

RiftSaw 2.2.0.CR2

User Guide

by Gary Brown, Kurt Stam, Heiko Braun, and Jeff Yu

1. Introduction	1
1.1. Overview	1
2. Administration	2
2.1. Overview	2
2.2. BPEL Console	2
2.2.1. Overview	2
2.2.2. Logging in	2
2.2.3. Deployed Process Definitions	3
2.2.4. Process Instances	4
2.2.5. Process Versions	4
2.2.6. Execution History	5
2.2.7. Instance Data and Execution Path	6
2.2.8. Retiring and Reactivating Process Definitions	7
2.3. BPEL Properties	8
2.4. Enable Execution Events	9
2.5. How to safely undeploy processes	10
3. Deploying BPEL Processes	11
3.1. Overview	11
3.2. Direct deployment to JBossAS server	11
3.3. Eclipse based Deployment	12
3.4. Changing Endpoint Configuration Properties	19
4. Web Service Configuration	20
4.1. Overview	20
4.2. Configuring a JAX-WS Handler	20
4.3. Apache CXF Configuration	21
4.3.1. Configuring the Server endpoint	21
4.3.2. Configuring the Client endpoint	23
4.4. Deployed Service WSDL	25
5. UDDI Integration	26
5.1. Overview	26
5.2. UDDI config properties	26
5.3. Default configurations	28
5.4. Other UDDI v3 Registries	28
5.5. UDDI Registry Entities and UDDI Seed Data	28
6. JBoss ESB Integration	29
6.1. Overview	29
6.2. Using the BPELInvoke ESB action	29
6.2.1. Fault Handling	31
6.2.2. SAML Support	32
7. RiftSaw Clustering Support	34
7.1. Overview	34
7.2. Installation	34
7.3. Deployment	35
7.4. Bpel Process Service Invocation	35

8. Database	36
8.1. Upgrade database schema	36
8.2. Database schema diagram	36

Introduction

1.1. Overview

This is the User Guide for the RiftSaw BPEL process engine.

RiftSaw provides a JBoss AS integration for the Apache ODE BPEL engine. For detailed information on executing BPEL processes within Apache ODE, we would refer the reader to the [Apache ODE](#) website and documentation.

In addition to the ability to run the Apache ODE engine within JBoss AS, the RiftSaw project also provides a GWT based administration console, replaces the Axis2 based transport with JBossWS (which can be configured to use Apache CXF), and provides tighter integration with JBossESB.

Administration

2.1. Overview

This section describes the administration capabilities associated with RiftSaw.

2.2. BPEL Console

2.2.1. Overview


This section provides an overview of the BPEL Console. The console provides the ability to view:

- The process definitions deployed to the BPEL engine
- The process instances executing in the BPEL engine
- The execution history of a process

2.2.2. Logging in

The BPEL console can be located using the URL: <http://localhost:8080/bpel-console>.

The first screen that is presented is the login screen:



The default username is *admin* with password *password*.

The Access Control mechanism used by the admin console is configured in the `$deployFolder/bpel-console/bpel-identity.sar/META-INF/jboss-service.xml`. The JAAS login module is initially set to use a property file based access mechanism, but can be replaced to use any appropriate alternative implementation.

The users for the default mechanism are configured in the property file `$deployFolder/bpel-console/bpel-identity.sar/bpel-users.properties`. The entries in this file represent *username=password*.


The user roles for the default mechanism are configured in the property file `$deployFolder/bpel-console/bpel-identity.sar/bpel-roles.properties`. The entries in this file represent *username=role*. The only role of interest currently is *administrator*.

2.2.3. Deployed Process Definitions

Once logged in, the 'Manage Instances' tab shows the currently deployed BPEL processes and their versions.

2.2.4. Process Instances

When a process definition is selected (*Open*), the list of active process instances for that process definition (and version) will be displayed in the bottom panel.




2.2.5. Process Versions

There is only one active version at a time. If you open a process definition the active version will be chosen by default. Sometimes you need to manage instances of a retired version (i.e. to terminate running instances). In that case the button *More - Change Version* allows you to select a different version of that process.


HelloGoodbye
http://www.jboss.org/bpel/examples

Version: 1 (active)



2.2.6. Execution History

The *Execution History* allows you to inspect the BPAF history data associated with a process. You can specify a timespan and whether or not you'd like to see failed/terminated instances included in the chart.



The chart has various keyboard bindings to browse through the displayed execution data:

Table 2.1.

Keyboard and Mouse Commands			
Up Arrow	Zoom In	Down Arrow	Zoom Out
Left Arrow	Half-Page Left	Right Arrow	Half-Page Right
Page-Up	Page Left	Page-Down	Page Right
TAB	Next Focus	Shift-TAB	Previous Focus
HOME	Max Zoom Out	ENTER	Max Zoom In to Focus

Keyboard and Mouse Commands

Mouse Drag	Scroll Chart	Shift Mouse Drag	Drag Select/Zoom
Mouse Wheel Up/Z	Zoom In	Mouse Wheel Down/X	Zoom Out
Backspace/Back Button	Back	Right Mouse Button	Context Menu
Left Click	Set Focus	Double Click	Max Zoom In to Focus

2.2.6.1. Logging configuration options

You need to explicitly enable history logging for a particular process through the 'deploy.xml' file and for the BPEL engine in general through the 'bpel.properties' file. In order to record history data, set the 'process-events' option for a particular process and make sure the 'org.jboss.soa.bpel.console.bpaf.BPAFLogAdapter' is enabled.


Process configuration example (deploy.xml):

```
<deploy xmlns="http://www.apache.org/ode/schemas/dd/2007/03"
  xmlns:bp1="http://www.jboss.org/bpel/examples"
  xmlns:intf="http://www.jboss.org/bpel/examples/wsdl">

  <process name="bp1:HelloGoodbye">
    <active>true</active>
    <process-events generate="all"/>
    <provide partnerLink="helloGoodbyePartnerLink">
      <service name="intf:HelloGoodbyeService" port="HelloGoodbyePort"/>
    </provide>
  </process>
</deploy>
```

2.2.7. Instance Data and Execution Path

When a process instance is selected, its details will be displayed in the *Execution Details* panel. The *Instance Data* button will also become enabled, allowing further detail about the process to be displayed. Likewise the *Execution Path* button opens the related instance execution graph.




2.2.8. Retiring and Reactivating Process Definitions

When a process definition is initially deployed (i.e. the first version of the process), it automatically becomes the active process definition. If that BPEL process definition is subsequently change and redeployed, then the previous version is *retired*, and the new version becomes the *active* version.

The only difference between an *active* and *retired* process definition is that a *retired* process definition can no longer create new process instances. However if there are current process instances associated with the *retired* process definition version, then these will continue to execute.

On some occasions, the administrator may wish to change which version of a process definition is considered the *active* version. Or they may simply want to *retire* the currently active process definition, so that no more process instances can be created, only allowing the already running process instances to continue until completed.

To change the status of a process definition, the administrator should select the *Runtime* tab from the lefthand panel, and then select the *Deployments* option. This will show the process definitions, their versions and their current status (active or retired).



To change a particular version from *retired* to *active*, simply select the *retired* version and press the *Activate* button in the bottom right.

To retire a currently active process definition, simply select the particular version and then press the *Retire* button in the bottom right.



Note

If you deploy a process definition version into RiftSaw server, the previous version of that process definition will go to 'retire' process automatically. Also please noted that if you undeploy a process, the associated endpoints get deactivated only if it doesn't have any previous version of that process.

2.3. BPEL Properties

When RiftSaw has been installed within the JBossAS environment, there is a property file located at `${JBossAS}/server/default/deploy/riftsaw.sar/bpel.properties`.

This property file contains a number of properties that are specific to ODE, and if interested in these properties, then you should refer to the ODE documentation. Only one point to note, the name of the property in this file maybe prefixed with *bpel.*, however in the ODE documentation the prefix would be *ode.*

This section will present the properties that are specific to RiftSaw.

Table 2.2. RiftSaw specific properties

<code>bpel.uddi.*</code>	These properties relate to the UDDI support, which is discussed in a subsequent chapter.
<code>bpel.jaxws.client.initializer.impl</code>	This property is automatically set upon installation, based on the JAXWS stack being used. This value should not be changed.
<code>bpel.ws.stableInterface</code> (default false)	This property determines whether the Web Service interface, associated with a BPEL process, will be updated when a new version of the BPEL process is deployed. The benefit of setting this to <i>false</i> is that changes to the WSDL will be made active with the BPEL process. However the issue is that during the transition between the interfaces, the web service will momentarily be unavailable - which may cause heavily used services to reject requests. By setting this value to <i>true</i> , then the web service will remain available while the BPEL process is updated, however any changes in the WSDL will not be made available.
<code>bpel.event.listeners</code> (<code>org.jboss.soa.bpel.console.bpaf.BPAFLogAdapter</code>)	Enables the BPAF logging used by the console to display the process execution history.

2.4. Enable Execution Events

Apache ODE generates a set of events on the process running. below is the steps for you to enable it. for more detail information on this, please consult the [ODE's Execution events guide](#).

- In `deploy.xml`, adding `<process-events generate="all"/>`.
- Adding following two properties in the `bpel.properties`.

```
bpel.event.listeners=org.apache.ode.bpel.common.evt.DebugBpelEventListener
debugeventlistener.dumpToStdOut=on
```

By registering the `DebugBpelEventListener`, it will use the log's `DEBUG` level for the `BpelEvent` information. By setting the `debugeventlistener.dumpToStdOut=on`, it will render the information on the console as well.

We've also provided an out-of-box example to demonstrate the adding a custom listener. Please see the `bpel_event_listener` example for more information.

2.5. How to safely undeploy processes

If an active process definition is undeployed, any active process instances will be terminated. If this is not desired behaviour, then the administrator should first retire the process definition (version) using the BPEL console, and then when all active process instances for the version have completed, it can be safely undeployed.

So below is the recommended steps for safely undeploy processes.

- Check if the processes have any active process instances.
- If it has unfinished process instances, retire the process in BPEL Console first.
- Undeploy the processes only if they don't have any active process instances.

Deploying BPEL Processes

3.1. Overview

This section outlines the mechanisms that can be used to deploy a BPEL process to RiftSaw BPEL engine running within a JBoss AS server.

3.2. Direct deployment to JBossAS server

The direct deployment approach is demonstrated using an *Ant* script in each of the quickstart examples. For example,

```
<!-- Import the base Ant build script... -->
<property file="../../install/deployment.properties" />

<property name="version" value="1" />

<property name="server.dir" value="${org_jboss_as_home}/server/${org_jboss_as_config}"/>
<property name="conf.dir" value="${server.dir}/conf"/>
<property name="deploy.dir" value="${server.dir}/deploy"/>
<property name="server.lib.dir" value="${server.dir}/lib"/>

<property name="sample.jar.name" value="${ant.project.name}-${version}.jar" />

<target name="deploy">
    <echo>Deploy ${ant.project.name}</echo>
    <jar basedir="bpel" destfile="${deploy.dir}/${sample.jar.name}" />
</target>

<target name="undeploy">
    <echo>Undeploy ${ant.project.name}</echo>
    <delete file="${deploy.dir}/${sample.jar.name}" />
</target>
```

This excerpt from the *Ant* build file for the *hello_world* quickstart example shows that deploying a RiftSaw BPEL process using *Ant* is very straightforward. The main points of interest are:

- It is necessary to identify the location of the JBoss AS server in which the BPEL process will be deployed. This is achieved in this example by referring to the `deployment.properties` file that has been configured in the RiftSaw distribution (install folder).
- If a versioned approach is being used, so that multiple versions of the same BPEL process may be deployed at one time, then the name of the archive (jar) containing the BPEL process (and associated artifacts) has a version number suffix. This would need to be manually incremented for each distinct version of the BPEL process being deployed.



Warning


Currently the version must be specified as a single integer value. Non-numeric values, such as versions expressed in a major.minor.incremental (maven style), will result in an exception when deployed to the server.

- The next step is to define the *deploy* target, which will create the BPEL process archive, using the contents of the *bpel* sub-folder in this case, and store it within the JBoss AS server's *deploy* folder.
- The final step is to define the *undeploy* target, which simply removes the BPEL process archive from the JBoss AS server's *deploy* folder.


3.3. Eclipse based Deployment

This section will explain how to deploy an Eclipse BPEL project to the RiftSaw BPEL engine running in a JBossAS server.


The first step is to create or import the Eclipse BPEL project. In this case we are going to import an existing project from the `${RiftSaw}/samples/quickstart/hello_world` folder. This can be achieved by selecting the *Import ...* menu item associated with the lefthand navigator panel in Eclipse, and then select the *General->Existing Projects into Workspace* entry and press the *Next* button.



Then press the *Browse* button and navigate to the *hello_world* quickstart folder. Once located, press the *Finish* button.




Once the project has been imported, you can inspect the contents, such as the BPEL process and WSDL description.



The next step is to create a server configuration for the JBoss AS environment in which the RiftSaw BPEL engine has previously been installed. From the Eclipse *Java EE* perspective, the *Server* tab should be visible in the lower region of the Eclipse window. If this view is not present, then go to the *Window->Show Views->Servers* menu item to open the view explicitly.


In the *Servers* view, right click and select the *New->Server* menu item.




Select the appropriate JBoss AS version, and then press *Finish*.

Before being able to deploy an example, we should start the new server. This can be achieved by right clicking on the server in the *Servers* tab, and selecting the *Start* menu item. The output from the server will be displayed in the *Console* tab.

Once the server has been started, right click on the server entry again, and select the *Add and Remove ...* menu item.




Select the *Quickstart_bpel_hello_world* project, press the *Add* button and then press the *Finish* button. This will cause the project to be deployed to the server.




Once the project has been deployed, it will show up as an entry below the server in the *Servers* tab.


The final step is to test the deployed BPEL process. In this example, we can do this using the *ant* script provided with the quickstart sample. Right click on the *build.xml* file in the root folder of the project, and select the *Run As->Ant Build ...* menu item. NOTE: It is important to select the menu item with the "...", as this provides a dialog window to enable you to select which *ant target* you wish to perform.



Deselect the *deploy* target, and select the *sendhello* target, before pressing the *Run* button. This was send a test 'hello' message to the server, and then display the response in the *Console* tab.




If you now want to update the BPEL process, select the *assignHelloMesg* node in the diagram, and select the *Properties* view. On the left of the view is a vertical list of tabs. Select the *Details* tab. Then select the "Expression to Variable" from the list, and update the *concat* function's second parameter - for example to add 'UPDATED' to the text.



Once the update has been saved, go to the *Server View* and select the *Full Publish* option from the menu associated with the *Quickstart_bpel_hello_world* project. This will cause the project to be re-deployed to the RiftSaw server.

The final step is to then re-run the 'sendhello' target within the *build.xml* file, to send a new request, and view the response. The response should now be modified according to the changes you made in the BPEL process.



If you expand the deployed project node, in the *Server View*, you will see that both of the deployed versions are shown. The older version is retained as there may still be BPEL process instances using that version of the process. If you right-click on each of the child nodes, you will see that it is also possible to undeploy the specific versions. However, if you explicitly undeploy a version, then any remaining active process instances for that version will be terminated.

You can then use the menu associated with the project, contained in the *Server View*, to undeploy the project (using the *Add and Remove ...* menu item) and finally use the menu associated with the server itself to *Stop* the server.

3.4. Changing Endpoint Configuration Properties

Apache ODE provides the means to customise certain properties, associated with a BPEL endpoint, by specifying the properties in a file with an extension of `.endpoint`.

For information on the properties that can be specified in this file, please see the Apache ODE documentation, located at: <http://ode.apache.org/endpoint-configuration.html>.



Note

RiftSaw currently only supports the following properties: *mex.timeout*

This section explains how to deploy these `.endpoint` files as part of a RiftSaw deployment.

Apache ODE supports two locations for finding these `.endpoint` files. A 'global' configuration folder, which by default is `ode/WEB-INF/conf`, and a process deployment specific location, which is `ode/WEB-INF/processes/$your_process`. Properties associated with the 'global' configuration override any property values provided in the process specific location.

RiftSaw currently does not support a 'global' configuration location, so it will only obtain the configuration files from the deployed BPEL bundle. More specifically, from the root location within the BPEL deployment unit, along side the BPEL deployment descriptor.

So, for example, if you place a file called `test.endpoint` in the `${RiftSaw}/samples/quickstart/hello_world/bpelContent` folder, with the following content:

```
# 3 minutes
mex.timeout=180000
```

then once deployed, the helloworld example could wait up to a maximum of 3 minutes to respond. To test this out, edit the `hello_world.bpel` and insert a `wait` activity before the response - similar to the following:

```
<wait>
  <for>'PT150S'</for>
</wait>
```

This will wait 2 minutes 30 seconds before responding, which is 30 seconds more than the default timeout, but still within the new timeout period specified within the `test.endpoint` file. If you then wish to try forcing the timeout to occur, simply increase the wait duration to 3 minutes 30 seconds, and resubmit the test message using the `ant sendhello` command.

Web Service Configuration

4.1. Overview

This section outlines the mechanisms that are available for configuring the web service stack used in providing the web service for a BPEL process, as well as invoking external web services from a BPEL process.

4.2. Configuring a JAX-WS Handler

JAX-WS is a standard Java API for client and server support of web services. The JAX-WS handler mechanism can be used by a client or server (i.e. the web service) to invoke a user specified class whenever a message (or fault) is sent or received. The handler is therefore installed into the message pipeline, and can manipulate the message header or body as required.

The handlers are usually installed either programmatically, or through a *HandlerChain* annotation on the Java interface representing the Web Service. However, in the case of a BPEL process deployed to RiftSaw, the JAX-WS service (representing the web service associated with the BPEL process) is dynamically created on deployment.

Therefore to associate the configuration of a JAX-WS handler chain with the Web Service dynamically created to support the BPEL process, the user must place a file called `jws_handler.xml` alongside the BPEL process definition and deployment descriptor.

The following provides an example of the XML configuration associated with the `jws_handler.xml` file. This particular example is used by the *service_handler* quickstart sample.

```
<handler-chains 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">
  <handler-chain>
    <handler>
      <handler-name>JAXWSHandler</handler-name>
      <handler-class>org.jboss.soa.bpel.examples.jaxws.JAXWSHandler</handler-class>
      <init-param>
        <param-name>TestParam</param-name>
        <param-value>TestValue</param-value>
      </init-param>
    </handler>
  </handler-chain>
</handler-chains>
```

The format of this file is the standard JAX-WS handler chain configuration. One or more handler elements can be specified, with each handler defining a name and class. The handler configuration can optionally provide initialization parameters that are passed to the *init* method on the handler implementation.



Note

This mechanism only installs JAX-WS handlers on the 'provider' web service. It is not currently possible to configure JAX-WS handlers for the client endpoints that invoke external web services from a BPEL process.

An example of this mechanism can be found in the *service_handler* quickstart sample.

4.3. Apache CXF Configuration

RiftSaw integrates with JBossWS, using the JAX-WS standard API, to support the following web service stacks: JBossWS native and Apache CXF. This section explains how RiftSaw deployed BPEL processes can include additional configuration specifically applicable to the Apache CXF web service stack - and is therefore only relevant if the JBossAS application server has been configured to use this stack. See the *Getting Started Guide* for information on how to switch to the Apache CXF stack when installing RiftSaw.

This section will explain how web service endpoints, whether server (i.e. representing the BPEL process) or client (i.e. being used to invoke external web services), are configured using the Apache CXF configuration format. It will also discuss reasons why you may wish to do this additional CXF specific configuration. However, for further information on how to configure CXF, and the features that it offers, the reader is referred to the Apache CXF website <http://cxf.apache.org>.

4.3.1. Configuring the Server endpoint

To create a CXF configuration that will be used by the RiftSaw web service provider (i.e. the server), it is simply a case of placing a file called `jbossws-cxf.xml` into the root folder of the BPEL deployment (along side the deployment descriptor).

This is the same filename as used by `jbossws-cxf`, when deploying a web service based on the use of JAXWS annotations. An example of the file content is:

```
<beans
  xmlns='http://www.springframework.org/schema/beans'
  xmlns:xsi='http://www.w3.org/2001/XMLSchema-instance'
  xmlns:beans='http://www.springframework.org/schema/beans'
  xmlns:jaxws='http://cxf.apache.org/jaxws'
  xsi:schemaLocation='http://cxf.apache.org/core
    http://cxf.apache.org/schemas/core.xsd
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
    http://cxf.apache.org/jaxws
    http://cxf.apache.org/schemas/jaxws.xsd'>

  <bean id="UsernameTokenSign_Request"
    class="org.apache.cxf.ws.security.wss4j.WSS4JInInterceptor">
    <constructor-arg>
      <map>
```

```

    <entry key="action" value="UsernameToken Timestamp Signature"/>
    <entry key="passwordType" value="PasswordDigest"/>
    <entry key="user" value="serverx509v1"/>
    <entry key="passwordCallbackClass"
          value="org.jboss.test.ws.jaxws.samples.wsse.ServerUsernamePasswordCallback"/>
    <entry key="signaturePropFile" value="etc/Server_SignVerf.properties"/>
    <entry key="signatureKeyIdentifier" value="DirectReference"/>
  </map>
</constructor-arg>
</bean>

<bean id="UsernameTokenSign_Response"
      class="org.apache.cxf.ws.security.wss4j.WSS4JOutInterceptor">
  <constructor-arg>
    <map>
      <entry key="action" value="UsernameToken Timestamp Signature"/>
      <entry key="passwordType" value="PasswordText"/>
      <entry key="user" value="serverx509v1"/>
      <entry key="passwordCallbackClass"
            value="org.jboss.test.ws.jaxws.samples.wsse.ServerUsernamePasswordCallback"/>
      <entry key="signaturePropFile" value="etc/Server_Decrypt.properties"/>
      <entry key="signatureKeyIdentifier" value="DirectReference"/>
      <entry key="signatureParts"
            value="{Element}{http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-
wssecurity-utility-1.0.xsd}Timestamp;{Element}{http://schemas.xmlsoap.org/soap/
envelope/}Body"/>
    </map>
  </constructor-arg>
</bean>

<jaxws:endpoint
  id='SecureHelloWorldWS'
  address='http://@jboss.bind.address@:8080/Quickstart_bpel_secure_serviceWS'
  implementor='@provider@'>
  <jaxws:inInterceptors>
    <ref bean="UsernameTokenSign_Request"/>
    <bean class="org.apache.cxf.binding.soap.saaj.SAAJInInterceptor"/>
  </jaxws:inInterceptors>
  <jaxws:outInterceptors>
    <ref bean="UsernameTokenSign_Response"/>
    <bean class="org.apache.cxf.binding.soap.saaj.SAAJOutInterceptor"/>
  </jaxws:outInterceptors>
</jaxws:endpoint>

</beans>

```

This example configures the web service to use username token and digital signature authentication.



Note

The `jaxws:endpoint` element has an attribute called *implementor* that defines the Java class implementing the JAXWS service. RiftSaw dynamically creates this class, and therefore

it is important that the attribute is set to the value `@provider@` to enable the dynamically created Java class to be correctly configured during deployment.

4.3.2. Configuring the Client endpoint

When configuring client endpoints, representing web services invoked by a BPEL process, the configuration is currently separated into different files on a per port basis - similar to the approach used by the Axis2 ODE integration.

The file name is of the form `jbossws-cxf-{portname_local_part}.xml`, where the `portname_local_part` represents the local part of the portname of the web service being invoked. For example, if the WSDL for the invoked web service is:

```
<definitions name='SecureHelloWorldWSService'
  targetNamespace='http://secure_invoke/helloworld' .... >
  <portType name='SecureHelloWorld'>
    ...
  </portType>
  <service name='SecureHelloWorldWSService'>
    <port name='SecureHelloWorldPort' ... >
      ...
    </port>
  </service>
</definitions>
```

then the CXF configuration file would be `jbossws-cxf-SecureHelloWorldPort.xml`.

The CXF configuration information within this file is associated with the CXF bus. For example:

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:cxf="http://cxf.apache.org/core"
  xmlns:wsa="http://cxf.apache.org/ws/addressing"
  xmlns:http="http://cxf.apache.org/transports/http/configuration"
  xmlns:wsm-policy="http://schemas.xmlsoap.org/ws/2005/02/rm/policy"
  xmlns:wsm-mgr="http://cxf.apache.org/ws/rm/manager"
  xmlns:beans='http://www.springframework.org/schema/beans'
  xmlns:jaxws='http://cxf.apache.org/jaxws'
  xmlns:ns1='http://secure_invoke/helloworld'
  xsi:schemaLocation="
    http://cxf.apache.org/core http://cxf.apache.org/schemas/core.xsd
    http://cxf.apache.org/transports/http/configuration http://cxf.apache.org/schemas/
configuration/http-conf.xsd
    http://schemas.xmlsoap.org/ws/2005/02/rm/policy http://schemas.xmlsoap.org/ws/2005/02/
rm/wsm-policy.xsd
    http://cxf.apache.org/ws/rm/manager http://cxf.apache.org/schemas/configuration/wsm-
manager.xsd
```

```

    http://www.springframework.org/schema/beans http://www.springframework.org/schema/beans/
    spring-beans.xsd
    http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd">

    <bean id="UsernameTokenSign_Request"
        class="org.apache.cxf.ws.security.wss4j.WSS4JOutInterceptor" >
        <constructor-arg>
            <map>
                <entry key="action" value="UsernameToken Timestamp Signature"/>
                <entry key="passwordType" value="PasswordDigest"/>
                <entry key="user" value="clientx509v1"/>
                <entry key="passwordCallbackClass"
                    value="org.jboss.test.ws.jaxws.samples.wsse.ClientUsernamePasswordCallback"/>
                <entry key="signaturePropFile" value="etc/Client_Sign.properties"/>
                <entry key="signatureKeyIdentifier" value="DirectReference"/>
                <entry key="signatureParts"
                    value="{Element}{http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-
wssecurity-utility-1.0.xsd}Timestamp;{Element}{http://schemas.xmlsoap.org/soap/
envelope/}Body"/>
            </map>
        </constructor-arg>
    </bean>

    <bean id="UsernameTokenSign_Response"
        class="org.apache.cxf.ws.security.wss4j.WSS4JInInterceptor" >
        <constructor-arg>
            <map>
                <entry key="action" value="UsernameToken Timestamp Signature"/>
                <entry key="passwordType" value="PasswordText"/>
                <entry key="user" value="serverx509v1"/>
                <entry key="passwordCallbackClass"
                    value="org.jboss.test.ws.jaxws.samples.wsse.ClientUsernamePasswordCallback"/>
                <entry key="signaturePropFile" value="etc/Client_Encrypt.properties"/>
                <entry key="signatureKeyIdentifier" value="DirectReference"/>
            </map>
        </constructor-arg>
    </bean>

    <cxf:bus>
        <cxf:outInterceptors>
            <ref bean="UsernameTokenSign_Request"/>
            <bean class="org.apache.cxf.binding.soap.saaj.SAAJOutInterceptor"/>
        </cxf:outInterceptors>
        <cxf:inInterceptors>
            <ref bean="UsernameTokenSign_Response"/>
            <bean class="org.apache.cxf.binding.soap.saaj.SAAJInInterceptor"/>
        </cxf:inInterceptors>
    </cxf:bus>

</beans>

```

This example configures the web service client to use username token and digital signature authentication.

4.4. Deployed Service WSDL

When a BPEL process is deployed, it results in a web service being created to represent the service endpoint, based on a WSDL description included with the deployed process. If the WSDL is displayed, using the endpoint URL of the service (with the *?wsdl* suffix), then the *<soap:address>* by default will be associated with the host bind address of the server (as defined in the *\${jboss.bind.address}* property).

To change this default address, for example if wanting to present the web service via a firewall that has a different host address, then you will need to refer to the JBossWS documentation.

As an example, jbossws-native-3.2.2.GA uses a configuration file (*\${jbossas}/server/default/deployers/jbossws.deployer/META-INF/stack-agnostic-jboss-beans.xml*) to define this information. The bean *WSServerConfig* has a property called *webServiceHost* that can be used to define the value to be used.

UDDI Integration

5.1. Overview

The integration of a UDDI client into the RiftSaw runtime codebase allows for the auto-registration of BPEL services to an UDDI registry upon deployment of the service. The registration process uses the jUDDI-3 client libraries which are capable of communicating to any UDDI v3 complaint registry.

Upon deployment both the Service and its BindingTemplate (EndPoint information) are registered, and a partnerLinkChannel is created for each partnerLink. UDDI lookup will obtain the WSDL Endpoint from the UDDI and access this URL to obtain the WSDL straight from the Service. Upon undeployment the BindingTemplate is removed from the UDDI Registry.

5.2. UDDI config properties

By default RiftSaw uses the jUDDI client libraries of JBossESB/SOA-P, and the client configuration is found in the `deploy/jbossesb.sar/esb.juddi.client.xml`. Both the name of the ClerkManager and the Clerk itself are specified in the *bpel.properties* file.

Table 5.1. The UDDI related properties in the *bpel.properties*

attribute	type (default)	description
<code>bpel.uddi.registration</code>	boolean (true)	If set to 'false', the UDDI integration is turned off. The RiftSaw installation process sets this value to 'true' only if the <code>jbossesb-registry.sar</code> is detected containing a jUDDI v3 registry. In all other case it is defaulted to false.
<code>bpel.webservice.secure</code>	boolean (false)	The UDDI Registration process registers an WSDL AccessPoint in the BindingTemplate for the BPEL Service it is registering. The BPEL server exposes the service WSDL Endpoint on the WS stack (Currently we support JBossWS and CXF), if your WS stack is configured to a use a secure (https) protocol, then you need to switch this setting to 'true'. Note that this setting is used during the registration process only.

attribute	type (default)	description
bpel.uddi.client.impl	String (org.jboss.soa.bpel.uddi.UDDIRegistryImpl)	Name of the class that implements the <i>org.jboss.soa.bpel.runtime.engine.ode.UDDIRegistry</i> interface.
bpel.uddi.clerk.config	String (not used by default)	Defines the path to the <i>bpel.uddi.client.xml</i> config file. This can be left commented out if you want to use the <i>jbossesb.sar/esb.uddi.client.xml</i> . However in that case a <i>bpel.uddi.clerk.manager</i> needs to be defined.
bpel.uddi.clerk.manager	String (esb-registry)	Defines the ClerkManager name that will be used if the <i>bpel.uddi.clerk.config</i> is left commented out. This value should correspond to the name of the manager in the <i>esb.juddi.client.xml</i> . Note that if the <i>bpel.uddi.clerk.config</i> is defined, the setting of the <i>bpel.uddi.clerk.manager</i> is ignored.
bpel.uddi.clerk	String (BPEL_clerk)	Defines the Clerk name that will be used. This value should correspond to the name of the clerk in the <i>esb.juddi.client.xml</i> . By default this is set to 'BPEL_clerk'.
bpel.uddi.lookup	boolean (true)	If set to true, the creating process of the partner channel will do a lookup by serviceName in the UDDI, and a WSDL Endpoint is retrieved. This WSDL Endpoint is then used to obtain the WSDL. This process makes it easier to move Endpoints around within your deployment, without having to update the partnerlink WSDL

attribute	type (default)	description
		files in your bpel deployments. Note that an it is still a requirement to deploy the initial partnerlink WSDL file for each partnerLink.

5.3. Default configurations

When RiftSaw is deployed to JBossAS-5.1.0, jUDDI v3 is not installed, and therefore the UDDI integration is turned off (`bpel.uddi.registration=false`).

When RiftSaw is deployed to SOA-P-5.0.0 (or JBossESB 4.8 or higher) UDDI integration is turned on and the `bpel.uddi.client.impl` is set to `org.jboss.soa.bpel.uddi.UDDIRegistrationImpl`. The `jbossesb.sar/esb.uddi.client.xml` is used, with manager name 'esb.registry'.

5.4. Other UDDI v3 Registries

Other UDDI v3 compliant registries can be used, however the UDDIv3 spec only requires communication using the UDDI WebServices. To set up SOAP based communication specify the JAXWS-Transport. At this point it makes sense to no longer use the `esb.uddi.client.xml`, but rather use your own `bpel.uddi.client.xml`. For more details please see the jUDDI v3 documentation.

5.5. UDDI Registry Entities and UDDI Seed Data

In the `esb.uddi.client.xml` a few properties are defined that are used by the Clerk at registration time. These settings of these values can be customized, however they must correspond to the UDDI seed data specified for the `jbossesb-registry.sar/juddi_custom_install_data`. So you will need to change it there as well.

The clerk is configured to use the `jbossesb` publisher and the `keyDomain` is set to "esb.jboss.org".

The `businessKey` is set to "redhat-jboss".

The `serviceDescription` is set to "BPEL Service deployed by Riftsaw".

The `bindingDescription` is set to "BPEL Endpoint deployed by Riftsaw".

Note that in SOA-P-5 the `jbossesb-registry.sar/esb.uddi.xml` contains a property `juddi.seed.always` which is set to "true". This means that that it is always trying to load the root seed data on startup of the server. It is recommended to turn this value to "false" once you are content with the UDDI Seed Data.

JBoss ESB Integration

6.1. Overview

This section outlines the support provided for the direct integration between RiftSaw and JBossESB.

Bi-directional loose integration is available through the use of web services. For example, an ESB action may invoke a BPEL process running within RiftSaw by invoking the appropriate web service represented by a WSDL interface. Similarly, a BPEL process can invoke an ESB managed service that is capable of presenting itself as a web service.

However this section will describe how integration between RiftSaw and JBossESB actions can be achieved without the use of web services (i.e. WSDL and SOAP).

6.2. Using the *BPELInvoke* ESB action

The *BPELInvoke* ESB action can be used within a *jboss-esb.xml* to request an invocation on a BPEL process running inside RiftSaw. The only constraints are that RiftSaw is installed within the same Java VM and that the requested BPEL process must have been deployed to the local RiftSaw engine.

The following example illustrates the *BPELInvoke* ESB action being used as part of the *bpel_helloworld* sample.

```
<action name="action2" class="org.jboss.soa.esb.actions.bpel.BPELInvoke">
  <property name="service" value="{http://www.jboss.org/bpel/examples/wSDL}HelloService"/>
  <property name="port" value="HelloPort" />
  <property name="operation" value="hello" />
  <property name="requestPartName" value="TestPart" />
  <property name="responsePartName" value="TestPart" />
</action>
```

The ESB action class is *org.jboss.soa.esb.actions.bpel.BPELInvoke*.

The properties for this ESB action are:

- service

This property is mandatory, and defines the service name registered in the WSDL associated with the deployed BPEL process.

- port

This property is optional, and defines the port name registered in the WSDL associated with the deployed BPEL process. This parameter is only required if port specific endpoint configuration information has been registered as part of the BPEL process deployment.

- operation

This property is mandatory, and represents the WSDL operation that is being invoked.

- requestPartName

This optional property can be used to define the WSDL message part that the inbound ESB message content should be mapped to. This property should be used where the ESB message does not already represent a multi-part message.

- responsePartName

This optional property can be used to extract the content of a response multi-part WSDL message, and place this in the ESB message being passed to the next ESB action in the pipeline. If this property is not defined, then the complete multi-part message value will be placed in the ESB message.

- abortOnFault

This optional property can be used to indicate whether a fault, generated during the invocation of a BPEL process, should be treated as a message (when the value of this property is 'false'), or as an exception that will abort the ESB service. The default value is 'true', causing the ESB service to abort.

This ESB action supports inbound messages with content defined as either:

- DOM

If the message content is a DOM document or element, then this can either be used as the complete multi-part message, or as the content of a message part defined using the *requestPartName* property.

If the message content is a DOM text node, then this can ONLY be used if a multi-part name has been defined in the *requestPartName* property.

- Java String

If the message content is a string representation of an XML document, then the *requestPartName* is optional. If not specified, then the document must represent the multipart message.


If the message content is a string that does not represent an XML document, then the *requestPartName* must be specified.

When the message content represents the complete multipart message, this must be defined as a top level element (whose name is irrelevant) with immediate child elements that represent each of the multiple parts of the message. Each of these elements must then have a single element/node, that represents the value of the named part.

```
<message>
  <TestPart>
    Hello World
  </TestPart>
</message>
```

This shows an example of a multipart message structure. The top element (i.e. *message*) is unimportant. The elements at the next level represent the part names - in this case there is only a single part, with name *TestPart*. The value of this part is defined as a text node, with value "Hello World". However this could have been an element representing the root node of a more complex XML value.

The following diagram illustrates the inter-relationship of the JBossESB *bpel_helloworld* quickstart and the RiftSaw BPEL process configuration files.



6.2.1. Fault Handling

The normal response from a WSDL operation will be returned from the *BPELInvoke* ESB action as a normal message and placed on the action pipeline ready for processing by the next ESB action, or alternatively if no further actions have been defined, then returned back to the service client.

Faults, associated with a WSDL operation, are handled slightly differently. Depending on configuration it is possible to receive the fault as an ESB message or for the fault to be treated as an exception which aborts the action pipeline. The configuration property used to determine which behaviour is used is called *abortOnFault*. The default value for this property is "true". As an example, from the loan fault quickstart sample,

```
<action name="action2" class="org.jboss.soa.esb.actions.bpel.BPELInvoke">
  <property name="service" value="{http://example.com/loan-approval/wsdl/}loanService"/>
  <property name="operation" value="request" />
  <property name="abortOnFault" value="true" />
</action>
```

A WSDL fault has two relevant pieces of information, the fault type (or code) and the fault details. These are both returned in specific parts of ESB message's body.

1. Fault code (as javax.xml.namespace.QName)

ESB message body part: *org.jboss.soa.esb.message.fault.detail.code*

This body part identifies the specific WSDL fault returned by the BPEL process, associated with the WSDL operation that was invoked.



Warning

The specific version of the QName used by the JBoss server is from the stax-api.jar found in the server's lib/endorsed directory. If the client does not also include this jar in a folder that is in its endorsed directories, then a class version exception will occur when this ESB message part is accessed.

2. Fault code (as textual representation of QName)

ESB message body part: *org.jboss.soa.bpel.message.fault.detail.code*

This body part will return the textual representation of the QName for the fault code. The textual representation is of the form "{namespace}localpart" and can be converted back into a QName using the *javax.xml.namespace.QName.valueOf(String)* method.

3. Fault details

ESB message body part: *org.jboss.soa.esb.message.fault.detail.detail*

This body part will contain the textual representation of the message content associated with the fault.

6.2.2. SAML Support

If the ESB service uses PicketLink to obtain a SAML token, then this assertion can be passed to the invoked BPEL process, using the *requestSAMLPartName* property.

```
<action name="action2" class="org.jboss.soa.esb.actions.bpel.BPELInvoke">
  <property name="service" value="{http://simple_invoke/helloworld}HelloHeaderWSService"/>
  <property name="operation" value="sayHi" />
  <property name="requestPartName" value="sayHello" />
  <property name="responsePartName" value="sayHelloResponse" />
  <property name="requestSAMLPartName" value="Security" />
</action>
```


The part name identified by the *requestSAMLPartName* must be defined as a WS-Security Security element, e.g.

```
<part name="Security" element="wsse:Security"
      xmlns:wsse="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-secext-1.0.xsd" />
```

RiftSaw Clustering Support

7.1. Overview

In order to make riftsaw to be able to work in a clustering environment, it has to be configured to a shared database for all the nodes that among the clustering environment, because riftsaw persists all of process states into database. For the frontend, it needs a load balancer to dispatch the soap message into nodes properly. Below is the picture of riftsaw server deployment architecture for clustering.



7.2. Installation

In the \$Riftsaw/install folder, if you set the `org.jboss.as.config=all` in the `deployment.properties`, it will deploy riftsaw-clustering libraries and files into JBoss AS. If you want to deploy the riftsaw clustering feature into other config, say like you copied 'all' into 'node1', you can run the following command, (which basically added '-Dclustering.support=true').

```
ant deploy -Dclustering.support=true
```

**Note**

Please be noted that the riftsaw clustering support depends on the JBoss AS Clustering's *HAPartitionService*, so please be sure that you have this service started if you use some customized config for JBoss AS.

7.3. Deployment

Deploying bpel artifact in clustering environment is different from deploying it into a single server. You have to copy your bpel artifact (say `hello_world.jar`) into `$JBossAS/server/$config (like all)/farm/`, any artifacts deployed in this folder will be copied into all nodes that in the clustering environment.

7.4. Bpel Process Service Invocation

For the bpel service that you deployed in the clustering environment, if you want to invoke that service, you should be specifying the *load balancer's url* instead of the soap address that you specified in the wsdl file.

Database

8.1. Upgrade database schema

The RiftSaw database schema has been changed between 2.0.0.Final and 2.1.0.Final, please refer to [this wiki page](#) for the upgrade detail information.

8.2. Database schema diagram

Below is the EER Diagram generated by Mysql Workbench tool.



Note

This diagram is from RiftSaw 2.1.0.Final database schemas.

