jbpm.kie.services.impl.RuntimeDataServiceImpl">
  <property name="commandService" ref="transactionCmdService"/>
  <property name="identityProvider" ref="identityProvider"/>
  <property name="taskService" ref="taskService"/>
</bean>

<!-- -- deployment service -->
<bean id="deploymentService"
class="org.jbpm.kie.services.impl.KModuleDeploymentService"
depends-on="entityManagerFactory" init-method="onInit">
  <property name="bpmn2Service" ref="definitionService"/>
  <property name="emf" ref="entityManagerFactory"/>
  <property name="managerFactory" ref="runtimeManagerFactory"/>
  <property name="identityProvider" ref="identityProvider"/>
  <property name="runtimeDataService" ref="runtimeDataService"/>
</bean>

<!-- process service -->
<bean id="processService"
class="org.jbpm.kie.services.impl.ProcessServiceImpl" depends-
on="deploymentService">
  <property name="dataService" ref="runtimeDataService"/>
  <property name="deploymentService" ref="deploymentService"/>
</bean>

<!-- user task service -->
```

```xml
<bean id="userTaskService"
class="org.jbpm.kie.services.impl.UserTaskServiceImpl" depends-
on="deploymentService">
  <property name="dataService" ref="runtimeDataService"/>
  <property name="deploymentService" ref="deploymentService"/>
</bean>

<!-- register runtime data service as listener on deployment
service so it can receive notification about deployed and
undeployed units -->
<bean id="data"
class="org.springframework.beans.factory.config.MethodInvokingFac
toryBean" depends-on="deploymentService">
  <property name="targetObject"
ref="deploymentService"></property>
  <property
name="targetMethod"><value>addListener</value></property>
  <property name="arguments">
  <list>
      <ref bean="runtimeDataService"/>
  </list>
  </property>
</bean>
```

CHAPTER 14. CDI INTEGRATION

14.1. CDI INTEGRATION

To make use of jbpm-kie-services in your system, you will need to provide some mbeans to satisfy all dependencies of the services. There are several mbeans that depend on actual scenarios.

» Entity manager and entity manager factory

» User group callback for human tasks

» Identity provider to pass authenticated user information to the services

When running in JEE environment, like JBoss Application Server, the mbean should satisfy all requirements of the jbpm-kie-services:

```
public class EnvironmentProducer {

    @PersistenceUnit(unitName = "org.jbpm.domain")
    private EntityManagerFactory emf;

    @Inject
    @Selectable
    private UserGroupCallback userGroupCallback;

    @Produces
    public EntityManagerFactory getEntityManagerFactory() {
        return this.emf;
    }

    @Produces
    @RequestScoped
    public EntityManager getEntityManager() {
        EntityManager em = emf.createEntityManager();
        return em;
    }

    public void close(@Disposes EntityManager em) {
        em.close();
    }

    @Produces
    public UserGroupCallback produceSelectedUserGroupCalback() {
        return userGroupCallback;
    }
    @Produces

    public IdentityProvider produceIdentityProvider {
        return new IdentityProvider() {
            // implement IdentityProvider
        };
    }
}
```

Then the **deployments/business-central.war/WEB-INF/beans.xml** file may be configured to change the current settings of the new **usergroupcallback** implementation:

```
<beans xmlns="http://java.sun.com/xml/ns/javaee"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
http://docs.jboss.org/cdi/beans_1_0.xsd">

<alternatives>

<class>org.jbpm.services.task.identity.JAASUserGroupCallbackImpl</class>
</alternatives>

</beans>
```

> **Note**
>
> **org.jbpm.services.task.identity.JAASUserGroupCallbackImpl** is just an example here to demonstrate the settings of the application server regardless of what it actually is (LDAP, DB, etc).

CHAPTER 15. PERSISTENCE

The runtime data of the Process Engine can be persisted in data stores. The persistence mechanism saves the data using marshalling: the runtime data is converted into a binary dataset and the dataset is saved in the data storage.

Note that persistence is not configured by default and the engine runs without persistence.

Custom marshalling mechanism

The runtime data is saved using marshalling (binary persistence). The marshalling mechanism is a custom serialization mechanism.

Red Hat JBoss BPM Suite will persist the following when persistence is configured:

» *Session state*: This includes the session ID, date of last modification, the session data that business rules would need for evaluation, state of timer jobs.

» *Process instance state*: This includes the process instance ID, process ID, date of last modification, date of last read access, process instance start date, runtime data (the execution status including the node being executed, variable values, etc.)and the eventtypes.

» *Work item runtime state*: This includes the work item ID, creation date, name, process instance ID, and the work item state itself.

Based on the persisted data, it is possible to restore the state of execution of all running process instances in case of failure or to temporarily remove running instances from memory and restore them later. By default, no persistence is configured.

To allow persistence, you need to add the jbpm-persistence jar files to the classpath of your application and configure the engine to use persistence. The engine automatically stores the runtime state in the storage when the engine reaches a safe point. Safe points are points where the process instance has paused. When a process instance invocation reaches a safe point in the engine, the engine stores any changes to the process instance as a snapshot of the process runtime data. However, when a process instance is completed, the persisted snapshot of process instance runtime data is automatically deleted.

If a failure occurs and you need to restore the engine runtime from the storage, the process instances are automatically restored and their execution resumes so there is no need to reload and trigger the process instances manually.

The runtime persistence data is to be considered internal to the engine. You should not access persisted runtime data or modify them directly as this might have unexpected side effects.

To obtain information about the current execution state, refer to the history log. Query the database for runtime data only if absolutely necessary.

15.1. SESSION

Sessions are persisted as **SessionInfo** entities. These persist the state of the runtime KIE session, and store the following data:

Field	Description	Nullable
id	primary key	false
lastmodificationdate	last saved to data store	N/A
rulesbytearray	binary dataset with session state (binary blob	false
startdate	session start	
optlock	version number used to lock value for optimistic locking	

15.2. PROCESS INSTANCE

Process instances are persisted as **ProcessInstanceInfo** entities, which persist the state of a process instance on runtime and store the following data:

Field	Description	Nullable
instanceid	primary key	false
lastmodificationdate	last saved to data store	N/A
lastreaddate	last read from data store	N/A
processid	ID of the process the instance is based on	false
processinstancebytearray	binary dataset with process instance state (binary blob)	false
startdate	Process instance start date	

Field	Description	Nullable
optlock	version number used lock value for optimistic locking	
state	Process instance state	false

ProcessInstanceInfo has a 1:N relationship to the **EventTypes** entity.

The **EventTypes** entity contains the following data:

Field	Description	Nullable
instanceid	reference to the Process instance (foreign key to the **processinstanceinfo**)	false
element	text field related to an event the Process instance has undergone	

Pessimistic Locking Support

The default locking mechanism for persistence of processes is *optimistic*. With multi-thread high concurrency to the same process instance, this locking strategy can result in bad performance.

With the release of the 6.1 version of Red Hat JBoss BPM Suite, this can be changed at runtime to allow the user to set locking on a per process basis and to allow it to be *pessimistic* (the change can be made at a per KIE Session level or Runtime Manager level as well and not just at the process level).

To set a process to use pessmistic locking, do this in the runtime environment:

```
import org.kie.api.runtime.Environment;
import org.kie.api.runtime.EnvironmentName;
import org.kie.api.runtime.manager.RuntimeManager;
import org.kie.api.runtime.manager.RuntimeManagerFactory;

...

// here env is an instance of org.kie.api.runtime.Environment
env.set(EnvironmentName.USE_PESSIMISTIC_LOCKING, true);
```

```
// now create your Runtime Manager using this enviornment
RuntimeManager manager =
RuntimeManagerFactory.Factory.get().newPerRequestRuntimeManager(env
ironment);
```

15.3. WORK ITEM

Work Items are persisted as **workiteminfo** entities, which persist the state of the particular work item instance on runtime and store the following data:

Field	Description	Nullable
workitemid	primary key	false
name	work item name	
processinstanceid	parent Process instance id	false
state	integer representing work item state	false
optlock	version number used lock value for optimistic locking	
workitembytearray	binary dataset with work item state (binary blob)	false
creationDate	timestampe on which the work item was created	false

15.4. PERSISTENCE CONFIGURATION

15.4.1. Persistence configuration

Although persistence is not used by default, the dependencies needed are available in the runtime directory as jar files.

Persistence is defined per session and you can define it either using the **JBPMHelper** class after you create a session or using the **JPAKnowledgeService** to create your session. The latter option provides more flexibility, while **JBPMHelper** has a method to create a session, and uses a configuration file to configure this session.

15.4.2. Configuring persistence using JBPMHelper

To configure persistence of your session using JBPMHelper, do the following:

1. Define your application to use an appropriate JBPMHelper session construtor:

   - **`KieSession ksession = JBPMHelper.newKieSession(kbase);`**

   - **`KieSession ksession = JBPMHelper.loadKieSession(kbase, sessionId);`**

2. Configure the persistence in the **`jBPM.properties`** file.

   > **Example 15.1. A sample jBPM.properties file with persistence for the in-memory H2 database**
   >
   > ```
   > # for creating a datasource
   > persistence.datasource.name=jdbc/jbpm-ds
   > persistence.datasource.user=sa
   > persistence.datasource.password=
   > persistence.datasource.url=jdbc:h2:tcp://localhost/~/jbpm-
   > db
   > persistence.datasource.driverClassName=org.h2.Driver
   >
   > # for configuring persistence of the session
   > persistence.enabled=true
   > persistence.persistenceunit.name=org.jbpm.persistence.jpa
   > persistence.persistenceunit.dialect=org.hibernate.dialect.H
   > 2Dialect
   >
   > # for configuring the human task service
   > taskservice.enabled=true
   > taskservice.datasource.name=org.jbpm.task
   > taskservice.transport=mina
   > taskservice.usergroupcallback=org.jbpm.task.service.DefaultU
   > serGroupCallbackImpl
   > ```

Any invocations on the session will now trigger the persistance process.

Make sure the datasource is up and running on engine start. If you are running the in-memory H2 database, you can start the database from your application using the **`JBPMHelper.startH2Server();`** method call and register it with the engine using **`JBPMHelper.setupDataSource();`** method call.

15.4.3. Configuring persistence using JPAKnowledgeService

To create your knowledge session and configure its persistence using JPAKnowledgeService, do the following:

1. Define your application to use the knowledge session created by JPAKnowledgeService:

a. Define the session based on a knowledge base, a knowledge session configuration, and an environment. The environment must contain a reference to your Entity Manager Factory:

```
// create the entity manager factory and register it in
the environment
EntityManagerFactory emf =
Persistence.createEntityManagerFactory(
"org.jbpm.persistence.jpa" );
Environment env = KnowledgeBaseFactory.newEnvironment();
env.set( EnvironmentName.ENTITY_MANAGER_FACTORY, emf );

// create a new knowledge session that uses JPA to store
the runtime state
KieSession ksession = JPAKnowledgeService.newKieSession(
kbase, null, env );
int sessionId = ksession.getId();

// invoke methods on your method here
ksession.startProcess( "MyProcess" );
ksession.dispose();
```

b. Define the session based on a specific session id.

```
// recreate the session from database using the sessionId
ksession = JPAKnowledgeService.loadKieSession(sessionId,
kbase, null, env );
```

2. Configure the persistence in the **META-INF/persistence.xml** file: configure JPA to use Hibernate and the respective database.

Information on how to configure data source on your application server should be available in the documentation delivered with the application server. For this information for JBoss Enterprise Application Platform, see the *Administration and Configuration Guide* for this product.

Example 15.2. A sample persistence.xml file with persistence for an H2 data source jdbc/jbpm-ds

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<persistence
  version="1.0"
  xsi:schemaLocation=
    "http://java.sun.com/xml/ns/persistence

http://java.sun.com/xml/ns/persistence/persistence_1_0.xsd
      http://java.sun.com/xml/ns/persistence/orm
      http://java.sun.com/xml/ns/persistence/orm_1_0.xsd"
  xmlns:orm="http://java.sun.com/xml/ns/persistence/orm"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="http://java.sun.com/xml/ns/persistence">
  <persistence-unit name="org.jbpm.persistence.jpa"
transaction-type="JTA">
```

```xml
<provider>org.hibernate.ejb.HibernatePersistence</provider>
    <jta-data-source>jdbc/jbpm-ds</jta-data-source>
    <mapping-file>META-INF/JBPMorm.xml</mapping-file>
    <class>org.drools.persistence.info.SessionInfo</class>

<class>org.jbpm.persistence.processinstance.ProcessInstanceI
nfo</class>
    <class>org.drools.persistence.info.WorkItemInfo</class>
    <properties>
      <property name="hibernate.dialect"
value="org.hibernate.dialect.H2Dialect"/>
      <property name="hibernate.max_fetch_depth"
value="3"/>
      <property name="hibernate.hbm2ddl.auto"
value="update"/>
      <property name="hibernate.show_sql" value="true"/>
      <property
name="hibernate.transaction.manager_lookup_class"

value="org.hibernate.transaction.BTMTransactionManagerLooku
p"/>
    </properties>
  </persistence-unit>
</persistence>
```

Any invocations on the session will now trigger the persistance process.

Make sure the datasource is up and running on engine start. If you are running the in-memory H2 database, you can start the database from your application using the **JBPMHelper. startH2Server();** method call and register it with the engine using **JBPMHelper. setupDataSource();** method call.

Sample data source configuration for simple Java environment

If you are running JBoss BPM Suite in a simple Java environment, your data source configuration will be similar to the following:

```java
PoolingDataSource ds = new PoolingDataSource();
ds.setUniqueName("jdbc/jbpm-ds");
ds.setClassName("bitronix.tm.resource.jdbc.lrc.LrcXADataSourc
e");
ds.setMaxPoolSize(3);
ds.setAllowLocalTransactions(true);
ds.getDriverProperties().put("user", "sa");
ds.getDriverProperties().put("password", "sasa");
ds.getDriverProperties().put("URL",
"jdbc:h2:tcp://localhost/~/jbpm-db");
ds.getDriverProperties().put("driverClassName",
"org.h2.Driver");
ds.init();
```

CHAPTER 16. TRANSACTIONS

16.1. TRANSACTIONS

The Process Engine supports JTA transactions: local transactions are only supported when using Spring. Pure local transactions are not supported.

By default, each method invocation is considered a transaction. To change this behavior, for example, to combine multiple commands into one transaction, you will need to specify transaction boundaries.

16.2. DEFINING TRANSACTIONS

To define a transaction, do the following:

1. Register the transaction manager in your environment:

   Example 16.1. Code with transaction manager registration

   ```
   // create the entity manager factory
   EntityManagerFactory emf =
   EntityManagerFactoryManager.get().getOrCreate("org.jbpm.per
   sistence.jpa");
   TransactionManager tm =
   TransactionManagerServices.getTransactionManager();
   Environment env = EnvironmentFactory.newEnvironment();
   env.set(EnvironmentName.ENTITY_MANAGER_FACTORY, emf);
   env.set(EnvironmentName.TRANSACTION_MANAGER, tm);

   // setup the runtime environment
   RuntimeEnvironment environment =
   RuntimeEnvironmentBuilder.Factory.get()
   .newDefaultBuilder()
   .addAsset(ResourceFactory.newClassPathResource("MyProcessDef
   inition.bpmn2"), ResourceType.BPMN2)

   .addEnvironmentEntry(EnvironmentName.TRANSACTION_MANAGER,
   tm)

   .addEnvironmentEntry(EnvironmentName.PERSISTENCE_CONTEXT_MA
   NAGER, new JpaProcessPersistenceContextManager(env))

   .addEnvironmentEntry(EnvironmentName.TASK_PERSISTENCE_CONTE
   XT_MANAGER, new JPATaskPersistenceContextManager(env))
       .get();
   ```

2. Initialize the KieSession:

   ```
   // get the KieSession
   RuntimeManager manager =
   ```

```
RuntimeManagerFactory.Factory.get().newPerProcessInstanceRunti
meManager(environment);
RuntimeEngine runtime =
manager.getRuntimeEngine(ProcessInstanceIdContext.get());
KieSession ksession = runtime.getKieSession();
```

3. Define the transaction manager in **jndi.properties**:

> **Example 16.2. Definition of Bitronix transaction manager in jndi.properties**
>
> ```
> java.naming.factory.initial=bitronix.tm.jndi.BitronixInitia
> lContextFactory
> ```

> **Using another transaction manager**
>
> To use a different JTA transaction manager, edit the **hibernate.transaction.manager_lookup_class**, the transaction manager property, in the **persistence.xml** file to load your transaction manager.
>
> For example, if choosing JBoss Transaction Manager:
>
> ```
> <property
> name="hibernate.transaction.manager_lookup_class"
> value="org.hibernate.transaction.JBossTransactionManage
> rLookup"/>
> ```

4. Define the start and the end of the transaction.

```
// start the transaction
UserTransaction ut =
InitialContext.doLookup("java:comp/UserTransaction");
ut.begin();

// perform multiple commands inside one transaction
ksession.insert( new Person( "John Doe" ) );
ksession.startProcess("MyProcess");

// commit the transaction
ut.commit();
```

16.3. CONTAINER MANAGED TRANSACTIONS

In cases where JBoss BPM Suite is embedded inside an application that is in a container that can manage transactions by itself (Container Managed Transactions - CMT), a special dedicated transaction manager is provided using the **org.jbpm.persistence.jta.ContainerManagerTransactionManager** class. This

is because the default implementation of the transaction manager in JBoss BPM Suite is based on the **UserTransaction** class getting the transaction status. However, some application servers in a CMT mode do not allow accessing the **UserTransaction** instance from JNDI.

Operations executed on this manager are all no-op because they cannot affect the underlying CMT. The **ContainerManagedTransactionManager** class expects that the transaction is always active (returning **ACTIVE** to the **getStatus()** method).

> **Note**
>
> Even though the container manages transactions, the container should be made aware of any exceptions that happen during process instance execution. Exceptions thrown by the engine should be propagated up to the container to properly rollback transactions.

Configuring the Transaction Manager

To configure and use the **ContainerManagedTransactionManager**, it needs to be inserted into the environment before you create or load a session:

```
Environment env = EnvironmentFactory.newEnvironment();
env.set(EnvironmentName.ENTITY_MANAGER_FACTORY, emf);
env.set(EnvironmentName.TRANSACTION_MANAGER, new
ContainerManagedTransactionManager());
env.set(EnvironmentName.PERSISTENCE_CONTEXT_MANAGER, new
JpaProcessPersistenceContextManager(env));
```

Next, setup the JPA Provider in your **persistence.xml** file. For example, if using IBM WebSphere:

```
<property name="hibernate.transaction.factory_class"
value="org.hibernate.transaction.CMTTransactionFactory"/>
<property name="hibernate.transaction.manager_lookup_class"
value="org.hibernate.transaction.WebSphereExtendedJTATransactionLook
up"/>
```

Disposing the KSession in a CMT

In a CMT, you should not dispose the ksession directly (by using the **dispose()** method). Doing so will cause exceptions on transaction completion as the Process Engine needs to clean up the state after the invocation has finished.

Instead, use the specialized class **org.jbpm.persistence.jta.ContainerManagedTransactionDisposeCommand**'s **execute()** method. Using this command ensures that the ksession will be disposed when the transaction is actually complete.

This method checks to see if the transaction is active. If it is, it delegates the actual disposal to the **afterDisposal** phase of the transaction instead of executing it directly. If there is no active transaction or if there is no active transaction, the ksession is disposed immediately.

CHAPTER 17. LOGGING

The logging mechanism allows you to store information about the execution of a process instance. It is provided by a special event listener that listens to the Process Engine for any relevant events to be logged, so that the information can be stored separately from other non-log information stored either in the server built-in database (H2) or a connected data source using JPA or Hibernate.

The **jbpm-audit** module provides the event listener and also allows you to store process-related information directly in a database using JPA or Hibernate. The data of the following entities is stored as follows:

- Process instance as **processinstancelog**

- Element instance as **nodeinstancelog**

- Variable instance as **variableinstancelog**

Table 17.1. Fields of the ProcessInstanceLog table

Field	Description	Nullable
id	The primary key of the log entity	No
end_date	The end date of the process instance	Yes
processid	The name (id) of the underlying process	Yes
processinstanceid	The id of the process instance	No
start_date	The start date of the process instance	Yes
status	The status of the process instance	Yes
parentProcessInstanceId	The process instance id of the parent process instance if applicable	Yes

Field	Description	Nullable
outcome	The outcome of the process instance (details on the process finish, such as error code)	Yes

Table 17.2. Fields of the NodeInstanceLog table

Field	Description	Nullable
id	The primary key of the log entity	No
log_date	The date of the event	Yes
nodeid	The node id of the underlying Process Element	Yes
nodeinstanceid	The id of the node instance	Yes
nodename	The name of the underlying node	Yes
processid	The id of the underlying process	Yes
processinstanceid	The id of the parent process instance	No
type	The type of the event (**0** = enter event, **1** = exit event)	No

Table 17.3. Fields of the VariableInstanceLog table

Field	Description	Nullable
id	The primary key of the log entity	No
log_date	The date of the event	Yes
processid	The name (id) of the underlying process	Yes
processinstanceid	The id of the process instance	No
value	The value of the variable at log time	Yes
variableid	The variable id as defined in the process definition	Yes
variableinstanceid	The id of the variable instance	Yes
outcome	The outcome of the process instance (details on the process finish, such as error code)	Yes

If neccessary, define your own data model of custom information and use the process event listeners to extract the information.

17.1. LOGGING EVENTS TO DATABASE

To log an event that occurs on runtime in a Process instance, an Element instance, or a variable instance, you need to do the following:

1. Map the Log classes to the data source, so that the given data source accepts the log entries. On Red Hat JBoss EAP, edit the data source properties in the `persistence.xml` file.

Example 17.1. The ProcessInstanceLog, NodeInstanceLog and VariableInstanceLog classes enabled for processInstanceDS

```xml
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<persistence  version="1.0"  xsi:schemaLocation=
    "http://java.sun.com/xml/ns/persistence

http://java.sun.com/xml/ns/persistence/persistence_1_0.xsd
      http://java.sun.com/xml/ns/persistence/orm
      http://java.sun.com/xml/ns/persistence/orm_1_0.xsd"
  xmlns:orm="http://java.sun.com/xml/ns/persistence/orm"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns="http://java.sun.com/xml/ns/persistence">

  <persistence-unit name="org.jbpm.persistence.jpa">

<provider>org.hibernate.ejb.HibernatePersistence</provider>
    <jta-data-source>jdbc/processInstanceDS</jta-data-source>
    <class>org.drools.persistence.info.SessionInfo</class>

<class>org.jbpm.persistence.processinstance.ProcessInstanceInfo
</class>
    <class>org.drools.persistence.info.WorkItemInfo</class>
    <class>org.jbpm.process.audit.ProcessInstanceLog</class>
    <class>org.jbpm.process.audit.NodeInstanceLog</class>
    <class>org.jbpm.process.audit.VariableInstanceLog</class>

    <properties>
      <property name="hibernate.dialect"
value="org.hibernate.dialect.H2Dialect"/>
      <property name="hibernate.max_fetch_depth" value="3"/>
      <property name="hibernate.hbm2ddl.auto"
value="update"/>
      <property name="hibernate.show_sql" value="true"/>
      <property
name="hibernate.transaction.manager_lookup_class"

value="org.hibernate.transaction.BTMTransactionManagerLookup"/
>
    </properties>
  </persistence-unit>
</persistence>
```

2. Register a logger on your Kie Session.

Example 17.2. Import the Loggers

```java
import org.jbpm.process.audit.AuditLogService;
import org.jbpm.process.audit.AuditLoggerFactory;
import org.jbpm.process.audit.AuditLoggerFactory.Type;
import org.jbpm.process.audit.JPAAuditLogService;
...
```

–

Example 17.3. Registering a Logger to a Kie Session

```
@PersistenceUnit(unitName = PERSISTENCE_UNIT_NAME)
  private EntityManagerFactory emf;

  private AuditLogService auditLogService;
@PostConstruct
  public void configure() {

  auditLogService = new JPAAuditLogService(emf);
  ((JPAAuditLogService)
auditLogService).setPersistenceUnitName(PERSISTENCE_UNIT_NAME)
;

  if( emf == null ) {
  ((JPAAuditLogService)
auditLogService).setPersistenceUnitName(PERSISTENCE_UNIT_NAME)
;
  }

  RuntimeEngine runtime =
singletonManager.getRuntimeEngine(EmptyContext.get());
  KieSession ksession = runtime.getKieSession();
  AuditLoggerFactory.newInstance(Type.JPA, ksession, null);

  }
```

3. Optionally, call the method **addFilter** on the logger to filter out irrelevant information. Only information accepted by all filters appears in the database.

4. Logger classes can be viewed in the Audit View:

```
<dependency>
 <groupId>org.jbpm</groupId>
 <artifactId>jbpm-audit</artifactId>
 <version>6.0.1.Final</version>
</dependency>
```

17.2. LOGBACK FUNCTIONALITY

Red Hat JBoss BPM Suite provides **logback** functionality for logging configuration.

Accordingly, everything configured is logged to the *Simple Logging Facade for Java* SLF4J, which delegates any log to Logback, Apache Commons Logging, Log4j or java.util.logging. Add a dependency to the logging adaptor for your logging framework of choice. If you are not using any logging framework yet, you can use Logback by adding this Maven dependency:

```
<dependency>
    <groupId>ch.qos.logback</groupId>
```

```
        <artifactId>logback-classic</artifactId>
        <version>1.x</version>
</dependency>
```

> **Note**
>
> **slf4j-nop** and **slf4j-simple** are ideal for a light environment.

17.3. CONFIGURING LOGGING

To configure the logging level of the packages, create a **logback.xml** file in **business-central.war/WEB-INF/classes/logback.xml**. To set the logging level of the **org.drools** package to "debug" for verbose logging, you would need to add the following line to the file:

```
<configuration>
  <logger name="org.drools" level="debug"/>

  ...
<configuration>
```

Similarly, you can configure logging for packages such as the following:

- **org.guvnor**

- **org.jbpm**

- **org.kie**

- **org.slf4j**

- **org.dashbuilder**

- **org.uberfire**

- **org.errai**

- etc…

If configuring with **log4j**, the **log4j.xml** can be located at **business-central.war/WEB-INF/classes/log4j.xml** and can be configured in the following way:

```
<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">

  <category name="org.drools">
    <priority value="debug" />
  </category>

  ...

</log4j:configuration>
```

Note

Additional logging can be configured in the individual container. To configure logging for JBoss Enterprise Application Platform, please refer to the Red Hat JBoss Enterprise Application Platform Administration and Configuration Guide.

17.4. MANAGING LOG FILES

Red Hat JBoss BPM Suite manages most of the required maintenance. Automatically cleaned runtime data includes:

» Process instances data, which is removed upon process instance completion.

» Work items data, which is removed upon work item completion.

» Task instances data, which is removed upon completion of a process to which given task belongs.

Runtime data, which may not be automatically cleaned, includes session information data. This depends on the selected runtime strategy:

» Singleton strategy ensures that session information runtime data will not be automatically removed.

» Per request strategy allows automatic removal when a given request terminates.

» Per process instances will be automatically removed when process instance mapped to a given session completes or is aborted.

Red Hat JBoss BPM Suite does not remove executor request and error information.

In order not to lose track of process instances, Red Hat JBoss BPM Suite offers audit data tables. These are used by default and keep track of the BPM Suite environment. JBoss BPM Suite offers two ways of how to manage and maintain the audit data tables:

» Automatic clean-up

» Manual clean-up

17.4.1. Automatic Clean-Up

Automatic clean-up uses the **LogCleanupCommand** executor command, which consists of logic to clean up all or selected data automatically. An advantage of the automatic clean-up method is the ability to schedule repeated clean-ups by using reoccurring job feature of the JBoss BPM Suite executor. This means that when one job completes, it provides information to the JBoss BPM Suite executor if and when the next instance of this job should be executed. By default, **LogCleanupCommand** is executed once a day but can be reconfigured to run on different intervals.

There are several important configuration options that can be used with the **LogCleanupCommand** command:

Table 17.4. LogCleanupCommand parameters table

Name	Description	Is Exclusive
SkipProcessLog	Indicates if the clean-up of process instances, node instances and variables log cleanup should be omitted (default: false)	No, can be used with other parameters
SkipTaskLog	Indicates if the task audit and the task event log clean-up should be omitted (default: false)	No, can be used with other parameters
SkipExecutorLog	Indicates if the JBoss BPM Suite executor entries clean-up should be omitted (default: false)	No, can be used with other parameters
SingleRun	Indicates if the job routine should run only once (default: false)	No, can be used with other parameters
NextRun	Sets a date for the next run. For example, 12h is set for jobs to be executed every 12 hours. If the option is not given, the next job will run 24 hours after the completion of the current job	Yes, cannot be used when the OlderThanPeriod parameter is used
OlderThan	Causes logs older than the given date to be removed. The date format is **YYYY-MM-DD**. Usually, this parameter is used for single run jobs	Yes, cannot be used when the NextRun parameter is used
OlderThanPeriod	Causes logs older than the given timer expression should be removed. For example, set 30d to remove logs older than 30 day from current time	No, can be used with other parameters

Name	Description	Is Exclusive
ForProcess	Specifies process definition ID for which logs should be removed	No, can be used with other parameters
ForDeployment	Specifies deployment ID for which logs should be removed	No, can be used with other parameters
EmfName	Persistence unit name that shall be used to perform operation deletion	N/A

Note

LogCleanupCommand does not remove any active instances, such as running process instances, task instances, or executor jobs.

Warning

While all audit tables have a time stamp, some may be missing other parameters, such as process id, or deployment id. For that reason, it is recommended to use the date parameter when managing the clean-up job routine.

17.4.2. Setting up Automatic Clean-up Job

To set up automatic clean-up job, do the following:

1. Open Business Central in your web browser (if running locally http://*localhost*:8080/business-central) and log in as a user with the **admin** role.

2. Go to **Deploy → Jobs**.

3. Click New Job in the top right hand corner of the page.

4. Enter a name, due date and time. Enter the following into the *Type* text field:

   ```
   org.jbpm.executor.commands.LogCleanupCommand
   ```

5. Click on **Add Parameter** if you wish to use parameters listed above. In the key section, enter a parameter name. In the value section, enter true or false, depending on the desired outcome.

6. Click **Create** to finalize the job creation wizard. You have successfully created an automatic clean-up job.

17.4.3. Manual Clean-Up

You may make use of audit API to do the clean-up manually with more control over parameters and thus more control over what will be removed. Audit API is divided into following areas:

- Process audit, which is used to clean up process, node and variables logs, accessible in the **jbpm-audit** module

- Task audit, which is used to clean up tasks and task events, accessible in the **jbpm-human-task-audit** module

- Executor jobs, which is used to clean up executor jobs and errors, accessible in the **jbpm-executor** module

Modules are sorted hierarchically and can be accessed as follows:

- **org.jbpm.process.audit.JPAAuditLogService**

- **org.jbpm.services.task.audit.service.TaskJPAAuditService**

- **org.jbpm.executor.impl.jpa.ExecutorJPAAuditService**

Several examples of manual clean-up follow:

Example 17.4. Removal of completed process instance logs

```
import org.jbpm.process.audit.JPAAuditLogService;
import
org.kie.internal.runtime.manager.audit.query.ProcessInstanceLogDelet
eBuilder;
import org.kie.api.runtime.process.ProcessInstance;

JPAAuditLogService auditService = new JPAAuditLogService(emf);
ProcessInstanceLogDeleteBuilder updateBuilder =
auditService.processInstanceLogDelete().status(ProcessInstance.STATE
_COMPLETED);
int result = updateBuilder.build().execute();
```

Example 17.5. Task audit logs removal for the org.jbpm:HR:1.0 deployment

```
import org.jbpm.services.task.audit.service.TaskJPAAuditService;
import org.kie.internal.task.query.AuditTaskDeleteBuilder;

TaskJPAAuditService auditService = new TaskJPAAuditService(emf);
```

```
AuditTaskDeleteBuilder updateBuilder =
auditService.auditTaskDelete().deploymentId("org.jbpm:HR:1.0");
int result = updateBuilder.build().execute();
```

Example 17.6. Executor error and requests removal

```
import org.jbpm.executor.impl.jpa.ExecutorJPAAuditService;
import
org.kie.internal.runtime.manager.audit.query.ErrorInfoDeleteBuilder
;
import
org.kie.internal.runtime.manager.audit.query.RequestInfoLogDeleteBu
ilder;

ExecutorJPAAuditService auditService = new
ExecutorJPAAuditService(emf);
ErrorInfoDeleteBuilder updateBuilder =
auditService.errorInfoLogDeleteBuilder().dateRangeEnd(new Date());
int result = updateBuilder.build().execute();

RequestInfoLogDeleteBuilder updateBuilder2 =
auditService.requestInfoLogDeleteBuilder().dateRangeEnd(new
Date());
result = updateBuilder.build().execute();
```

Note

When removing executor entries, ensure that the error information is removed before the request information because of constraints setup on database.

Warning

Parts of the code utilize internal API. While this does not have any direct impact on the functionality of our product, internal API is subject to change and Red Hat cannot guarantee backward compatibility.

CHAPTER 18. LOCALIZATION AND CUSTOMIZATION

18.1. AVAILABLE LANGUAGES

The Red Hat JBoss BPM Suite web user interface can be viewed in multiple languages:

- United States English (**en_US**)
- Spanish (**es_ES**)
- Japanese (**ja_JP**)
- Simplified Chinese (**zh_CN**)
- Traditional Chinese (**zh_TW**)
- Portuguese (**pt_BR**)
- French (**fr_CA**)
- German (**de_DE**)

> **Note**
>
> If a language is not specified, US English is used by default.

> **Note**
>
> Dashbuilder does not support Traditional Chinese (**zh_TW**).

18.2. CHANGING LANGUAGE SETTINGS

Changing the User Interface Language in Business Central

By default, Business Central uses the system locale. If you need to change it, then append the required locale code at the end of the Business Central URL. For example, the following URL will set the locale to Portuguese (pt_BR).

```
http://localhost:8080/business-central/?locale=pt_BR
```

Changing the User Interface Language in Dashbuilder

To change the user interface language in dashbuilder, do the following:

1. Log into the dashbuilder after the server has been successfully started by navigating to http://*localhost*:8080/dashbuilder in a web browser.

2. Select the language of your choice by clicking on the available locales on the top center of the dashbuilder user interface to change the language.

Setting a Default User Interface Language in Dashbuilder

Following is an example to set the default user interface language in dashbuilder:

Procedure: Setting the Default Language as French

1. Navigate to **jboss-eap-6.4/standalone/configuration** and define the following in the **standalone.xml** file.

   ```
   <system-properties>
     <property
   name="org.jboss.dashboard.LocaleManager.installedLocaleIds"
   value="en,es,de,fr,ja,pt,zh"/>
     <property
   name="org.jboss.dashboard.LocaleManager.defaultLocaleId"
   value="fr"/>
   </system-properties>
   ```

2. The default user interface language of the dashbuilder is now set to French.

Defining the Installed Locales in Dashbuilder

Following is an example to define the installed locales in dashbuilder:

Procedure: Defining the Installed Locale

1. Navigate to **jboss-eap-6.4/standalone/configuration** and define the following in the **standalone.xml** file:

   ```
   <system-properties>
     <property
   name="org.jboss.dashboard.LocaleManager.installedLocaleIds"
   value="en,es,de,fr,ja,pt"/>
      <property
   name="org.jboss.dashboard.LocaleManager.defaultLocaleId"
   value="fr"/>
   </system-properties>
   ```

In this example, the Chinese language (zh) has been removed from the list of installed locales so users will not be able to switch the dashbuilder to Chinese. Dashbuilder will show content in French, which is the default locale. Users will be able to select other languages that are defined (en, es, de, ja, pt) in this file.

> **Note**
>
> Within Business Central, the application server does not need to be restarted after changing locale if you append the "locale" parameter to the URL of Business Central. However, with Dashbuilder, the application server should be restarted after the configuration files have been changed.

18.3. RUNNING THE JVM WITH UTF-8 ENCODING

Red Hat JBoss BPM Suite is designed to work with UTF-8 encoding. If a different encoding system is being used by the JVM, unexpected errors might occur.

To ensure UTF-8 is used by the JVM, use the JVM option **`-Dfile.encoding=UTF-8`**.

PART IV. EXECUTION

CHAPTER 19. PROCESS EXECUTION SERVER CONFIGURATION

19.1. ASSIGNMENT RULES

Assignment rules are rules executed automatically when a Human Task is created or completed. This mechanism can be used, for example, to assign a Human Task automatically to a particular user of a group or prevent a user from completing a Task if data is missing.

19.1.1. Defining Assignment Rules

To define assignment rules, do the following:

1. Create a file that will contain the rule definition on the Business Central classpath (the recommended location is *DEPLOY_DIR*/**standalone/deployments/business-central.war/WEB-INF/classes/**):

   » **default-add-task.drl** with the rules to be checked when the Human Task is created

   » **default-complete-task.drl** with the rules to be checked when the Human Task is completed

2. Define the rules in the file.

Example 19.1. The default-add-task.drl Content

```
package defaultPackage

import org.kie.api.task.model.Task;
import org.kie.api.task.model.User;
import org.kie.api.task.model.Status;
import org.kie.api.task.model.PeopleAssignments;
import org.jbpm.services.task.rule.TaskServiceRequest;
import
org.jbpm.services.task.exception.PermissionDeniedException;
import org.jbpm.services.task.impl.model.*;
import java.util.HashMap;
import java.util.List;


global TaskServiceRequest request;

rule "Don't allow Mary to complete task when rejected"
    when
        $task : Task()
        $actualOwner : User( id == 'mary') from
$task.getTaskData().getActualOwner()
        $params : HashMap(this["approved"] == false)
    then
```

```
                request.setAllowed(false);

request.setExceptionClass(PermissionDeniedException.class);
        request.addReason("Mary is not allowed to complete task
with approved false");
end
```

If the potential owners of a Human Task will contain the user **Mary**, the task will be automatically assigned to the user **mary**.

Example 19.2. The default-complete-task.drl Content

```
package defaultPackage

import org.kie.api.task.model.Task;
import org.kie.api.task.model.User;
import org.kie.api.task.model.Status;
import org.kie.api.task.model.PeopleAssignments;
import org.jbpm.services.task.rule.TaskServiceRequest;
import
org.jbpm.services.task.exception.PermissionDeniedException;
import org.jbpm.services.task.impl.model.*;
import java.util.HashMap;
import java.util.List;


global TaskServiceRequest request;

rule "Don't allow Mary to complete task when rejected"
    when
        $task : Task()
        $actualOwner : User( id == 'mary') from
$task.getTaskData().getActualOwner()
        $params : HashMap(this["approved"] == false)
    then
        request.setAllowed(false);

request.setExceptionClass(PermissionDeniedException.class);
        request.addReason("Mary is not allowed to complete task
without approval.");
end
```

If the potential owners of a Human Task will contain the user **Mary**, the task will be automatically assigned to the user **mary**.

19.2. MAIL SESSION

Mail session defines the mail server properties that are used for sending emails if required by the application, such as, escalation or notification mechanisms (refer to the *Red Hat JBoss BPM Suite User Guide*).

19.2.1. Setting up Mail Session

To set up the mail session for your execution engine, do the following:

1. Open the respective profile configuration file (**standalone.xml** or **host.xml**) for editing.

2. Add the mail session to the **urn:jboss:domain:mail:1.1** subsystem.

   Example 19.3. New Mail Session on localhost

   ```
   <subsystem xmlns="urn:jboss:domain:mail:1.1">
     <!-- omitted code -->

     <mail-session jndi-name="mail/jbpmMailSession"
   debug="true" from="bpms@company.com">
       <smtp-server outbound-socket-binding-ref="bpmsMail"/>
     </mail-session>
   </subsystem>
   ```

3. Define the session outbound socket in the profile configuration file.

   Example 19.4. Outbound Socket Definition

   ```
   <outbound-socket-binding name="bpmsMail">
     <remote-destination host="localhost" port="12345"/>
   </outbound-socket-binding>
   ```

PART V. MONITORING

CHAPTER 20. PROCESS MONITORING

20.1. JBOSS OPERATIONS NETWORK

A JBoss Operations Network plug-in can be used to monitor rules sessions for Red Hat JBoss BPM Suite.

Due to a limitation of passing the JVM monitoring arguments via the Maven command line, all `com.sun.management.jmxremote.*` parameters must be passed to the Red Hat JBoss BPM Suite application via the `pom.xml` configuration file.

Please refer to the **JBoss Operations Network** *Installation Guide* for installation instructions for the JBoss ON server.

20.2. INSTALLING THE JBOSS BRMS PLUG-IN INTO JBOSS ON

Red Hat JBoss BRMS plug-in for JBoss Operations Network can be installed by either copying the plug-in JAR files to the JBoss Operations Network plug-in directory or through the JBoss Operations Network GUI.

The following procedure guides the user to copy the plug-in JAR files to the JBoss Operations Network plug-in directory:

Procedure: Copying the JBoss BRMS plug-in JAR files

1. Extract the JBoss BRMS plug-in pack archive to a temporary location. This creates a subdirectory with the name **jon-plugin-pack-brms-bpms-3.3.0.GA**. For example:

   ```
   [root@server rhq-agent]# unzip jon-plugin-pack-brms-bpms-
   3.3.0.GA.zip -d /tmp
   ```

2. Copy the extracted JBoss BRMS plug-in JAR files from the **jon-plugin-pack-brms-bpms-3.2.0.GA/** directory to the JBoss ON server plug-in directory. For example:

   ```
   [root@server rhq-agent]# cp /tmp/jon-plugin-pack-brms-bpms-
   3.3.0.GA/*.jar /opt/jon/jon-server-3.3.0.GA1/plugins
   ```

3. Start the JBoss Operations Network server to update the JBoss BRMS plug-in.

To upload the JBoss BRMS plug-in through the JBoss Operations Network GUI, follow this procedure:

Procedure: Uploading the JBoss BRMS plug-in through GUI

1. Start the JBoss Operations Network Server and Log in to access the GUI.

2. In the top navigation of the GUI, open the **Administration** menu.

3. In the **Configuration** area on the left, select the **Server Plugins** link.

4. At the bottom of the list of loaded server plug-ins, click the **Upload a plugin** button and choose the BRMS plugin.

5. The JBoss BRMS plug-in for JBoss Operations Network is now uploaded.

20.3. MONITORING KIE BASES AND KIE SESSIONS

In order for JBoss Operations Network to monitor KieBases and KieSessions, MBeans must be enabled.

MBeans can be enabled either by passing the parameter: **-kie.mbeans = enabled** or via the API:

```
KieBaseConfiguration kbconf =
KieServices.Factory.get().newKieBaseConfiguration();
kbconf.setOption(MBeansOption.ENABLED);
```

> **Note**
>
> **Kie Services** have been implemented for JBoss BRMS 6; for JBoss BRMS 5, **Drools Services** was the naming convention used and it had different measurements on sessions. For example, **activation → match** renaming occured in the updated version.

Please refer to the *JBoss Operations Network Resource Monitoring and Operations Reference* guide for information on importing Kie Sessions into the Inventory View for monitoring purposes.

CHAPTER 21. MANAGING SECURITY FOR RED HAT JBOSS BPM SUITE DASHBUILDER

21.1. ACCESSING RED HAT JBOSS BPM SUITE DASHBUILDER

Dashbuilder is the Red Hat JBoss BPM Suite web-based user interface for Business Activity Monitoring. To access the Dashbuilder from Business Central, go to **Dashboards → Process & Task Dashboards**.

The displayed dashboard provides statistics on runtime data selected on the left. You can create your own dashboard in the Dashbuilder. To do so, display the Dashbuilder by clicking **Dashboards → Business Dashboards**.

21.2. MANAGING SECURITY

To manage security, you can define custom authorization policies to grant or deny access to workspace, page, or panel instances per role.

Defined below is a list of the available roles for Dashbuilder:

- `admin`: Administrates the Red Hat JBoss BPM Suite system. Has full access rights to make any changes necessary. Also has the ability to add and remove users from the system.

- `developer`: Implements code required for process to work. Mostly uses the JBDS connection to view processes, but may use the web tool occasionally.:

- `analyst`: Responsible for creating and designing processes into the system. Creates process flows and handles process change requests. Needs to test processes that they create. Also creates forms and dashboards.

- `user`: Daily user of the system to take actions on business tasks that are required for the processes to continue forward. Works primarily with the task lists.

- `manager`: Viewer of the system that is interested in statistics around the business processes and their performance, business indicators, and other reporting of the system and people who interact with the system.

Thanks to the permissions system, you can build a workspace structure with several pages, menus, and panels and define what pages and panels within a page will be visible for each role. You can also define special types of users and give them restricted access to certain tooling features, or even restricted access to a page subset.
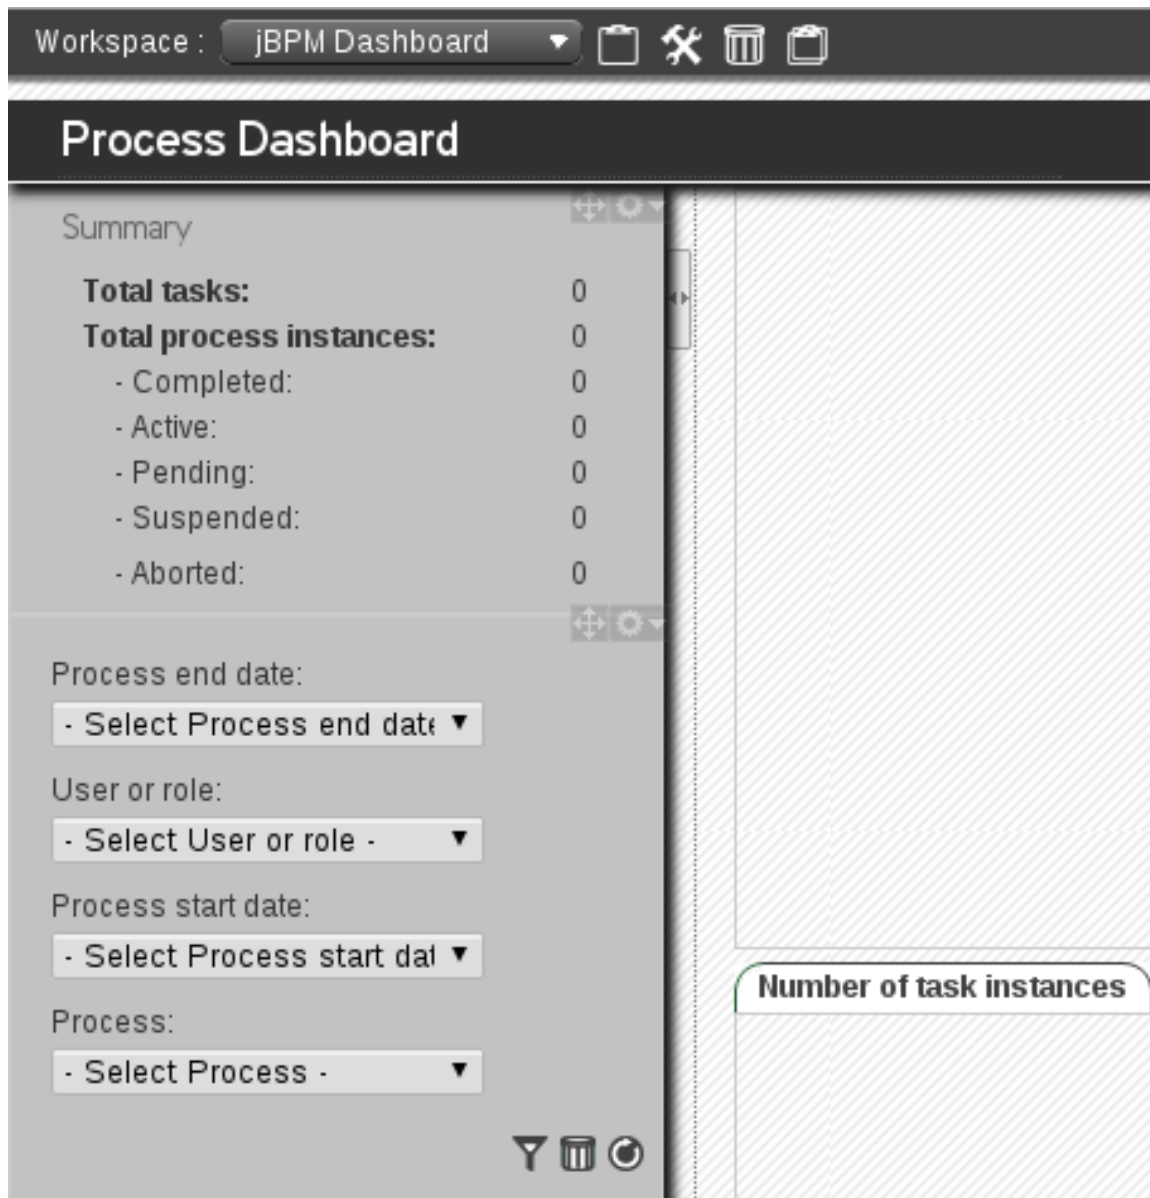
21.3. WORKSPACE PERMISSIONS

Procedure: Accessing Workspace Permissions

1. Log into Business Dashboards from Business Central (as described in the Accessing Red Hat JBoss BPM Suite Dashbuilder topic).

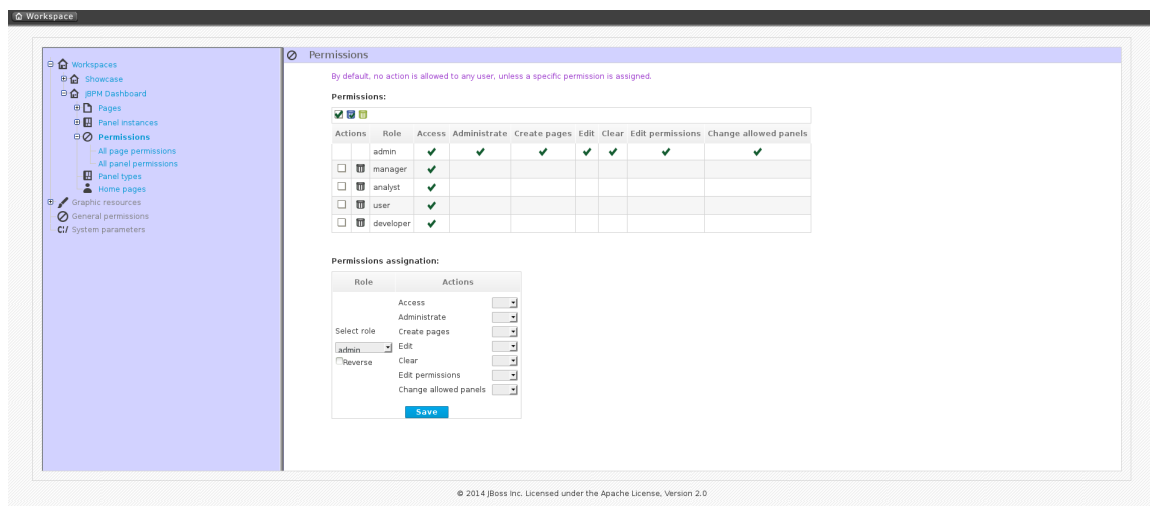2. Select the appropriate Dashboard from the Wokspace drop-down:

   Figure 21.1. Dashbuilder Workspace

3. Click the ⚒ button to access the Workspace Dashboard.

4. Click the *Permissions* label to view the permission management screen.

Figure 21.2. Permissions Screen

Under the **Permissions** section is a list of allowed actions that are applied to the selected role:

» **Access**: Permission to login into the application.

» **Administrate**: Permission to access the toolbar and system configuration features.

» **Create pages**: Ability to create new project pages.

» **Edit**: Permission to change the workspace properties.

» **Clear**: Ability to delete the workspace.

» **Edit permissions**: Ability to grant/deny permissions.

» **Change allowed panels**: Permission to restrict the type of panels that can be used in this workspace.

To assign a permission, you must select the target role and the list of actions allowed over the selected resource:

Figure 21.3. Permissions assignment

» _Target roles_ (who): What user will be granted/denied with the permissions defined.

» _Allowed actions_: Depending on the type of the resource we can enable/disable what the user can do on this resource.

» _Reverse_ (optional): When we have a set of roles and we want to grant/deny a permission to all the roles but one.

Note

By default, the full set of permissions go to the role `admin`. This makes it easy to create a user that can do everything as long as the role `admin` is assigned.

21.4. PAGE PERMISSIONS

1. To access **Page permissions**, locate the **Pages** drop-down under the jBPM Dashboard (or whatever Dashboard you selected).

2. After expanding **Pages**, expand the **Process dashboard** option.

3. Select the **Page permissions** option.

Figure 21.4. Page Permissions
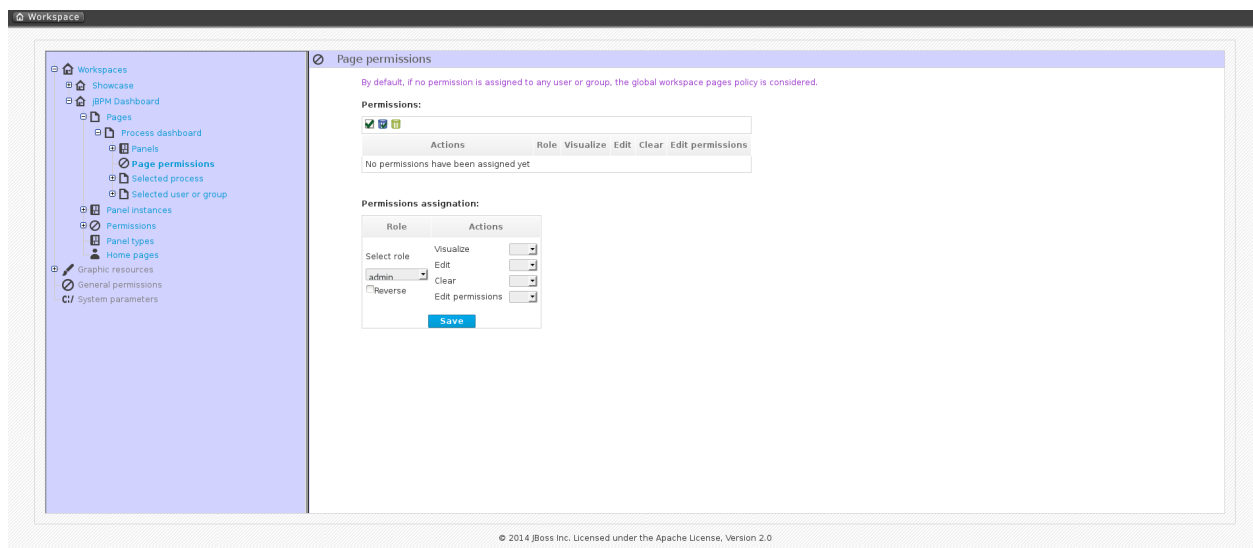
Under the **Permissions** section is a list of allowed actions that are applied to the selected role:

▸ **Visualize**: Permission to make the page visible.

▸ **Edit**: Ability to change the page properties.

▸ **Clear**: Ability to delete the page.

▸ **Edit permissions**: Ability to grant/deny permissions for the page.

21.5. PANEL PERMISSIONS

1. To access the **Panel permissions** page, expand the **Panel instances** option under the jBPM Dashboard (or whatever Dashboard you are using).

2. Expand the **Dashboard** option and then expand the **Process dashboard**.

3. Expand the **Panels** choice and select the appropriate process.

4. Open the **Panel permissions** page.

Below is a screenshot of the permission management screen for a given panel (in this example, the Process dashboard):

Figure 21.5. Panel permissions configuration screen
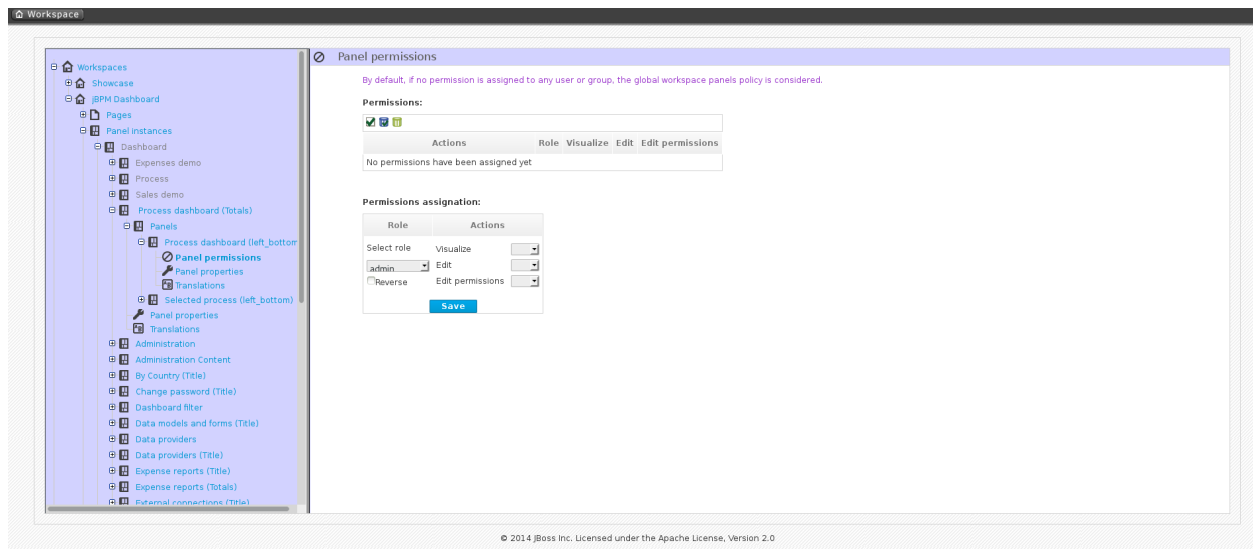
Allowed actions are the following:

➤ **Visualize**: Make the panel visible.

➤ **Edit**: Change the panel properties.

➤ **Edit permissions**: Ability to grant/deny permissions for the panel.

APPENDIX A. REVISION HISTORY

Note that revision numbers relate to the edition of this manual, not to version numbers of Red Hat JBoss BRMSBPM Suite.

Revision 6.3.0-11	**Thu Sep 15 2016**	**Tomas Radej**
Rebuilt.		
Revision 6.3.0-10	**Thu Sep 15 2016**	**Tomas Radej**
Published the AsciiDoc version of the docs.		
Revision 6.3.0-8	**Thu Jul 14 2016**	**Tomas Radej**
Updated documentation with release 6.3.1.		
Revision 6.3.0-7	**Thu Jun 2 2016**	**Marek Czernek**
Releasing the newest documentation.		
Revision 6.3.0-6	**Thu May 5 2016**	**Tomas Radej**
Fixed linking errors.		
Revision 6.3.0-5	**Thu May 5 2016**	**Tomas Radej**
Built with live links.		
Revision 6.3.0-4	**Thu May 5 2016**	**Tomas Radej**
Bootstrapping links.		
Revision 6.3.0-3	**Thu Apr 28 2016**	**Tomas Radej**
All books rebuilt.		
Revision 6.3.0-2	**Thu Apr 28 2016**	**Tomas Radej**
All books rebuilt.		
Revision 6.3.0-1	**Thu Apr 28 2016**	**Tomas Radej**
Initial build for release 6.3.0 of JBoss BPM SuiteJBoss BRMS.		