



## **Red Hat JBoss Fuse 6.0**

# **Configuring Web Service Endpoints**

Deploy service endpoints



# Red Hat JBoss Fuse 6.0 Configuring Web Service Endpoints

---

Deploy service endpoints

JBoss A-MQ Docs Team

Content Services

[fuse-docs-support@redhat.com](mailto:fuse-docs-support@redhat.com)

## Legal Notice

Copyright © 2013 Red Hat.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This guide describes how to create Apache CXF endpoints in Red Hat JBoss Fuse.

---

## Table of Contents

<b>CHAPTER 1. CONFIGURING JAX-WS ENDPOINTS</b> .....	<b>4</b>
1.1. CONFIGURING SERVICE PROVIDERS	4
1.2. CONFIGURING CONSUMER ENDPOINTS	13
<b>CHAPTER 2. CONFIGURING THE HTTP TRANSPORT</b> .....	<b>17</b>
2.1. CONFIGURING A CONSUMER	17
2.2. CONFIGURING A SERVICE PROVIDER	24
2.3. USING THE HTTP TRANSPORT IN DECOUPLED MODE	29
<b>CHAPTER 3. USING SOAP OVER JMS</b> .....	<b>34</b>
3.1. BASIC CONFIGURATION	34
3.2. JMS URIS	36
3.3. WSDL EXTENSIONS	38
<b>CHAPTER 4. USING GENERIC JMS</b> .....	<b>43</b>
4.1. USING THE JMS CONFIGURATION BEAN	43
4.2. USING WSDL TO CONFIGURE JMS	49
4.3. USING A NAMED REPLY DESTINATION	54
<b>CHAPTER 5. APACHE CXF LOGGING</b> .....	<b>55</b>
5.1. OVERVIEW OF APACHE CXF LOGGING	55
5.2. SIMPLE EXAMPLE OF USING LOGGING	56
5.3. DEFAULT LOGGING CONFIGURATION FILE	57
5.4. ENABLING LOGGING AT THE COMMAND LINE	60
5.5. LOGGING FOR SUBSYSTEMS AND SERVICES	60
5.6. LOGGING MESSAGE CONTENT	62
<b>CHAPTER 6. DEPLOYING WS-ADDRESSING</b> .....	<b>65</b>
6.1. INTRODUCTION TO WS-ADDRESSING	65
6.2. WS-ADDRESSING INTERCEPTORS	65
6.3. ENABLING WS-ADDRESSING	66
6.4. CONFIGURING WS-ADDRESSING ATTRIBUTES	67
<b>CHAPTER 7. ENABLING RELIABLE MESSAGING</b> .....	<b>69</b>
7.1. INTRODUCTION TO WS-RM	69
7.2. WS-RM INTERCEPTORS	70
7.3. ENABLING WS-RM	71
7.4. CONFIGURING WS-RM	74
7.5. CONFIGURING WS-RM PERSISTENCE	82
<b>CHAPTER 8. ENABLING HIGH AVAILABILITY</b> .....	<b>85</b>
8.1. INTRODUCTION TO HIGH AVAILABILITY	85
8.2. ENABLING HA WITH STATIC FAILOVER	85
8.3. CONFIGURING HA WITH STATIC FAILOVER	87
<b>CHAPTER 9. ENABLING HIGH AVAILABILITY IN FUSE FABRIC</b> .....	<b>89</b>
9.1. LOAD BALANCING CLUSTER	89
9.2. FAILOVER CLUSTER	102
<b>CHAPTER 10. PACKAGING AN APPLICATION</b> .....	<b>105</b>
CREATING A BUNDLE	105
REQUIRED BUNDLE	105
REQUIRED PACKAGES	105
EXAMPLE	106

<b>CHAPTER 11. DEPLOYING AN APPLICATION</b> .....	<b>107</b>
OVERVIEW	107
HOT DEPLOYMENT	107
DEPLOYING FROM THE CONSOLE	107
REFRESHING AN APPLICATION	107
STOPPING AN APPLICATION	107
UNINSTALLING AN APPLICATION	108
<b>APPENDIX A. APACHE CXF BINDING IDS</b> .....	<b>109</b>
<b>APPENDIX B. USING THE MAVEN OSGI TOOLING</b> .....	<b>110</b>
B.1. SETTING UP A RED HAT JBOSS FUSE OSGI PROJECT	110
B.2. CONFIGURING THE BUNDLE PLUG-IN	113
<b>APPENDIX C. CONDUITS</b> .....	<b>118</b>
OVERVIEW	118
CONDUIT LIFE-CYCLE	118
CONDUIT WEIGHT	118
<b>INDEX</b> .....	<b>119</b>



# CHAPTER 1. CONFIGURING JAX-WS ENDPOINTS

## Abstract

JAX-WS endpoints are configured using one of three Spring configuration elements. The correct element depends on what type of endpoint you are configuring and which features you wish to use. For consumers you use the `jaxws:client` element. For service providers you can use either the `jaxws:endpoint` element or the `jaxws:server` element.

The information used to define an endpoint is typically defined in the endpoint's contract. You can use the configuration element's to override the information in the contract. You can also use the configuration elements to provide information that is not provided in the contract.



## NOTE

When dealing with endpoints developed using a Java-first approach it is likely that the SEI serving as the endpoint's contract is lacking information about the type of binding and transport to use.

You must use the configuration elements to activate advanced features such as WS-RM. This is done by providing child elements to the endpoint's configuration element.

## 1.1. CONFIGURING SERVICE PROVIDERS

Apache CXF has two elements that can be used to configure a service provider:

- [Section 1.1.1, “Using the jaxws:endpoint Element”](#)
- [Section 1.1.2, “Using the jaxws:server Element”](#)

The differences between the two elements are largely internal to the runtime. The `jaxws:endpoint` element injects properties into the `org.apache.cxf.jaxws.EndpointImpl` object created to support a service endpoint. The `jaxws:server` element injects properties into the `org.apache.cxf.jaxws.support.JaxWsServerFactoryBean` object created to support the endpoint. The `EndpointImpl` object passes the configuration data to the `JaxWsServerFactoryBean` object. The `JaxWsServerFactoryBean` object is used to create the actual service object. Because either configuration element will configure a service endpoint, you can choose based on the syntax you prefer.

### 1.1.1. Using the jaxws:endpoint Element

#### Overview

The `jaxws:endpoint` element is the default element for configuring JAX-WS service providers. Its attributes and children specify all of the information needed to instantiate a service provider. Many of the attributes map to information in the service's contract. The children are used to configure interceptors and other advanced features.

#### Identifying the endpoint being configured



For the runtime to apply the configuration to the proper service provider, it must be able to identify it. The basic means for identifying a service provider is to specify the class that implements the endpoint. This is done using the `jaxws:endpoint` element's `implementor` attribute.

For instances where different endpoints share a common implementation, it is possible to provide different configuration for each endpoint. There are two approaches for distinguishing a specific endpoint in configuration:

- a combination of the `serviceName` attribute and the `endpointName` attribute

The `serviceName` attribute specifies the `wsdl:service` element defining the service's endpoint. The `endpointName` attribute specifies the specific `wsdl:port` element defining the service's endpoint. Both attributes are specified as QNames using the format `ns:name`. `ns` is the namespace of the element and `name` is the value of the element's `name` attribute.

### TIP

If the `wsdl:service` element only has one `wsdl:port` element, the `endpointName` attribute can be omitted.

- the `name` attribute

The `name` attribute specifies the QName of the specific `wsdl:port` element defining the service's endpoint. The QName is provided in the format `{ns}localPart`. `ns` is the namespace of the `wsdl:port` element and `localPart` is the value of the `wsdl:port` element's `name` attribute.

## Attributes

The attributes of the `jaxws:endpoint` element configure the basic properties of the endpoint. These properties include the address of the endpoint, the class that implements the endpoint, and the bus that hosts the endpoint.

Table 1.1, “Attributes for Configuring a JAX-WS Service Provider Using the `jaxws:endpoint` Element” describes the attribute of the `jaxws:endpoint` element.

**Table 1.1. Attributes for Configuring a JAX-WS Service Provider Using the `jaxws:endpoint` Element**

Attribute	Description
<code>id</code>	Specifies a unique identifier that other configuration elements can use to refer to the endpoint.
<code>implementor</code>	Specifies the class implementing the service. You can specify the implementation class using either the class name or an ID reference to a Spring bean configuring the implementation class. This class must be on the classpath.

Attribute	Description
<b>implementorClass</b>	Specifies the class implementing the service. This attribute is useful when the value provided to the <b>implementor</b> attribute is a reference to a bean that is wrapped using Spring AOP.
<b>address</b>	Specifies the address of an HTTP endpoint. This value overrides the value specified in the services contract.
<b>wSDLLocation</b>	Specifies the location of the endpoint's WSDL contract. The WSDL contract's location is relative to the folder from which the service is deployed.
<b>endpointName</b>	Specifies the value of the service's <b>wSDL:port</b> element's <b>name</b> attribute. It is specified as a QName using the format <i>ns:name</i> where <i>ns</i> is the namespace of the <b>wSDL:port</b> element.
<b>serviceName</b>	Specifies the value of the service's <b>wSDL:service</b> element's <b>name</b> attribute. It is specified as a QName using the format <i>ns:name</i> where <i>ns</i> is the namespace of the <b>wSDL:service</b> element.
<b>publish</b>	Specifies if the service should be automatically published. If this is set to <b>false</b> , the developer must explicitly publish the endpoint.
<b>bus</b>	Specifies the ID of the Spring bean configuring the bus used to manage the service endpoint. This is useful when configuring several endpoints to use a common set of features.
<b>bindingUri</b>	Specifies the ID of the message binding the service uses. A list of valid binding IDs is provided in <a href="#">Appendix A, Apache CXF Binding IDs</a> .
<b>name</b>	Specifies the stringified QName of the service's <b>wSDL:port</b> element. It is specified as a QName using the format <i>{ns}localPart</i> . <i>ns</i> is the namespace of the <b>wSDL:port</b> element and <i>localPart</i> is the value of the <b>wSDL:port</b> element's <b>name</b> attribute.
<b>abstract</b>	Specifies if the bean is an abstract bean. Abstract beans act as parents for concrete bean definitions and are not instantiated. The default is <b>false</b> . Setting this to <b>true</b> instructs the bean factory not to instantiate the bean.

Attribute	Description
<b>depends-on</b>	Specifies a list of beans that the endpoint depends on being instantiated before it can be instantiated.
<b>createdFromAPI</b>	<p>Specifies that the user created that bean using Apache CXF APIs, such as <code>Endpoint.publish()</code> or <code>Service.getPort()</code>.</p> <p>The default is <b>false</b>.</p> <p>Setting this to <b>true</b> does the following:</p> <ul style="list-style-type: none"> <li>• Changes the internal name of the bean by appending <code>.jaxws-endpoint</code> to its id</li> <li>• Makes the bean abstract</li> </ul>
<b>publishedEndpointUrl</b>	The URL that is placed in the <b>address</b> element of the generated WSDL. If this value is not specified, the value of the <b>address</b> attribute is used. This attribute is useful when the "public" URL is not be the same as the URL on which the service is deployed.

In addition to the attributes listed in [Table 1.1, “Attributes for Configuring a JAX-WS Service Provider Using the `jaxws:endpoint` Element”](#), you might need to use multiple `xmlns:shortName` attributes to declare the namespaces used by the `endpointName` and `serviceName` attributes.

## Example

[Example 1.1, “Simple JAX-WS Endpoint Configuration”](#) shows the configuration for a JAX-WS endpoint that specifies the address where the endpoint is published. The example assumes that you want to use the defaults for all other values or that the implementation has specified values in the annotations.

### Example 1.1. Simple JAX-WS Endpoint Configuration

```
<beans ...
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  ...
  schemaLocation="...
    http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd
    ..."
  <jaxws:endpoint id="example"
    implementor="org.apache.cxf.example.DemoImpl"
    address="http://localhost:8080/demo" />
</beans>
```

[Example 1.2, “JAX-WS Endpoint Configuration with a Service Name”](#) shows the configuration for a JAX-WS endpoint whose contract contains two service definitions. In this case, you must specify which service definition to instantiate using the `serviceName` attribute.

### Example 1.2. JAX-WS Endpoint Configuration with a Service Name

```
<beans ...
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  ...
  schemaLocation="...
    http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd
    ...">

  <jaxws:endpoint id="example2"
    implementor="org.apache.cxf.example.DemoImpl"
    serviceName="samp:demoService2"
    xmlns:samp="http://org.apache.cxf/wsdl/example" />

</beans>
```

The `xmlns:samp` attribute specifies the namespace in which the WSDL `service` element is defined.

## 1.1.2. Using the `jaxws:server` Element

### Overview

The `jaxws:server` element is an element for configuring JAX-WS service providers. It injects the configuration information into the `org.apache.cxf.jaxws.support.JaxWsServerFactoryBean`. This is a Apache CXF specific object. If you are using a pure Spring approach to building your services, you will not be forced to use Apache CXF specific APIs to interact with the service.

The attributes and children of the `jaxws:server` element specify all of the information needed to instantiate a service provider. The attributes specify the information that is required to instantiate an endpoint. The children are used to configure interceptors and other advanced features.

### Identifying the endpoint being configured

In order for the runtime to apply the configuration to the proper service provider, it must be able to identify it. The basic means for identifying a service provider is to specify the class that implements the endpoint. This is done using the `jaxws:server` element's `serviceBean` attribute.

For instances where different endpoint's share a common implementation, it is possible to provide different configuration for each endpoint. There are two approaches for distinguishing a specific endpoint in configuration:

- a combination of the `serviceName` attribute and the `endpointName` attribute

The `serviceName` attribute specifies the `wsdl:service` element defining the service's endpoint. The `endpointName` attribute specifies the specific `wsdl:port` element defining the service's endpoint. Both attributes are specified as QNames using the format `ns:name`. `ns` is the namespace of the element and `name` is the value of the element's `name` attribute.

**TIP**

If the `wsdl:service` element only has one `wsdl:port` element, the `endpointName` attribute can be omitted.

- the `name` attribute

The `name` attribute specifies the QName of the specific `wsdl:port` element defining the service's endpoint. The QName is provided in the format `{ns}localPart`. `ns` is the namespace of the `wsdl:port` element and `localPart` is the value of the `wsdl:port` element's `name` attribute.

**Attributes**

The attributes of the `jaxws:server` element configure the basic properties of the endpoint. These properties include the address of the endpoint, the class that implements the endpoint, and the `bus` that hosts the endpoint.

Table 1.2, “Attributes for Configuring a JAX-WS Service Provider Using the `jaxws:server` Element” describes the attribute of the `jaxws:server` element.

**Table 1.2. Attributes for Configuring a JAX-WS Service Provider Using the `jaxws:server` Element**

Attribute	Description
<code>id</code>	Specifies a unique identifier that other configuration elements can use to refer to the endpoint.
<code>serviceBean</code>	Specifies the class implementing the service. You can specify the implementation class using either the class name or an ID reference to a Spring bean configuring the implementation class. This class must be on the classpath.
<code>serviceClass</code>	Specifies the class implementing the service. This attribute is useful when the value provided to the <code>implementor</code> attribute is a reference to a bean that is wrapped using Spring AOP.
<code>address</code>	Specifies the address of an HTTP endpoint. This value will override the value specified in the services contract.
<code>wsdlLocation</code>	Specifies the location of the endpoint's WSDL contract. The WSDL contract's location is relative to the folder from which the service is deployed.
<code>endpointName</code>	Specifies the value of the service's <code>wsdl:port</code> element's <code>name</code> attribute. It is specified as a QName using the format <code>ns:name</code> , where <code>ns</code> is the namespace of the <code>wsdl:port</code> element.

Attribute	Description
<b>serviceName</b>	Specifies the value of the service's <code>wsdl:service</code> element's <code>name</code> attribute. It is specified as a QName using the format <code>ns:name</code> , where <code>ns</code> is the namespace of the <code>wsdl:service</code> element.
<b>start</b>	Specifies if the service should be automatically published. If this is set to <code>false</code> , the developer must explicitly publish the endpoint.
<b>bus</b>	Specifies the ID of the Spring bean configuring the bus used to manage the service endpoint. This is useful when configuring several endpoints to use a common set of features.
<b>bindingId</b>	Specifies the ID of the message binding the service uses. A list of valid binding IDs is provided in <a href="#">Appendix A, Apache CXF Binding IDs</a> .
<b>name</b>	Specifies the stringified QName of the service's <code>wsdl:port</code> element. It is specified as a QName using the format <code>{ns}localPart</code> , where <code>ns</code> is the namespace of the <code>wsdl:port</code> element and <code>localPart</code> is the value of the <code>wsdl:port</code> element's <code>name</code> attribute.
<b>abstract</b>	Specifies if the bean is an abstract bean. Abstract beans act as parents for concrete bean definitions and are not instantiated. The default is <code>false</code> . Setting this to <code>true</code> instructs the bean factory not to instantiate the bean.
<b>depends-on</b>	Specifies a list of beans that the endpoint depends on being instantiated before the endpoint can be instantiated.
<b>createdFromAPI</b>	<p>Specifies that the user created that bean using Apache CXF APIs, such as <code>Endpoint.publish()</code> or <code>Service.getPort()</code>.</p> <p>The default is <code>false</code>.</p> <p>Setting this to <code>true</code> does the following:</p> <ul style="list-style-type: none"> <li>• Changes the internal name of the bean by appending <code>.jaxws-endpoint</code> to its id</li> <li>• Makes the bean abstract</li> </ul>

In addition to the attributes listed in [Table 1.2, “Attributes for Configuring a JAX-WS Service Provider Using the `jaxws:server` Element”](#), you might need to use multiple `xmlns:shortName` attributes to declare the namespaces used by the `endpointName` and `serviceName` attributes.

## Example

[Example 1.3, “Simple JAX-WS Server Configuration”](#) shows the configuration for a JAX-WS endpoint that specifies the address where the endpoint is published.

### Example 1.3. Simple JAX-WS Server Configuration

```
<beans ...
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  ...
  schemaLocation="...
    http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd
    ..."
  <jaxws:server id="exampleServer"
    serviceBean="org.apache.cxf.example.DemoImpl"
    address="http://localhost:8080/demo" />
</beans>
```

## 1.1.3. Adding Functionality to Service Providers

### Overview

The `jaxws:endpoint` and the `jaxws:server` elements provide the basic configuration information needed to instantiate a service provider. To add functionality to your service provider or to perform advanced configuration you must add child elements to the configuration.

Child elements allow you to do the following:

- [Change the endpoint's logging behavior](#)
- [Add interceptors to the endpoint's messaging chain](#)
- [Enable WS-Addressing features](#)
- [Enable reliable messaging](#)

### Elements

[Table 1.3, “Elements Used to Configure JAX-WS Service Providers”](#) describes the child elements that `jaxws:endpoint` supports.

**Table 1.3. Elements Used to Configure JAX-WS Service Providers**

Element	Description
<code>jaxws:handlers</code>	Specifies a list of JAX-WS <code>Handler</code> implementations for processing messages.

Element	Description
<code>jaxws:inInterceptors</code>	Specifies a list of interceptors that process inbound requests. For more information see " <a href="#">Developing Apache CXF Interceptors</a> ".
<code>jaxws:inFaultInterceptors</code>	Specifies a list of interceptors that process inbound fault messages. For more information see " <a href="#">Developing Apache CXF Interceptors</a> ".
<code>jaxws:outInterceptors</code>	Specifies a list of interceptors that process outbound replies. For more information see " <a href="#">Developing Apache CXF Interceptors</a> ".
<code>jaxws:outFaultInterceptors</code>	Specifies a list of interceptors that process outbound fault messages. For more information see " <a href="#">Developing Apache CXF Interceptors</a> ".
<code>jaxws:binding</code>	Specifies a bean configuring the message binding used by the endpoint. Message bindings are configured using implementations of the <code>org.apache.cxf.binding.BindingFactory</code> interface. <sup>[a]</sup>
<code>jaxws:dataBinding</code> <sup>[b]</sup>	Specifies the class implementing the data binding used by the endpoint. This is specified using an embedded bean definition.
<code>jaxws:executor</code>	Specifies a Java executor that is used for the service. This is specified using an embedded bean definition.
<code>jaxws:features</code>	Specifies a list of beans that configure advanced features of Apache CXF. You can provide either a list of bean references or a list of embedded beans.
<code>jaxws:invoker</code>	Specifies an implementation of the <code>org.apache.cxf.service.Invoker</code> interface used by the service. <sup>[c]</sup>
<code>jaxws:properties</code>	Specifies a Spring map of properties that are passed along to the endpoint. These properties can be used to control features like enabling MTOM support.
<code>jaxws:serviceFactory</code>	Specifies a bean configuring the <code>JaxWsServiceFactoryBean</code> object used to instantiate the service.



Element	Description
[a]	The SOAP binding is configured using the <code>soap:soapBinding</code> bean.
[b]	The <code>jaxws:endpoint</code> element does not support the <code>jaxws:dataBinding</code> element.
[c]	The <b>Invoker</b> implementation controls how a service is invoked. For example, it controls whether each request is handled by a new instance of the service implementation or if state is preserved across invocations.

## 1.2. CONFIGURING CONSUMER ENDPOINTS

### Overview

JAX-WS consumer endpoints are configured using the `jaxws:client` element. The element's attributes provide the basic information necessary to create a consumer.

To add other functionality, like WS-RM, to the consumer you add children to the `jaxws:client` element. Child elements are also used to configure the endpoint's logging behavior and to inject other properties into the endpoint's implementation.

### Basic Configuration Properties

The attributes described in [Table 1.4, “Attributes Used to Configure a JAX-WS Consumer”](#) provide the basic information necessary to configure a JAX-WS consumer. You only need to provide values for the specific properties you want to configure. Most of the properties have sensible defaults, or they rely on information provided by the endpoint's contract.

**Table 1.4. Attributes Used to Configure a JAX-WS Consumer**

Attribute	Description
<code>address</code>	Specifies the HTTP address of the endpoint where the consumer will make requests. This value overrides the value set in the contract.
<code>bindingId</code>	Specifies the ID of the message binding the consumer uses. A list of valid binding IDs is provided in <a href="#">Appendix A, Apache CXF Binding IDs</a> .
<code>bus</code>	Specifies the ID of the Spring bean configuring the bus managing the endpoint.
<code>endpointName</code>	Specifies the value of the <code>wSDL:port</code> element's <code>name</code> attribute for the service on which the consumer is making requests. It is specified as a QName using the format <code>ns:name</code> , where <code>ns</code> is the namespace of the <code>wSDL:port</code> element.

Attribute	Description
<b>serviceName</b>	Specifies the value of the <code>wSDL:service</code> element's <code>name</code> attribute for the service on which the consumer is making requests. It is specified as a QName using the format <code>ns:name</code> where <code>ns</code> is the namespace of the <code>wSDL:service</code> element.
<b>username</b>	Specifies the username used for simple username/password authentication.
<b>password</b>	Specifies the password used for simple username/password authentication.
<b>serviceClass</b>	Specifies the name of the service endpoint interface(SEI).
<b>wSDLLocation</b>	Specifies the location of the endpoint's WSDL contract. The WSDL contract's location is relative to the folder from which the client is deployed.
<b>name</b>	Specifies the stringified QName of the <code>wSDL:port</code> element for the service on which the consumer is making requests. It is specified as a QName using the format <code>{ns}localPart</code> , where <code>ns</code> is the namespace of the <code>wSDL:port</code> element and <code>localPart</code> is the value of the <code>wSDL:port</code> element's <code>name</code> attribute.
<b>abstract</b>	Specifies if the bean is an abstract bean. Abstract beans act as parents for concrete bean definitions and are not instantiated. The default is <b>false</b> . Setting this to <b>true</b> instructs the bean factory not to instantiate the bean.
<b>depends-on</b>	Specifies a list of beans that the endpoint depends on being instantiated before it can be instantiated.
<b>createdFromAPI</b>	<p>Specifies that the user created that bean using Apache CXF APIs like <code>Service.getPort()</code>.</p> <p>The default is <b>false</b>.</p> <p>Setting this to <b>true</b> does the following:</p> <ul style="list-style-type: none"> <li>• Changes the internal name of the bean by appending <code>.jaxws-client</code> to its id</li> <li>• Makes the bean abstract</li> </ul>

In addition to the attributes listed in [Table 1.4, “Attributes Used to Configure a JAX-WS Consumer”](#), it might be necessary to use multiple `xmlns:shortName` attributes to declare the namespaces used by the `endpointName` and the `serviceName` attributes.

## Adding functionality

To add functionality to your consumer or to perform advanced configuration, you must add child elements to the configuration.

Child elements allow you to do the following:

- [Change the endpoint's logging behavior](#)
- [Add interceptors to the endpoint's messaging chain](#)
- [Enable WS-Addressing features](#)
- [Enable reliable messaging](#)

[Table 1.5, “Elements For Configuring a Consumer Endpoint”](#) describes the child element's you can use to configure a JAX-WS consumer.

**Table 1.5. Elements For Configuring a Consumer Endpoint**

Element	Description
<code>jaxws:binding</code>	Specifies a bean configuring the message binding used by the endpoint. Message bindings are configured using implementations of the <code>org.apache.cxf.binding.BindingFactory</code> interface. <sup>[a]</sup>
<code>jaxws:dataBinding</code>	Specifies the class implementing the data binding used by the endpoint. You specify this using an embedded bean definition. The class implementing the JAXB data binding is <code>org.apache.cxf.jaxb.JAXBDataBinding</code> .
<code>jaxws:features</code>	Specifies a list of beans that configure advanced features of Apache CXF. You can provide either a list of bean references or a list of embedded beans.
<code>jaxws:handlers</code>	Specifies a list of <code>JAX-WSHandler</code> implementations for processing messages.
<code>jaxws:inInterceptors</code>	Specifies a list of interceptors that process inbound responses. For more information see <a href="#">"Developing Apache CXF Interceptors"</a> .
<code>jaxws:inFaultInterceptors</code>	Specifies a list of interceptors that process inbound fault messages. For more information see <a href="#">"Developing Apache CXF Interceptors"</a> .

Element	Description
<code>jaxws:outInterceptors</code>	Specifies a list of interceptors that process outbound requests. For more information see <a href="#">"Developing Apache CXF Interceptors"</a> .
<code>jaxws:outFaultInterceptors</code>	Specifies a list of interceptors that process outbound fault messages. For more information see <a href="#">"Developing Apache CXF Interceptors"</a> .
<code>jaxws:properties</code>	Specifies a map of properties that are passed to the endpoint.
<code>jaxws:conduitSelector</code>	Specifies an <code>org.apache.cxf.endpoint.ConduitSelector</code> implementation for the client to use. A <code>ConduitSelector</code> implementation will override the default process used to select the <code>Conduit</code> object that is used to process outbound requests.
[a] The SOAP binding is configured using the <code>soap:soapBinding</code> bean.	

## Example

[Example 1.4, "Simple Consumer Configuration"](#) shows a simple consumer configuration.

### Example 1.4. Simple Consumer Configuration

```
<beans ...
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  ...
  schemaLocation="...
    http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd
    ...">
  <jaxws:client id="bookClient"
    serviceClass="org.apache.cxf.demo.BookClientImpl"
    address="http://localhost:8080/books"/>
  ...
</beans>
```

## CHAPTER 2. CONFIGURING THE HTTP TRANSPORT

### Abstract

The Apache CXF HTTP transport is highly configurable.

## 2.1. CONFIGURING A CONSUMER

HTTP consumer endpoints can specify a number of HTTP connection attributes including whether the endpoint automatically accepts redirect responses, whether the endpoint can use chunking, whether the endpoint will request a keep-alive, and how the endpoint interacts with proxies. In addition to the HTTP connection properties, an HTTP consumer endpoint can specify how it is secured.

A consumer endpoint can be configured using two mechanisms:

- [Configuration](#)
- [WSDL](#)

### 2.1.1. Using Configuration

#### Namespace

The elements used to configure an HTTP consumer endpoint are defined in the namespace `http://cxf.apache.org/transport/http/configuration`. It is commonly referred to using the prefix `http-conf`. In order to use the HTTP configuration elements you must add the lines shown in [Example 2.1, “HTTP Consumer Configuration Namespace”](#) to the `beans` element of your endpoint's configuration file. In addition, you must add the configuration elements' namespace to the `xsi:schemaLocation` attribute.

#### Example 2.1. HTTP Consumer Configuration Namespace

```
<beans ...
  xmlns:http-
conf="http://cxf.apache.org/transport/http/configuration"
  ...
  xsi:schemaLocation="...

http://cxf.apache.org/transport/http/configuration

http://cxf.apache.org/schemas/configuration/http-conf.xsd
  ...">
```

#### The conduit element

You configure an HTTP consumer endpoint using the `http-conf:conduit` element and its children. The `http-conf:conduit` element takes a single attribute, `name`, that specifies the WSDL `port` element corresponding to the endpoint. The value for the `name` attribute takes the form `portQName.http-conduit`. [Example 2.2, “http-conf:conduit Element”](#) shows the `http-`

`conf:conduit` element that would be used to add configuration for an endpoint that is specified by the WSDL fragment `<port binding="widgetSOAPBinding" name="widgetSOAPPort">` when the endpoint's target namespace is `http://widgets.widgetvendor.net`.

### Example 2.2. `http-conf:conduit` Element

```
...
  <http-conf:conduit name="
{http://widgets/widgetvendor.net}widgetSOAPPort.http-conduit">
    ...
  </http-conf:conduit>
...

```

The `http-conf:conduit` element has child elements that specify configuration information. They are described in [Table 2.1, “Elements Used to Configure an HTTP Consumer Endpoint”](#).

**Table 2.1. Elements Used to Configure an HTTP Consumer Endpoint**

Element	Description
<code>http-conf:client</code>	Specifies the HTTP connection properties such as timeouts, keep-alive requests, content types, etc. See <a href="#">the section called “The client element”</a> .
<code>http-conf:authorization</code>	Specifies the parameters for configuring the basic authentication method that the endpoint uses preemptively.  The preferred approach is to supply a <a href="#">Basic Authentication Supplier</a> object.
<code>http-conf:proxyAuthorization</code>	Specifies the parameters for configuring basic authentication against outgoing HTTP proxy servers.
<code>http-conf:tlsClientParameters</code>	Specifies the parameters used to configure SSL/TLS.
<code>http-conf:basicAuthSupplier</code>	Specifies the bean reference or class name of the object that supplies the basic authentication information used by the endpoint, either preemptively or in response to a <b>401</b> HTTP challenge.
<code>http-conf:trustDecider</code>	Specifies the bean reference or class name of the object that checks the HTTP(S) <code>URLConnection</code> object to establish trust for a connection with an HTTPS service provider before any information is transmitted.

### The client element

The `http-conf:client` element is used to configure the non-security properties of a consumer endpoint's HTTP connection. Its attributes, described in [Table 2.2, “HTTP Consumer Configuration Attributes”](#), specify the connection's properties.

Table 2.2. HTTP Consumer Configuration Attributes

Attribute	Description
<b>ConnectionTimeout</b>	<p>Specifies the amount of time, in milliseconds, that the consumer attempts to establish a connection before it times out. The default is <b>30000</b>.</p> <p><b>0</b> specifies that the consumer will continue to send the request indefinitely.</p>
<b>ReceiveTimeout</b>	<p>Specifies the amount of time, in milliseconds, that the consumer will wait for a response before it times out. The default is <b>30000</b>.</p> <p><b>0</b> specifies that the consumer will wait indefinitely.</p>
<b>AutoRedirect</b>	<p>Specifies if the consumer will automatically follow a server issued redirection. The default is <b>false</b>.</p>
<b>MaxRetransmits</b>	<p>Specifies the maximum number of times a consumer will retransmit a request to satisfy a redirect. The default is <b>-1</b> which specifies that unlimited retransmissions are allowed.</p>
<b>AllowChunking</b>	<p>Specifies whether the consumer will send requests using chunking. The default is <b>true</b> which specifies that the consumer will use chunking when sending requests.</p> <p>Chunking cannot be used if either of the following are true:</p> <ul style="list-style-type: none"> <li>• <b>http-conf:basicAuthSupplier</b> is configured to provide credentials preemptively.</li> <li>• <b>AutoRedirect</b> is set to <b>true</b>.</li> </ul> <p>In both cases the value of <b>AllowChunking</b> is ignored and chunking is disallowed.</p>
<b>Accept</b>	<p>Specifies what media types the consumer is prepared to handle. The value is used as the value of the HTTP Accept property. The value of the attribute is specified using multipurpose internet mail extensions (MIME) types.</p>

Attribute	Description
<b>AcceptLanguage</b>	<p>Specifies what language (for example, American English) the consumer prefers for the purpose of receiving a response. The value is used as the value of the HTTP AcceptLanguage property.</p> <p>Language tags are regulated by the International Organization for Standards (ISO) and are typically formed by combining a language code, determined by the ISO-639 standard, and country code, determined by the ISO-3166 standard, separated by a hyphen. For example, en-US represents American English.</p>
<b>AcceptEncoding</b>	<p>Specifies what content encodings the consumer is prepared to handle. Content encoding labels are regulated by the Internet Assigned Numbers Authority (IANA). The value is used as the value of the HTTP AcceptEncoding property.</p>
<b>ContentType</b>	<p>Specifies the media type of the data being sent in the body of a message. Media types are specified using multipurpose internet mail extensions (MIME) types. The value is used as the value of the HTTP ContentType property. The default is <code>text/xml</code>.</p> <p>For web services, this should be set to <code>text/xml</code>. If the client is sending HTML form data to a CGI script, this should be set to <code>application/x-www-form-urlencoded</code>. If the HTTP POST request is bound to a fixed payload format (as opposed to SOAP), the content type is typically set to <code>application/octet-stream</code>.</p>
<b>Host</b>	<p>Specifies the Internet host and port number of the resource on which the request is being invoked. The value is used as the value of the HTTP Host property.</p> <p>This attribute is typically not required. It is only required by certain DNS scenarios or application designs. For example, it indicates what host the client prefers for clusters (that is, for virtual servers mapping to the same Internet protocol (IP) address).</p>
<b>Connection</b>	<p>Specifies whether a particular connection is to be kept open or closed after each request/response dialog. There are two valid values:</p> <ul style="list-style-type: none"> <li>● <b>Keep-Alive</b> (default) – Specifies that the consumer wants the connection kept open after the initial request/response sequence. If the server honors it, the connection is kept open until the consumer closes it.</li> <li>● <b>close</b> – Specifies that the connection to the server is closed after each request/response sequence.</li> </ul>



Attribute	Description
<b>CacheControl</b>	Specifies directives about the behavior that must be adhered to by caches involved in the chain comprising a request from a consumer to a service provider. See <a href="#">Section 2.1.3, “Consumer Cache Control Directives”</a> .
<b>Cookie</b>	Specifies a static cookie to be sent with all requests.
<b>BrowserType</b>	Specifies information about the browser from which the request originates. In the HTTP specification from the World Wide Web consortium (W3C) this is also known as the <i>user-agent</i> . Some servers optimize based on the client that is sending the request.
<b>Referer</b>	<p>Specifies the URL of the resource that directed the consumer to make requests on a particular service. The value is used as the value of the HTTP Referer property.</p> <p>This HTTP property is used when a request is the result of a browser user clicking on a hyperlink rather than typing a URL. This can allow the server to optimize processing based upon previous task flow, and to generate lists of back-links to resources for the purposes of logging, optimized caching, tracing of obsolete or mistyped links, and so on. However, it is typically not used in web services applications.</p> <p>If the <b>AutoRedirect</b> attribute is set to <b>true</b> and the request is redirected, any value specified in the <b>Referer</b> attribute is overridden. The value of the HTTP Referer property is set to the URL of the service that redirected the consumer’s original request.</p>
<b>DecoupledEndpoint</b>	<p>Specifies the URL of a decoupled endpoint for the receipt of responses over a separate provider-&gt;consumer connection. For more information on using decoupled endpoints see, <a href="#">Section 2.3, “Using the HTTP Transport in Decoupled Mode”</a>.</p> <p>You must configure both the consumer endpoint and the service provider endpoint to use WS-Addressing for the decoupled endpoint to work.</p>
<b>ProxyServer</b>	Specifies the URL of the proxy server through which requests are routed.
<b>ProxyServerPort</b>	Specifies the port number of the proxy server through which requests are routed.

Attribute	Description
<b>ProxyServerType</b>	Specifies the type of proxy server used to route requests. Valid values are: <ul style="list-style-type: none"> <li>• <b>HTTP</b>(default)</li> <li>• <b>SOCKS</b></li> </ul>

## Example

[Example 2.3, “HTTP Consumer Endpoint Configuration”](#) shows the configuration of an HTTP consumer endpoint that wants to keep its connection to the provider open between requests, that will only retransmit requests once per invocation, and that cannot use chunking streams.

### Example 2.3. HTTP Consumer Endpoint Configuration

```
<beans xmlns="http://www.springframework.org/schema/beans"
      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xmlns:http-conf="http://cxf.apache.org/transports/http/configuration"
      xsi:schemaLocation="http://cxf.apache.org/transports/http/configuration
http://cxf.apache.org/schemas/configuration/http-conf.xsd
                        http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd">

  <http-conf:conduit name="
{http://apache.org/hello_world_soap_http}SoapPort.http-conduit">
    <http-conf:client Connection="Keep-Alive"
                      MaxRetransmits="1"
                      AllowChunking="false" />
  </http-conf:conduit>
</beans>
```

## More information

For more information on HTTP conduits see [Appendix C, Conduits](#).

## 2.1.2. Using WSDL

### Namespace

The WSDL extension elements used to configure an HTTP consumer endpoint are defined in the namespace `http://cxf.apache.org/transports/http/configuration`. It is commonly referred to using the prefix `http-conf`. In order to use the HTTP configuration elements you must add the line shown in [Example 2.4, “HTTP Consumer WSDL Element's Namespace”](#) to the `definitions` element of your endpoint's WSDL document.

**Example 2.4. HTTP Consumer WSDL Element's Namespace**

```
<definitions ...
  xmlns:http-
  conf="http://cxf.apache.org/transport/http/configuration"
```

**The client element**

The `http-conf:client` element is used to specify the connection properties of an HTTP consumer in a WSDL document. The `http-conf:client` element is a child of the WSDL `port` element. It has the same attributes as the `client` element used in the configuration file. The attributes are described in [Table 2.2, “HTTP Consumer Configuration Attributes”](#).

**Example**

[Example 2.5, “WSDL to Configure an HTTP Consumer Endpoint”](#) shows a WSDL fragment that configures an HTTP consumer endpoint to specify that it does not interact with caches.

**Example 2.5. WSDL to Configure an HTTP Consumer Endpoint**

```
<service ... >
  <port ... >
    <soap:address ... />
    <http-conf:client CacheControl="no-cache" />
  </port>
</service>
```

**2.1.3. Consumer Cache Control Directives**

[Table 2.3, “http-conf:client Cache Control Directives”](#) lists the cache control directives supported by an HTTP consumer.

**Table 2.3. http-conf:client Cache Control Directives**

Directive	Behavior
no-cache	Caches cannot use a particular response to satisfy subsequent requests without first revalidating that response with the server. If specific response header fields are specified with this value, the restriction applies only to those header fields within the response. If no response header fields are specified, the restriction applies to the entire response.
no-store	Caches must not store either any part of a response or any part of the request that invoked it.
max-age	The consumer can accept a response whose age is no greater than the specified time in seconds.

Directive	Behavior
max-stale	The consumer can accept a response that has exceeded its expiration time. If a value is assigned to max-stale, it represents the number of seconds beyond the expiration time of a response up to which the consumer can still accept that response. If no value is assigned, the consumer can accept a stale response of any age.
min-fresh	The consumer wants a response that is still fresh for at least the specified number of seconds indicated.
no-transform	Caches must not modify media type or location of the content in a response between a provider and a consumer.
only-if-cached	Caches should return only responses that are currently stored in the cache, and not responses that need to be reloaded or revalidated.
cache-extension	Specifies additional extensions to the other cache directives. Extensions can be informational or behavioral. An extended directive is specified in the context of a standard directive, so that applications not understanding the extended directive can adhere to the behavior mandated by the standard directive.

## 2.2. CONFIGURING A SERVICE PROVIDER

HTTP service provider endpoints can specify a number of HTTP connection attributes including if it will honor keep alive requests, how it interacts with caches, and how tolerant it is of errors in communicating with a consumer.

A service provider endpoint can be configured using two mechanisms:

- [Configuration](#)
- [WSDL](#)

### 2.2.1. Using Configuration

#### Namespace

The elements used to configure an HTTP provider endpoint are defined in the namespace `http://cxf.apache.org/transports/http/configuration`. It is commonly referred to using the prefix `http-conf`. In order to use the HTTP configuration elements you must add the lines shown in [Example 2.6, “HTTP Provider Configuration Namespace”](#) to the `beans` element of your endpoint's configuration file. In addition, you must add the configuration elements' namespace to the `xsi:schemaLocation` attribute.

#### Example 2.6. HTTP Provider Configuration Namespace

```

<beans ...
  xmlns:http-
conf="http://cxf.apache.org/transports/http/configuration"
  ...
  xsi:schemaLocation="...

http://cxf.apache.org/transports/http/configuration

http://cxf.apache.org/schemas/configuration/http-conf.xsd
  ...">

```

## The destination element

You configure an HTTP service provider endpoint using the `http-conf:destination` element and its children. The `http-conf:destination` element takes a single attribute, `name`, that specifies the WSDL `port` element that corresponds to the endpoint. The value for the `name` attribute takes the form `portQName.http-destination`. [Example 2.7, “http-conf:destination Element”](#) shows the `http-conf:destination` element that is used to add configuration for an endpoint that is specified by the WSDL fragment `<port binding="widgetSOAPBinding" name="widgetSOAPPort">` when the endpoint's target namespace is `http://widgets.widgetvendor.net`.

### Example 2.7. http-conf:destination Element

```

...
  <http-conf:destination name="
{http://widgets/widgetvendor.net}widgetSOAPPort.http-destination">
  ...
  </http-conf:destination>
...

```

The `http-conf:destination` element has a number of child elements that specify configuration information. They are described in [Table 2.4, “Elements Used to Configure an HTTP Service Provider Endpoint”](#).

**Table 2.4. Elements Used to Configure an HTTP Service Provider Endpoint**

Element	Description
<code>http-conf:server</code>	Specifies the HTTP connection properties. See <a href="#">the section called “The server element”</a> .
<code>http-conf:contextMatchStrategy</code>	Specifies the parameters that configure the context match strategy for processing HTTP requests.
<code>http-conf:fixedParameterOrder</code>	Specifies whether the parameter order of an HTTP request handled by this destination is fixed.

## The server element

The `http-conf:server` element is used to configure the properties of a service provider endpoint's HTTP connection. Its attributes, described in [Table 2.5, “HTTP Service Provider Configuration Attributes”](#), specify the connection's properties.

**Table 2.5. HTTP Service Provider Configuration Attributes**

Attribute	Description
<b>ReceiveTimeout</b>	<p>Sets the length of time, in milliseconds, the service provider attempts to receive a request before the connection times out. The default is <b>30000</b>.</p> <p><b>0</b> specifies that the provider will not timeout.</p>
<b>SuppressClientSendErrors</b>	<p>Specifies whether exceptions are to be thrown when an error is encountered on receiving a request. The default is <b>false</b>; exceptions are thrown on encountering errors.</p>
<b>SuppressClientReceiveErrors</b>	<p>Specifies whether exceptions are to be thrown when an error is encountered on sending a response to a consumer. The default is <b>false</b>; exceptions are thrown on encountering errors.</p>
<b>HonorKeepAlive</b>	<p>Specifies whether the service provider honors requests for a connection to remain open after a response has been sent. The default is <b>false</b>; keep-alive requests are ignored.</p>
<b>RedirectURL</b>	<p>Specifies the URL to which the client request should be redirected if the URL specified in the client request is no longer appropriate for the requested resource. In this case, if a status code is not automatically set in the first line of the server response, the status code is set to <b>302</b> and the status description is set to <b>Object Moved</b>. The value is used as the value of the HTTP RedirectURL property.</p>
<b>CacheControl</b>	<p>Specifies directives about the behavior that must be adhered to by caches involved in the chain comprising a response from a service provider to a consumer. See <a href="#">Section 2.2.3, “Service Provider Cache Control Directives”</a>.</p>
<b>ContentLocation</b>	<p>Sets the URL where the resource being sent in a response is located.</p>
<b>ContentType</b>	<p>Specifies the media type of the information being sent in a response. Media types are specified using multipurpose internet mail extensions (MIME) types. The value is used as the value of the HTTP ContentType location.</p>

Attribute	Description
<b>ContentEncoding</b>	<p>Specifies any additional content encodings that have been applied to the information being sent by the service provider. Content encoding labels are regulated by the Internet Assigned Numbers Authority (IANA). Possible content encoding values include <b>zip</b>, <b>gzip</b>, <b>compress</b>, <b>deflate</b>, and <b>identity</b>. This value is used as the value of the HTTP ContentEncoding property.</p> <p>The primary use of content encodings is to allow documents to be compressed using some encoding mechanism, such as zip or gzip. Apache CXF performs no validation on content codings. It is the user's responsibility to ensure that a specified content coding is supported at application level.</p>
<b>ServerType</b>	<p>Specifies what type of server is sending the response. Values take the form <i>program-name/version</i>; for example, <b>Apache/1.2.5</b>.</p>

## Example

[Example 2.8, “HTTP Service Provider Endpoint Configuration”](#) shows the configuration for an HTTP service provider endpoint that honors keep-alive requests and suppresses all communication errors.

### Example 2.8. HTTP Service Provider Endpoint Configuration

```
<beans xmlns="http://www.springframework.org/schema/beans"
      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xmlns:http-conf="http://cxf.apache.org/transports/http/configuration"
      xsi:schemaLocation="http://cxf.apache.org/transports/http/configuration
http://cxf.apache.org/schemas/configuration/http-conf.xsd
                        http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd">
  <http-conf:destination name="
{http://apache.org/hello_world_soap_http}SoapPort.http-destination">
    <http-conf:server SuppressClientSendErrors="true"
                      SuppressClientReceiveErrors="true"
                      HonorKeepAlive="true" />
  </http-conf:destination>
</beans>
```

## 2.2.2. Using WSDL

### Namespace

The WSDL extension elements used to configure an HTTP provider endpoint are defined in the namespace `http://cxf.apache.org/transports/http/configuration`. It is commonly referred to using the prefix `http-conf`. To use the HTTP configuration elements you must add the line shown in [Example 2.9, “HTTP Provider WSDL Element's Namespace”](#) to the `definitions` element of your endpoint's WSDL document.

#### Example 2.9. HTTP Provider WSDL Element's Namespace

```
<definitions ...
  xmlns:http-
  conf="http://cxf.apache.org/transports/http/configuration"
```

### The server element

The `http-conf:server` element is used to specify the connection properties of an HTTP service provider in a WSDL document. The `http-conf:server` element is a child of the WSDL `port` element. It has the same attributes as the `server` element used in the configuration file. The attributes are described in [Table 2.5, “HTTP Service Provider Configuration Attributes”](#).

#### Example

[Example 2.10, “WSDL to Configure an HTTP Service Provider Endpoint”](#) shows a WSDL fragment that configures an HTTP service provider endpoint specifying that it will not interact with caches.

#### Example 2.10. WSDL to Configure an HTTP Service Provider Endpoint

```
<service ... >
  <port ... >
    <soap:address ... />
    <http-conf:server CacheControl="no-cache" />
  </port>
</service>
```

### 2.2.3. Service Provider Cache Control Directives

[Table 2.6, “http-conf:server Cache Control Directives”](#) lists the cache control directives supported by an HTTP service provider.

Table 2.6. `http-conf:server` Cache Control Directives

Directive	Behavior
no-cache	Caches cannot use a particular response to satisfy subsequent requests without first revalidating that response with the server. If specific response header fields are specified with this value, the restriction applies only to those header fields within the response. If no response header fields are specified, the restriction applies to the entire response.



Directive	Behavior
public	Any cache can store the response.
private	Public ( <i>shared</i> ) caches cannot store the response because the response is intended for a single user. If specific response header fields are specified with this value, the restriction applies only to those header fields within the response. If no response header fields are specified, the restriction applies to the entire response.
no-store	Caches must not store any part of the response or any part of the request that invoked it.
no-transform	Caches must not modify the media type or location of the content in a response between a server and a client.
must-revalidate	Caches must revalidate expired entries that relate to a response before that entry can be used in a subsequent response.
proxy-revalidate	Does the same as must-revalidate, except that it can only be enforced on shared caches and is ignored by private unshared caches. When using this directive, the public cache directive must also be used.
max-age	Clients can accept a response whose age is no greater than the specified number of seconds.
s-max-age	Does the same as max-age, except that it can only be enforced on shared caches and is ignored by private unshared caches. The age specified by s-max-age overrides the age specified by max-age. When using this directive, the proxy-revalidate directive must also be used.
cache-extension	Specifies additional extensions to the other cache directives. Extensions can be informational or behavioral. An extended directive is specified in the context of a standard directive, so that applications not understanding the extended directive can adhere to the behavior mandated by the standard directive.

## 2.3. USING THE HTTP TRANSPORT IN DECOUPLED MODE

### Overview

In normal HTTP request/response scenarios, the request and the response are sent using the same HTTP connection. The service provider processes the request and responds with a response containing the appropriate HTTP status code and the contents of the response. In the case of a successful request, the HTTP status code is set to **200**.

In some instances, such as when using WS-RM or when requests take an extended period of time to execute, it makes sense to decouple the request and response message. In this case the service providers sends the consumer a **202 Accepted** response to the consumer over the back-channel of the HTTP connection on which the request was received. It then processes the request and sends the response back to the consumer using a new decoupled server->client HTTP connection. The consumer runtime receives the incoming response and correlates it with the appropriate request before returning to the application code.

## Configuring decoupled interactions

Using the HTTP transport in decoupled mode requires that you do the following:

1. Configure the consumer to use WS-Addressing.  
See [the section called “Configuring an endpoint to use WS-Addressing”](#).
2. Configure the consumer to use a decoupled endpoint.  
See [the section called “Configuring the consumer”](#).
3. Configure any service providers that the consumer interacts with to use WS-Addressing.  
See [the section called “Configuring an endpoint to use WS-Addressing”](#).

## Configuring an endpoint to use WS-Addressing

Specify that the consumer and any service provider with which the consumer interacts use WS-Addressing.

You can specify that an endpoint uses WS-Addressing in one of two ways:

- Adding the `wsa:UsingAddressing` element to the endpoint's WSDL `port` element as shown in [Example 2.11, “Activating WS-Addressing using WSDL”](#).

### Example 2.11. Activating WS-Addressing using WSDL

```
...
<service name="WidgetSOAPService">
  <port name="WidgetSOAPPort" binding="tns:WidgetSOAPBinding">
    <soap:address="http://widgetvendor.net/widgetSeller" />
    <wsa:UsingAddressing
xmlns:wsa="http://www.w3.org/2005/02/addressing/wsd1"/>
  </port>
</service>
...
```

- Adding the WS-Addressing policy to the endpoint's WSDL `port` element as shown in [Example 2.12, “Activating WS-Addressing using a Policy”](#).

### Example 2.12. Activating WS-Addressing using a Policy

```
...
<service name="WidgetSOAPService">
  <port name="WidgetSOAPPort" binding="tns:WidgetSOAPBinding">
```

```

    <soap:address="http://widgetvendor.net/widgetSeller" />
    <wsp:Policy xmlns:wsp="http://www.w3.org/2006/07/ws-policy">
      <wsam:Addressing
xmlns:wsam="http://www.w3.org/2007/02/addressing/metadata">
        <wsp:Policy/>
      </wsam:Addressing>
    </wsp:Policy>
  </port>
</service>
...

```

**NOTE**

The WS-Addressing policy supersedes the `wsa:UsingAddressing` WSDL element.

**Configuring the consumer**

Configure the consumer endpoint to use a decoupled endpoint using the `DecoupledEndpoint` attribute of the `http-conf:conduit` element.

[Example 2.13, “Configuring a Consumer to Use a Decoupled HTTP Endpoint”](#) shows the configuration for setting up the endpoint defined in [Example 2.11, “Activating WS-Addressing using WSDL”](#) to use a decoupled endpoint. The consumer now receives all responses at `http://widgetvendor.net/widgetSellerInbox`.

**Example 2.13. Configuring a Consumer to Use a Decoupled HTTP Endpoint**

```

<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:http="http://cxf.apache.org/transports/http/configuration"

  xsi:schemaLocation="http://cxf.apache.org/transports/http/configuration
http://cxf.apache.org/schemas/configuration/http-conf.xsd
  http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd">

  <http:conduit name="
{http://widgetvendor.net/services}WidgetSOAPPort.http-conduit">
    <http:client
DecoupledEndpoint="http://widgetvendor.net:9999/decoupled_endpoint" />
  </http:conduit>
</beans>

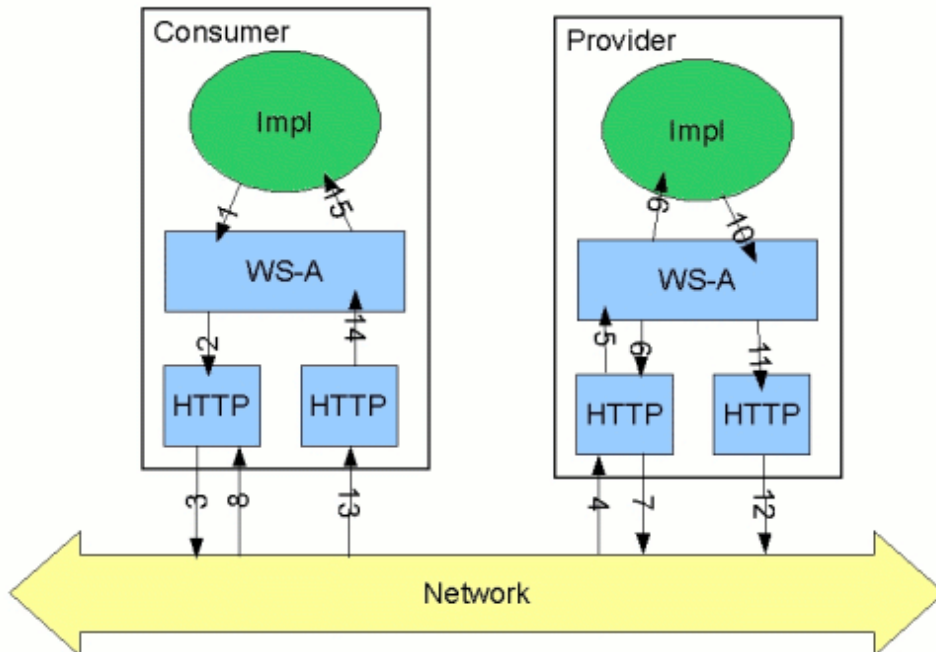
```

**How messages are processed**

Using the HTTP transport in decoupled mode adds extra layers of complexity to the processing of HTTP messages. While the added complexity is transparent to the implementation level code in an application, it might be important to understand what happens for debugging reasons.

Figure 2.1, “Message Flow in for a Decoupled HTTP Transport” shows the flow of messages when using HTTP in decoupled mode.

Figure 2.1. Message Flow in for a Decoupled HTTP Transport



A request starts the following process:

1. The consumer implementation invokes an operation and a request message is generated.
2. The WS-Addressing layer adds the WS-A headers to the message.

When a decoupled endpoint is specified in the consumer's configuration, the address of the decoupled endpoint is placed in the WS-A ReplyTo header.

3. The message is sent to the service provider.
4. The service provider receives the message.
5. The request message from the consumer is dispatched to the provider's WS-A layer.
6. Because the WS-A ReplyTo header is not set to anonymous, the provider sends back a message with the HTTP status code set to **202**, acknowledging that the request has been received.
7. The HTTP layer sends a **202 Accepted** message back to the consumer using the original connection's back-channel.

8. The consumer receives the **202 Accepted** reply on the back-channel of the HTTP connection used to send the original message.

When the consumer receives the **202 Accepted** reply, the HTTP connection closes.

9. The request is passed to the service provider's implementation where the request is processed.
10. When the response is ready, it is dispatched to the WS-A layer.
11. The WS-A layer adds the WS-Addressing headers to the response message.
12. The HTTP transport sends the response to the consumer's decoupled endpoint.
13. The consumer's decoupled endpoint receives the response from the service provider.
14. The response is dispatched to the consumer's WS-A layer where it is correlated to the proper request using the WS-A `RelatesTo` header.
15. The correlated response is returned to the client implementation and the invoking call is unblocked.

## CHAPTER 3. USING SOAP OVER JMS

### Abstract

Apache CXF implements the W3C standard SOAP/JMS transport. This standard is intended to provide a more robust alternative to SOAP/HTTP services. Apache CXF applications using this transport should be able to interoperate with applications that also implement the SOAP/JMS standard. The transport is configured directly in an endpoint's WSDL.

### 3.1. BASIC CONFIGURATION

#### Overview

The [SOAP over JMS protocol](#) is defined by the World Wide Web Consortium(W3C) as a way of providing a more reliable transport layer to the customary SOAP/HTTP protocol used by most services. The Apache CXF implementation is fully compliant with the specification and should be compatible with any framework that is also compliant.

This transport uses JNDI to find the JMS destinations. When an operation is invoked, the request is packaged as a SOAP message and sent in the body of a JMS message to the specified destination.

To use the SOAP/JMS transport:

1. Specify that the transport type is SOAP/JMS.
2. Specify the target destination using a JMS URI.
3. Optionally, configure the JNDI connection.
4. Optionally, add additional JMS configuration.

#### Specifying the JMS transport type

You configure a SOAP binding to use the JMS transport when specifying the WSDL binding. You set the `soap:binding` element's `transport` attribute to `http://www.w3.org/2010/soapjms/`.

[Example 3.1, "SOAP over JMS binding specification"](#) shows a WSDL binding that uses SOAP/JMS.

##### Example 3.1. SOAP over JMS binding specification

```
<wsdl:binding ... >
  <soap:binding style="document"
                transport="http://www.w3.org/2010/soapjms/" />
  ...
</wsdl:binding>
```

#### Specifying the target destination

You specify the address of the JMS target destination when specifying the WSDL port for the endpoint. The address specification for a SOAP/JMS endpoint uses the same `soap:address` element and attribute as a SOAP/HTTP endpoint. The difference is the address specification. JMS endpoints use a

JMS URI as defined in the [URI Scheme for JMS 1.0](#). [Example 3.2, “JMS URI syntax”](#) shows the syntax for a JMS URI.

### Example 3.2. JMS URI syntax

```
jms:variant:destination?options
```

[Table 3.1, “JMS URI variants”](#) describes the available variants for the JMS URI.

**Table 3.1. JMS URI variants**

Variant	Description
jndi	Specifies that the destination is a JNDI name for the target destination. When using this variant, you must provide the configuration for accessing the JNDI provider.
topic	Specifies that the destination is the name of the topic to be used as the target destination. The string provided is passed into <code>Session.createTopic()</code> to create a representation of the destination.
queue	Specifies that the destination is the name of the queue to be used as the target destination. The string provided is passed into <code>Session.createQueue()</code> to create a representation of the destination.

The *options* portion of a JMS URI are used to configure the transport and are discussed in [Section 3.2, “JMS URIs”](#).

[Example 3.3, “SOAP/JMS endpoint address”](#) shows the WSDL port entry for a SOAP/JMS endpoint whose target destination is looked up using JNDI.

### Example 3.3. SOAP/JMS endpoint address

```
<wsdl:port ... >
  ...
  <soap:address
location="jms:jndi:dynamicQueues/test.cxf.jmstransport.queue" />
</wsdl:port>
```

For working with SOAP/JMS services in Java see [chapter “Using SOAP over JMS”](#) in [“Developing Applications Using JAX-WS”](#).

## Configuring JNDI and the JMS transport

The SOAP/JMS provides several ways to configure the JNDI connection and the JMS transport:

- [Using the JMS URI](#)
- [Using WSDL extensions](#)

## 3.2. JMS URIS

### Overview

When using SOAP/JMS, a JMS URI is used to specify the endpoint's target destination. The JMS URI can also be used to configure JMS connection by appending one or more options to the URI. These options are detailed in the IETF standard, [URI Scheme for Java Message Service 1.0](#). They can be used to configure the JNDI system, the reply destination, the delivery mode to use, and other JMS properties.

### Syntax

As shown in [Example 3.2, “JMS URI syntax”](#), you can append one or more options to the end of a JMS URI by separating them from the destination's address with a question mark(?). Multiple options are separated by an ampersand(&). [Example 3.4, “Syntax for JMS URI options”](#) shows the syntax for using multiple options in a JMS URI.

#### Example 3.4. Syntax for JMS URI options

```
jmsAddress?option1=value1&option2=value2&...optionN=valueN
```

### JMS properties

[Table 3.2, “JMS properties settable as URI options”](#) shows the URI options that affect the JMS transport layer.

**Table 3.2. JMS properties settable as URI options**

Property	Default	Description
<code>deliveryMode</code>	<b>PERSISTENT</b>	Specifies whether to use JMS <b>PERSISTENT</b> or <b>NON_PERSISTENT</b> message semantics. In the case of <b>PERSISTENT</b> delivery mode, the JMS broker stores messages in persistent storage before acknowledging them; whereas <b>NON_PERSISTENT</b> messages are kept in memory only.



Property	Default	Description
<b>replyToName</b>		<p>Explicitly specifies the reply destination to appear in the <b>JMSReplyTo</b> header. Setting this property is recommended for applications that have request-reply semantics because the JMS provider will assign a temporary reply queue if one is not explicitly set.</p> <p>The value of this property has an interpretation that depends on the variant specified in the JMS URI:</p> <ul style="list-style-type: none"> <li>• <b>jndi</b> variant—the JNDI name of the destination</li> <li>• <b>queue</b> or <b>topic</b> variants—the actual name of the destination</li> </ul>
<b>priority</b>	<b>4</b>	Specifies the JMS message priority, which ranges from 0 (lowest) to 9 (highest).
<b>timeToLive</b>	<b>0</b>	Time (in milliseconds) after which the message will be discarded by the JMS provider. A value of <b>0</b> represents an infinite lifetime (the default).

## JNDI properties

Table 3.3, “JNDI properties settable as URI options” shows the URI options that can be used to configure JNDI for this endpoint.

Table 3.3. JNDI properties settable as URI options

Property	Description
<b>jndiConnectionFactoryName</b>	Specifies the JNDI name of the JMS connection factory.
<b>jndiInitialContextFactory</b>	Specifies the fully qualified Java class name of the JNDI provider (which must be of <b>javax.jms.InitialContextFactory</b> type). Equivalent to setting the <b>java.naming.factory.initial</b> Java system property.

Property	Description
<code>jndiURL</code>	Specifies the URL that initializes the JNDI provider. Equivalent to setting the <code>java.naming.provider.url</code> Java system property.

## Additional JNDI properties

The properties, `java.naming.factory.initial` and `java.naming.provider.url`, are standard properties, which are required to initialize any JNDI provider. Sometimes, however, a JNDI provider might support custom properties in addition to the standard ones. In this case, you can set an arbitrary JNDI property by setting a URI option of the form `jndi-PropertyName`.

For example, if you were using SUN's LDAP implementation of JNDI, you could set the JNDI property, `java.naming.factory.control`, in a JMS URI as shown in [Example 3.5, “Setting a JNDI property in a JMS URI”](#).

### Example 3.5. Setting a JNDI property in a JMS URI

```
jms:queue:FOO.BAR?jndi-
java.naming.factory.control=com.sun.jndi.ldap.ResponseControlFactory
```

## Example

If the JMS provider is *not* already configured, it is possible to provide the requisite JNDI configuration details in the URI using options (see [Table 3.3, “JNDI properties settable as URI options”](#)). For example, to configure an endpoint to use the Apache ActiveMQ JMS provider and connect to the queue called `test.cxf.jmstransport.queue`, use the URI shown in [Example 3.6, “JMS URI that configures a JNDI connection”](#).

### Example 3.6. JMS URI that configures a JNDI connection

```
jms:jndi:dynamicQueues/test.cxf.jmstransport.queue
?
jndiInitialContextFactory=org.apache.activemq.jndi.ActiveMQInitialContextFactory
&jndiConnectionFactoryName=ConnectionFactory
&jndiURL=tcp://localhost:61616
```

## 3.3. WSDL EXTENSIONS

### Overview

You can specify the basic configuration of the JMS transport by inserting WSDL extension elements into the contract, either at binding scope, service scope, or port scope. The WSDL extensions enable you to specify the properties for bootstrapping a JNDI `InitialContext`, which can then be used to

look up JMS destinations. You can also set some properties that affect the behavior of the JMS transport layer.

## SOAP/JMS namespace

the SOAP/JMS WSDL extensions are defined in the `http://www.w3.org/2010/soapjms/` namespace. To use them in your WSDL contracts add the following setting to the `wSDL:definitions` element:

```
<wSDL:definitions ...
  xmlns:soapjms="http://www.w3.org/2010/soapjms/"
  ... >
```

## WSDL extension elements

Table 3.4, “SOAP/JMS WSDL extension elements” shows all of the WSDL extension elements you can use to configure the JMS transport.

Table 3.4. SOAP/JMS WSDL extension elements

Element	Default	Description
<code>soapjms:jndiInitialContextFactory</code>		Specifies the fully qualified Java class name of the JNDI provider. Equivalent to setting the <code>java.naming.factory.initial</code> Java system property.
<code>soapjms:jndiURL</code>		Specifies the URL that initializes the JNDI provider. Equivalent to setting the <code>java.naming.provider.url</code> Java system property.
<code>soapjms:jndiContextParameter</code>		Enables you to specify an additional property for creating the JNDI <code>InitialContext</code> . Use the <code>name</code> and <code>value</code> attributes to specify the property.
<code>soapjms:jndiConnectionFactoryName</code>		Specifies the JNDI name of the JMS connection factory.

Element	Default	Description
<code>soapjms:deliveryMode</code>	<b>PERSISTENT</b>	Specifies whether to use JMS <b>PERSISTENT</b> or <b>NON_PERSISTENT</b> message semantics. In the case of <b>PERSISTENT</b> delivery mode, the JMS broker stores messages in persistent storage before acknowledging them; whereas <b>NON_PERSISTENT</b> messages are kept in memory only.
<code>soapjms:replyToName</code>		<p>Explicitly specifies the reply destination to appear in the <b>JMSReplyTo</b> header. Setting this property is recommended for SOAP invocations that have request-reply semantics. If this property is not set the JMS provider allocates a temporary queue with an automatically generated name.</p> <p>The value of this property has an interpretation that depends on the variant specified in the JMS URI, as follows:</p> <ul style="list-style-type: none"> <li>• <b>jndi</b> variant—the JNDI name of the destination.</li> <li>• <b>queue</b> or <b>topic</b> variants—the actual name of the destination.</li> </ul>
<code>soapjms:priority</code>	<b>4</b>	Specifies the JMS message priority, which ranges from 0 (lowest) to 9 (highest).
<code>soapjms:timeToLive</code>	<b>0</b>	Time, in milliseconds, after which the message will be discarded by the JMS provider. A value of <b>0</b> represents an infinite lifetime.

## Configuration scopes

The WSDL elements placement in the WSDL contract effect the scope of the configuration changes on the endpoints defined in the contract. The SOAP/JMS WSDL elements can be placed as children of either the `wsdl:binding` element, the `wsdl:service` element, or the `wsdl:port` element. The parent of the SOAP/JMS elements determine which of the following scopes the configuration is placed into.

### Binding scope

You can configure the JMS transport at the *binding scope* by placing extension elements inside the `wsdl:binding` element. Elements in this scope define the default configuration for all endpoints that use this binding. Any settings in the binding scope can be overridden at the service scope or the port scope.

### Service scope

You can configure the JMS transport at the *service scope* by placing extension elements inside a `wsdl:service` element. Elements in this scope define the default configuration for all endpoints in this service. Any settings in the service scope can be overridden at the port scope.

### Port scope

You can configure the JMS transport at the *port scope* by placing extension elements inside a `wsdl:port` element. Elements in the port scope define the configuration for this port. They override any defaults defined at the service scope or at the binding scope.

## Example

[Example 3.7, “WSDL contract with SOAP/JMS configuration”](#) shows a WSDL contract for a SOAP/JMS service. It configures the JNDI layer in the binding scope, the message delivery details in the service scope, and the reply destination in the port scope.

### Example 3.7. WSDL contract with SOAP/JMS configuration

```

<wsd:definitions ...
  1  xmlns:soapjms="http://www.w3.org/2010/soapjms/"
    ... >
    ...
    <wsdl:binding name="JMSSGreeterPortBinding"
type="tns:JMSSGreeterPortType">
    ...
  2  <soapjms:jndiInitialContextFactory>
        org.apache.activemq.jndi.ActiveMQInitialContextFactory
    </soapjms:jndiInitialContextFactory>
    <soapjms:jndiURL>tcp://localhost:61616</soapjms:jndiURL>
    <soapjms:jndiConnectionFactoryName>
        ConnectionFactory
    </soapjms:jndiConnectionFactoryName>
    ...
    </wsdl:binding>
    ...
    <wsdl:service name="JMSSGreeterService">
    ...
  3  <soapjms:deliveryMode>NON_PERSISTENT</soapjms:deliveryMode>
        <soapjms:timeToLive>60000</soapjms:timeToLive>
    ...
    <wsdl:port binding="tns:JMSSGreeterPortBinding" name="GreeterPort">
  4  <soap:address
location="jms:jndi:dynamicQueues/test.cxf.jmstransport.queue" />
  5  <soapjms:replyToName>
        dynamicQueues/greeterReply.queue
    </soapjms:replyToName>
    ...
    </wsdl:port>

```

```
    ...  
  </wsdl:service>  
  ...  
</wsdl:definitions>
```

The WSDL in [Example 3.7, “WSDL contract with SOAP/JMS configuration”](#) does the following:

- 1 Declare the namespace for the SOAP/JMS extensions.
- 2 Configure the JNDI connections in the binding scope.
- 3 Configure the JMS delivery style to non-persistent and each message to live for one minute.
- 4 Specify the target destination.
- 5 Configure the JMS transport so that reply messages are delivered on the `greeterReply.queue` queue.

## CHAPTER 4. USING GENERIC JMS

### Abstract

Apache CXF provides a generic implementation of a JMS transport. The generic JMS transport is not restricted to using SOAP messages and allows for connecting to any application that uses JMS.

The Apache CXF generic JMS transport can connect to any JMS provider and work with applications that exchange JMS messages with bodies of either `TextMessage` or `ByteMessage`.

There are two ways to enable and configure the JMS transport:

- [JMS configuration bean](#)
- [WSDL](#)

### 4.1. USING THE JMS CONFIGURATION BEAN

#### Overview

To simplify JMS configuration and make it more powerful, Apache CXF uses a single JMS configuration bean to configure JMS endpoints. The bean is implemented by the `org.apache.cxf.transport.jms.JMSConfiguration` class. It can be used to either configure endpoint's directly or to configure the JMS conduits and destinations.

#### Configuration namespace

The JMS configuration bean uses the [Spring p-namespace](#) to make the configuration as simple as possible. To use this namespace you need to declare it in the configuration's root element as shown in [Example 4.1, “Declaring the Spring p-namespace”](#).

#### Example 4.1. Declaring the Spring p-namespace

```
<beans ...
  xmlns:p="http://www.springframework.org/schema/p"
  ... >
  ...
</beans>
```

#### Specifying the configuration

You specify the JMS configuration by defining a bean of class `org.apache.cxf.transport.jms.JMSConfiguration`. The properties of the bean provide the configuration settings for the transport.

[Table 4.1, “General JMS Configuration Properties”](#) lists properties that are common to both providers and consumers.

#### Table 4.1. General JMS Configuration Properties

Property	Default	Description
<b>connectionFactory-ref</b>		Specifies a reference to a bean that defines a JMS <b>ConnectionFactory</b> .
<b>wrapInSingleConnectionFactory</b>	<b>true</b>	Specifies whether to wrap the <b>ConnectionFactory</b> with a Spring <b>SingleConnectionFactory</b> . Doing so can improve the performance of the JMS transport when the specified connection factory does not pool connections.
<b>reconnectOnException</b>	<b>false</b>	Specifies whether to create a new connection in the case of an exception. This property is only used when wrapping the connection factory with a Spring <b>SingleConnectionFactory</b> .
<b>targetDestination</b>		Specifies the JNDI name or provider specific name of a destination.
<b>replyDestination</b>		Specifies the JMS name of the JMS destinations where replies are sent. This attribute allows you to use a user defined destination for replies. For more details see <a href="#">Section 4.3, “Using a Named Reply Destination”</a> .
<b>destinationResolver</b>		Specifies a reference to a Spring <b>DestinationResolver</b> . This allows you to define how destination names are resolved. By default a <b>DynamicDestinationResolver</b> is used. It resolves destinations using the JMS providers features. If you reference a <b>JndiDestinationResolver</b> you can resolve the destination names using JNDI.



Property	Default	Description
<b>transactionManager</b>		Specifies a reference to a Spring transaction manager. This allows the service to participate in JTA Transactions.
<b>taskExecutor</b>		Specifies a reference to a Spring <b>TaskExecutor</b> . This is used in listeners to decide how to handle incoming messages. By default the transport uses the Spring <b>SimpleAsyncTaskExecutor</b> .
<b>useJms11</b>	<b>false</b>	Specifies whether JMS 1.1 features are available.
<b>messageIdEnabled</b>	<b>true</b>	Specifies whether the JMS transport wants the JMS broker to provide message IDs. Setting this to <b>false</b> causes the endpoint to call its message producer's <b>setDisableMessageID()</b> method with a value of <b>true</b> . The JMS broker is then given a hint that it does not need to generate message IDs or add them to the messages from the endpoint. The JMS broker can choose to accept the hint or ignore it.
<b>messageTimestampEnabled</b>	<b>true</b>	Specifies whether the JMS transport wants the JMS broker to provide message time stamps. Setting this to <b>false</b> causes the endpoint to call its message producer's <b>setDisableMessageTimestamp()</b> method with a value of <b>true</b> . The JMS broker is then given a hint that it does not need to generate time stamps or add them to the messages from the endpoint. The JMS broker can choose to accept the hint or ignore it.

Property	Default	Description
<b>cacheLevel</b>	<b>3</b>	Specifies the level of caching allowed by the listener. Valid values are <b>0</b> (CACHE_NONE), <b>1</b> (CACHE_CONNECTION), <b>2</b> (CACHE_SESSION), <b>3</b> (CACHE_CONSUMER), <b>4</b> (CACHE_AUTO).
<b>pubSubNoLocal</b>	<b>false</b>	Specifies whether to receive messages produced from the same connection.
<b>receiveTimeout</b>	<b>0</b>	Specifies, in milliseconds, the amount of time to wait for response messages. <b>0</b> means wait indefinitely.
<b>explicitQosEnabled</b>	<b>false</b>	Specifies whether the QoS settings like priority, persistence, and time to live are explicitly set for each message or if they are allowed to use default values.
<b>deliveryMode</b>	<b>1</b>	Specifies if a message is persistent. The two values are: <ul style="list-style-type: none"> <li><b>1(NON_PERSISTENT)</b>— messages will be kept memory</li> <li><b>2(PERSISTENT)</b>— messages will be persisted to disk</li> </ul>
<b>priority</b>	<b>4</b>	Specifies the message's priority for the messages. JMS priority values can range from 0 to 9. The lowest priority is 0 and the highest priority is 9.
<b>timeToLive</b>	<b>0</b>	Specifies, in milliseconds, the message will be available after it is sent. 0 specifies an infinite time to live.
<b>sessionTransacted</b>	<b>false</b>	Specifies if JMS transactions are used.

Property	Default	Description
<code>concurrentConsumers</code>	<code>1</code>	Specifies the minimum number of concurrent consumers created by the listener.
<code>maxConcurrentConsumers</code>	<code>1</code>	Specifies the maximum number of concurrent consumers by listener.
<code>messageSelector</code>		Specifies the string value of the selector. For more information on the syntax used to specify message selectors, see the JMS 1.1 specification.
<code>subscriptionDurable</code>	<code>false</code>	Specifies whether the server uses durable subscriptions.
<code>durableSubscriptionName</code>		Specifies the string used to register the durable subscription.
<code>messageType</code>	<code>text</code>	Specifies how the message data will be packaged as a JMS message. <code>text</code> specifies that the data will be packaged as a <code>TextMessage</code> . <code>binary</code> specifies that the data will be packaged as an <code>ByteMessage</code> .
<code>pubSubDomain</code>	<code>false</code>	Specifies whether the target destination is a topic.
<code>jmsProviderTibcoEms</code>	<code>false</code>	Specifies if your JMS provider is Tibco EMS. This causes the principal in the security context to be populated from the <code>JMS_TIBCO_SENDER</code> header.
<code>useMessageIDAsCorrelationID</code>	<code>false</code>	Specifies whether JMS will use the message ID to correlate messages. If not, the client will set a generated correlation ID.

As shown in [Example 4.2, “JMS configuration bean”](#), the bean's properties are specified as attributes to the `bean` element. They are all declared in the Spring `p` namespace.

#### Example 4.2. JMS configuration bean

```
<bean id="jmsConfig"
      class="org.apache.cxf.transport.jms.JMSConfiguration"
```

```
p:connectionFactory-ref="connectionFactory"
p:targetDestination="dynamicQueues/greeter.request.queue"
p:pubSubDomain="false" />
```

## Applying the configuration to an endpoint

The `JMSConfiguration` bean can be applied directly to both server and client endpoints using the Apache CXF features mechanism. To do so:

1. Set the endpoint's `address` attribute to `jms://`.
2. Add a `jaxws:feature` element to the endpoint's configuration.
3. Add a bean of type `org.apache.cxf.transport.jms.JMSConfigFeature` to the feature.
4. Set the bean element's `p:jmsConfig-ref` attribute to the ID of the `JMSConfiguration` bean.

[Example 4.3, “Adding JMS configuration to a JAX-WS client”](#) shows a JAX-WS client that uses the JMS configuration from [Example 4.2, “JMS configuration bean”](#).

### Example 4.3. Adding JMS configuration to a JAX-WS client

```
<jaxws:client id="CustomerService"
  xmlns:customer="http://customerservice.example.com/"
  serviceName="customer:CustomerServiceService"
  endpointName="customer:CustomerServiceEndpoint"
  address="jms://"
  serviceClass="com.example.customerservice.CustomerService">
  <jaxws:features>
    <bean class="org.apache.cxf.transport.jms.JMSConfigFeature"
      p:jmsConfig-ref="jmsConfig"/>
  </jaxws:features>
</jaxws:client>
```

## Applying the configuration to the transport

The `JMSConfiguration` bean can be applied to JMS conduits and JMS destinations using the `jms:jmsConfig-ref` element. The `jms:jmsConfig-ref` element's value is the ID of the `JMSConfiguration` bean.

[Example 4.4, “Adding JMS configuration to a JMS conduit”](#) shows a JMS conduit that uses the JMS configuration from [Example 4.2, “JMS configuration bean”](#).

### Example 4.4. Adding JMS configuration to a JMS conduit

```
<jms:conduit name="
{http://cxf.apache.org/jms_conf_test>HelloWorldQueueBinMsgPort.jms-
conduit">
```

```

...
<jms:jmsConfig-ref>jmsConf</jms:jmsConfig-ref>
</jms:conduit>

```

## 4.2. USING WSDL TO CONFIGURE JMS

The WSDL extensions for defining a JMS endpoint are defined in the namespace `http://cxf.apache.org/transports/jms`. In order to use the JMS extensions you will need to add the line shown in [Example 4.5, “JMS WSDL extension namespace”](#) to the definitions element of your contract.

### Example 4.5. JMS WSDL extension namespace

```
xmlns:jms="http://cxf.apache.org/transports/jms"
```

### 4.2.1. Basic JMS configuration

#### Overview

The JMS address information is provided using the `jms:address` element and its child, the `jms:JMSNamingProperties` element. The `jms:address` element’s attributes specify the information needed to identify the JMS broker and the destination. The `jms:JMSNamingProperties` element specifies the Java properties used to connect to the JNDI service.



#### IMPORTANT

Information specified using the JMS feature will override the information in the endpoint’s WSDL file.

#### Specifying the JMS address

The basic configuration for a JMS endpoint is done by using a `jms:address` element as the child of your service’s `port` element. The `jms:address` element used in WSDL is identical to the one used in the configuration file. Its attributes are listed in [Table 4.2, “JMS endpoint attributes”](#).

Table 4.2. JMS endpoint attributes

Attribute	Description
<code>destinationStyle</code>	Specifies if the JMS destination is a JMS queue or a JMS topic.
<code>jndiConnectionFactoryName</code>	Specifies the JNDI name bound to the JMS connection factory to use when connecting to the JMS destination.
<code>jmsDestinationName</code>	Specifies the JMS name of the JMS destination to which requests are sent.

Attribute	Description
<b>jmsReplyDestinationName</b>	Specifies the JMS name of the JMS destinations where replies are sent. This attribute allows you to use a user defined destination for replies. For more details see <a href="#">Section 4.3, “Using a Named Reply Destination”</a> .
<b>jndiDestinationName</b>	Specifies the JNDI name bound to the JMS destination to which requests are sent.
<b>jndiReplyDestinationName</b>	Specifies the JNDI name bound to the JMS destinations where replies are sent. This attribute allows you to use a user defined destination for replies. For more details see <a href="#">Section 4.3, “Using a Named Reply Destination”</a> .
<b>connectionUserName</b>	Specifies the user name to use when connecting to a JMS broker.
<b>connectionPassword</b>	Specifies the password to use when connecting to a JMS broker.

The `jms:address` WSDL element uses a `jms:JMSNamingProperties` child element to specify additional information needed to connect to a JNDI provider.

### Specifying JNDI properties

To increase interoperability with JMS and JNDI providers, the `jms:address` element has a child element, `jms:JMSNamingProperties`, that allows you to specify the values used to populate the properties used when connecting to the JNDI provider. The `jms:JMSNamingProperties` element has two attributes: `name` and `value`. `name` specifies the name of the property to set. `value` attribute specifies the value for the specified property. `jms:JMSNamingProperties` element can also be used for specification of provider specific properties.

The following is a list of common JNDI properties that can be set:

1. `java.naming.factory.initial`
2. `java.naming.provider.url`
3. `java.naming.factory.object`
4. `java.naming.factory.state`
5. `java.naming.factory.url.pkgs`
6. `java.naming.dns.url`
7. `java.naming.authoritative`
8. `java.naming.batchsize`

9. `java.naming.referral`
10. `java.naming.security.protocol`
11. `java.naming.security.authentication`
12. `java.naming.security.principal`
13. `java.naming.security.credentials`
14. `java.naming.language`
15. `java.naming.applet`

For more details on what information to use in these attributes, check your JNDI provider's documentation and consult the Java API reference material.

## Example

[Example 4.6, "JMS WSDL port specification"](#) shows an example of a JMS WSDL port specification.

### Example 4.6. JMS WSDL port specification

```
<service name="JMSService">
  <port binding="tns:Greeter_SOAPBinding" name="SoapPort">
    <jms:address jndiConnectionFactoryName="ConnectionFactory"
jndiDestinationName="dynamicQueues/test.Celtix.jmstransport" >
      <jms:JMSNamingProperty name="java.naming.factory.initial"
value="org.activemq.jndi.ActiveMQInitialContextFactory" />
      <jms:JMSNamingProperty name="java.naming.provider.url"
value="tcp://localhost:61616" />
    </jms:address>
  </port>
</service>
```

## 4.2.2. JMS client configuration

### Overview

JMS consumer endpoints specify the type of messages they use. JMS consumer endpoint can use either a `JMS ByteMessage` or a `JMS TextMessage`.

When using an `ByteMessage` the consumer endpoint uses a `byte[]` as the method for storing data into and retrieving data from the JMS message body. When messages are sent, the message data, including any formatting information, is packaged into a `byte[]` and placed into the message body before it is placed on the wire. When messages are received, the consumer endpoint will attempt to unmarshal the data stored in the message body as if it were packed in a `byte[]`.

When using a `TextMessage`, the consumer endpoint uses a `string` as the method for storing and retrieving data from the message body. When messages are sent, the message information, including any format-specific information, is converted into a `string` and placed into the JMS message body.

When messages are received the consumer endpoint will attempt to unmarshal the data stored in the JMS message body as if it were packed into a string.

When native JMS applications interact with Apache CXF consumers, the JMS application is responsible for interpreting the message and the formatting information. For example, if the Apache CXF contract specifies that the binding used for a JMS endpoint is SOAP, and the messages are packaged as `TextMessage`, the receiving JMS application will get a text message containing all of the SOAP envelope information.

## Specifying the message type

The type of messages accepted by a JMS consumer endpoint is configured using the optional `jms:client` element. The `jms:client` element is a child of the WSDL `port` element and has one attribute:

**Table 4.3. JMS Client WSDL Extensions**

<p><code>messageType</code></p>	<p>Specifies how the message data will be packaged as a JMS message. <code>text</code> specifies that the data will be packaged as a <code>TextMessage</code>. <code>binary</code> specifies that the data will be packaged as an <code>ByteMessage</code>.</p>
---------------------------------	---

## Example

[Example 4.7, “WSDL for a JMS consumer endpoint”](#) shows the WSDL for configuring a JMS consumer endpoint.

### Example 4.7. WSDL for a JMS consumer endpoint

```
<service name="JMSService">
  <port binding="tns:Greeter_SOAPBinding" name="SoapPort">
    <jms:address jndiConnectionFactoryName="ConnectionFactory"
jndiDestinationName="dynamicQueues/test.Celtix.jmstransport" >
      <jms:JMSNamingProperty name="java.naming.factory.initial"
value="org.activemq.jndi.ActiveMQInitialContextFactory" />
      <jms:JMSNamingProperty name="java.naming.provider.url"
value="tcp://localhost:61616" />
    </jms:address>
    <jms:client messageType="binary" />
  </port>
</service>
```

## 4.2.3. JMS provider configuration

### Overview

JMS provider endpoints have a number of behaviors that are configurable. These include:

- how messages are correlated



- the use of durable subscriptions
- if the service uses local JMS transactions
- the message selectors used by the endpoint

## Specifying the configuration

Provider endpoint behaviors are configured using the optional `jms:server` element. The `jms:server` element is a child of the WSDL `wsdl:port` element and has the following attributes:

**Table 4.4. JMS provider endpoint WSDL extensions**

Attribute	Description
<code>useMessageIDAsCorrelationID</code>	Specifies whether JMS will use the message ID to correlate messages. The default is <code>false</code> .
<code>durableSubscriberName</code>	Specifies the name used to register a durable subscription.
<code>messageSelector</code>	Specifies the string value of a message selector to use. For more information on the syntax used to specify message selectors, see the JMS 1.1 specification.
<code>transactional</code>	Specifies whether the local JMS broker will create transactions around message processing. The default is <code>false</code> . <sup>[a]</sup>

[a] Currently, setting the `transactional` attribute to `true` is not supported by the runtime.

## Example

**Example 4.8, “WSDL for a JMS provider endpoint”** shows the WSDL for configuring a JMS provider endpoint.

### Example 4.8. WSDL for a JMS provider endpoint

```
<service name="JMSService">
  <port binding="tns:Greeter_SOAPBinding" name="SoapPort">
    <jms:address jndiConnectionFactoryName="ConnectionFactory"
jndiDestinationName="dynamicQueues/test.Celtix.jmstransport" >
      <jms:JMSNamingProperty name="java.naming.factory.initial"
value="org.activemq.jndi.ActiveMQInitialContextFactory" />
      <jms:JMSNamingProperty name="java.naming.provider.url"
value="tcp://localhost:61616" />
    </jms:address>
    <jms:server messageSelector="cxf_message_selector"
useMessageIDAsCorrelationID="true"
```

```

transactional="true"
durableSubscriberName="cxf_subscriber" />
</port>
</service>

```

## 4.3. USING A NAMED REPLY DESTINATION

### Overview

By default, Apache CXF endpoints using JMS create a temporary queue for sending replies back and forth. If you prefer to use named queues, you can configure the queue used to send replies as part of an endpoint's JMS configuration.

### Setting the reply destination name

You specify the reply destination using either the `jmsReplyDestinationName` attribute or the `jndiReplyDestinationName` attribute in the endpoint's JMS configuration. A client endpoint will listen for replies on the specified destination and it will specify the value of the attribute in the `ReplyTo` field of all outgoing requests. A service endpoint will use the value of the `jndiReplyDestinationName` attribute as the location for placing replies if there is no destination specified in the request's `ReplyTo` field.

### Example

[Example 4.9, “JMS Consumer Specification Using a Named Reply Queue”](#) shows the configuration for a JMS client endpoint.

#### Example 4.9. JMS Consumer Specification Using a Named Reply Queue

```

<jms:conduit name="
{http://cxf.apache.org/jms_endpt}HelloWorldJMSPort.jms-conduit">
  <jms:address destinationStyle="queue"
    jndiConnectionFactoryName="myConnectionFactory"
    jndiDestinationName="myDestination"
    jndiReplyDestinationName="myReplyDestination" >
    <jms:JMSNamingProperty name="java.naming.factory.initial"
value="org.apache.cxf.transport.jms.MyInitialContextFactory" />
    <jms:JMSNamingProperty name="java.naming.provider.url"
value="tcp://localhost:61616" />
  </jms:address>
</jms:conduit>

```

## CHAPTER 5. APACHE CXF LOGGING

### Abstract

This chapter describes how to configure logging in the Apache CXF runtime.

## 5.1. OVERVIEW OF APACHE CXF LOGGING

### Overview

Apache CXF uses the Java logging utility, `java.util.logging`. Logging is configured in a logging configuration file that is written using the standard `java.util.Properties` format. To run logging on an application, you can specify logging programmatically or by defining a property at the command that points to the logging configuration file when you start the application.

### Default logging.properties file

Apache CXF comes with a default `logging.properties` file, which is located in your `InstallDir/etc` directory. This file configures both the output destination for the log messages and the message level that is published. The default configuration sets the loggers to print message flagged with the **WARNING** level to the console. You can either use the default file without changing any of the configuration settings or you can change the configuration settings to suit your specific application.

### Logging feature

Apache CXF includes a logging feature that can be plugged into your client or your service to enable logging. [Example 5.1, “Configuration for Enabling Logging”](#) shows the configuration to enable the logging feature.

#### Example 5.1. Configuration for Enabling Logging

```
<jaxws:endpoint...>
  <jaxws:features>
    <bean class="org.apache.cxf.feature.LoggingFeature"/>
  </jaxws:features>
</jaxws:endpoint>
```

For more information, see [Section 5.6, “Logging Message Content”](#).

### Where to begin?

To run a simple example of logging follow the instructions outlined in a [Section 5.2, “Simple Example of Using Logging”](#).

For more information on how logging works in Apache CXF, read this entire chapter.

### More information on `java.util.logging`

The `java.util.logging` utility is one of the most widely used Java logging frameworks. There is a lot of information available online that describes how to use and extend this framework. As a starting point, however, the following documents gives a good overview of `java.util.logging`:

- <http://java.sun.com/j2se/1.5.0/docs/guide/logging/overview.html>
- <http://java.sun.com/j2se/1.5.0/docs/api/java/util/logging/package-summary.html>

## 5.2. SIMPLE EXAMPLE OF USING LOGGING

### Changing the log levels and output destination

To change the log level and output destination of the log messages in the `wsdl_first` sample application, complete the following steps:

1. Run the sample server as described in the *Running the demo using javasection* of the `README.txt` file in the `InstallDir/samples/wsdl_first` directory. Note that the `server start` command specifies the default `logging.properties` file, as follows:

Platform	Command
Windows	<pre>start java - Djava.util.logging.config.file=% CXF_HOME%\etc\logging.properties demo.hw.server.Server</pre>
UNIX	<pre>java - Djava.util.logging.config.file=\$ CXF_HOME/etc/logging.properties demo.hw.server.Server &amp;</pre>

The default `logging.properties` file is located in the `InstallDir/etc` directory. It configures the Apache CXF loggers to print `WARNING` level log messages to the console. As a result, you see very little printed to the console.

2. Stop the server as described in the `README.txt` file.
3. Make a copy of the default `logging.properties` file, name it `mylogging.properties` file, and save it in the same directory as the default `logging.properties` file.
4. Change the global logging level and the console logging levels in your `mylogging.properties` file to `INFO` by editing the following lines of configuration:

```
.level= INFO
java.util.logging.ConsoleHandler.level = INFO
```

5. Restart the server using the following command:

Platform	Command
Windows	<pre>start java - Djava.util.logging.config.file=% CXF_HOME%\etc\mylogging.properti es demo.hw.server.Server</pre>
UNIX	<pre>java - Djava.util.logging.config.file=\$ CXF_HOME/etc/mylogging.propertie s demo.hw.server.Server &amp;</pre>

Because you configured the global logging and the console logger to log messages of level **INFO**, you see a lot more log messages printed to the console.

### 5.3. DEFAULT LOGGING CONFIGURATION FILE

The default logging configuration file, `logging.properties`, is located in the `InstallDir/etc` directory. It configures the Apache CXF loggers to print **WARNING** level messages to the console. If this level of logging is suitable for your application, you do not have to make any changes to the file before using it. You can, however, change the level of detail in the log messages. For example, you can change whether log messages are sent to the console, to a file or to both. In addition, you can specify logging at the level of individual packages.



#### NOTE

This section discusses the configuration properties that appear in the default `logging.properties` file. There are, however, many other `java.util.logging` configuration properties that you can set. For more information on the `java.util.logging` API, see the `java.util.logging` javadoc at: <http://java.sun.com/j2se/1.5/docs/api/java/util/logging/package-summary.html>.

#### 5.3.1. Configuring Logging Output

The Java logging utility, `java.util.logging`, uses handler classes to output log messages. [Table 5.1, “Java.util.logging Handler Classes”](#) shows the handlers that are configured in the default `logging.properties` file.

Table 5.1. Java.util.logging Handler Classes

Handler Class	Outputs to
<code>ConsoleHandler</code>	Outputs log messages to the console
<code>FileHandler</code>	Outputs log messages to a file



#### IMPORTANT

The handler classes must be on the system classpath in order to be installed by the Java VM when it starts. This is done when you set the Apache CXF environment.

## Configuring the console handler

[Example 5.2, “Configuring the Console Handler”](#) shows the code for configuring the console logger.

### Example 5.2. Configuring the Console Handler

```
handlers= java.util.logging.ConsoleHandler
```

The console handler also supports the configuration properties shown in [Example 5.3, “Console Handler Properties”](#).

### Example 5.3. Console Handler Properties

```
1 java.util.logging.ConsoleHandler.level = WARNING  
  java.util.logging.ConsoleHandler.formatter =  
2 java.util.logging.SimpleFormatter
```

The configuration properties shown in [Example 5.3, “Console Handler Properties”](#) can be explained as follows:

- 1 The console handler supports a separate log level configuration property. This allows you to limit the log messages printed to the console while the global logging setting can be different (see [Section 5.3.2, “Configuring Logging Levels”](#)). The default setting is **WARNING**.
- 2 Specifies the `java.util.logging` formatter class that the console handler class uses to format the log messages. The default setting is the `java.util.logging.SimpleFormatter`.

## Configuring the file handler

[Example 5.4, “Configuring the File Handler”](#) shows code that configures the file handler.

### Example 5.4. Configuring the File Handler

```
handlers= java.util.logging.FileHandler
```

The file handler also supports the configuration properties shown in [Example 5.5, “File Handler Configuration Properties”](#).

### Example 5.5. File Handler Configuration Properties

```
1 java.util.logging.FileHandler.pattern = %h/java%u.log  
2 java.util.logging.FileHandler.limit = 50000  
3 java.util.logging.FileHandler.count = 1  
  java.util.logging.FileHandler.formatter =  
4 java.util.logging.XMLFormatter
```

The configuration properties shown in [Example 5.5, “File Handler Configuration Properties”](#) can be explained as follows:

- 1 Specifies the location and pattern of the output file. The default setting is your home directory.
- 2 Specifies, in bytes, the maximum amount that the logger writes to any one file. The default setting is `50000`. If you set it to zero, there is no limit on the amount that the logger writes to any one file.
- 3 Specifies how many output files to cycle through. The default setting is `1`.
- 4 Specifies the `java.util.logging` formatter class that the file handler class uses to format the log messages. The default setting is the `java.util.logging.XMLFormatter`.

### Configuring both the console handler and the file handler

You can set the logging utility to output log messages to both the console and to a file by specifying the console handler and the file handler, separated by a comma, as shown in [Example 5.6, “Configuring Both Console Logging and File Logging”](#).

#### Example 5.6. Configuring Both Console Logging and File Logging

```
handlers= java.util.logging.FileHandler,  
java.util.logging.ConsoleHandler
```

## 5.3.2. Configuring Logging Levels

### Logging levels

The `java.util.logging` framework supports the following levels of logging, from the least verbose to the most verbose:

- SEVERE
- WARNING
- INFO
- CONFIG
- FINE
- FINER
- FINEST

### Configuring the global logging level

To configure the types of event that are logged across all loggers, configure the global logging level as shown in [Example 5.7, “Configuring Global Logging Levels”](#).

#### Example 5.7. Configuring Global Logging Levels

```
-
```

```
.level= WARNING
```

### Configuring logging at an individual package level

The `java.util.logging` framework supports configuring logging at the level of an individual package. For example, the line of code shown in [Example 5.8, “Configuring Logging at the Package Level”](#) configures logging at a `SEVERE` level on classes in the `com.xyz.foo` package.

#### Example 5.8. Configuring Logging at the Package Level

```
com.xyz.foo.level = SEVERE
```

## 5.4. ENABLING LOGGING AT THE COMMAND LINE

### Overview

You can run the logging utility on an application by defining a `java.util.logging.config.file` property when you start the application. You can either specify the default `logging.properties` file or a `logging.properties` file that is unique to that application.

### Specifying the log configuration file on application start-up

To specify logging on application start-up add the flag shown in [Example 5.9, “Flag to Start Logging on the Command Line”](#) when starting the application.

#### Example 5.9. Flag to Start Logging on the Command Line

```
-Djava.util.logging.config.file=myfile
```

## 5.5. LOGGING FOR SUBSYSTEMS AND SERVICES

You can use the `com.xyz.foo.level` configuration property described in [the section called “Configuring logging at an individual package level”](#) to set fine-grained logging for specified Apache CXF logging subsystems.

### Apache CXF logging subsystems

[Table 5.2, “Apache CXF Logging Subsystems”](#) shows a list of available Apache CXF logging subsystems.

Table 5.2. Apache CXF Logging Subsystems

Subsystem	Description
<code>org.apache.cxf.aegis</code>	Aegis binding



Subsystem	Description
<code>org.apache.cxf.binding.coloc</code>	colocated binding
<code>org.apache.cxf.binding.http</code>	HTTP binding
<code>org.apache.cxf.binding.jbi</code>	JBI binding
<code>org.apache.cxf.binding.object</code>	Java Object binding
<code>org.apache.cxf.binding.soap</code>	SOAP binding
<code>org.apache.cxf.binding.xml</code>	XML binding
<code>org.apache.cxf.bus</code>	Apache CXF bus
<code>org.apache.cxf.configuration</code>	configuration framework
<code>org.apache.cxf.endpoint</code>	server and client endpoints
<code>org.apache.cxf.interceptor</code>	interceptors
<code>org.apache.cxf.jaxws</code>	Front-end for JAX-WS style message exchange, JAX-WS handler processing, and interceptors relating to JAX-WS and configuration
<code>org.apache.cxf.jbi</code>	JBI container integration classes
<code>org.apache.cxf.jca</code>	JCA container integration classes
<code>org.apache.cxf.js</code>	JavaScript front-end
<code>org.apache.cxf.transport.http</code>	HTTP transport
<code>org.apache.cxf.transport.https</code>	secure version of HTTP transport, using HTTPS
<code>org.apache.cxf.transport.jbi</code>	JBI transport
<code>org.apache.cxf.transport.jms</code>	JMS transport
<code>org.apache.cxf.transport.local</code>	transport implementation using local file system
<code>org.apache.cxf.transport.servlet</code>	HTTP transport and servlet implementation for loading JAX-WS endpoints into a servlet container
<code>org.apache.cxf.ws.addressing</code>	WS-Addressing implementation

Subsystem	Description
<code>org.apache.cxf.ws.policy</code>	WS-Policy implementation
<code>org.apache.cxf.ws.rm</code>	WS-ReliableMessaging (WS-RM) implementation
<code>org.apache.cxf.ws.security.wss4j</code>	WSS4J security implementation

## Example

The WS-Addressing sample is contained in the *InstallDir/samples/ws\_addressing* directory. Logging is configured in the `logging.properties` file located in that directory. The relevant lines of configuration are shown in [Example 5.10, “Configuring Logging for WS-Addressing”](#).

### Example 5.10. Configuring Logging for WS-Addressing

```
java.util.logging.ConsoleHandler.formatter =
demos.ws_addressing.common.ConciseFormatter
...
org.apache.cxf.ws.addressing.soap.MAPCodec.level = INFO
```

The configuration in [Example 5.10, “Configuring Logging for WS-Addressing”](#) enables the snooping of log messages relating to WS-Addressing headers, and displays them to the console in a concise form.

For information on running this sample, see the `README.txt` file located in the *InstallDir/samples/ws\_addressing* directory.

## 5.6. LOGGING MESSAGE CONTENT

You can log the content of the messages that are sent between a service and a consumer. For example, you might want to log the contents of SOAP messages that are being sent between a service and a consumer.

### Configuring message content logging

To log the messages that are sent between a service and a consumer, and vice versa, complete the following steps:

1. [Add the logging feature to your endpoint's configuration.](#)
2. [Add the logging feature to your consumer's configuration.](#)
3. [Configure the logging system log `INFO` level messages.](#)

### Adding the logging feature to an endpoint

Add the logging feature your endpoint's configuration as shown in [Example 5.11, “Adding Logging to Endpoint Configuration”](#).

### Example 5.11. Adding Logging to Endpoint Configuration

```

<jaxws:endpoint ...>
  <jaxws:features>
    <bean class="org.apache.cxf.feature.LoggingFeature"/>
  </jaxws:features>
</jaxws:endpoint>

```

The example XML shown in [Example 5.11, “Adding Logging to Endpoint Configuration”](#) enables the logging of SOAP messages.

## Adding the logging feature to a consumer

Add the logging feature your client's configuration as shown in [Example 5.12, “Adding Logging to Client Configuration”](#).

### Example 5.12. Adding Logging to Client Configuration

```

<jaxws:client ...>
  <jaxws:features>
    <bean class="org.apache.cxf.feature.LoggingFeature"/>
  </jaxws:features>
</jaxws:client>

```

The example XML shown in [Example 5.12, “Adding Logging to Client Configuration”](#) enables the logging of SOAP messages.

## Set logging to log INFO level messages

Ensure that the `logging.properties` file associated with your service is configured to log `INFO` level messages, as shown in [Example 5.13, “Setting the Logging Level to INFO”](#).

### Example 5.13. Setting the Logging Level to INFO

```

.level= INFO
java.util.logging.ConsoleHandler.level = INFO

```

## Logging SOAP messages

To see the logging of SOAP messages modify the `wsdl_first` sample application located in the `InstallDir/samples/wsdl_first` directory, as follows:

1. Add the `jaxws:features` element shown in [Example 5.14, “Endpoint Configuration for Logging SOAP Messages”](#) to the `cxf.xml` configuration file located in the `wsdl_first` sample's directory:

### Example 5.14. Endpoint Configuration for Logging SOAP Messages

```

<jaxws:endpoint name="
{http://apache.org/hello_world_soap_http}SoapPort"
  createdFromAPI="true">

```

```

<jaxws:properties>
  <entry key="schema-validation-enabled" value="true" />
</jaxws:properties>
<jaxws:features>
  <bean class="org.apache.cxf.feature.LoggingFeature"/>
</jaxws:features>
</jaxws:endpoint>

```

- The sample uses the default `logging.properties` file, which is located in the `InstallDir/etc` directory. Make a copy of this file and name it `mylogging.properties`.
- In the `mylogging.properties` file, change the logging levels to `INFO` by editing the `.level` and the `java.util.logging.ConsoleHandler.level` configuration properties as follows:

```

.level= INFO
java.util.logging.ConsoleHandler.level = INFO

```

- Start the server using the new configuration settings in both the `cx.xml` file and the `mylogging.properties` file as follows:

Platform	Command
Windows	<pre>start java - Djava.util.logging.config.file=% CXF_HOME%\etc\mylogging.properti es demo.hw.server.Server</pre>
UNIX	<pre>java - Djava.util.logging.config.file=\$ CXF_HOME/etc/mylogging.propertie s demo.hw.server.Server &amp;</pre>

- Start the hello world client using the following command:

Platform	Command
Windows	<pre>java - Djava.util.logging.config.file=% CXF_HOME%\etc\mylogging.properti es demo.hw.client.Client .\wsdl\hello_world.wsdl</pre>
UNIX	<pre>java - Djava.util.logging.config.file=\$ CXF_HOME/etc/mylogging.propertie s demo.hw.client.Client ./wsdl/hello_world.wsdl</pre>

The SOAP messages are logged to the console.

## CHAPTER 6. DEPLOYING WS-ADDRESSING

### Abstract

Apache CXF supports WS-Addressing for JAX-WS applications. This chapter explains how to deploy WS-Addressing in the Apache CXF runtime environment.

## 6.1. INTRODUCTION TO WS-ADDRESSING

### Overview

WS-Addressing is a specification that allows services to communicate addressing information in a transport neutral way. It consists of two parts:

- A structure for communicating a reference to a Web service endpoint
- A set of Message Addressing Properties (MAP) that associate addressing information with a particular message

### Supported specifications

Apache CXF supports both the WS-Addressing 2004/08 specification and the WS-Addressing 2005/03 specification.

### Further information

For detailed information on WS-Addressing, see the 2004/08 submission at <http://www.w3.org/Submission/ws-addressing/>.

## 6.2. WS-ADDRESSING INTERCEPTORS

### Overview

In Apache CXF, WS-Addressing functionality is implemented as interceptors. The Apache CXF runtime uses interceptors to intercept and work with the raw messages that are being sent and received. When a transport receives a message, it creates a message object and sends that message through an interceptor chain. If the WS-Addressing interceptors are added to the application's interceptor chain, any WS-Addressing information included with a message is processed.

### WS-Addressing Interceptors

The WS-Addressing implementation consists of two interceptors, as described in [Table 6.1, “WS-Addressing Interceptors”](#).

**Table 6.1. WS-Addressing Interceptors**

Interceptor	Description

Interceptor	Description
<code>org.apache.cxf.ws.addressing.MAPAggregator</code>	A logical interceptor responsible for aggregating the Message Addressing Properties (MAPs) for outgoing messages.
<code>org.apache.cxf.ws.addressing.soap.MAPCodec</code>	A protocol-specific interceptor responsible for encoding and decoding the Message Addressing Properties (MAPs) as SOAP headers.

## 6.3. ENABLING WS-ADDRESSING

### Overview

To enable WS-Addressing the WS-Addressing interceptors must be added to the inbound and outbound interceptor chains. This is done in one of the following ways:

- [Apache CXF Features](#)
- [RMAssertion and WS-Policy Framework](#)
- [Using Policy Assertion in a WS-Addressing Feature](#)

### Adding WS-Addressing as a Feature

WS-Addressing can be enabled by adding the WS-Addressing feature to the client and the server configuration as shown in [Example 6.1, “client.xml—Adding WS-Addressing Feature to Client Configuration”](#) and [Example 6.2, “server.xml—Adding WS-Addressing Feature to Server Configuration”](#) respectively.

#### Example 6.1. client.xml—Adding WS-Addressing Feature to Client Configuration

```
<?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:jaxws="http://cxf.apache.org/jaxws"
  xmlns:wsa="http://cxf.apache.org/ws/addressing"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd">

  <jaxws:client ...>
    <jaxws:features>
      <wsa:addressing/>
    </jaxws:features>
  </jaxws:client>
</beans>
```

#### Example 6.2. server.xml—Adding WS-Addressing Feature to Server Configuration

```

<?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:jaxws="http://cxf.apache.org/jaxws"
       xmlns:wsa="http://cxf.apache.org/ws/addressing"
       xsi:schemaLocation="
http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd">

    <jaxws:endpoint ...>
        <jaxws:features>
            <wsa:addressing/>
        </jaxws:features>
    </jaxws:endpoint>
</beans>

```

## 6.4. CONFIGURING WS-ADDRESSING ATTRIBUTES

### Overview

The Apache CXF WS-Addressing feature element is defined in the namespace `http://cxf.apache.org/ws/addressing`. It supports the two attributes described in [Table 6.2](#), “WS-Addressing Attributes”.

**Table 6.2. WS-Addressing Attributes**

Attribute Name	Value
<code>allowDuplicates</code>	A boolean that determines if duplicate MessageIDs are tolerated. The default setting is <code>true</code> .
<code>usingAddressingAdvisory</code>	A boolean that indicates if the presence of the <code>UsingAddressing</code> element in the WSDL is advisory only; that is, its absence does not prevent the encoding of WS-Addressing headers.

### Configuring WS-Addressing attributes

Configure WS-Addressing attributes by adding the attribute and the value you want to set it to the WS-Addressing feature in your server or client configuration file. For example, the following configuration extract sets the `allowDublicates` attribute to `false` on the server endpoint:

```

<beans ... xmlns:wsa="http://cxf.apache.org/ws/addressing" ...>
    <jaxws:endpoint ...>
        <jaxws:features>
            <wsa:addressing allowDublicates="false"/>
        </jaxws:features>
    </jaxws:endpoint>
</beans>

```

## Using a WS-Policy assertion embedded in a feature

In [Example 6.3, “Using the Policies to Configure WS-Addressing”](#) an addressing policy assertion to enable non-anonymous responses is embedded in the `policies` element.

### Example 6.3. Using the Policies to Configure WS-Addressing

```
<?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:wsa="http://cxf.apache.org/ws/addressing"
  xmlns:wsp="http://www.w3.org/2006/07/ws-policy"
  xmlns:policy="http://cxf.apache.org/policy-config"
  xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-
wss-wssecurity-utility-1.0.xsd"
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  xsi:schemaLocation="
http://www.w3.org/2006/07/ws-policy http://www.w3.org/2006/07/ws-
policy.xsd
http://cxf.apache.org/ws/addressing
http://cxf.apache.org/schema/ws/addressing.xsd
http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd
http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd">

  <jaxws:endpoint name="
{http://cxf.apache.org/greeter_control}GreeterPort"
    createdFromAPI="true">
    <jaxws:features>
      <policy:policies>
        <wsp:Policy
xmlns:wsam="http://www.w3.org/2007/02/addressing/metadata">
          <wsam:Addressing>
            <wsp:Policy>
              <wsam:NonAnonymousResponses/>
            </wsp:Policy>
          </wsam:Addressing>
        </wsp:Policy>
      </policy:policies>
    </jaxws:features>
  </jaxws:endpoint>
</beans>
```



## CHAPTER 7. ENABLING RELIABLE MESSAGING

### Abstract

Apache CXF supports WS-Reliable Messaging(WS-RM). This chapter explains how to enable and configure WS-RM in Apache CXF.

## 7.1. INTRODUCTION TO WS-RM

### Overview


WS-ReliableMessaging (WS-RM) is a protocol that ensures the reliable delivery of messages in a distributed environment. It enables messages to be delivered reliably between distributed applications in the presence of software, system, or network failures.

For example, WS-RM can be used to ensure that the correct messages have been delivered across a network exactly once, and in the correct order.

### How WS-RM works

WS-RM ensures the reliable delivery of messages between a source and a destination endpoint. The source is the initial sender of the message and the destination is the ultimate receiver, as shown in [Figure 7.1, “Web Services Reliable Messaging”](#).

Figure 7.1. Web Services Reliable Messaging



The flow of WS-RM messages can be described as follows:

1. The RM source sends a `CreateSequence` protocol message to the RM destination. This contains a reference for the endpoint that receives acknowledgements (the `wsm:AcksTo` endpoint).
2. The RM destination sends a `CreateSequenceResponse` protocol message back to the RM source. This message contains the sequence ID for the RM sequence session.

3. The RM source adds an RM **Sequence** header to each message sent by the application source. This header contains the sequence ID and a unique message ID.
4. The RM source transmits each message to the RM destination.
5. The RM destination acknowledges the receipt of the message from the RM source by sending messages that contain the RM **SequenceAcknowledgement** header.
6. The RM destination delivers the message to the application destination in an exactly-once-in-order fashion.
7. The RM source retransmits a message that it has not yet received an acknowledgement.

The first retransmission attempt is made after a base retransmission interval. Successive retransmission attempts are made, by default, at exponential back-off intervals or, alternatively, at fixed intervals. For more details, see [Section 7.4, “Configuring WS-RM”](#).

This entire process occurs symmetrically for both the request and the response message; that is, in the case of the response message, the server acts as the RM source and the client acts as the RM destination.

## WS-RM delivery assurances

WS-RM guarantees reliable message delivery in a distributed environment, regardless of the transport protocol used. Either the source or the destination endpoint logs an error if reliable delivery can not be assured.

## Supported specifications

Apache CXF supports the 2005/02 version of the WS-RM specification, which is based on the WS-Addressing 2004/08 specification.

## Further information

For detailed information on WS-RM, see the specification at <http://specs.xmlsoap.org/ws/2005/02/rm/ws-reliablemessaging.pdf>.

## 7.2. WS-RM INTERCEPTORS

### Overview

In Apache CXF, WS-RM functionality is implemented as interceptors. The Apache CXF runtime uses interceptors to intercept and work with the raw messages that are being sent and received. When a transport receives a message, it creates a message object and sends that message through an interceptor chain. If the application's interceptor chain includes the WS-RM interceptors, the application can participate in reliable messaging sessions. The WS-RM interceptors handle the collection and aggregation of the message chunks. They also handle all of the acknowledgement and retransmission logic.

### Apache CXF WS-RM Interceptors

The Apache CXF WS-RM implementation consists of four interceptors, which are described in [Table 7.1, “Apache CXF WS-ReliableMessaging Interceptors”](#).

Table 7.1. Apache CXF WS-ReliableMessaging Interceptors

Interceptor	Description
<code>org.apache.cxf.ws.rm.RMOutInterceptor</code>	<p>Deals with the logical aspects of providing reliability guarantees for outgoing messages.</p> <p>Responsible for sending the <b>CreateSequence</b> requests and waiting for their <b>CreateSequenceResponse</b> responses.</p> <p>Also responsible for aggregating the sequence properties—ID and message number—for an application message.</p>
<code>org.apache.cxf.ws.rm.RMInInterceptor</code>	<p>Responsible for intercepting and processing RM protocol messages and <b>SequenceAcknowledgement</b> messages that are piggybacked on application messages.</p>
<code>org.apache.cxf.ws.rm.soap.RMSoapInterceptor</code>	<p>Responsible for encoding and decoding the reliability properties as SOAP headers.</p>
<code>org.apache.cxf.ws.rm.RetransmissionInterceptor</code>	<p>Responsible for creating copies of application messages for future resending.</p>

## Enabling WS-RM

The presence of the WS-RM interceptors on the interceptor chains ensures that WS-RM protocol messages are exchanged when necessary. For example, when intercepting the first application message on the outbound interceptor chain, the **RMOutInterceptor** sends a **CreateSequence** request and waits to process the original application message until it receives the **CreateSequenceResponse** response. In addition, the WS-RM interceptors add the sequence headers to the application messages and, on the destination side, extract them from the messages. It is not necessary to make any changes to your application code to make the exchange of messages reliable.

For more information on how to enable WS-RM, see [Section 7.3, “Enabling WS-RM”](#).

## Configuring WS-RM Attributes

You control sequence demarcation and other aspects of the reliable exchange through configuration. For example, by default Apache CXF attempts to maximize the lifetime of a sequence, thus reducing the overhead incurred by the out-of-band WS-RM protocol messages. To enforce the use of a separate sequence per application message configure the WS-RM source’s sequence termination policy (setting the maximum sequence length to 1).

For more information on configuring WS-RM behavior, see [Section 7.4, “Configuring WS-RM”](#).

## 7.3. ENABLING WS-RM

### Overview

To enable reliable messaging, the WS-RM interceptors must be added to the interceptor chains for both inbound and outbound messages and faults. Because the WS-RM interceptors use WS-Addressing, the WS-Addressing interceptors must also be present on the interceptor chains.

You can ensure the presence of these interceptors in one of two ways:

- **Explicitly**, by adding them to the dispatch chains using Spring beans
- **Implicitly**, using WS-Policy assertions, which cause the Apache CXF runtime to transparently add the interceptors on your behalf.

## Spring beans—explicitly adding interceptors

To enable WS-RM add the WS-RM and WS-Addressing interceptors to the Apache CXF bus, or to a consumer or service endpoint using Spring bean configuration. This is the approach taken in the WS-RM sample that is found in the *InstallDir/samples/ws\_rm* directory. The configuration file, *ws\_rm.cxf*, shows the WS-RM and WS-Addressing interceptors being added one-by-one as Spring beans (see [Example 7.1, “Enabling WS-RM Using Spring Beans”](#)).

### Example 7.1. Enabling WS-RM Using Spring Beans

```
<?xml version="1.0" encoding="UTF-8"?>
1 <beans xmlns="http://www.springframework.org/schema/beans"
      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:schemaLocation="http://www.springframework.org/schema/
      beans http://www.springframework.org/schema/beans/spring-beans.xsd">
2   <bean id="mapAggregator"
      class="org.apache.cxf.ws.addressing.MAPAggregator"/>
      <bean id="mapCodec"
      class="org.apache.cxf.ws.addressing.soap.MAPCodec"/>
3   <bean id="rmLogicalOut"
      class="org.apache.cxf.ws.rm.RMOutInterceptor">
      <property name="bus" ref="cxf"/>
    </bean>
      <bean id="rmLogicalIn" class="org.apache.cxf.ws.rm.RMInInterceptor">
      <property name="bus" ref="cxf"/>
    </bean>
      <bean id="rmCodec"
      class="org.apache.cxf.ws.rm.soap.RMSoapInterceptor"/>
      <bean id="cxf" class="org.apache.cxf.bus.CXFBusImpl">
4       <property name="inInterceptors">
          <list>
            <ref bean="mapAggregator"/>
            <ref bean="mapCodec"/>
            <ref bean="rmLogicalIn"/>
            <ref bean="rmCodec"/>
          </list>
        </property>
5       <property name="inFaultInterceptors">
          <list>
            <ref bean="mapAggregator"/>
            <ref bean="mapCodec"/>
            <ref bean="rmLogicalIn"/>
            <ref bean="rmCodec"/>
          </list>
        </property>
    </bean>
  </beans>
```

```

6      <property name="outInterceptors">
          <list>
            <ref bean="mapAggregator"/>
            <ref bean="mapCodec"/>
            <ref bean="rmLogicalOut"/>
            <ref bean="rmCodec"/>
          </list>
        </property>
7      <property name="outFaultInterceptors">
          <list>
            <ref bean="mapAggregator">
            <ref bean="mapCodec"/>
            <ref bean="rmLogicalOut"/>
            <ref bean="rmCodec"/>
          </list>
        </property>
      </bean>
    </beans>

```

The code shown in [Example 7.1, “Enabling WS-RM Using Spring Beans”](#) can be explained as follows:

- 1 A Apache CXF configuration file is a Spring XML file. You must include an opening Spring **beans** element that declares the namespaces and schema files for the child elements that are encapsulated by the **beans** element.
- 2 Configures each of the WS-Addressing interceptors—**MAPAggregator** and **MAPCodec**. For more information on WS-Addressing, see [Chapter 6, Deploying WS-Addressing](#).
- 3 Configures each of the WS-RM interceptors—**RMOutInterceptor**, **RMInInterceptor**, and **RMSoapInterceptor**.
- 4 Adds the WS-Addressing and WS-RM interceptors to the interceptor chain for inbound messages.
- 5 Adds the WS-Addressing and WS-RM interceptors to the interceptor chain for inbound faults.
- 6 Adds the WS-Addressing and WS-RM interceptors to the interceptor chain for outbound messages.
- 7 Adds the WS-Addressing and WS-RM interceptors to the interceptor chain for outbound faults.

## WS-Policy framework—implicitly adding interceptors

The WS-Policy framework provides the infrastructure and APIs that allow you to use WS-Policy. It is compliant with the November 2006 draft publications of the [Web Services Policy 1.5—Framework](#) and [Web Services Policy 1.5—Attachment](#) specifications.

To enable WS-RM using the Apache CXF WS-Policy framework, do the following:

1. Add the policy feature to your client and server endpoint. [Example 7.2, “Configuring WS-RM using WS-Policy”](#) shows a reference bean nested within a `jaxws:feature` element. The reference bean specifies the **AddressingPolicy**, which is defined as a separate element within the same configuration file.

**Example 7.2. Configuring WS-RM using WS-Policy**

```

<jaxws:client>
  <jaxws:features>
    <ref bean="AddressingPolicy"/>
  </jaxws:features>
</jaxws:client>
<wsp:Policy wsu:Id="AddressingPolicy"
xmlns:wsam="http://www.w3.org/2007/02/addressing/metadata">
  <wsam:Addressing>
    <wsp:Policy>
      <wsam:NonAnonymousResponses/>
    </wsp:Policy>
  </wsam:Addressing>
</wsp:Policy>

```

2. Add a reliable messaging policy to the `wsdl:service` element—or any other WSDL element that can be used as an attachment point for policy or policy reference elements—to your WSDL file, as shown in [Example 7.3, “Adding an RM Policy to Your WSDL File”](#).

**Example 7.3. Adding an RM Policy to Your WSDL File**

```

<wsp:Policy wsu:Id="RM"
xmlns:wsp="http://www.w3.org/2006/07/ws-policy"
xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-
wss-wssecurity-utility-1.0.xsd">
  <wsam:Addressing
xmlns:wsam="http://www.w3.org/2007/02/addressing/metadata">
    <wsp:Policy/>
  </wsam:Addressing>
  <wsrmp:RMAssertion
xmlns:wsrmp="http://schemas.xmlsoap.org/ws/2005/02/rm/policy">
    <wsrmp:BaseRetransmissionInterval Milliseconds="10000"/>
  </wsrmp:RMAssertion>
</wsp:Policy>
...
<wsdl:service name="ReliableGreeterService">
  <wsdl:port binding="tns:GreeterSOAPBinding"
name="GreeterPort">
    <soap:address
location="http://localhost:9020/SoapContext/GreeterPort"/>
    <wsp:PolicyReference URI="#RM"
xmlns:wsp="http://www.w3.org/2006/07/ws-policy"/>
  </wsdl:port>
</wsdl:service>

```

## 7.4. CONFIGURING WS-RM

You can configure WS-RM by:

- Setting Apache CXF-specific attributes that are defined in the Apache CXF WS-RM manager namespace, `http://cxf.apache.org/ws/rm/manager`.

- Setting standard WS-RM policy attributes that are defined in the `http://schemas.xmlsoap.org/ws/2005/02/rm/policy` namespace.

### 7.4.1. Configuring Apache CXF-Specific WS-RM Attributes

#### Overview

To configure the Apache CXF-specific attributes, use the `rmManager` Spring bean. Add the following to your configuration file:

- The `http://cxf.apache.org/ws/rm/manager` namespace to your list of namespaces.
- An `rmManager` Spring bean for the specific attribute that you want to configure.

[Example 7.4, “Configuring Apache CXF-Specific WS-RM Attributes”](#) shows a simple example.

#### Example 7.4. Configuring Apache CXF-Specific WS-RM Attributes

```
<beans xmlns="http://www.springframework.org/schema/beans"
      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xmlns:wsm-mgr="http://cxf.apache.org/ws/rm/manager"
      xsi:schemaLocation="http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
http://cxf.apache.org/ws/rm/manager
http://cxf.apache.org/schemas/configuration/wsmr-manager.xsd">
  ...
  <wsm-mgr:rmManager>
  <!--
    ...Your configuration goes here
  -->
  </wsm-mgr:rmManager>
```

#### Children of the `rmManager` Spring bean

[Table 7.2, “Children of the `rmManager` Spring Bean”](#) shows the child elements of the `rmManager` Spring bean, defined in the `http://cxf.apache.org/ws/rm/manager` namespace.

Table 7.2. Children of the `rmManager` Spring Bean

Element	Description
<code>RMAssertion</code>	An element of type <code>RMAssertion</code>
<code>deliveryAssurance</code>	An element of type <code>DeliveryAssuranceType</code> that describes the delivery assurance that should apply
<code>sourcePolicy</code>	An element of type <code>SourcePolicyType</code> that allows you to configure details of the RM source

Element	Description
<b>destinationPolicy</b>	An element of type <code>DestinationPolicyType</code> that allows you to configure details of the RM destination

### Example

For an example, see [the section called “Maximum unacknowledged messages threshold”](#).

## 7.4.2. Configuring Standard WS-RM Policy Attributes

### Overview

You can configure standard WS-RM policy attributes in one of the following ways:

- [RMAssertion in rmManager Spring bean](#)
- [Policy within a feature](#)
- [WSDL file](#)
- [External attachment](#)

### WS-Policy RMAssertion Children

[Table 7.3, “Children of the WS-Policy RMAssertion Element”](#) shows the elements defined in the `http://schemas.xmlsoap.org/ws/2005/02/rm/policy` namespace:

**Table 7.3. Children of the WS-Policy RMAssertion Element**

Name	Description
<b>InactivityTimeout</b>	Specifies the amount of time that must pass without receiving a message before an endpoint can consider an RM sequence to have been terminated due to inactivity.
<b>BaseRetransmissionInterval</b>	Sets the interval within which an acknowledgement must be received by the RM Source for a given message. If an acknowledgement is not received within the time set by the <b>BaseRetransmissionInterval</b> , the RM Source will retransmit the message.
<b>ExponentialBackoff</b>	Indicates the retransmission interval will be adjusted using the commonly known exponential backoff algorithm (Tanenbaum).  For more information, see <i>Computer Networks</i> , Andrew S. Tanenbaum, Prentice Hall PTR, 2003.



Name	Description
<b>AcknowledgementInterval</b>	In WS-RM, acknowledgements are sent on return messages or sent stand-alone. If a return message is not available to send an acknowledgement, an RM Destination can wait for up to the acknowledgement interval before sending a stand-alone acknowledgement. If there are no unacknowledged messages, the RM Destination can choose not to send an acknowledgement.

### More detailed reference information

For more detailed reference information, including descriptions of each element's sub-elements and attributes, please refer to <http://schemas.xmlsoap.org/ws/2005/02/rm/wsrp-policy.xsd>.

### RMAssertion in rmManager Spring bean

You can configure standard WS-RM policy attributes by adding an `RMAssertion` within a Apache CXF `rmManager` Spring bean. This is the best approach if you want to keep all of your WS-RM configuration in the same configuration file; that is, if you want to configure Apache CXF-specific attributes and standard WS-RM policy attributes in the same file.

For example, the configuration in [Example 7.5, “Configuring WS-RM Attributes Using an RMAssertion in an rmManager Spring Bean”](#) shows:

- A standard WS-RM policy attribute, `BaseRetransmissionInterval`, configured using an `RMAssertion` within an `rmManager` Spring bean.
- An Apache CXF-specific RM attribute, `intraMessageThreshold`, configured in the same configuration file.

#### Example 7.5. Configuring WS-RM Attributes Using an RMAssertion in an rmManager Spring Bean

```
<beans xmlns:wsrm-
policy="http://schemas.xmlsoap.org/ws/2005/02/rm/policy"
xmlns:wsrm-mgr="http://cxf.apache.org/ws/rm/manager"
...>
<wsrm-mgr:rmManager id="org.apache.cxf.ws.rm.RMManager">
  <wsrm-policy:RMAssertion>
    <wsrm-policy:BaseRetransmissionInterval Milliseconds="4000" />
  </wsrm-policy:RMAssertion>
  <wsrm-mgr:destinationPolicy>
    <wsrm-mgr:acksPolicy intraMessageThreshold="0" />
  </wsrm-mgr:destinationPolicy>
</wsrm-mgr:rmManager>
</beans>
```

### Policy within a feature

You can configure standard WS-RM policy attributes within features, as shown in [Example 7.6](#), “Configuring WS-RM Attributes as a Policy within a Feature”.

### Example 7.6. Configuring WS-RM Attributes as a Policy within a Feature

```
<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:wsa="http://cxf.apache.org/ws/addressing"
  xmlns:wsp="http://www.w3.org/2006/07/ws-policy"
  xmlns:wsu="http://docs.oasis-open.org/wss/2004/01/oasis-200401-
wss-wssecurity-utility-1.0.xsd"
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  xsi:schemaLocation="
http://www.w3.org/2006/07/ws-policy http://www.w3.org/2006/07/ws-
policy.xsd
http://cxf.apache.org/ws/addressing
http://cxf.apache.org/schema/ws/addressing.xsd
http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd
http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd">
  <jaxws:endpoint name="
{http://cxf.apache.org/greeter_control}GreeterPort"
createdFromAPI="true">
    <jaxws:features>
      <wsp:Policy>
        <wsrm:RMAssertion
xmlns:wsrm="http://schemas.xmlsoap.org/ws/2005/02/rm/policy">
          <wsrm:AcknowledgementInterval Milliseconds="200"
/>
        </wsrm:RMAssertion>
        <wsam:Addressing
xmlns:wsam="http://www.w3.org/2007/02/addressing/metadata">
          <wsp:Policy>
            <wsam:NonAnonymousResponses/>
          </wsp:Policy>
        </wsam:Addressing>
      </wsp:Policy>
    </jaxws:features>
  </jaxws:endpoint>
</beans>
```

### WSDL file

If you use the WS-Policy framework to enable WS-RM, you can configure standard WS-RM policy attributes in a WSDL file. This is a good approach if you want your service to interoperate and use WS-RM seamlessly with consumers deployed to other policy-aware Web services stacks.

For an example, see [the section called “WS-Policy framework—implicitly adding interceptors”](#) where the base retransmission interval is configured in the WSDL file.

### External attachment

You can configure standard WS-RM policy attributes in an external attachment file. This is a good approach if you cannot, or do not want to, change your WSDL file.

[Example 7.7, “Configuring WS-RM in an External Attachment”](#) shows an external attachment that enables both WS-A and WS-RM (base retransmission interval of 30 seconds) for a specific EPR.

#### Example 7.7. Configuring WS-RM in an External Attachment

```
<attachments xmlns:wsp="http://www.w3.org/2006/07/ws-policy"
xmlns:wsa="http://www.w3.org/2005/08/addressing">
  <wsp:PolicyAttachment>
    <wsp:AppliesTo>
      <wsa:EndpointReference>

<wsa:Address>http://localhost:9020/SoapContext/GreeterPort</wsa:Address>
      </wsa:EndpointReference>
    </wsp:AppliesTo>
    <wsp:Policy>
      <wsam:Addressing
xmlns:wsam="http://www.w3.org/2007/02/addressing/metadata">
        <wsp:Policy/>
      </wsam:Addressing>
      <wsrmp:RMAssertion
xmlns:wsrmp="http://schemas.xmlsoap.org/ws/2005/02/rm/policy">
        <wsrmp:BaseRetransmissionInterval
Milliseconds="30000"/>
      </wsrmp:RMAssertion>
    </wsp:Policy>
  </wsp:PolicyAttachment>
</attachments>/
```

### 7.4.3. WS-RM Configuration Use Cases

#### Overview

This subsection focuses on configuring WS-RM attributes from a use case point of view. Where an attribute is a standard WS-RM policy attribute, defined in the <http://schemas.xmlsoap.org/ws/2005/02/rm/policy> namespace, only the example of setting it in an `RMAssertion` within an `rmManager` Spring bean is shown. For details of how to set such attributes as a policy within a feature; in a WSDL file, or in an external attachment, see [Section 7.4.2, “Configuring Standard WS-RM Policy Attributes”](#).

The following use cases are covered:

- [Base retransmission interval](#)
- [Exponential backoff for retransmission](#)
- [Acknowledgement interval](#)
- [Maximum unacknowledged messages threshold](#)
- [Maximum length of an RM sequence](#)

- [Message delivery assurance policies](#)

## Base retransmission interval

The `BaseRetransmissionInterval` element specifies the interval at which an RM source retransmits a message that has not yet been acknowledged. It is defined in the <http://schemas.xmlsoap.org/ws/2005/02/rm/wsrmpolicy.xsd> schema file. The default value is 3000 milliseconds.

[Example 7.8, “Setting the WS-RM Base Retransmission Interval”](#) shows how to set the WS-RM base retransmission interval.

### Example 7.8. Setting the WS-RM Base Retransmission Interval

```
<beans xmlns:wsrm-
policy="http://schemas.xmlsoap.org/ws/2005/02/rm/policy
...>
<wsrm-mgr:rmManager id="org.apache.cxf.ws.rm.RMManager">
  <wsrm-policy:RMAssertion>
    <wsrm-policy:BaseRetransmissionInterval Milliseconds="4000"/>
  </wsrm-policy:RMAssertion>
</wsrm-mgr:rmManager>
</beans>
```

## Exponential backoff for retransmission

The `ExponentialBackoff` element determines if successive retransmission attempts for an unacknowledged message are performed at exponential intervals.

The presence of the `ExponentialBackoff` element enables this feature. An exponential backoff ratio of 2 is used by default.

[Example 7.9, “Setting the WS-RM Exponential Backoff Property”](#) shows how to set the WS-RM exponential backoff for retransmission.

### Example 7.9. Setting the WS-RM Exponential Backoff Property

```
<beans xmlns:wsrm-
policy="http://schemas.xmlsoap.org/ws/2005/02/rm/policy
...>
<wsrm-mgr:rmManager id="org.apache.cxf.ws.rm.RMManager">
  <wsrm-policy:RMAssertion>
    <wsrm-policy:ExponentialBackoff="4"/>
  </wsrm-policy:RMAssertion>
</wsrm-mgr:rmManager>
</beans>
```

## Acknowledgement interval

The `AcknowledgementInterval` element specifies the interval at which the WS-RM destination sends asynchronous acknowledgements. These are in addition to the synchronous acknowledgements

that it sends on receipt of an incoming message. The default asynchronous acknowledgement interval is 0 milliseconds. This means that if the `AcknowledgementInterval` is not configured to a specific value, acknowledgements are sent immediately (that is, at the first available opportunity).

Asynchronous acknowledgements are sent by the RM destination only if both of the following conditions are met:

- The RM destination is using a non-anonymous `wstrm:acksTo` endpoint.
- The opportunity to piggyback an acknowledgement on a response message does not occur before the expiry of the acknowledgement interval.

[Example 7.10, “Setting the WS-RM Acknowledgement Interval”](#) shows how to set the WS-RM acknowledgement interval.

#### Example 7.10. Setting the WS-RM Acknowledgement Interval

```
<beans xmlns:wstrm-
policy="http://schemas.xmlsoap.org/ws/2005/02/rm/policy
...>
<wstrm-mgr:rmManager id="org.apache.cxf.ws.rm.RMManager">
  <wstrm-policy:RMAssertion>
    <wstrm-policy:AcknowledgementIntervalMilliseconds="2000"/>
  </wstrm-policy:RMAssertion>
</wstrm-mgr:rmManager>
</beans>
```

#### Maximum unacknowledged messages threshold

The `maxUnacknowledged` attribute sets the maximum number of unacknowledged messages that can accrue per sequence before the sequence is terminated.

[Example 7.11, “Setting the WS-RM Maximum Unacknowledged Message Threshold”](#) shows how to set the WS-RM maximum unacknowledged messages threshold.

#### Example 7.11. Setting the WS-RM Maximum Unacknowledged Message Threshold

```
<beans xmlns:wstrm-mgr="http://cxf.apache.org/ws/rm/manager
...>
<wstrm-mgr:reliableMessaging>
  <wstrm-mgr:sourcePolicy>
    <wstrm-mgr:sequenceTerminationPolicy maxUnacknowledged="20" />
  </wstrm-mgr:sourcePolicy>
</wstrm-mgr:reliableMessaging>
</beans>
```

#### Maximum length of an RM sequence

The `maxLength` attribute sets the maximum length of a WS-RM sequence. The default value is 0, which means that the length of a WS-RM sequence is unbound.

When this attribute is set, the RM endpoint creates a new RM sequence when the limit is reached, and after receiving all of the acknowledgements for the previously sent messages. The new message is sent using a new sequence.

[Example 7.12, “Setting the Maximum Length of a WS-RM Message Sequence”](#) shows how to set the maximum length of an RM sequence.

#### Example 7.12. Setting the Maximum Length of a WS-RM Message Sequence

```
<beans xmlns:wsmr-mgr="http://cxf.apache.org/ws/rm/manager
...>
<wsmr-mgr:reliableMessaging>
  <wsmr-mgr:sourcePolicy>
    <wsmr-mgr:sequenceTerminationPolicy maxLength="100" />
  </wsmr-mgr:sourcePolicy>
</wsmr-mgr:reliableMessaging>
</beans>
```

### Message delivery assurance policies

You can configure the RM destination to use the following delivery assurance policies:

- **AtMostOnce** – The RM destination delivers the messages to the application destination only once. If a message is delivered more than once an error is raised. It is possible that some messages in a sequence may not be delivered.
- **AtLeastOnce** – The RM destination delivers the messages to the application destination at least once. Every message sent will be delivered or an error will be raised. Some messages might be delivered more than once.
- **InOrder** – The RM destination delivers the messages to the application destination in the order that they are sent. This delivery assurance can be combined with the **AtMostOnce** or **AtLeastOnce** assurances.

[Example 7.13, “Setting the WS-RM Message Delivery Assurance Policy”](#) shows how to set the WS-RM message delivery assurance.

#### Example 7.13. Setting the WS-RM Message Delivery Assurance Policy

```
<beans xmlns:wsmr-mgr="http://cxf.apache.org/ws/rm/manager
...>
<wsmr-mgr:reliableMessaging>
  <wsmr-mgr:deliveryAssurance>
    <wsmr-mgr:AtLeastOnce />
  </wsmr-mgr:deliveryAssurance>
</wsmr-mgr:reliableMessaging>
</beans>
```


## 7.5. CONFIGURING WS-RM PERSISTENCE

## Overview

The Apache CXF WS-RM features already described in this chapter provide reliability for cases such as network failures. WS-RM persistence provides reliability across other types of failure such as an RM source or an RM destination crash.

WS-RM persistence involves storing the state of the various RM endpoints in persistent storage. This enables the endpoints to continue sending and receiving messages when they are reincarnated.

Apache CXF enables WS-RM persistence in a configuration file. The default WS-RM persistence store is JDBC-based. For convenience, Apache CXF includes Derby for out-of-the-box deployment. In addition, the persistent store is also exposed using a Java API.



### IMPORTANT

WS-RM persistence is supported for oneway calls only, and it is disabled by default.

## How it works

Apache CXF WS-RM persistence works as follows:

- At the RM source endpoint, an outgoing message is persisted before transmission. It is evicted from the persistent store after the acknowledgement is received.
- After a recovery from crash, it recovers the persisted messages and retransmits until all the messages have been acknowledged. At that point, the RM sequence is closed.
- At the RM destination endpoint, an incoming message is persisted, and upon a successful store, the acknowledgement is sent. When a message is successfully dispatched, it is evicted from the persistent store.
- After a recovery from a crash, it recovers the persisted messages and dispatches them. It also brings the RM sequence to a state where new messages are accepted, acknowledged, and delivered.

## Enabling WS-persistence

To enable WS-RM persistence, you must specify the object implementing the persistent store for WS-RM. You can develop your own or you can use the JDBC based store that comes with Apache CXF.

The configuration shown in [Example 7.14, “Configuration for the Default WS-RM Persistence Store”](#) enables the JDBC-based store that comes with Apache CXF.

### Example 7.14. Configuration for the Default WS-RM Persistence Store

```
<bean id="RMTxStore"
class="org.apache.cxf.ws.rm.persistence.jdbc.RMTxStore"/>
<wsrm-mgr:rmManager id="org.apache.cxf.ws.rm.RMManager">
  <property name="store" ref="RMTxStore"/>
</wsrm-mgr:rmManager>
```

## Configuring WS-persistence

The JDBC-based store that comes with Apache CXF supports the properties shown in [Table 7.4, “JDBC Store Properties”](#).

**Table 7.4. JDBC Store Properties**

Attribute Name	Type	Default Setting
driverClassName	String	<code>org.apache.derby.jdbc.EmbeddedDriver</code>
userName	String	null
passWord	String	null
url	String	<code>jdbc:derby:rmdb;create=true</code>

The configuration shown in [Example 7.15, “Configuring the JDBC Store for WS-RM Persistence”](#) enables the JDBC-based store that comes with Apache CXF, while setting the `driverClassName` and `url` to non-default values.

#### Example 7.15. Configuring the JDBC Store for WS-RM Persistence

```
<bean id="RMTxStore"
class="org.apache.cxf.ws.rm.persistence.jdbc.RMTxStore">
  <property name="driverClassName" value="com.acme.jdbc.Driver"/>
  <property name="url" value="jdbc:acme:rmdb;create=true"/>
</bean>
```



## CHAPTER 8. ENABLING HIGH AVAILABILITY

### Abstract

This chapter explains how to enable and configure high availability in the Apache CXF runtime.

## 8.1. INTRODUCTION TO HIGH AVAILABILITY

### Overview

Scalable and reliable applications require high availability to avoid any single point of failure in a distributed system. You can protect your system from single points of failure using *replicated services*.

A replicated service is comprised of multiple instances, or *replicas*, of the same service. Together these act as a single logical service. Clients invoke requests on the replicated service, and Apache CXF delivers the requests to one of the member replicas. The routing to a replica is transparent to the client.

### HA with static failover

Apache CXF supports high availability (HA) with static failover in which replica details are encoded in the service WSDL file. The WSDL file contains multiple ports, and can contain multiple hosts, for the same service. The number of replicas in the cluster remains static as long as the WSDL file remains unchanged. Changing the cluster size involves editing the WSDL file.

## 8.2. ENABLING HA WITH STATIC FAILOVER

### Overview

To enable HA with static failover, you must do the following:

1. [Encode replica details in your service WSDL file](#)
2. [Add the clustering feature to your client configuration](#)

### Encode replica details in your service WSDL file

You must encode the details of the replicas in your cluster in your service WSDL file. [Example 8.1, “Enabling HA with Static Failover–WSDL File”](#) shows a WSDL file extract that defines a service cluster of three replicas.

#### Example 8.1. Enabling HA with Static Failover–WSDL File

```
1 <wsdl:service name="ClusteredService">
2   <wsdl:port binding="tns:Greeter_SOAPBinding" name="Replica1">
      <soap:address
location="http://localhost:9001/SoapContext/Replica1"/>
    </wsdl:port>
3   <wsdl:port binding="tns:Greeter_SOAPBinding" name="Replica2">
      <soap:address
```

```

location="http://localhost:9002/SoapContext/Replica2"/>
  </wsdl:port>

  4 <wsdl:port binding="tns:Greeter_SOAPBinding" name="Replica3">
      <soap:address
location="http://localhost:9003/SoapContext/Replica3"/>
      </wsdl:port>

</wsdl:service>

```

The WSDL extract shown in [Example 8.1, “Enabling HA with Static Failover–WSDL File”](#) can be explained as follows:

- 1 Defines a service, **ClusterService**, which is exposed on three ports:
  1. **Replica1**
  2. **Replica2**
  3. **Replica3**
- 2 Defines **Replica1** to expose the **ClusterService** as a SOAP over HTTP endpoint on port **9001**.
- 3 Defines **Replica2** to expose the **ClusterService** as a SOAP over HTTP endpoint on port **9002**.
- 4 Defines **Replica3** to expose the **ClusterService** as a SOAP over HTTP endpoint on port **9003**.

## Add the clustering feature to your client configuration

In your client configuration file, add the clustering feature as shown in [Example 8.2, “Enabling HA with Static Failover–Client Configuration”](#).

### Example 8.2. Enabling HA with Static Failover–Client Configuration

```

<?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:jaxws="http://cxf.apache.org/jaxws"
  xmlns:clustering="http://cxf.apache.org/clustering"
  xsi:schemaLocation="http://cxf.apache.org/jaxws
  http://cxf.apache.org/schemas/jaxws.xsd
  http://www.springframework.org/schema/beans
  http://www.springframework.org/schema/beans/spring-beans.xsd">

  <jaxws:client name="
{http://apache.org/hello_world_soap_http}Replica1"
      createdFromAPI="true">
    <jaxws:features>
      <clustering:failover/>
    </jaxws:features>
  </jaxws:client>

  <jaxws:client name="

```

```

{http://apache.org/hello_world_soap_http}Replica2"
    createdFromAPI="true">
    <jaxws:features>
      <clustering:failover/>
    </jaxws:features>
  </jaxws:client>

  <jaxws:client name="
{http://apache.org/hello_world_soap_http}Replica3"
    createdFromAPI="true">
    <jaxws:features>
      <clustering:failover/>
    </jaxws:features>
  </jaxws:client>

</beans>

```

## 8.3. CONFIGURING HA WITH STATIC FAILOVER

### Overview

By default, HA with static failover uses a sequential strategy when selecting a replica service if the original service with which a client is communicating becomes unavailable, or fails. The sequential strategy selects a replica service in the same sequential order every time it is used. Selection is determined by Apache CXF's internal service model and results in a deterministic failover pattern.

### Configuring a random strategy

You can configure HA with static failover to use a random strategy instead of the sequential strategy when selecting a replica. The random strategy selects a random replica service each time a service becomes unavailable, or fails. The choice of failover target from the surviving members in a cluster is entirely random.

To configure the random strategy, add the configuration shown in [Example 8.3, “Configuring a Random Strategy for Static Failover”](#) to your client configuration file.

#### Example 8.3. Configuring a Random Strategy for Static Failover

```

<beans ...>
  ❶ <bean id="Random"
    class="org.apache.cxf.clustering.RandomStrategy"/>

    <jaxws:client name="
{http://apache.org/hello_world_soap_http}Replica3"
    createdFromAPI="true">
    <jaxws:features>
      <clustering:failover>
        ❷ <clustering:strategy>
          <ref bean="Random"/>
        </clustering:strategy>
      </clustering:failover>
    </jaxws:features>
  </jaxws:client>

```

```
| | </beans>
```

The configuration shown in [Example 8.3, “Configuring a Random Strategy for Static Failover”](#) can be explained as follows:

- 1 Defines a **Random** bean and implementation class that implements the random strategy.
- 2 Specifies that the random strategy is used when selecting a replica.

## CHAPTER 9. ENABLING HIGH AVAILABILITY IN FUSE FABRIC

### Abstract

When all of your servers and clients are deployed within the same fabric, you can use an alternative mechanism for implementing high availability cluster, which works by exploiting the fabric registry. Because all the parts of the application must be deployed on the same fabric, this mechanism is suitable for deployment on a LAN.

## 9.1. LOAD BALANCING CLUSTER

### 9.1.1. Introduction to Load Balancing

#### Overview

The fabric load balancing mechanism exploits the fact that fabric provides a distributed fabric registry, which is accessible to all of the container in the fabric. This makes it possible to use the fabric registry as a discovery mechanism for locating WS endpoints in the fabric. By storing all of the endpoint addresses belonging to a particular cluster under the same registry node, any WS clients in the fabric can easily discover the location of the endpoints in the cluster.

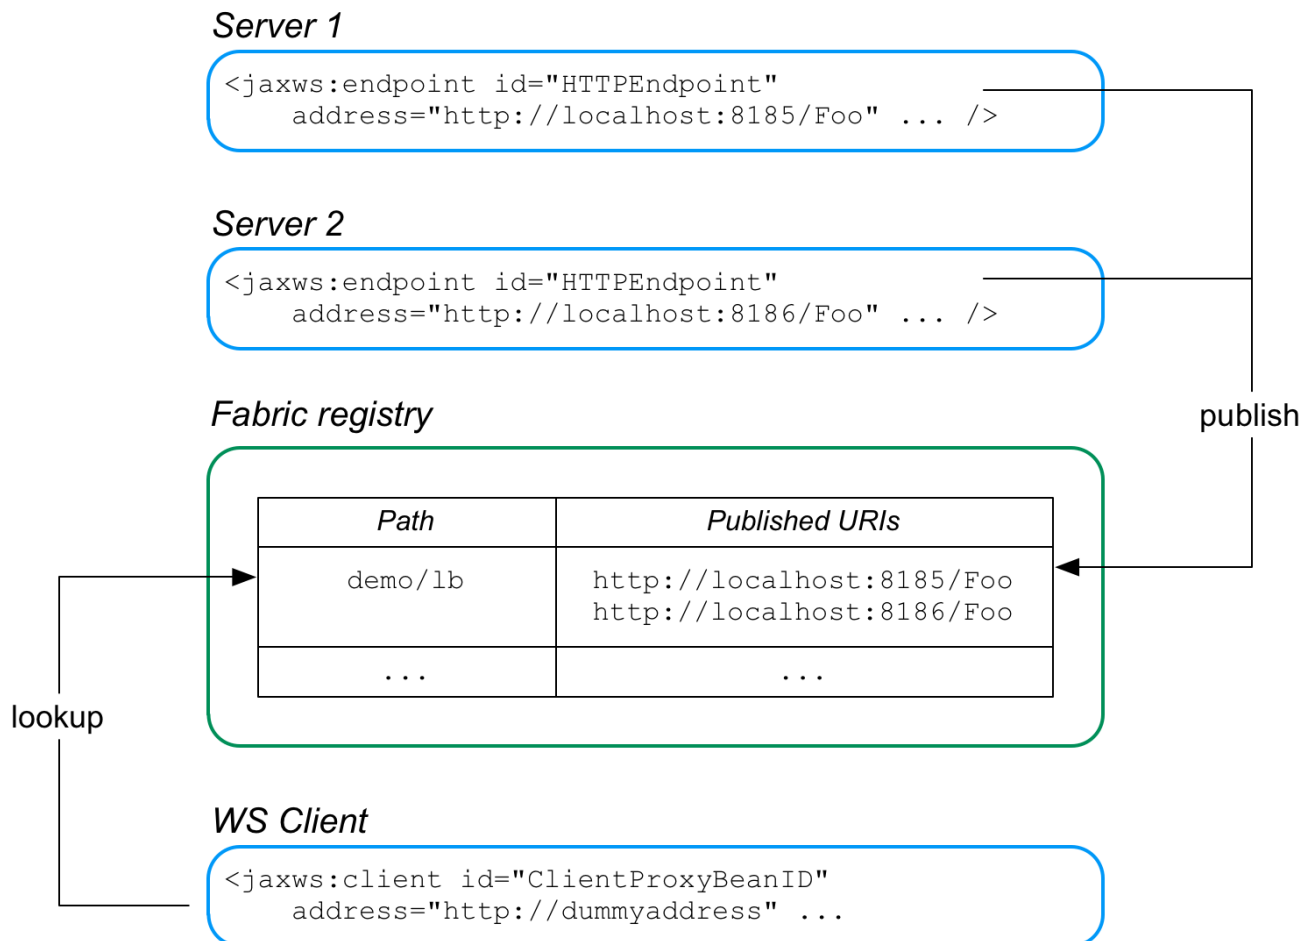
#### Fuse Fabric

A fabric is a distributed collection of containers that share a common database of configuration settings (the *fabric registry*). Every container in the fabric has a fabric agent deployed in it, which manages the container and redeploys applications to the container whenever a new profile is assigned to the container (a profile is the basic deployment unit in a fabric).

#### Load-balancing cluster

[Figure 9.1, “Fabric Load Balancing for Apache CXF”](#) gives an overview of the fabric load balancing mechanism for Apache CXF endpoints.

Figure 9.1. Fabric Load Balancing for Apache CXF



In this example, two WS servers are created, with the URIs, `http://localhost:8185/Foo` and `http://localhost:8186/Foo`. For both of these servers, the load balancer feature is configured to store the cluster endpoints under the path, `demo/lb`, in the fabric registry.

Now, when the WS client starts, it is configured to look up the cluster path, `demo/lb`, in the fabric registry. Because the `demo/lb` path is associated with multiple endpoint addresses, fabric implements a random load balancing algorithm to choose one of the available URIs to connect to.

### FabricLoadBalancerFeature

The fabric load balancer feature is implemented by the following class:

```
org.fusesource.fabric.cxf.FabricLoadBalancerFeature
```

The `FabricLoadBalancerFeature` class exposes the following bean properties:

#### fabricPath

This property specifies a node in the fabric registry (specified relative to the base node, `/fabric/cxf/endpoints`) that is used to store the data for a particular endpoint cluster.

#### zkClient

A proxy reference to the OSGi service exposed by the fabric agent (of type, `org.linkedin.zookeeper.client.IZKClient`).

**maximumConnectionTimeout**

The maximum length of time to attempt to connect to the fabric agent, specified in milliseconds. The default is 10000 (10 seconds).

**connectionRetryTime**

How long to wait between connection attempts, specified in milliseconds. The default is 100.

**loadBalanceStrategy**

By implementing a bean of type `org.fusesource.fabric.cxf.LoadBalanceStrategy` and setting this property, you can customise the load balancing algorithm used by the load balancing feature.

**Prerequisites**

To use the fabric load balancer feature in your application, your project must satisfy the following prerequisites:

- [the section called “Maven dependency”](#).
- [the section called “OSGi package import”](#).
- [the section called “Fabric deployment”](#).
- [the section called “Required feature”](#).

**Maven dependency**

The fabric load balancer feature requires the `fabric-cxf` Maven artifact. Add the following dependency to your project's POM file:

```
<dependency>
  <groupId>org.fusesource.fabric</groupId>
  <artifactId>fabric-cxf</artifactId>
  <version>6.0.0.redhat-024</version>
</dependency>
```

**OSGi package import**

If you are packaging your project as an OSGi bundle, you must add `org.fusesource.fabric.cxf` to the list of imported packages. For example, using the Maven bundle plug-in, you can specify this package import by adding `org.fusesource.fabric.cxf` to the comma-separated list in the `Import-Package` element, as follows:

```
<plugin>
  <groupId>org.apache.felix</groupId>
  <artifactId>maven-bundle-plugin</artifactId>
  <version>2.2.0</version>
  <extensions>>true</extensions>
  <configuration>
    <instructions>
      <Bundle-SymbolicName>${project.artifactId}</Bundle-
SymbolicName>
```

```

        <Import-Package>
            ...
    org.fusesource.fabric.cxf,
            *
        </Import-Package>
        ...
    </instructions>
</configuration>
</plugin>

```

## Fabric deployment

When you come to deploy your application into a Red Hat JBoss Fuse container, you *must* deploy it into a fabric. The fabric load balancer feature is not supported in a standalone container.

## Required feature

The fabric load balancer requires the `fabric-cxf` Apache Karaf feature to be installed in the container. In the context of a fabric, this means you must add the `fabric-cxf` feature to the relevant deployment profile. For example, if you are using the `cxf-lb-server` profile to deploy a load-balancing WS server, you can add the `fabric-cxf` feature by entering the following console command:

```
JBossFuse:karaf@root> profile-edit -f fabric-cxf cxf-lb-server
```

## 9.1.2. Configure the Server

### Overview

To configure a WS server to use fabric load balancing, you must configure a fabric load balancer feature and install it in the default Apache CXF bus instance. This section describes how to configure the load balancer feature in Spring XML and in blueprint XML.

### Prerequisites

For the basic prerequisites to build a fabric load-balancing WS server, see [the section called “Prerequisites”](#).

### Spring XML

The following fragment from a Spring XML file shows how to add the fabric load balancer feature, `FabricLoadBalancerFeature`, to an Apache CXF bus. Any Apache CXF endpoints subsequently created on this bus will automatically have the load-balancer feature enabled.

```

<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
    ...
    xmlns:osgi="http://www.springframework.org/schema/osgi"
    ...
    xmlns:cxfcore="http://cxf.apache.org/core"
>
    ...
    <!-- Reference the fabric agent -->

```



```

<osgi:reference id="IZKClient"
  interface="org.linkedin.zookeeper.client.IZKClient" />

<!-- Configure the Fabric load balancer feature -->
<bean id="fabricLoadBalancerFeature"
  class="org.fusesource.fabric.cxf.FabricLoadBalancerFeature">
  <property name="zkClient" ref="IZKClient" />
  <property name="fabricPath" value="ZKPath" />
</bean>

<!-- Add the feature to the bus -->
<cxfc core:bus>
  <cxfc core:features>
    <ref bean="fabricLoadBalancerFeature" />
  </cxfc core:features>
</cxfc core:bus>
  ...
</beans>

```

The following beans are used to install the fabric load-balancer feature:

#### **IZKClient OSGi service reference**

The **IZKClient** reference is a proxy of the local fabric agent, which it accesses through the **org.linkedin.zookeeper.client.IZKClient** interface. This reference is needed in order to integrate the load balancer feature with the underlying fabric.

#### **FabricLoadBalancerFeature bean**

The **FabricLoadBalancerFeature** bean is initialised with the following properties:

##### **zkClient**

A reference to the fabric agent proxy, **IZKClient**.

##### **fabricPath**

The path of a node in the fabric registry, where the cluster data is stored (for example, the addresses of the endpoints in the load-balancing cluster). The node path is specified relative to the base node, **/fabric/cxf/endpoints**.

#### **Apache CXF bus**

The **cxfc core:bus** element installs the fabric load balancer feature in the default bus instance.

#### **Blueprint XML**

The following fragment from a blueprint XML file shows how to add the fabric load balancer feature, **FabricLoadBalancerFeature**, to an Apache CXF bus. Any Apache CXF endpoints subsequently created on this bus will automatically have the load-balancer feature enabled.

```

<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  ...
  xmlns:cxf="http://cxf.apache.org/blueprint/core"
  ...

```

```

>
...
<!-- Reference the fabric agent -->
<reference id="org.linkedin.zookeeper.client.IZKClient"
          interface="org.linkedin.zookeeper.client.IZKClient" />

<!-- Create the Fabric load balancer feature -->
<bean id="fabricLoadBalancerFeature"
      class="org.fusesource.fabric.cxf.FabricLoadBalancerFeature">
  <property name="zkClient"
    ref="org.linkedin.zookeeper.client.IZKClient" />
  <property name="fabricPath" value="ZKPath" />
</bean>

<!-- setup the feature on the bus to help publish the services to the
fabric-->
<cxf:bus bus="cxf">
  <cxf:features>
    <ref component-id="fabricLoadBalancerFeature"/>
  </cxf:features>
</cxf:bus>
...
</blueprint>

```

The following beans are used to install the fabric load-balancer feature:

#### IZKClient reference

The **IZKClient** reference is a proxy of the local fabric agent, which it accesses through the `org.linkedin.zookeeper.client.IZKClient` interface. This reference is needed in order to integrate the load balancer feature with the underlying fabric.

#### FabricLoadBalancerFeature bean

The **FabricLoadBalancerFeature** bean is initialised with the following properties:

##### zkClient

A reference to the fabric agent proxy, **IZKClient**.

##### fabricPath

The path of a node in the fabric registry, where the cluster data is stored (for example, the addresses of the endpoints in the load-balancing cluster). The node path is specified relative to the base node, `/fabric/cxf/endpoints`.

#### Apache CXF bus

The `cxf:bus` element installs the fabric load balancer feature in the default bus instance.

#### Example using Spring XML

[Example 9.1, “WS Server with Fabric Load Balancer Feature”](#) shows a complete Spring XML example of a WS endpoint configured to use the fabric load balancing feature.

**Example 9.1. WS Server with Fabric Load Balancer Feature**

```

<?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:jaxws="http://cxf.apache.org/jaxws"
  xmlns:osgi="http://www.springframework.org/schema/osgi"
  xmlns:osgix="http://www.springframework.org/schema/osgi-compendium"
  xmlns:ctx="http://www.springframework.org/schema/context"
  xmlns:cxfcore="http://cxf.apache.org/core"
  xsi:schemaLocation="
    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
    http://www.springframework.org/schema/osgi
    http://www.springframework.org/schema/osgi/spring-osgi.xsd
    http://www.springframework.org/schema/context
    http://www.springframework.org/schema/context/spring-
context.xsd
    http://www.springframework.org/schema/osgi
    http://www.springframework.org/schema/osgi/spring-osgi.xsd
    http://www.springframework.org/schema/osgi-compendium
    http://www.springframework.org/schema/osgi-
compendium/spring-osgi-compendium.xsd
    http://cxf.apache.org/core
    http://cxf.apache.org/schemas/core.xsd
">

  <!-- Configuration Admin entry -->
  <osgix:cm-properties id="cmProps"
    persistent-id="org.fusesource.example.fabric.lb">
    <prop key="portNumber">8191</prop>
  </osgix:cm-properties>

  <!-- placeholder configurer -->
  <ctx:property-placeholder properties-ref="cmProps" />

  <!-- Create the WS endpoint -->
  <jaxws:endpoint id="HTTPEndpoint"
    implementor="org.fusesource.example.PersonImpl"
    address="http://localhost:${portNumber}/PersonServiceCF"/>

  <!-- Reference the fabric agent -->
  <osgi:reference id="IZKClient"
    interface="org.linkedin.zookeeper.client.IZKClient" />

  <!-- Configure the Fabric load balancer feature -->
  <bean id="fabricLoadBalancerFeature"
    class="org.fusesource.fabric.cxf.FabricLoadBalancerFeature">
    <property name="zkClient" ref="IZKClient" />
    <property name="fabricPath" value="demo/lb" />
  </bean>

```

```

        <!-- Add the feature to the bus -->
        <cxfc core:bus>
            <cxfc core:features>
                <ref bean="fabricLoadBalancerFeature" />
            </cxfc core:features>
        </cxfc core:bus>
    </beans>

```

The preceding Spring XML configuration consists of the following main sections:

- *Setting up OSGi Configuration Admin*—the `osgi:cm-properties` element and the `ctx:property-placeholder` element set up the OSGi Configuration Admin service, enabling you to substitute the WS endpoint's IP port number using the placeholder, `${portNumber}`.

Use of the OSGi Configuration Admin service is optional, but it provides a convenient way of creating a cluster of WS servers listening on different ports. With the configuration shown here, you can deploy the *same* server bundle into different container instances, using the OSGi Configuration Admin service to customise the IP port of each deployed server instance.

- *Creating the WS endpoint*—create the WS endpoint in the usual way, using the `jaxws:endpoint` element. By default, this endpoint is automatically associated with the default bus instance, which has load balancing enabled. The only unusual aspect of this endpoint definition is that the OSGi Configuration Admin service is used to define the port number in the endpoint's address (through the `${portNumber}` placeholder).
- *Enabling the fabric load balancing feature*—the fabric load balancing feature is installed in the default bus instance, as previously described. In this example, the `fabricPath` property is set to the value, `demo/lb`.

## Creating fabric profiles for the example

To deploy the WS endpoint in a fabric, you need to create the appropriate profiles. Because you want to create a cluster of WS endpoints, it makes sense to define multiple profiles, where each WS endpoint gets its own profile. For example, you could define the following set of profiles for deployment:

### `cxf-lb-server`

A base profile, containing details common to all of the cluster endpoints.

### `cxf-lb-server-8185`

A profile that inherits from `cxf-lb-server` and sets the `portNumber` property to 8185.

### `cxf-lb-server-8186`

A profile that inherits from `cxf-lb-server` and sets the `portNumber` property to 8186.

Assuming that the example server bundle (containing the WS endpoint implementation) is stored in the local Maven repository and has the Maven coordinates, `org.fusesource.example/cxf-lb-server/1.0-SNAPSHOT`, you can create the `cxf-lb-server` base profile by entering the following console commands:

```
JBossFuse:karaf@root> profile-create --parents cxf cxf-lb-server
JBossFuse:karaf@root> profile-edit -f fabric-cxf cxf-lb-server
JBossFuse:karaf@root> profile-edit -b mvn:org.fusesource.example/cxf-lb-
server/1.0-SNAPSHOT cxf-lb-server
```

You can then create the **cxf-lb-server-8185** profile and the **cxf-lb-server-8186** profile as follows:

```
JBossFuse:karaf@root> profile-create --parents cxf-lb-server cxf-lb-
server-8185
JBossFuse:karaf@root> profile-create --parents cxf-lb-server cxf-lb-
server-8186
JBossFuse:karaf@root> profile-edit -p
org.fusesource.example.fabric.lb/portNumber=8185 cxf-lb-server-8185
JBossFuse:karaf@root> profile-edit -p
org.fusesource.example.fabric.lb/portNumber=8186 cxf-lb-server-8186
```

The **cxf-lb-server-8185** profile and the **cxf-lb-server-8186** profile are now ready to be deployed to the container of your choice, using the **fabric:container-change-profile** command.

### 9.1.3. Configure the Client

#### Overview

To configure a WS client to use fabric load balancing, you must install the fabric load balancer feature directly in the client proxy instance. This section describes how to configure the load balancer feature in Spring XML, blueprint XML, and by programming in Java.

#### Prerequisites

For the basic prerequisites to build a fabric load-balancing WS client, see [the section called “Prerequisites”](#).

#### Spring XML

The following fragment from a Spring XML file shows how to add the fabric load balancer feature, **FabricLoadBalancerFeature**, directly into a WS client proxy instance.

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  ...
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  xmlns:osgi="http://www.springframework.org/schema/osgi"
  ...
>
<!-- Create a client proxy, with load balancing enabled -->
<jaxws:client id="ClientProxyBeanID"
  address="http://dummyaddress"
  serviceClass="SEI">
  <jaxws:features>
    <ref bean="fabricLoadBalancerFeature" />
  </jaxws:features>
```

```

</jaxws:client>
...
<!-- Reference the fabric agent -->
<osgi:reference id="IZKClient"
    interface="org.linkedin.zookeeper.client.IZKClient" />

<!-- Configure the Fabric load balancer feature -->
<bean id="fabricLoadBalancerFeature"
    class="org.fusesource.fabric.cxf.FabricLoadBalancerFeature">
    <property name="zkClient" ref="IZKClient" />
    <property name="fabricPath" value="ZKPath" />
</bean>
...
</beans>

```

The fabric load balancer feature is installed directly into the WS client proxy by inserting it as a child of the `jaxws:features` element (or, as in this case, by inserting a bean reference to the actual instance). The following beans are used to initialise the fabric load-balancer feature:

#### IZKClient OSGi service reference

The `IZKClient` reference is a proxy of the local fabric agent, which it accesses through the `org.linkedin.zookeeper.client.IZKClient` interface. This reference is needed in order to integrate the load balancer feature with the underlying fabric.

#### FabricLoadBalancerFeature bean

The `FabricLoadBalancerFeature` bean is initialised with the following properties:

##### zkClient

A reference to the fabric agent proxy, `IZKClient`.

##### fabricPath

The path of a node in the fabric registry, where the cluster data is stored (for example, the addresses of the endpoints in the load-balancing cluster). The node path is specified relative to the base node, `/fabric/cxf/endpoints`.

## Blueprint XML

The following fragment from a blueprint XML file shows how to add the fabric load balancer feature, `FabricLoadBalancerFeature`, directly into a WS client proxy instance.

```

<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
    ...
    xmlns:jaxws="http://cxf.apache.org/blueprint/jaxws"
    xmlns:cxf="http://cxf.apache.org/blueprint/core"
    ...
>
<!-- Create a client proxy, with load balancing enabled -->
<jaxws:client id="ClientProxyBeanID"
    address="http://dummyaddress"
    serviceClass="SEI">
    <jaxws:features>

```

```

        <ref component-id="fabricLoadBalancerFeature" />
    </jaxws:features>
</jaxws:client>
...
<reference id="org.linkedin.zookeeper.client.IZKClient"
    interface="org.linkedin.zookeeper.client.IZKClient" />

<!-- Create the Fabric load balancer feature -->
<bean id="fabricLoadBalancerFeature"
    class="org.fusesource.fabric.cxf.FabricLoadBalancerFeature">
    <property name="zkClient"
ref="org.linkedin.zookeeper.client.IZKClient" />
    <property name="fabricPath" value="ZKPath" />
</bean>
...
</blueprint>

```

The fabric load balancer feature is installed directly into the WS client proxy by inserting it as a child of the `jaxws:features` element (or, as in this case, by inserting a bean reference to the actual instance). The following beans are used to initialise the fabric load-balancer feature:

#### IZKClient reference

The `IZKClient` reference is a proxy of the local fabric agent, which it accesses through the `org.linkedin.zookeeper.client.IZKClient` interface. This reference is needed in order to integrate the load balancer feature with the underlying fabric.

#### FabricLoadBalancerFeature bean

The `FabricLoadBalancerFeature` bean is initialised with the following properties:

##### zkClient

A reference to the fabric agent proxy, `IZKClient`.

##### fabricPath

The path of a node in the fabric registry, where the cluster data is stored (for example, the addresses of the endpoints in the load-balancing cluster). The node path is specified relative to the base node, `/fabric/cxf/endpoints`.

#### Java

As an alternative to using XML configuration, you can enable the fabric load balancing feature on the client side by programming directly in Java. The following example shows how to enable fabric load balancing on a proxy for the `Hello` Web service.

```

// Java
package org.fusesource.fabric.demo.cxf.client;

import org.apache.cxf.feature.AbstractFeature;
import org.apache.cxf.frontend.ClientProxyFactoryBean;
import org.apache.cxf.jaxws.JaxWsProxyFactoryBean;
import org.fusesource.fabric.cxf.FabricLoadBalancerFeature;
import org.fusesource.fabric.demo.cxf.Hello;

```

```

import java.util.ArrayList;
import java.util.List;

public class Client {

    private Hello hello;

    public void initializeHelloProxy() {
        // The feature will try to create a zookeeper client itself
        // by checking the system property of zookeeper.url
        FabricLoadBalancerFeature feature = new
FabricLoadBalancerFeature();
        // Feature will use this path to locate the service
        feature.setFabricPath("demo/lb");

        ClientProxyFactoryBean clientFactory = new
JaxWsProxyFactoryBean();
        clientFactory.setServiceClass(ClientProxyFactoryBean.class);
        // The address is not the actual address that the client will
access
        clientFactory.setAddress("http://dummyaddress");

        List<AbstractFeature> features = new ArrayList<AbstractFeature>();
        features.add(feature);
        // we need to setup the feature on the client factory
        clientFactory.setFeatures(features);

        // Create the proxy of Hello
        hello = clientFactory.create(Hello.class);
    }

    public static void main(String args[]) {
        initializeHelloProxy();
        ...
    }
}

```

In this example, the `fabricPath` property is set to the value, `demo/lb` (which matches the example value used by the server in [Example 9.1, “WS Server with Fabric Load Balancer Feature”](#)).

The address that the client proxy accesses is set to a dummy value, `http://dummyaddress`, because this value is not used. When the client is initialized, the load balancer feature substitutes the address value retrieved from the `demo/lb` node of the fabric registry.

### Example using Spring XML

[Example 9.2, “Client Proxy with Fabric Load Balancer Feature”](#) shows a detailed Spring XML example of how to configure a WS client proxy with the fabric load balancer feature.

#### Example 9.2. Client Proxy with Fabric Load Balancer Feature

```

<?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:jaxws="http://cxf.apache.org/jaxws"
xmlns:osgi="http://www.springframework.org/schema/osgi"
xmlns:cxfcore="http://cxf.apache.org/core"
xsi:schemaLocation="
    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
    http://www.springframework.org/schema/osgi
    http://www.springframework.org/schema/osgi/spring-osgi.xsd
    http://www.springframework.org/schema/osgi
    http://www.springframework.org/schema/osgi/spring-osgi.xsd
    http://cxf.apache.org/core
    http://cxf.apache.org/schemas/core.xsd
">

<!-- Create a client proxy, with load balancing enabled -->
<jaxws:client id="personServiceProxy"
    address="http://dummyaddress"
    serviceClass="org.fusesource.example.Person">
    <jaxws:features>
        <ref bean="fabricLoadBalancerFeature" />
    </jaxws:features>
</jaxws:client>

<!-- Inject the client proxy into a bean... -->
<!-- (not shown) -->

<!-- Reference the fabric agent -->
<osgi:reference id="IZKClient"
    interface="org.linkedin.zookeeper.client.IZKClient" />

<!-- Configure the Fabric load balancer feature -->
<bean id="fabricLoadBalancerFeature"
    class="org.fusesource.fabric.cxf.FabricLoadBalancerFeature">
    <property name="zkClient" ref="IZKClient" />
    <property name="fabricPath" value="demo/lb" />
</bean>

</beans>

```

## Creating a fabric profile for the client

To deploy the WS client in a fabric, you need to create a profile for it. Assuming that the example client bundle is stored in the local Maven repository and has the Maven coordinates, **org.fusesource.example/cxf-lb-client/1.0-SNAPSHOT**, you can create the **cxf-lb-client** profile by entering the following console commands:

```

JBossFuse:karaf@root> profile-create --parents cxf cxf-lb-client
JBossFuse:karaf@root> profile-edit -f fabric-cxf cxf-lb-client
JBossFuse:karaf@root> profile-edit -b mvn:org.fusesource.example/cxf-lb-
client/1.0-SNAPSHOT cxf-lb-client

```

The `cxf-lb-client` profile is now ready to be deployed to the container of your choice, using the `fabric:container-change-profile` command.

## 9.2. FAILOVER CLUSTER

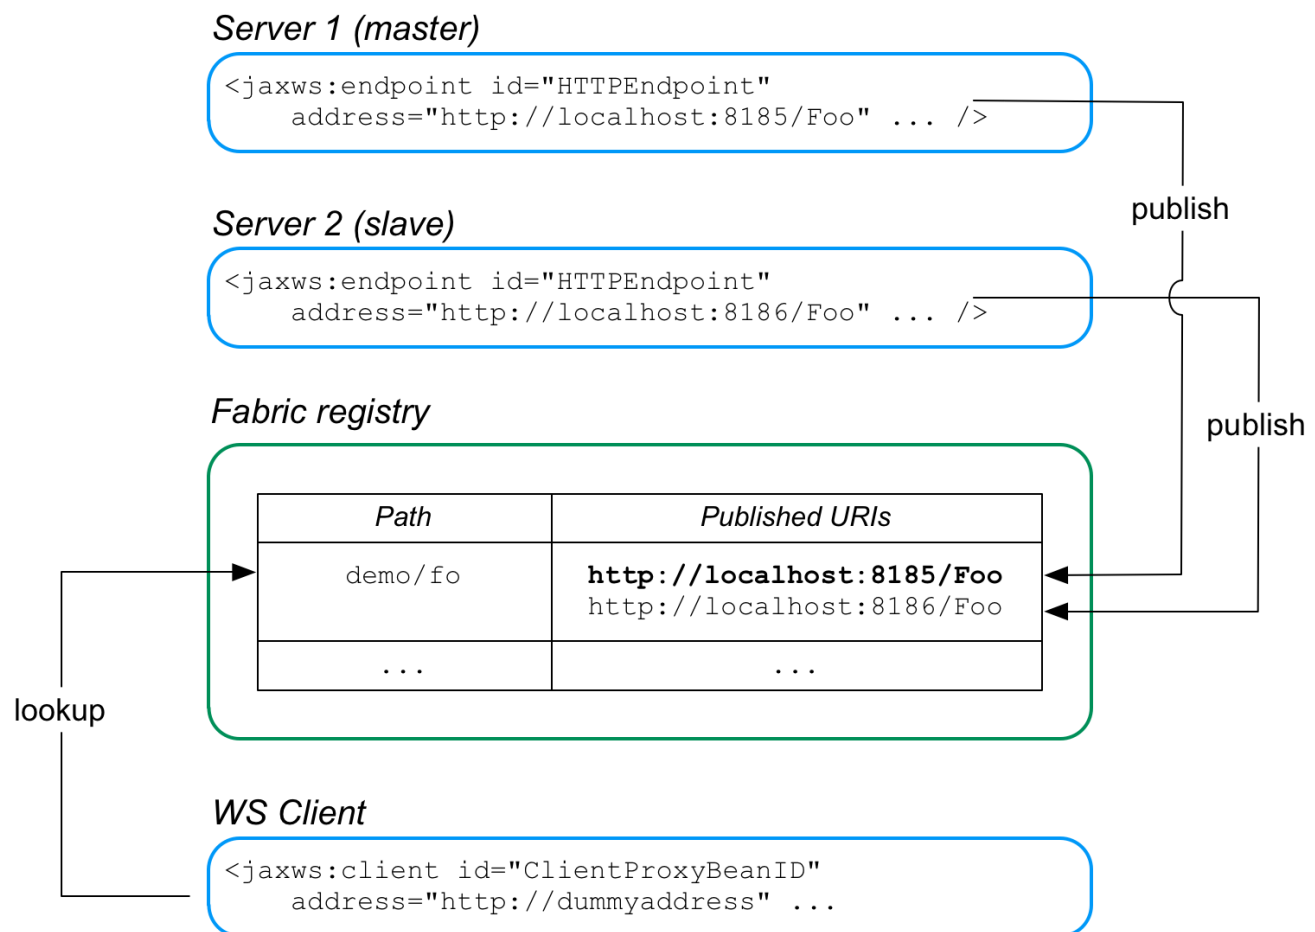
### Overview

A failover cluster in Fuse Fabric is based on an ordered list of WS endpoints that are registered under a particular node in the fabric registry. A client detects the failure of a master endpoint by catching the exception that occurs when it tries to make an invocation. When that happens, the client automatically moves to the next available endpoint in the cluster.

### Failover cluster

Figure 9.2, “Fabric Failover for Apache CXF” gives an overview of the fabric failover mechanism for Apache CXF endpoints.

Figure 9.2. Fabric Failover for Apache CXF



In this example, two WS servers are created, with the URIs, `http://localhost:8185/Foo` and `http://localhost:8186/Foo`. In both servers, the failover feature is configured to store the cluster endpoints under the path, `demo/fo`, in the fabric registry. The cluster endpoints stored under `demo/fo` are *ordered*. The first endpoint in the cluster is the *master* and all of the other endpoints are *slaves*.

The failover algorithm works as follows:

1. When the WS client starts, it is configured to look up the cluster path, `demo/fo`, in the fabric registry. The failover feature initially returns the *first* address registered under `demo/fo` (the master).
2. At some point, the master server could fail. The client determines whether the master has failed by catching the exception that occurs when it tries to make an invocation: if the caught exception matches one of the exceptions in a specified list (by default, just the `java.io.IOException`), the master is deemed to have failed and the client now ignores the corresponding address entry under `demo/fo`.
3. The client selects the *next* address entry under `demo/fo` and attempts to connect to that server. Assuming that this server is healthy, it is effectively the new master.
4. At some point in the future, if the failed old master is restarted successfully, it creates a new address entry under `demo/fo` *after* the existing entries, and is then available to clients, in case the other server (or servers) fail.

## FabricFailOverFeature

The fabric failover feature is implemented by the following class:

```
org.fusesource.fabric.cxf.FabricFailOverFeature
```

The `FabricFailOverFeature` class exposes the following bean properties:

### fabricPath

This property specifies a node in the fabric registry (specified relative to the base node, `/fabric/cxf/endpoints`) that is used to store the data for a particular endpoint cluster.

### zkClient

A proxy reference to the OSGi service exposed by the fabric agent (of type, `org.linkedin.zookeeper.client.IZKClient`).

### maximumConnectionTimeout

The maximum length of time to attempt to connect to the fabric agent, specified in milliseconds. The default is 10000 (10 seconds).

### connectionRetryTime

How long to wait between connection attempts, specified in milliseconds. The default is 100.

### exceptions

A semicolon-separated list of exceptions that signal to the client that a server has failed. If not set, this property defaults to `java.io.IOException`.

For example, you could set the `exceptions` property to a value like the following:

```
java.io.IOException;javax.xml.ws.soap.SOAPFaultException
```

## Spring XML

The configuration of WS servers and WS clients in the failover case is similar to the load balancing case (see [Section 9.1.2, “Configure the Server”](#) and [Section 9.1.3, “Configure the Client”](#)), except that instead of instantiating and referencing a `FabricLoadBalancerFeature` bean, you must instantiate and reference a `FabricFailOverFeature` bean.

For example, in Spring XML you can create a `FabricFailOverFeature` bean instance as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  ...
  xmlns:osgi="http://www.springframework.org/schema/osgi"
  ...
  xmlns:cxfc="http://cxf.apache.org/core"
>
  ...
  <!-- Reference the fabric agent -->
  <osgi:reference id="IZKClient"
    interface="org.linkedin.zookeeper.client.IZKClient" />

  <!-- Configure the Fabric load balancer feature -->
  <bean id="failoverFeature"
    class="org.fusesource.fabric.cxf.FabricFailOverFeature">
    <property name="zkClient" ref="IZKClient" />
    <property name="fabricPath" value="ZKPath" />
  </bean>
  ...
</beans>
```

Remember to customise the value of the `fabricPath` property and to reference the appropriate bean ID (`failoverFeature` in the preceding example).

## Blueprint XML

In blueprint XML you can create a `FabricFailOverFeature` bean instance as follows:

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  ...
  xmlns:cxfc="http://cxf.apache.org/blueprint/core"
  ...
>
  ...
  <!-- Reference the fabric agent -->
  <reference id="org.linkedin.zookeeper.client.IZKClient"
    interface="org.linkedin.zookeeper.client.IZKClient" />

  <!-- Create the Fabric load balancer feature -->
  <bean id="failoverFeature"
    class="org.fusesource.fabric.cxf.FabricFailOverFeature">
    <property name="zkClient"
ref="org.linkedin.zookeeper.client.IZKClient" />
    <property name="fabricPath" value="ZKPath" />
  </bean>
  ...
</blueprint>
```

## CHAPTER 10. PACKAGING AN APPLICATION

### Abstract

Applications must be packed as an OSGi bundle before they can be deployed into Red Hat JBoss Fuse. You will not need to include any Apache CXF specific packages in your bundle. The Apache CXF packages are included in JBoss Fuse. You need to ensure you import the required packages when building your bundle.

### CREATING A BUNDLE

To deploy a Apache CXF application into Red Hat JBoss Fuse, you need to package it as an OSGi bundle. There are several tools available for assisting in the process. JBoss Fuse uses the Maven bundle plug-in whose use is described in [Appendix B, Using the Maven OSGi Tooling](#).

### REQUIRED BUNDLE

The Apache CXF runtime components are included in JBoss Fuse as an OSGi bundle called `org.apache.cxf.cxf-bundle`. This bundle needs to be installed in the JBoss Fuse container before your application's bundle can be started.

To inform the container of this dependency, you use the OSGi manifest's Required-Bundle property.

### REQUIRED PACKAGES

In order for your application to use the Apache CXF components, you need to import their packages into the application's bundle. Because of the complex nature of the dependencies in Apache CXF, you cannot rely on the Maven bundle plug-in, or the `bnd` tool, to automatically determine the needed imports. You will need to explicitly declare them.

You need to import the following packages into your bundle:

- `javax.jws`
- `javax.wsdl`
- `META-INF.cxf`
- `META-INF.cxf.osgi`
- `org.apache.cxf.bus`
- `org.apache.cxf.bus.spring`
- `org.apache.cxf.bus.resource`
- `org.apache.cxf.configuration.spring`
- `org.apache.cxf.resource`
- `org.apache.servicemix.cxf.transport.http_osgi`
- `org.springframework.beans.factory.config`

## EXAMPLE

[Example 10.1, “Apache CXF Application Manifest”](#) shows a manifest for a Apache CXF application's OSGi bundle.

### Example 10.1. Apache CXF Application Manifest

```
Manifest-Version: 1.0
Built-By: FinnMcCumial
Created-By: Apache Maven Bundle Plugin
Bundle-License: http://www.apache.org/licenses/LICENSE-2.0.txt
Import-Package: javax.jws, javax.wsdl, META-INF.cxf, META-INF.cxf.osgi,
org.apache.cxf.bus, org.apache.cxf.bus.spring, org.apache.bus.resource,
org.apache.cxf.configuration.spring, org.apache.cxf.resource,
org.apache.servicemix.cxf.transport.http_cxf,
org.springframework.beans.factory.config
Bnd-LastModified: 1222079507224
Bundle-Version: 4.0.0.fuse
Bundle-Name: Fuse CXF Example
Bundle-Description: This is a sample CXF manifest.
Build-Jdk: 1.5.0_08
Private-Package: org.apache.servicemix.examples.cxf
Required-Bundle: org.apache.cxf.cxf-bundle
Bundle-ManifestVersion: 2
Bundle-SymbolicName: cxf-wsdl-first-osgi
Tool: Bnd-0.0.255
```

# CHAPTER 11. DEPLOYING AN APPLICATION

## Abstract

Red Hat JBoss Fuse will automatically install and deploy your application. You can also manually control the state of your application using the console.

## OVERVIEW

There are two ways to deploy your application into Red Hat JBoss Fuse:

1. Rely on the hot deployment mechanism.
2. Use the console.

You can also start and stop a deployed application using the console.

## HOT DEPLOYMENT

The easiest way to deploy an application is to place it in the hot deployment folder. By default, the hot deployment folder is `InstallDir/deploy`. Any bundle placed in this folder is installed into the container. If its dependencies can be resolved, the bundle is activated.

Once the bundle is installed in the container, you can manage it using the console.

## DEPLOYING FROM THE CONSOLE

The easiest way to deploy an application from the console is to install it and start it in one step. This is done using the `osgi install -s` command. It takes the location of the bundle as a URI. So the command:

```
servicemix>osgi install -s file:/home/finn/ws/widgetapp.jar
```

Installs and attempts to start the bundle `widgetapp.jar` which is located in `/home/finn/ws`.

You can use the `osgi install` command without the `-s` flag. That will install the bundle without attempting to start it. You will then have to manually start the bundle using the `osgi start` command.

The `osgi start` command uses the bundle ID to determine which bundle to activate. <sup>[1]</sup>

## REFRESHING AN APPLICATION

If you make changes to your application and want to redeploy it, you can do so by replacing the installed bundle with a new version and using the `osgi refresh` command. This command instructs the container to stop the running instance of your application, reload the bundle, and restart it.

The `osgi refresh` command uses a bundle ID to determine which bundle to refresh. <sup>[1]</sup>

## STOPPING AN APPLICATION

If you want to temporarily deactivate your application you can use the `osgi stop` command. The `osgi stop` moves your application's bundle from the active state to the resolved state. This means that it can be easily restarted using the `osgi start` command.

The `osgi stop` command uses a bundle ID to determine which bundle to stop. [\[1\]](#)

## UNINSTALLING AN APPLICATION

When you want to permanently remove an application from the container you need to uninstall it. Bundles can only be installed when they are not active. This means that you have to stop your application using the `osgi stop` command before trying to uninstall it.

Once the application's bundle is stopped, you can use the `osgi uninstall` command to remove the bundle from the container. This does not delete the physical bundle. It just removes the bundle from the container's list of installed bundles.

The `osgi stop` command uses a bundle ID to determine which bundle to uninstall. [\[1\]](#)

---

[1] You can get a list of the bundle IDs using the `osgi list` command.



## APPENDIX A. APACHE CXF BINDING IDS

Table A.1. Binding IDs for Message Bindings

Binding	ID
CORBA	<a href="http://cxf.apache.org/bindings/corba">http://cxf.apache.org/bindings/corba</a>
HTTP/REST	<a href="http://apache.org/cxf/binding/http">http://apache.org/cxf/binding/http</a>
SOAP 1.1	<a href="http://schemas.xmlsoap.org/wsdl/soap/http">http://schemas.xmlsoap.org/wsdl/soap/http</a>
SOAP 1.1 w/ MTOM	<a href="http://schemas.xmlsoap.org/wsdl/soap/http?mtom=true">http://schemas.xmlsoap.org/wsdl/soap/http?mtom=true</a>
SOAP 1.2	<a href="http://www.w3.org/2003/05/soap/bindings/HTTP/">http://www.w3.org/2003/05/soap/bindings/HTTP/</a>
SOAP 1.2 w/ MTOM	<a href="http://www.w3.org/2003/05/soap/bindings/HTTP/?mtom=true">http://www.w3.org/2003/05/soap/bindings/HTTP/?mtom=true</a>
XML	<a href="http://cxf.apache.org/bindings/xformat">http://cxf.apache.org/bindings/xformat</a>

## APPENDIX B. USING THE MAVEN OSGI TOOLING

### Abstract

Manually creating a bundle, or a collection of bundles, for a large project can be cumbersome. The Maven bundle plug-in makes the job easier by automating the process and providing a number of shortcuts for specifying the contents of the bundle manifest.

The Red Hat JBoss Fuse OSGi tooling uses the [Maven bundle plug-in](#) from Apache Felix. The bundle plug-in is based on the [bnd](#) tool from Peter Kriens. It automates the construction of OSGi bundle manifests by introspecting the contents of the classes being packaged in the bundle. Using the knowledge of the classes contained in the bundle, the plug-in can calculate the proper values to populate the `Import-Packages` and the `Export-Package` properties in the bundle manifest. The plug-in also has default values that are used for other required properties in the bundle manifest.

To use the bundle plug-in, do the following:

1. [Add](#) the bundle plug-in to your project's POM file.
2. [Configure](#) the plug-in to correctly populate your bundle's manifest.

### B.1. SETTING UP A RED HAT JBOSS FUSE OSGI PROJECT

#### Overview

A Maven project for building an OSGi bundle can be a simple single level project. It does not require any sub-projects. However, it does require that you do the following:

1. [Add](#) the bundle plug-in to your POM.
2. [Instruct](#) Maven to package the results as an OSGi bundle.



#### NOTE

There are several Maven archetypes you can use to set up your project with the appropriate settings.

#### Directory structure

A project that constructs an OSGi bundle can be a single level project. It only requires that you have a top-level POM file and a `src` folder. As in all Maven projects, you place all Java source code in the `src/java` folder, and you place any non-Java resources in the `src/resources` folder.

Non-Java resources include Spring configuration files, JBI endpoint configuration files, and WSDL contracts.



#### NOTE

Red Hat JBoss Fuse OSGi projects that use Apache CXF, Apache Camel, or another Spring configured bean also include a `beans.xml` file located in the `src/resources/META-INF/spring` folder.

## Adding a bundle plug-in

Before you can use the bundle plug-in you must add a dependency on Apache Felix. After you add the dependency, you can add the bundle plug-in to the plug-in portion of the POM.

[Example B.1, “Adding an OSGi bundle plug-in to a POM”](#) shows the POM entries required to add the bundle plug-in to your project.

### Example B.1. Adding an OSGi bundle plug-in to a POM


```

...
<dependencies>
  1 <dependency>
    <groupId>org.apache.felix</groupId>
    <artifactId>org.osgi.core</artifactId>
    <version>1.0.0</version>
  </dependency>
...
</dependencies>
...
<build>
  <plugins>
    2 <plugin>
      <groupId>org.apache.felix</groupId>
      <artifactId>maven-bundle-plugin</artifactId>
      <configuration>
        <instructions>
          3 <Bundle-SymbolicName>${pom.artifactId}</Bundle-SymbolicName>
          4 <Import-Package>*,org.apache.camel.osgi</Import-Package>
          5 <Private-Package>org.apache.servicemix.examples.camel</Private-Package>
        </instructions>
      </configuration>
    </plugin>
  </plugins>
</build>
...

```

The entries in [Example B.1, “Adding an OSGi bundle plug-in to a POM”](#) do the following:

- 1 Adds the dependency on Apache Felix
- 2 Adds the bundle plug-in to your project
- 3 Configures the plug-in to use the project's artifact ID as the bundle's symbolic name
- 4 Configures the plug-in to include all Java packages imported by the bundled classes; also imports the `org.apache.camel.osgi` package
- 5 Configures the plug-in to bundle the listed class, but not to include them in the list of exported packages

**NOTE**

Edit the configuration to meet the requirements of your project.

For more information on configuring the bundle plug-in, see [Section B.2, “Configuring the Bundle Plug-In”](#).

**Activating a bundle plug-in**

To have Maven use the bundle plug-in, instruct it to package the results of the project as a bundle. Do this by setting the POM file's `packaging` element to `bundle`.

**Useful Maven archetypes**

There are several Maven archetypes to generate a project that is preconfigured to use the bundle plug-in:

- [the section called “Spring OSGi archetype”](#)
- [the section called “Apache CXF code-first archetype”](#)
- [the section called “Apache CXF wsdl-first archetype”](#)
- [the section called “Apache Camel archetype”](#)

**Spring OSGi archetype**

The Spring OSGi archetype creates a generic project for building an OSGi project using Spring DM, as shown:

```
org.springframework.osgi/spring-bundle-osgi-archetype/1.1.2
```

You invoke the archetype using the following command:

```
mvn archetype:create -DarchetypeGroupId=org.springframework.osgi -
DarchetypeArtifactId=spring-osgi-bundle-archetype -DarchetypeVersion=1.12
-DgroupId=groupId -DartifactId=artifactId -Dversion=version
```

**Apache CXF code-first archetype**

The Apache CXF code-first archetype creates a project for building a service from Java, as shown:

```
org.apache.servicemix.tooling/servicemix-osgi-cxf-code-first-
archetype/2008.01.0.3-fuse
```

You invoke the archetype using the following command:

```
mvn archetype:create -DarchetypeGroupId=org.apache.servicemix.tooling -
DarchetypeArtifactId=spring-osgi-bundle-archetype -
DarchetypeVersion=2008.01.0.3-fuse -DgroupId=groupId -
DartifactId=artifactId -Dversion=version
```

## Apache CXF wsdl-first archetype

The Apache CXF wsdl-first archetype creates a project for creating a service from WSDL, as shown:

```
org.apache.servicemix.tooling/servicemix-osgi-cxf-wsdl-first-archetype/2008.01.0.3-fuse
```

You invoke the archetype using the following command:

```
mvn archetype:create -DarchetypeGroupId=org.apache.servicemix.tooling -DarchetypeArtifactId=servicemix-osgi-cxf-wsdl-first-archetype -DarchetypeVersion=2008.01.0.3-fuse -DgroupId=groupId -DartifactId=artifactId -Dversion=version
```

## Apache Camel archetype

The Apache Camel archetype creates a project for building a route that is deployed into JBoss Fuse, as shown:

```
org.apache.servicemix.tooling/servicemix-osgi-camel-archetype/2008.01.0.3-fuse
```

You invoke the archetype using the following command:

```
mvn archetype:create -DarchetypeGroupId=org.apache.servicemix.tooling -DarchetypeArtifactId=servicemix-osgi-camel-archetype -DarchetypeVersion=2008.01.0.3-fuse -DgroupId=groupId -DartifactId=artifactId -Dversion=version
```

## B.2. CONFIGURING THE BUNDLE PLUG-IN

### Overview

A bundle plug-in requires very little information to function. All of the required properties use default settings to generate a valid OSGi bundle.

While you can create a valid bundle using just the default values, you will probably want to modify some of the values. You can specify most of the properties inside the plug-in's `instructions` element.

### Configuration properties

Some of the commonly used configuration properties are:

- [Bundle-SymbolicName](#)
- [Bundle-Name](#)
- [Bundle-Version](#)
- [Export-Package](#)
- [Private-Package](#)

- [Import-Package](#)

## Setting a bundle's symbolic name

By default, the bundle plug-in sets the value for the `Bundle-SymbolicName` property to `groupId + ". " + artifactId`, with the following exceptions:

- If `groupId` has only one section (no dots), the first package name with classes is returned.

For example, if the group ID is `commons-logging:commons-logging`, the bundle's symbolic name is `org.apache.commons.logging`.

- If `artifactId` is equal to the last section of `groupId`, then `groupId` is used.

For example, if the POM specifies the group ID and artifact ID as `org.apache.maven:maven`, the bundle's symbolic name is `org.apache.maven`.

- If `artifactId` starts with the last section of `groupId`, that portion is removed.

For example, if the POM specifies the group ID and artifact ID as `org.apache.maven:maven-core`, the bundle's symbolic name is `org.apache.maven.core`.

To specify your own value for the bundle's symbolic name, add a `Bundle-SymbolicName` child in the plug-in's `instructions` element, as shown in [Example B.2](#).

### Example B.2. Setting a bundle's symbolic name

```
<plugin>
  <groupId>org.apache.felix</groupId>
  <artifactId>maven-bundle-plugin</artifactId>
  <configuration>
    <instructions>
      <Bundle-SymbolicName>${project.artifactId}</Bundle-SymbolicName>
      ...
    </instructions>
  </configuration>
</plugin>
```

## Setting a bundle's name

By default, a bundle's name is set to `${project.name}`.

To specify your own value for the bundle's name, add a `Bundle-Name` child to the plug-in's `instructions` element, as shown in [Example B.3](#).

### Example B.3. Setting a bundle's name

```
<plugin>
  <groupId>org.apache.felix</groupId>
  <artifactId>maven-bundle-plugin</artifactId>
  <configuration>
    <instructions>
```

```

    <Bundle-Name>JoeFred</Bundle-Name>
    ...
  </instructions>
</configuration>
</plugin>

```

## Setting a bundle's version

By default, a bundle's version is set to `${project.version}`. Any dashes ( - ) are replaced with dots ( . ) and the number is padded up to four digits. For example, `4.2-SNAPSHOT` becomes `4.2.0.SNAPSHOT`.

To specify your own value for the bundle's version, add a `Bundle-Version` child to the plug-in's `instructions` element, as shown in [Example B.4](#).

### Example B.4. Setting a bundle's version


```

<plugin>
  <groupId>org.apache.felix</groupId>
  <artifactId>maven-bundle-plugin</artifactId>
  <configuration>
    <instructions>
      <Bundle-Version>1.0.3.1</Bundle-Version>
      ...
    </instructions>
  </configuration>
</plugin>

```

## Specifying exported packages

By default, the OSGi manifest's `Export-Package` list is populated by all of the packages in your local Java source code (under `src/main/java`), *except* for the default package, `.`, and any packages containing `.impl` or `.internal`.



### IMPORTANT

If you use a `Private-Package` element in your plug-in configuration and you do not specify a list of packages to export, the default behavior includes only the packages listed in the `Private-Package` element in the bundle. No packages are exported.

The default behavior can result in very large packages and in exporting packages that should be kept private. To change the list of exported packages you can add an `Export-Package` child to the plug-in's `instructions` element.

The `Export-Package` element specifies a list of packages that are to be included in the bundle and that are to be exported. The package names can be specified using the `*` wildcard symbol. For example, the entry `com.fuse.demo.*` includes all packages on the project's classpath that start with `com.fuse.demo`.

You can specify packages to be excluded by prefixing the entry with `!`. For example, the entry `!com.fuse.demo.private` excludes the package `com.fuse.demo.private`.

When excluding packages, the order of entries in the list is important. The list is processed in order from the beginning and any subsequent contradicting entries are ignored.


For example, to include all packages starting with `com.fuse.demo` except the package `com.fuse.demo.private`, list the packages using:

```
!com.fuse.demo.private,com.fuse.demo.*
```

However, if you list the packages using `com.fuse.demo.*,!com.fuse.demo.private`, then `com.fuse.demo.private` is included in the bundle because it matches the first pattern.

## Specifying private packages

If you want to specify a list of packages to include in a bundle *without* exporting them, you can add a **Private-Package** instruction to the bundle plug-in configuration. By default, if you do not specify a **Private-Package** instruction, all packages in your local Java source are included in the bundle.



### IMPORTANT

If a package matches an entry in both the **Private-Package** element and the **Export-Package** element, the **Export-Package** element takes precedence. The package is added to the bundle and exported.

The **Private-Package** element works similarly to the **Export-Package** element in that you specify a list of packages to be included in the bundle. The bundle plug-in uses the list to find all classes on the project's classpath that are to be included in the bundle. These packages are packaged in the bundle, but not exported (unless they are also selected by the **Export-Package** instruction).

[Example B.5](#) shows the configuration for including a private package in a bundle

#### Example B.5. Including a private package in a bundle

```
<plugin>
  <groupId>org.apache.felix</groupId>
  <artifactId>maven-bundle-plugin</artifactId>
  <configuration>
    <instructions>
      <Private-Package>org.apache.cxf.wsdlFirst.impl</Private-Package>
      ...
    </instructions>
  </configuration>
</plugin>
```


## Specifying imported packages

By default, the bundle plug-in populates the OSGi manifest's **Import-Package** property with a list of all the packages referred to by the contents of the bundle.



While the default behavior is typically sufficient for most projects, you might find instances where you want to import packages that are not automatically added to the list. The default behavior can also result in unwanted packages being imported.

To specify a list of packages to be imported by the bundle, add an **Import -Package** child to the plug-in's **instructions** element. The syntax for the package list is the same as for the **Export -Package** element and the **Private -Package** element.



## IMPORTANT

When you use the **Import -Package** element, the plug-in does not automatically scan the bundle's contents to determine if there are any required imports. To ensure that the contents of the bundle are scanned, you must place an **\*** as the last entry in the package list.

**Example B.6** shows the configuration for specifying the packages imported by a bundle

### Example B.6. Specifying the packages imported by a bundle

```
<plugin>
  <groupId>org.apache.felix</groupId>
  <artifactId>maven-bundle-plugin</artifactId>
  <configuration>
    <instructions>
      <Import -Package>javax.jws,
        javax.wsdl,
        org.apache.cxf.bus,
        org.apache.cxf.bus.spring,
        org.apache.cxf.bus.resource,
        org.apache.cxf.configuration.spring,
        org.apache.cxf.resource,
        org.springframework.beans.factory.config,
        *
      </Import -Package>
      ...
    </instructions>
  </configuration>
</plugin>
```

## More information

For more information on configuring a bundle plug-in, see:

- ["Managing OSGi Dependencies"](#)
- [Apache Felix documentation](#)
- [Peter Kriens' aQute Software Consultancy web site](#)

## APPENDIX C. CONDUITS

### Abstract

Conduits are a low-level piece of the transport architecture that are used to implement outbound connections. Their behavior and life-cycle can effect system performance and processing load.

### OVERVIEW

Conduits manage the client-side, or outbound, transport details in the Apache CXF runtime. They are responsible for opening ports, establishing outbound connections, sending messages, and listening for any responses between an application and a single external endpoint. If an application connects to multiple endpoints, it will have one conduit instance for each endpoint.

Each transport type implements its own conduit using the `Conduit` interface. This allows for a standardized interface between the application level functionality and the transports.

In general, you only need to worry about the conduits being used by your application when configuring the client-side transport details. The underlying semantics of how the runtime handles conduits is, generally, not something a developer needs to worry about.

However, there are cases when an understanding of conduit's can prove helpful:

- Implementing a custom transport
- Advanced application tuning to manage limited resources

### CONDUIT LIFE-CYCLE

Conduits are managed by the client implementation object. Once created, a conduit lives for the duration of the client implementation object. The conduit's life-cycle is:

1. When the client implementation object is created, it is given a reference to a `ConduitSelector` object.
2. When the client needs to send a message is request's a reference to a conduit from the conduit selector.

If the message is for a new endpoint, the conduit selector creates a new conduit and passes it to the client implementation. Otherwise, it passes the client a reference to the conduit for the target endpoint.

3. The conduit sends messages when needed.
4. When the client implementation object is destroyed, all of the conduits associated with it are destroyed.

### CONDUIT WEIGHT

The weight of a conduit object depends on the transport implementation. HTTP conduits are extremely light weight. JMS conduits are heavy because they are associated with the `JMS Session` object and one or more `JMSListenerContainer` objects.

---

# INDEX

## A

**AcknowledgementInterval**, [Acknowledgement interval](#)

application source, [How WS-RM works](#)

**AtLeastOnce**, [Message delivery assurance policies](#)

**AtMostOnce**, [Message delivery assurance policies](#)

## B

**BaseRetransmissionInterval**, [Base retransmission interval](#)

**Bundle-Name**, [Setting a bundle's name](#)

**Bundle-SymbolicName**, [Setting a bundle's symbolic name](#)

**Bundle-Version**, [Setting a bundle's version](#)

### bundles

exporting packages, [Specifying exported packages](#)

importing packages, [Specifying imported packages](#)

name, [Setting a bundle's name](#)

private packages, [Specifying private packages](#)

symbolic name, [Setting a bundle's symbolic name](#)

version, [Setting a bundle's version](#)

## C

### configuration

HTTP consumer connection properties, [The client element](#)

HTTP consumer endpoint, [Using Configuration](#)

HTTP service provider connection properties, [The server element](#)

HTTP service provider endpoint, [Using Configuration](#)

**CreateSequence**, [How WS-RM works](#)

**CreateSequenceResponse**, [How WS-RM works](#)

## D

**driverClassName**, [Configuring WS-persistence](#)

## E

**ExponentialBackoff**, [Exponential backoff for retransmission](#)

**Export-Package**, [Specifying exported packages](#)

## H

### high availability

client configuration, [Add the clustering feature to your client configuration](#)

configuring random strategy, [Configuring a random strategy](#)

configuring static failover, [Overview](#)

enabling static failover, [Overview](#)

static failover, [HA with static failover](#)

`http-conf:authorization`, [The conduit element](#)

`http-conf:basicAuthSupplier`, [The conduit element](#)

`http-conf:client`, [The client element](#)

Accept, [The client element](#)

AcceptEncoding, [The client element](#)

AcceptLanguage, [The client element](#)

AllowChunking, [The client element](#)

AutoRedirect, [The client element](#)

BrowserType, [The client element](#)

CacheControl, [The client element](#), [Consumer Cache Control Directives](#)

Connection, [The client element](#)

ConnectionTimeout, [The client element](#)

ContentType, [The client element](#)

Cookie, [The client element](#)

DecoupledEndpoint, [The client element](#), [Configuring the consumer](#)

Host, [The client element](#)

MaxRetransmits, [The client element](#)

ProxyServer, [The client element](#)

ProxyServerPort, [The client element](#)

ProxyServerType, [The client element](#)

ReceiveTimeout, [The client element](#)

Referer, [The client element](#)

`http-conf:conduit`, [The conduit element](#)

name attribute, [The conduit element](#)

`http-conf:contextMatchStrategy`, [The destination element](#)

---

**http-conf:destination**, [The destination element](#)

**name attribute**, [The destination element](#)

**http-conf:fixedParameterOrder**, [The destination element](#)

**http-conf:proxyAuthorization**, [The conduit element](#)

**http-conf:server**, [The destination element](#), [The server element](#)

**CacheControl**, [The server element](#), [Service Provider Cache Control Directives](#)

**ContentEncoding**, [The server element](#)

**ContentLocation**, [The server element](#)

**ContentType**, [The server element](#)

**HonorKeepAlive**, [The server element](#)

**ReceiveTimeout**, [The server element](#)

**RedirectURL**, [The server element](#)

**ServerType**, [The server element](#)

**SuppressClientReceiveErrors**, [The server element](#)

**SuppressClientSendErrors**, [The server element](#)

**http-conf:tlsClientParameters**, [The conduit element](#)

**http-conf:trustDecider**, [The conduit element](#)

## I

**Import-Package**, [Specifying imported packages](#)

**InOrder**, [Message delivery assurance policies](#)

## J

**jaxws:binding**, [Elements](#), [Adding functionality](#)

**jaxws:client**

**abstract**, [Basic Configuration Properties](#)

**address**, [Basic Configuration Properties](#)

**bindingId**, [Basic Configuration Properties](#)

**bus**, [Basic Configuration Properties](#)

**createdFromAPI**, [Basic Configuration Properties](#)

**depends-on**, [Basic Configuration Properties](#)

**endpointName**, [Basic Configuration Properties](#)

**name**, [Basic Configuration Properties](#)

**password**, [Basic Configuration Properties](#)

`serviceClass`, [Basic Configuration Properties](#)

`serviceName`, [Basic Configuration Properties](#)

`username`, [Basic Configuration Properties](#)

`wSDLLocation`, [Basic Configuration Properties](#)

`jaxws:conduitSelector`, [Adding functionality](#)

`jaxws:dataBinding`, [Elements](#), [Adding functionality](#)

`jaxws:endpoint`

`abstract`, [Attributes](#)

`address`, [Attributes](#)

`bindingUri`, [Attributes](#)

`bus`, [Attributes](#)

`createdFromAPI`, [Attributes](#)

`depends-on`, [Attributes](#)

`endpointName`, [Attributes](#)

`id`, [Attributes](#)

`implementor`, [Attributes](#)

`implementorClass`, [Attributes](#)

`name`, [Attributes](#)

`publish`, [Attributes](#)

`publishedEndpointUrl`, [Attributes](#)

`serviceName`, [Attributes](#)

`wSDLLocation`, [Attributes](#)

`jaxws:executor`, [Elements](#)

`jaxws:features`, [Elements](#), [Adding functionality](#)

`jaxws:handlers`, [Elements](#), [Adding functionality](#)

`jaxws:inFaultInterceptors`, [Elements](#), [Adding functionality](#)

`jaxws:inInterceptors`, [Elements](#), [Adding functionality](#)

`jaxws:invoker`, [Elements](#)

`jaxws:outFaultInterceptors`, [Elements](#), [Adding functionality](#)

`jaxws:outInterceptors`, [Elements](#), [Adding functionality](#)

`jaxws:properties`, [Elements](#), [Adding functionality](#)

`jaxws:server`

`abstract`, [Attributes](#)

---

address, [Attributes](#)

bindingId, [Attributes](#)

bus, [Attributes](#)

createdFromAPI, [Attributes](#)

depends-on, [Attributes](#)

endpointName, [Attributes](#)

id, [Attributes](#)

name, [Attributes](#)

publish, [Attributes](#)

serviceBean, [Attributes](#)

serviceClass, [Attributes](#)

serviceName, [Attributes](#)

wsdlLocation, [Attributes](#)

jaxws:serviceFactory, [Elements](#)

JMS

specifying the message type, [Specifying the message type](#)

JMS destination

specifying, [Specifying the JMS address](#)

jms:address, [Specifying the JMS address](#)

connectionPassword attribute, [Specifying the JMS address](#)

connectionUserName attribute, [Specifying the JMS address](#)

destinationStyle attribute, [Specifying the JMS address](#)

jmsDestinationName attribute, [Specifying the JMS address](#)

jmsReplyDestinationName attribute, [Using a Named Reply Destination](#)

jmsReplyDestinationName attribute, [Specifying the JMS address](#)

jndiConnectionFactoryName attribute, [Specifying the JMS address](#)

jndiDestinationName attribute, [Specifying the JMS address](#)

jndiReplyDestinationName attribute, [Specifying the JMS address](#) , [Using a Named Reply Destination](#)

jms:client, [Specifying the message type](#)

messageType attribute, [Specifying the message type](#)

jms:JMSNamingProperties, [Specifying JNDI properties](#)

**jms:server**, [Specifying the configuration](#)

**durableSubscriberName**, [Specifying the configuration](#)

**messageSelector**, [Specifying the configuration](#)

**transactional**, [Specifying the configuration](#)

**useMessageIDAsCorrealationID**, [Specifying the configuration](#)

**JMSConfiguration**, [Specifying the configuration](#)

**JNDI**

**specifying the connection factory**, [Specifying the JMS address](#)

**M**

**Maven archetypes**, [Useful Maven archetypes](#)

**Maven tooling**

**adding the bundle plug-in**, [Adding a bundle plug-in](#)

**maxLength**, [Maximum length of an RM sequence](#)

**maxUnacknowledged**, [Maximum unacknowledged messages threshold](#)

**N**

**named reply destination**

**specifying in WSDL**, [Specifying the JMS address](#)

**using**, [Using a Named Reply Destination](#)

**O**

**osgi install**, [Deploying from the console](#)

**osgi refresh**, [Refreshing an application](#)

**osgi start**, [Deploying from the console](#)

**osgi stop**, [Stopping an application](#)

**osgi uninstall**, [Uninstalling an application](#)

**P**

**passWord**, [Configuring WS-persistence](#)

**Private-Package**, [Specifying private packages](#)

**R**

**random strategy**, [Configuring a random strategy](#)

**replicated services**, [Overview](#)



---

**RMAssertion**, [WS-Policy RMAssertion Children](#)

## S

**Sequence**, [How WS-RM works](#)

**SequenceAcknowledgment**, [How WS-RM works](#)

**static failover**, [HA with static failover](#)

    configuring, [Overview](#)

    enabling, [Overview](#)

## U

**userName**, [Configuring WS-persistence](#)

## W

**WS-Addressing**

    using, [Configuring an endpoint to use WS-Addressing](#)

**WS-RM**

**AcknowledgementInterval**, [Acknowledgement interval](#)

**AtLeastOnce**, [Message delivery assurance policies](#)

**AtMostOnce**, [Message delivery assurance policies](#)

**BaseRetransmissionInterval**, [Base retransmission interval](#)

    configuring, [Configuring WS-RM](#)

    destination, [How WS-RM works](#)

    driverClassName, [Configuring WS-persistence](#)

    enabling, [Enabling WS-RM](#)

**ExponentialBackoff**, [Exponential backoff for retransmission](#)

    externalL attachment, [External attachment](#)

    initial sender, [How WS-RM works](#)

**InOrder**, [Message delivery assurance policies](#)

    interceptors, [Apache CXF WS-RM Interceptors](#)

    maxLength, [Maximum length of an RM sequence](#)

    maxUnacknowledged, [Maximum unacknowledged messages threshold](#)

    password, [Configuring WS-persistence](#)

    rmManager, [Children of the rmManager Spring bean](#)

    source, [How WS-RM works](#)

ultimate receiver, [How WS-RM works](#)

url, [Configuring WS-persistence](#)

userName, [Configuring WS-persistence](#)

wsam:Addressing, [Configuring an endpoint to use WS-Addressing](#)

#### WSDL extensors

jms:address (see jms:address)

jms:client (see jms:client)

jms:JMSNamingProperties (see jms:JMSNamingProperties)

jms:server (see jms:server)

wstrm:AcksTo, [How WS-RM works](#)

wswa:UsingAddressing, [Configuring an endpoint to use WS-Addressing](#)