

Progress | Artix.

PROGRESS[®] ARTIX[®]

Java Router, Defining Routes

Version 5.6, August 2011

Progress Software

Publication date 11 Aug 2011

Copyright © 2001-2011 Progress Software Corporation and/or its subsidiaries or affiliates.

Legal Notices

These materials and all Progress software products are copyrighted and all rights are reserved by Progress Software Corporation. The information in these materials is subject to change without notice, and Progress Software Corporation assumes no responsibility for any errors that may appear therein. The references in these materials to specific platforms supported are subject to change.

Actional, Apama, Artix, Business Empowerment, DataDirect (and design), DataDirect Connect, DataDirect Connect64, DataDirect Technologies, DataDirect XML Converters, DataDirect XQuery, DataXtend, Dynamic Routing Architecture, EdgeXtend, Empowerment Center, Fathom, Fuse Mediation Router, Fuse Message Broker, Fuse Services Framework, IntelliStream, IONA, Making Software Work Together, Mindreef, ObjectStore, OpenEdge, Orbix, PeerDirect, POSSENET, Powered by Progress, PowerTier, Progress, Progress DataXtend, Progress Dynamics, Progress Business Empowerment, Progress Empowerment Center, Progress Empowerment Program, Progress OpenEdge, Progress Profiles, Progress Results, Progress Software Developers Network, Progress Sonic, ProVision, PS Select, Savvion, SequelLink, Shadow, SOAPscope, SOAPStation, Sonic, Sonic ESB, SonicMQ, Sonic Orchestration Server, SpeedScript, Stylus Studio, Technical Empowerment, WebSpeed, Xcalia (and design), and Your Software, Our Technology-Experience the Connection are registered trademarks of Progress Software Corporation or one of its affiliates or subsidiaries in the U.S. and/or other countries. AccelEvent, Apama Dashboard Studio, Apama Event Manager, Apama Event Modeler, Apama Event Store, Apama Risk Firewall, AppsAlive, AppServer, ASPen, ASP-in-a-Box, BusinessEdge, Business Making Progress, Cache-Forward, CloudEdge, DataDirect Spy, DataDirect SupportLink, Fuse, FuseSource, Future Proof, GVAC, High Performance Integration, ObjectStore Inspector, ObjectStore Performance Expert, OpenAccess, Orbacus, Pantero, POSSE, ProDataSet, Progress Arcade, Progress CloudEdge, Progress Control Tower, Progress ESP Event Manager, Progress ESP Event Modeler, Progress Event Engine, Progress RFID, Progress RPM, Progress Software Business Making Progress, PSE Pro, SectorAlliance, SeeThinkAct, Shadow z/Services, Shadow z/Direct, Shadow z/Events, Shadow z/Presentation, Shadow Studio, SmartBrowser, SmartComponent, SmartDataBrowser, SmartDataObjects, SmartDataView, SmartDialog, SmartFolder, SmartFrame, SmartObjects, SmartPanel, SmartQuery, SmartViewer, SmartWindow, Sonic Business Integration Suite, Sonic Process Manager, Sonic Collaboration Server, Sonic Continuous Availability Architecture, Sonic Database Service, Sonic Workbench, Sonic XML Server, The Brains Behind BAM, WebClient, and Who Makes Progress are trademarks or service marks of Progress Software Corporation and/or its subsidiaries or affiliates in the U.S. and other countries. Java is a registered trademark of Oracle and/or its affiliates. Any other marks contained herein may be trademarks of their respective owners.

Third Party Acknowledgements -- See [Third Party Acknowledgements on page 15](#).

Table of Contents

Preface	11
Open Source Project Resources	12
Document Conventions	13
Third Party Acknowledgements	15
Defining Routes in Java DSL	17
Implementing a RouteBuilder Class	18
Basic Java DSL Syntax	20
Processors	24
Languages for Expressions and Predicates	30
Transforming Message Content	35
Defining Routes in XML	41
Using the Router Schema in an XML File	42
Defining a Basic Route in XML	44
Processors	45
Languages for Expressions and Predicates	52
Transforming Message Content	54
Basic Principles of Route Building	57
Pipeline Processing	58
Multiple Inputs	63
Exception Handling	68
Bean Integration	70

List of Figures

1. Local Routing Rules	21
2. Processor Modifying an In Message	58
3. Processor Creating an Out Message	59
4. Sample Pipeline for InOnly Exchanges	59
5. Sample Pipeline for InOut Exchanges	60
6. Example of Interceptor Chaining	61
7. Pipeline Alternative to Interceptor Chaining	62
8. Processing Multiple Inputs with Segmented Routes	63
9. Processing Multiple Inputs with a Content Enricher	66

List of Tables

1. Properties for Simple Language	31
2. Transformation Methods from the ProcessorType Class	36
3. Methods from the Builder Class	37
4. Modifier Methods from the ValueBuilder Class	38
5. Elements for Expression and Predicate Languages	52
6. Basic Bean Annotations	73
7. Expression Language Annotations	74

List of Examples

1. Implementation of a RouteBuilder Class	18
2. Implementing a Custom Processor Class	29
3. Simple Transformation of Incoming Messages	35
4. Specifying the Router Schema Location	42
5. Router Schema in a Spring Configuration File	42
6. Basic Route in XML	44

Preface

Open Source Project Resources	12
Document Conventions	13
Third Party Acknowledgements	15

Open Source Project Resources

Apache Incubator CXF

Web site: <http://cxf.apache.org/>

User's list: <user@cxf.apache.org>

Apache Tomcat

Web site: <http://tomcat.apache.org/>

User's list: <users@tomcat.apache.org>

Apache ActiveMQ

Web site: <http://activemq.apache.org/>

User's list: <users@activemq.apache.org>

Apache Camel

Web site: <http://camel.apache.org/>

User's list: <users@camel.apache.org>

Document Conventions

Typographical conventions

This book uses the following typographical conventions:

<code>fixed width</code>	<p>Fixed width (Courier font) in normal text represents portions of code and literal names of items such as classes, functions, variables, and data structures. For example, text might refer to the <code>javax.xml.ws.Endpoint</code> class.</p> <p>Constant width paragraphs represent code examples or information a system displays on the screen. For example:</p> <pre>import java.util.logging.Logger;</pre>
<i>Fixed width italic</i>	<p>Fixed width italic words or characters in code and commands represent variable values you must supply, such as arguments to commands or path names for your particular system. For example:</p> <pre>% cd /users/<i>YourUserName</i></pre>
<i>Italic</i>	<p>Italic words in normal text represent <i>emphasis</i> and introduce <i>new terms</i>.</p>
Bold	<p>Bold words in normal text represent graphical user interface components such as menu commands and dialog boxes. For example: the User Preferences dialog.</p>

Keying conventions

This book uses the following keying conventions:

No prompt	When a command's format is the same for multiple platforms, the command prompt is not shown.
%	A percent sign represents the UNIX command shell prompt for a command that does not require root privileges.
#	A number sign represents the UNIX command shell prompt for a command that requires root privileges.
>	The notation > represents the MS-DOS or Windows command prompt.

...	Horizontal or vertical ellipses in format and syntax descriptions indicate that material has been eliminated to simplify a discussion.
[]	Brackets enclose optional items in format and syntax descriptions.
{ }	Braces enclose a list from which you must choose an item in format and syntax descriptions.
	In format and syntax descriptions, a vertical bar separates items in a list of choices enclosed in {} (braces).

Admonition conventions

This book uses the following conventions for admonitions:

	Notes display information that may be useful, but not critical.
	Tips provide hints about completing a task or using a tool. They may also provide information about workarounds to possible problems.
	Important notes display information that is critical to the task at hand.
	Cautions display information about likely errors that can be encountered. These errors are unlikely to cause damage to your data or your systems.
	Warnings display information about errors that may cause damage to your systems. Possible damage from these errors include system failures and loss of data.

Third Party Acknowledgements

Progress Artix ESB v5.6 incorporates Apache Commons Codec v1.2 from The Apache Software Foundation. Such technology is subject to the following terms and conditions: The Apache Software License, Version 1.1 - Copyright (c) 2001-2003 The Apache Software Foundation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The end-user documentation included with the redistribution, if any, must include the following acknowledgement: "This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>)."
Alternately, this acknowledgement may appear in the software itself, if and wherever such third-party acknowledgements normally appear.
4. The names "Apache", "The Jakarta Project", "Commons", and "Apache Software Foundation" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact apache@apache.org.
5. Products derived from this software may not be called "Apache", "Apache" nor may "Apache" appear in their name without prior written permission of the Apache Software Foundation.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org/>.

Progress Artix ESB v5.6 incorporates Jcraft JSCH v0.1.44 from Jcraft. Such technology is subject to the following terms and conditions: Copyright (c) 2002-2010 Atsuhiko Yamanaka, JCraft, Inc. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met: 1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer. 2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution. 3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED ``AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL JCRAFT, INC. OR ANY CONTRIBUTORS TO THIS SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Defining Routes in Java DSL

You can define routing rules in Java, using a domain specific language (DSL). The routing rules represent the core of a router application and Java DSL is currently the most flexible way to define them.

Implementing a RouteBuilder Class	18
Basic Java DSL Syntax	20
Processors	24
Languages for Expressions and Predicates	30
Transforming Message Content	35

Implementing a RouteBuilder Class

Overview

In Java Router, you define routes by implementing a `RouteBuilder` class. You must override a single method, `RouteBuilder.configure()`, and in this method define the routing rules you want to associate with the `RouteBuilder`. The rules themselves are defined using a *Domain Specific Language* (DSL), which is implemented as a Java API.

You can define as many `RouteBuilder` classes as you like in a router application. Ultimately, each `RouteBuilder` class must get instantiated once and registered with the `CamelContext` object. Normally, however, the lifecycle of the `RouteBuilder` objects is managed automatically by the container in which you deploy the router. The core task for a router developer is simply to implement one or more `RouteBuilder` classes.

RouteBuilder class

The `org.apache.camel.builder.RouteBuilder` class is the base class for implementing your own route builder types. It defines an abstract method, `configure()`, that you *must* override in your derived implementation class. In addition, `RouteBuilder` also defines methods that are used to initiate the routing rules (for example, `from()`, `intercept()`, and `exception()`).

Implementing a RouteBuilder

[Example 1 on page 18](#) shows an example of a simpler `RouteBuilder` implementation. You need only define a single method, `configure()`, which contains a list of routing rules (one Java statement for each rule).

Example 1. Implementation of a RouteBuilder Class

```
import org.apache.camel.builder.RouteBuilder;

public class MyRouteBuilder extends RouteBuilder {

    public void configure() {
        // Define routing rules here:
        from("file:src/data?noop=true").to("file:target/messages");

        // More rules can be included, in you like.
        // ...
    }
}
```

Where the rule of the form `from(URL1).to(URL2)` instructs the router to read messages from the file system located in directory, `src/data`, and send them

to files located in the directory, `target/messages`. The option, `?noop=true`, specifies that the source messages are *not* to be deleted from the `src/data` directory.

Basic Java DSL Syntax

What is a DSL?

A Domain Specific Language (DSL) is essentially a mini-language designed for a special purpose. The DSL is not required to be logically complete; it need only have enough expressive power to describe problems adequately in the chosen domain.

Typically, a DSL does *not* require a dedicated parser, interpreter, or compiler. You can piggyback a DSL on top of an existing object-oriented host language by observing that it is possible to map an API in a host language to a specialized language syntax: that is, a sequence of commands in the DSL maps to a chain of method invocations in the host language. For example, a sequence of commands in some hypothetical DSL that might look like this:

```
command01;  
command02;  
command03;
```

Can be mapped to a chain of Java invocations, like this:

```
command01().command02().command03()
```

You could even define blocks, for example:

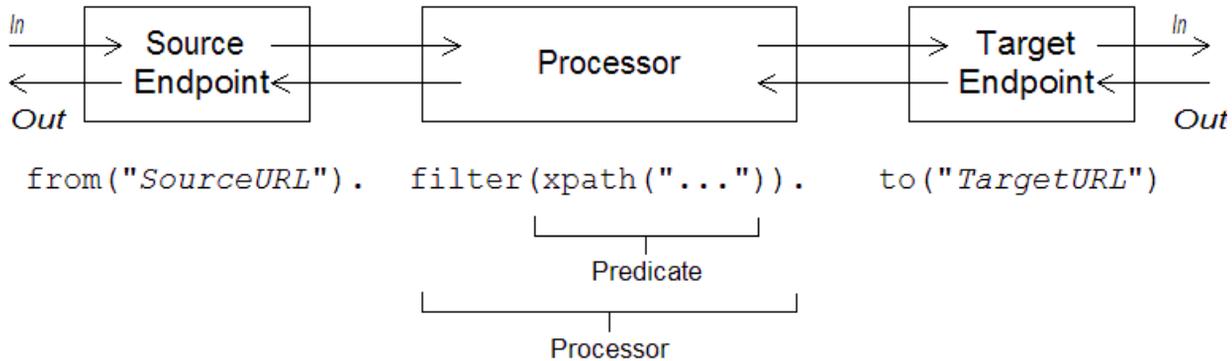
```
command01().startBlock().command02().command03().endBlock()
```

The syntax of the DSL is implicitly defined by the type system of the specialized API. For example, the return type of a method determines which methods can legally be invoked next (equivalent to the next command in the DSL).

Router rule syntax

The Java Router defines a *router DSL* for defining routing rules. You can use this DSL to define rules in the body of a `RouteBuilder.configure()` implementation. [Figure 1 on page 21](#) shows an overview of the basic syntax for defining local routing rules.

Figure 1. Local Routing Rules



A local rule always starts with a `from("EndpointURL")` method, which specifies the source of messages for the routing rule. You can then add an arbitrarily long chain of processors to the rule (for example, `filter()`), finishing off the rule with a `to("EndpointURL")` method, which specifies the target for the messages that pass through the rule. It is not always necessary to end a rule with `to()`, however. There are alternative ways of specifying the message target in a rule.



Note

It is also possible to define a global routing rule, by starting the rule with a special processor type (such as `intercept()`, `exception()`, `errorHandler()`, and so on). This kind of rule lies outside the scope of the *Getting Started* guide.

Sources and targets

A local rule always starts by defining a source endpoint, using `from("EndpointURL")`, and typically (but not always) ends by defining a target endpoint, using `to("EndpointURL")`. The endpoint URLs, `EndpointURL`, can use any of the components configured at deploy time. For example, you could use a file endpoint, `file:MyMessageDirectory`, a CXF endpoint, `cxf:MyServiceName`, or an ActiveMQ endpoint,

`activemq:queue:MyQName`. For a complete list of component types, see <http://camel.apache.org/components.html>¹.

Processors

A *processor* is a method that can access and modify the stream of messages passing through a rule. If a message is part of a remote procedure call (*InOut* call), the processor can potentially act on the messages flowing in *both* directions: on the request messages, flowing from source to target, and on the reply messages, flowing from target back to source (see [Message exchanges on page 22](#)). Processors can take *expression* or *predicate* arguments, that modify their behavior. For example, the rule shown in [Figure 1 on page 21](#) includes a `filter()` processor that takes an `xpath()` predicate as its argument.

Expressions and predicates

Expressions (evaluating to strings or other data types) and predicates (evaluating to true or false) occur frequently as arguments to the built-in processor types. You do not have to worry much about which type to pass to an expression argument, because they are usually automatically converted to the type you need. For example, you can usually just pass a string into an expression argument. Predicate expressions are useful for defining conditional behaviour in a route. For example, the following filter rule propagates *In* messages, only if the `foo` header is equal to the value `bar`:

```
from("seda:a").filter(header("foo").isEqualTo("bar")).to("seda:b");
```

Where the filter is qualified by the predicate, `header("foo").isEqualTo("bar")`. To construct more sophisticated predicates and expressions, based on the message content, you can use one of the expression and predicate languages (see [Languages for Expressions and Predicates on page 30](#)).

Message exchanges

When a router rule is activated, it can process messages passing in either direction: that is, from source to target or from target back to source. For example, if a router rule is mediating a remote procedure call (RPC), the rule would process requests, replies, and faults. How do you manage message correlation in this case? One of the most effective and straightforward ways is to use a *message exchange* object as the basis for processing messages. Java Router uses message exchange objects (of `org.apache.camel.Exchange` type) in its API for processing router rules.

¹ <http://activemq.apache.org/camel/components.html>

The basic idea of the message exchange is that, instead of accessing requests, replies, and faults separately, you encapsulate the correlated messages inside a single object (an `Exchange` object). Message correlation now becomes trivial from the perspective of a processor, because correlated messages are encapsulated in a single `Exchange` object and processors gain access to messages through the `Exchange` object.

Using an `Exchange` object makes it easy to generalize message processing to different kinds of *message exchange pattern*. For example, an asynchronous protocol might define a message exchange pattern that consists of a single message that flows from the source to the target (an *In* message). An RPC protocol, on the other hand, might define a message exchange pattern that consists of a request message correlated with either a reply or fault message. Currently, Java Router supports the following message exchange patterns:

- `InOnly`
- `RobustInOnly`
- `InOut`
- `InOptionalOut`
- `OutOnly`
- `RobustOutOnly`
- `OutIn`
- `OutOptionalIn`

Where these message exchange patterns are represented by constants in the enumeration type, `org.apache.camel.ExchangePattern`.

Processors

Overview

To enable the router to do something more interesting than simply connecting a source endpoint to a target endpoint, you can add *processors* to your route. A processor is a command you can insert into a routing rule in order to perform arbitrary processing of the messages that flow through the rule. Java Router provides a wide variety of different processors, as follows:

- [Filter on page 24.](#)
- [Choice on page 25.](#)
- [Pipeline on page 25](#)
- [Recipient list on page 25.](#)
- [Splitter on page 26.](#)
- [Aggregator on page 26.](#)
- [Resequencer on page 26.](#)
- [Throttler on page 27.](#)
- [Delayer on page 27.](#)
- [Load balancer on page 28](#)
- [Custom processor on page 28.](#)

Filter

The `filter()` processor can be used to prevent uninteresting messages from reaching the target endpoint. It takes a single predicate argument: if the predicate is true, the message exchange is allowed through to the target; if the predicate is false, the message exchange is blocked. For example, the following filter blocks a message exchange, unless the incoming message contains a header, `foo`, with value equal to `bar`:

```
from("SourceURL").filter(header("foo").isEqualTo("bar")).to("TargetURL");
```

Choice

The `choice()` processor is a conditional statement that is used to route incoming messages to alternative targets. The alternative targets are each preceded by a `when()` method, which takes a predicate argument. If the predicate is true, the following target is selected, otherwise processing proceeds to the next `when()` method in the rule. For example, the following `choice()` processor directs incoming messages to either `Target1`, `Target2`, or `Target3`, depending on the values of `Predicate1` and `Predicate2`:

```
from("SourceURL").choice().when(Predicate1).to("Target1")
                        .when(Predicate2).to("Target2")
                        .otherwise().to("Target3");
```

Pipeline

The `pipeline()` processor is used to link together a chain of targets, where the output of one target is fed into the input of the next target in the pipeline (analogous to the UNIX `pipe` command). The `pipeline()` method takes an arbitrary number of endpoint arguments, which specify the sequence of endpoints in the pipeline. For example, to pass messages from `SourceURL` to `Target1` to `Target2` to `Target3` in a pipeline, you could use the following rule:

```
from("SourceURL").pipeline("Target1","Target2","Target3");
```

Recipient list

If you want the messages from a source endpoint, `SourceURL`, to be sent to more than one target, there are two alternative approaches you can use. One approach is to invoke the `to()` method with multiple target endpoints (static recipient list), for example:

```
from("SourceURL").to("Target1","Target2","Target3");
```

The alternative approach is to invoke the `recipientList()` processor, which takes a list of recipients as its argument (dynamic recipient list). The advantage of the `recipientList()` processor is that the list of recipients can be calculated at runtime. For example, the following rule generates a recipient list by reading the contents of the `recipientListHeader` from the incoming message:

```
from("SourceURL").recipientList(header("recipientListHeader").tokenize(", "));
```

Splitter

The `splitter()` processor is used to split a message into parts, which are then processed as separate messages. The `splitter()` method takes a list argument, where each item in the list represents a message part that is to be re-sent as a separate message. For example, the following rule splits the body of an incoming message into separate lines and then sends each line to the target in a separate message:

```
from("SourceURL").splitter(bodyAs(String.class).tokenize("\n")).to("TargetURL");
```

Aggregator

The `aggregator()` processor is used to aggregate related incoming messages into a single message. In order to distinguish which messages are eligible to be aggregated together, you need to define a *correlation key* for the aggregator. The correlation key is normally derived from a field in the message (for example, a header field). Messages that have the same correlation key value are eligible to be aggregated together. You can also optionally specify an aggregation algorithm to the `aggregator()` processor (the default algorithm is to pick the latest message with a given value of the correlation key and to discard the older messages with that correlation key value).

For example, if you are monitoring a data stream that reports stock prices in real time, you might only be interested in the *latest* price of each stock symbol. In this case, you could configure an aggregator to transmit only the latest price for a given stock and discard the older (out-of-date) price notifications. The following rule implements this functionality, where the correlation key is read from the `stockSymbol` header and the default aggregator algorithm is used:

```
from("SourceURL").aggregator(header("stockSymbol")).to("TargetURL");
```

Resequencer

A `resequencer()` processor is used to re-arrange the order in which incoming messages are transmitted. The `resequencer()` method takes a sequence number as its argument (where the sequence number is calculated from the contents of a field in the incoming message). Naturally, before you can start re-ordering messages, you need to wait until a certain number of messages have been received from the source. There are a couple of different ways to specify how long the `resequencer()` processor should wait before attempting

to re-order the accumulated messages and forward them to the target, as follows:

- *Batch resequencing*—(the default) wait until a specified number of messages have accumulated before starting to re-order and forward messages. This processing option is specified by invoking `resequencer().batch()`. For example, the following resequencing rule would re-order messages based on the `timeOfDay` header, waiting until at least 300 messages have accumulated or 4000 ms have elapsed since the last message received.

```
from("SourceURL").resequencer(header("timeOfDay")).batch(new
BatchResequencerConfig(300, 4000L)).to("TargetURL");
```

- *Stream resequencing*—transmit messages as soon as they arrive *unless* the resequencer detects a gap in the incoming message stream (missing sequence numbers), in which case the resequencer waits until the missing messages arrive and then forwards the messages in the correct order. To avoid the resequencer blocking forever, you can specify a timeout (default is 1000 ms), after which time the message sequence is transmitted with unresolved gaps. For example, the following resequencing rule detects gaps in the message stream by monitoring the value of the `sequenceNumber` header, where the maximum buffer size is limited to 5000 and the timeout is specified to be 4000 ms:

```
from("SourceURL").resequencer(header("sequenceNum
ber")).stream(new StreamResequencerConfig(5000,
4000L)).to("TargetURL");
```

Throttler

The `throttler()` processor is used to ensure that a target endpoint does not get overloaded. The throttler works by limiting the number of messages that can pass through per second. If the incoming messages exceed the specified rate, the throttler accumulates excess messages in a buffer and transmits them more slowly to the target endpoint. For example, to limit the rate of throughput to 100 messages per second, you can define the following rule:

```
from("SourceURL").throttler(100).to("TargetURL");
```

Delayer

The `delayer()` processor is used to hold up messages for a specified length of time. The delay can either be relative (wait a specified length of time after

receipt of the incoming message) or absolute (wait until a specific time). For example, to add a delay of 2 seconds before transmitting received messages, you can use the following rule:

```
from("SourceURL").delayer(2000).to("TargetURL");
```

To wait until the absolute time specified in the `processAfter` header, you can use the following rule:

```
from("SourceURL").delayer(header("processAfter")).to("TargetURL");
```

The `delayer()` method is overloaded, such that an integer is interpreted as a relative delay and an expression (for example, a string) is interpreted as an absolute delay.

Load balancer

The `loadBalance()` processor is used to load balance message exchanges over a list of target endpoints. It is possible to customize the load balancing strategy. For example, to load balance incoming messages exchanges using a round robin algorithm (each endpoint in the target list is tried in sequence), you can use the following rule:

```
from("SourceURL").loadBalance().roundRobin().to("TargetURL_01",  
"TargetURL_02", "TargetURL_03");
```

Alternatively, you can customize the load balancing algorithm by implementing your own `LoadBalancer` class, as follows:

```
public class MyLoadBalancer implements org.apache.camel.pro  
cessor.loadbalancer.LoadBalancer {  
    ...  
};  
  
from("SourceURL").loadBalance().setLoadBalancer(new MyLoadBal  
ancer())  
    .to("TargetURL_01", "TargetURL_02", "TargetURL_03");
```

Custom processor

If none of the standard processors described here provide the functionality you need, you can always define your own custom processor. To create a custom processor, define a class that implements the `org.apache.camel.Processor` interface and override the `process()` method in this class. For example, the following custom processor, `MyProcessor`, removes the header named `foo` from incoming messages:

Example 2. Implementing a Custom Processor Class

```
public class MyProcessor implements org.apache.camel.Processor
{
    public void process(org.apache.camel.Exchange exchange)
    {
        inMessage = exchange.getIn();
        if (inMessage != null) {
            inMessage.removeHeader("foo");
        }
    }
};
```

To insert the custom processor into a router rule, invoke the `process()` method, which provides a generic mechanism for inserting processors into rules. For example, the following rule invokes the processor defined in [Example 2 on page 29](#):

```
org.apache.camel.Processor myProc = new MyProcessor();
from("SourceURL").process(myProc).to("TargetURL");
```

Languages for Expressions and Predicates

Overview

To provide you with greater flexibility when parsing and processing messages, Java Router supports language plug-ins for various scripting languages. For example, if an incoming message is formatted as XML, it is relatively easy to extract the contents of particular XML elements or attributes from the message using a language such as XPath. The Java Router implements script builder classes, which encapsulate the imported languages. Each languages is accessed through a static method that takes a script expression as its argument, processes the current message using that script, and then returns an expression or a predicate. In order to be usable as an expression or a predicate, the script builder classes implement the following interfaces:

```
org.apache.camel.Expression<E>
org.apache.camel.Predicate<E>
```

In addition to this, the `ScriptBuilder` class (which wraps scripting languages such as JavaScript, and so on) inherits from the following interface:

```
org.apache.camel.Processor
```

Which implies that the languages associated with the `ScriptBuilder` class can also be used as message processors (see [Custom processor on page 28](#)).

Simple

The simple language is a very limited expression language that is built into the router core. This language can be useful, if you need to eliminate dependancies on third-party libraries whilst testing. Otherwise, you should use one of the other languages. To use the simple language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.language.simple.SimpleLanguage.simple;
```

The simple language provides various elementary expressions that return different parts of a message exchange. For example, the expression, `simple("header.timeOfDay")`, would return the contents of a header called `timeOfDay` from the incoming message. You can also construct predicates by testing expressions for equality. For example, the predicate, `simple("header.timeOfDay = '14:30'")`, tests whether the `timeOfDay` header in the incoming message is equal to `14:30`. [Table 1 on page 31](#) shows the list of elementary expressions supported by the simple language.

Table 1. Properties for Simple Language

Elementary Expression	Description
<code>body</code>	Access the body of the incoming message.
<code>out.body</code>	Access the body of the outgoing message.
<code>header.HeaderName</code>	Access the contents of the <i>HeaderName</i> header from the incoming message.
<code>out.header.HeaderName</code>	Access the contents of the <i>HeaderName</i> header from the outgoing message.
<code>property.PropertyName</code>	Access the <i>PropertyName</i> property on the exchange.

XPath

The `xpath()` static method parses message content using the XPath language (to learn about XPath, see the W3 Schools tutorial, <http://www.w3schools.com/xpath/default.asp>). To use the XPath language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.builder.xml.XPathBuilder.xpath;
```

You can pass an XPath expression to `xpath()` as a string argument. The XPath expression implicitly acts on the message content and returns a node set as its result. Depending on the context, the return value is interpreted either as a predicate (where an empty node set is interpreted as false) or an expression. For example, if you are processing an XML message with the following content:

```
<person user="paddington">
  <firstName>Paddington</firstName>
  <lastName>Bear</lastName>
  <city>London</city>
</person>
```

You could choose which target endpoint to route the message to, based on the content of the `city` element, using the following rule:

```
from("file:src/data?noop=true") .
  choice() .
    when(xpath("/person/city = 'London'")) .to("file:target/messages/uk") .
    otherwise() .to("file:target/messages/others");
```

Where the return value of `xpath()` is treated as a predicate in this example.

XQuery

The `xquery()` static method parses message content using the XQuery language (to learn about XQuery, see the W3 Schools tutorial, <http://www.w3schools.com/xquery/default.asp>). XQuery is a superset of the XPath language; hence, any valid XPath expression is also a valid XQuery expression. To use the XQuery language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.builder.saxon.XQueryBuilder.xquery;
```

You can pass an XQuery expression to `xquery()` in several different ways. For simple expressions, you can pass the XQuery expressions as a string, `java.lang.String`. For longer XQuery expressions, on the other hand, you might prefer to store the expression in a file, which you can then reference by passing a `java.io.File` argument or a `java.net.URL` argument to the overloaded `xquery()` method. The XQuery expression implicitly acts on the message content and returns a node set as its result. Depending on the context, the return value is interpreted either as a predicate (where an empty node set is interpreted as false) or an expression.

JoSQL

The `sql()` static method enables you to call on the JoSQL (SQL for Java objects) language to evaluate predicates and expressions in Java Router. JoSQL employs a SQL-like query syntax to perform selection and ordering operations on data from in-memory Java objects—JoSQL is *not* a database, however. In the JoSQL syntax, each Java object instance is treated like a table row and each object method is treated like a column name. Using this syntax, it is possible to construct powerful statements for extracting and compiling data from collections of Java objects. For details, see <http://josql.sourceforge.net/>.

To use the JoSQL language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.builder.sql.SqlBuilder.sql;
```

OGNL

The `ognl()` static method enables you to call on OGNL (Object Graph Navigation Language) expressions, which can then be used as predicates and expressions in a router rule. For details, see <http://www.ognl.org/>.

To use the OGNL language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.language.ognl.OgnlExpression.ognl;
```

EL

The `el()` static method enables you to call on the Unified Expression Language (EL) to construct predicates and expressions in a router rule. The EL was originally specified as part of the JSP 2.1 standard (JSR-245), but is now available as a standalone language. Java Router integrates with JUEL (<http://juel.sourceforge.net/>), which is an open source implementation of the EL language.

To use the EL language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.language.juel.JuelExpression.el;
```

Groovy

The `groovy()` static method enables you to call on the Groovy scripting language to construct predicates and expressions in a route. To use the Groovy language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.builder.camel.script.ScriptBuilder.*;
```

JavaScript

The `javascript()` static method enables you to call on the JavaScript scripting language to construct predicates and expressions in a route. To use the JavaScript language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.builder.camel.script.ScriptBuilder.*;
```

PHP

The `php()` static method enables you to call on the PHP scripting language to construct predicates and expressions in a route. To use the PHP language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.builder.camel.script.Script
Builder.*;
```

Python

The `python()` static method enables you to call on the Python scripting language to construct predicates and expressions in a route. To use the Python language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.builder.camel.script.Script
Builder.*;
```

Ruby

The `ruby()` static method enables you to call on the Ruby scripting language to construct predicates and expressions in a route. To use the Ruby language in your application code, include the following import statement in your Java source files:

```
import static org.apache.camel.builder.camel.script.Script
Builder.*;
```

Bean

You can also use Java beans to evaluate predicates and expressions. For example, to evaluate the predicate on a filter using the `isGoldCustomer()` method on the bean instance, `myBean`, you can use a rule like the following:

```
from("SourceURL")
    .filter().method("myBean", "isGoldCustomer")
    .to("TargetURL");
```

A discussion of bean integration in Java Router is beyond the scope of this *Defining Routes* guide. For details, see <http://activemq.apache.org/camel/bean-language.html>.

Transforming Message Content

Overview

Java Router supports a variety of approaches to transforming message content. In addition to a simple native API for modifying message content, Java Router supports integration with several different third-party libraries and transformation standards. The following kinds of transformation are discussed in this section:

- [Simple transformations on page 35](#).
- [Marshalling and unmarshalling on page 39](#).

Simple transformations

The Java DSL has a built-in API that enables you to perform simple transformations on incoming and outgoing messages. For example, the rule shown in [Example 3 on page 35](#) would append the text, `World!`, to the end of the incoming message body.

Example 3. Simple Transformation of Incoming Messages

```
from("SourceURL").setBody(body().append(" World!")).to("TargetURL");
```

Where the `setBody()` command replaces the content of the incoming message's body. You can use the following API classes to perform simple transformations of the message content in a router rule:

- `org.apache.camel.model.ProcessorType`
- `org.apache.camel.builder.Builder`
- `org.apache.camel.builder.ValueBuilder`

ProcessorType class

The `org.apache.camel.model.ProcessorDefinition` class defines the DSL commands you can insert directly into a router rule—for example, the `setBody()` command in [Example 3 on page 35](#). [Table 2 on page 36](#) shows the `ProcessorDefinition` methods that are relevant to transforming message content:

Table 2. Transformation Methods from the ProcessorType Class

Method	Description
Type convertBodyTo(Class type)	Converts the IN message body to the specified type.
Type convertFaultBodyTo(Class type)	Converts the FAULT message body to the specified type.
Type convertOutBodyTo(Class type)	Converts the OUT message body to the specified type.
Type removeFaultHeader(String name)	Adds a processor which removes the header on the FAULT message.
Type removeHeader(String name)	Adds a processor which removes the header on the IN message.
Type removeOutHeader(String name)	Adds a processor which removes the header on the OUT message.
Type removeProperty(String name)	Adds a processor which removes the exchange property.
ExpressionClause<ProcessorType<Type>> setBody()	Adds a processor which sets the body on the IN message.
ExpressionClause<ProcessorType<Type>> setBody()	Adds a processor which sets the body on the IN message.
Type setFaultBody(Expression expression)	Adds a processor which sets the body on the FAULT message.
Type setFaultHeader(String name, Expression expression)	Adds a processor which sets the header on the FAULT message.
ExpressionClause<ProcessorType<Type>> setHeader(String name)	Adds a processor which sets the header on the IN message.
Type setHeader(String name, Expression expression)	Adds a processor which sets the header on the IN message.
ExpressionClause<ProcessorType<Type>> setOutHeader(String name)	Adds a processor which sets the header on the OUT message.
Type setOutHeader(String name, Expression expression)	Adds a processor which sets the header on the OUT message.
ExpressionClause<ProcessorType<Type>> setProperty(String name)	Adds a processor which sets the exchange property.

Method	Description
Type <code>setProperty(String name, Expression expression)</code>	Adds a processor which sets the exchange property.

Builder class

The `org.apache.camel.builder.Builder` class provides access to message content in contexts where expressions or predicates are expected. In other words, `Builder` methods are typically invoked in the *arguments* of DSL commands—for example, the `body()` command in [Example 3 on page 35](#). [Table 3 on page 37](#) summarizes the static methods available in the `Builder` class.

Table 3. Methods from the Builder Class

Method	Description
<code>static <E extends Exchange> ValueBuilder<E> body()</code>	Returns a predicate and value builder for the inbound body on an exchange.
<code>static <E extends Exchange, T> ValueBuilder<E> bodyAs(Class<T> type)</code>	Returns a predicate and value builder for the inbound message body as a specific type.
<code>static <E extends Exchange> ValueBuilder<E> constant(Object value)</code>	Returns a constant expression.
<code>static <E extends Exchange> ValueBuilder<E> faultBody()</code>	Returns a predicate and value builder for the fault body on an exchange.
<code>static <E extends Exchange, T> ValueBuilder<E> faultBodyAs(Class<T> type)</code>	Returns a predicate and value builder for the fault message body as a specific type.
<code>static <E extends Exchange> ValueBuilder<E> header(String name)</code>	Returns a predicate and value builder for headers on an exchange.
<code>static <E extends Exchange> ValueBuilder<E> outBody()</code>	Returns a predicate and value builder for the outbound body on an exchange.
<code>static <E extends Exchange> ValueBuilder<E> outBody()</code>	Returns a predicate and value builder for the outbound message body as a specific type.

Method	Description
<code>static <E extends Exchange> ValueBuilder<E> systemProperty(String name)</code>	Returns an expression for the given system property.
<code>static <E extends Exchange> ValueBuilder<E> systemProperty(String name, String defaultValue)</code>	Returns an expression for the given system property.

ValueBuilder class

The `org.apache.camel.builder.ValueBuilder` class enables you to modify values returned by the `Builder` methods. In other words, the methods in `ValueBuilder` provide a simple way of modifying message content. [Table 4 on page 38](#) summarizes the methods available in the `ValueBuilder` class. That is, the table shows only the methods that are used to modify the value they are invoked on (for full details, see the *API Reference* documentation).

Table 4. Modifier Methods from the ValueBuilder Class

Method	Description
<code>ValueBuilder<E> append(Object value)</code>	Appends the string evaluation of this expression with the given value.
<code>ValueBuilder<E> convertTo(Class type)</code>	Converts the current value to the given type using the registered type converters.
<code>ValueBuilder<E> convertToString()</code>	Converts the current value a <code>String</code> using the registered type converters.
<code>ValueBuilder<E> regexReplaceAll(String regex, Expression<E> replacement)</code>	Replaces all occurrences of the regular expression with the given replacement.
<code>ValueBuilder<E> regexReplaceAll(String regex, String replacement)</code>	Replaces all occurrences of the regular expression with the given replacement.
<code>ValueBuilder<E> regexTokenize(String regex)</code>	Tokenizes the string conversion of this expression using the given regular expression.
<code>ValueBuilder<E> tokenize()</code>	

Method	Description
<code>ValueBuilder<E> tokenize(String token)</code>	Tokenizes the string conversion of this expression using the given token separator.

Marshalling and unmarshalling

You can convert between low-level and high-level message formats using the following commands:

- `marshal()`—convert a high-level data format to a low-level data format.
- `unmarshal()`—convert a low-level data format to a high-level data format.

Java Router supports marshalling and unmarshalling of the following data formats:

- *Java serialization*—enables you to convert a Java object to a blob of binary data. For this data format, unmarshalling converts a binary blob to a Java object and marshalling converts a Java object to a binary blob. For example, to read a serialized Java object from an endpoint, *SourceURL*, and convert it to a Java object, you could use the following rule:

```
from("SourceURL").unmarshal().serialization()
.<FurtherProcessing>.to("TargetURL");
```

- *JAXB*—provides a mapping between XML schema types and Java types (see <https://jaxb.dev.java.net/>). For JAXB, unmarshalling converts an XML data type to a Java object and marshalling converts a Java object to an XML data type. Before you can use JAXB data formats, you must compile your XML schema using a JAXB compiler in order to generate the Java classes that represent the XML data types in the schema. This is called *binding* the schema. After you have bound the schema, you can define a rule to unmarshal XML data to a Java object, using code like the following:

```
org.apache.camel.spi.DataFormat jaxb = new
org.apache.camel.model.dataformat.JaxbDataFormat("Generated
PackageName");

from("SourceURL").unmarshal(jaxb)
.<FurtherProcessing>.to("TargetURL");
```

Where *GeneratedPackagename* is the name of the Java package generated by the JAXB compiler, which contains the Java classes representing your XML schema.

- *XMLBeans*—provides an alternative mapping between XML schema types and Java types (see <http://xmlbeans.apache.org/>). For XMLBeans, unmarshalling converts an XML data type to a Java object and marshalling converts a Java object to an XML data type. For example, to unmarshal XML data to a Java object using XMLBeans, you can use code like the following:

```
from("SourceURL").unmarshal().xmlBeans()  
    .<FurtherProcessing>.to("TargetURL");
```

- *XStream*—provides another mapping between XML types and Java types (see <http://xstream.codehaus.org/>). XStream is a serialization library (like Java serialization), enabling you to convert any Java object to XML. For XStream, unmarshalling converts an XML data type to a Java object and marshalling converts a Java object to an XML data type. For example, to unmarshal XML data to a Java object using XStream, you can use code like the following:

```
from("SourceURL").unmarshal().xstream()  
    .<FurtherProcessing>.to("TargetURL");
```

Defining Routes in XML

You can define routing rules in XML. This approach is not as flexible as Java DSL, but has the advantage that it is easy to reconfigure the routing rules at runtime.

Using the Router Schema in an XML File	42
Defining a Basic Route in XML	44
Processors	45
Languages for Expressions and Predicates	52
Transforming Message Content	54

Using the Router Schema in an XML File

Overview

The root element of the router schema is `camelContext`, which is defined in the XML namespace, `http://camel.apache.org/schema/spring`. Router configurations are typically embedded in other XML configuration files (for example, in a Spring configuration file). In general, whenever a router configuration is embedded in another configuration file, you need to specify the location of the router schema (so that the router configuration can be parsed). For example, [Example 4 on page 42](#) shows how to embed the router configuration, `camelContext`, in an arbitrary document, `DocRootElement`.

Example 4. Specifying the Router Schema Location

```
<DocRootElement ...
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="
    http://camel.apache.org/schema/spring ht
tp://camel.apache.org/schema/spring/camel-spring.xsd">

  <camelContext id="camel" xmlns="http://act
ivemq.apache.org/camel/schema/spring">
    <!-- Define your routing rules here -->
  </camelContext>
</DocRootElement>
```

Where the schema location is specified to be

`http://camel.apache.org/schema/spring/camel-spring.xsd`, which gives the location of the schema on the Apache Web site. This location always contains the latest, most up-to-date version of the XML schema. .

[Example 5 on page 42](#) shows an example of embedding a router configuration, `camelContext`, in a Spring configuration file.

Example 5. Router Schema in a Spring Configuration File

```
<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 ht
tp://www.springframework.org/schema/beans/spring-beans.xsd
    http://camel.apache.org/schema/spring ht
tp://camel.apache.org/schema/spring/camel-spring.xsd">

  <camelContext id="camel" xmlns="ht
tp://camel.apache.org/schema/spring">
    <!-- Define your routing rules in here -->
```

```
</camelContext>  
  <!-- Other Spring configuration -->  
  <!-- ... -->  
</beans>
```

Defining a Basic Route in XML

Basic concepts

In order to understand how to build a route using XML, you need to understand some of the basic concepts of the routing language—for example, sources and targets, processors, expressions and predicates, and message exchanges. For definitions and explanations of these concepts see [Basic Java DSL Syntax on page 20](#).

Example of a basic route

[Example 6 on page 44](#) shows an example of a basic route in XML, which connects a source endpoint, *SourceURL*, directly to a destination endpoint, *TargetURL*.

Example 6. Basic Route in XML

```
<camelContext id="CamelContextID" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <to uri="TargetURL"/>
  </route>
</camelContext>
```

Where *CamelContextID* is an arbitrary, unique identifier for the Camel context. The route is defined by a `route` element and there can be multiple `route` elements under the `camelContext` element.

Processors

Overview

To enable the router to do something more interesting than simply connecting a source endpoint to a target endpoint, you can add *processors* to your route. A processor is a command you can insert into a routing rule in order to perform arbitrary processing of the messages that flow through the rule. Java Router provides a wide variety of different processors, as follows:

- [Filter on page 45](#).
- [Choice on page 46](#).
- [Recipient list on page 46](#).
- [Splitter on page 47](#).
- [Aggregator on page 47](#).
- [Resequencer on page 48](#).
- [Throttler on page 49](#).
- [Delayer on page 50](#).

Filter

The `filter` processor can be used to prevent uninteresting messages from reaching the target endpoint. It takes a single predicate argument: if the predicate is true, the message exchange is allowed through to the target; if the predicate is false, the message exchange is blocked. For example, the following filter blocks a message exchange, unless the incoming message contains a header, `foo`, with value equal to `bar`:

```
<camelContext id="filterRoute" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <filter>
      <simple>header.foo = 'bar'</simple>
      <to uri="TargetURL"/>
    </filter>
  </route>
</camelContext>
```

```
</route>
</camelContext>
```

Choice

The `choice` processor is a conditional statement that is used to route incoming messages to alternative targets. The alternative targets are each enclosed in a `when` element, which takes a predicate argument. If the predicate is true, the current target is selected, otherwise processing proceeds to the next `when` element in the rule. For example, the following `choice()` processor directs incoming messages to either `Target1`, `Target2`, or `Target3`, depending on the values of the predicates:

```
<camelContext id="buildSimpleRouteWithChoice" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <choice>
      <when>
        <!-- First predicate -->
        <simple>header.foo = 'bar'</simple>
        <to uri="Target1"/>
      </when>
      <when>
        <!-- Second predicate -->
        <simple>header.foo = 'manchu'</simple>
        <to uri="Target2"/>
      </when>
      <otherwise>
        <to uri="Target3"/>
      </otherwise>
    </choice>
  </route>
</camelContext>
```

Recipient list

If you want the messages from a source endpoint, `SourceURL`, to be sent to more than one target, there are two alternative approaches you can use. One approach is to include multiple `to` elements in the route, for example:

```
<camelContext id="staticRecipientList" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <to uri="Target1"/>
    <to uri="Target2"/>
    <to uri="Target3"/>
  </route>
</camelContext>
```

```
</route>
</camelContext>
```

The alternative approach is to add a `recipientList` element, which takes a list of recipients as its argument (dynamic recipient list). The advantage of using the `recipientList` element is that the list of recipients can be calculated at runtime. For example, the following rule generates a recipient list by reading the contents of the `recipientListHeader` from the incoming message:

```
<camelContext id="dynamicRecipientList" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <recipientList>
      <!-- Requires XPath 2.0 -->
      <xpath>tokenize(/headers/recipientListHeader, "\s+")</xpath>
    </recipientList>
  </route>
</camelContext>
```

Splitter

The `splitter` processor is used to split a message into parts, which are then processed as separate messages. The `splitter` element must contain an expression that returns a list, where each item in the list represents a message part that is to be re-sent as a separate message. For example, the following rule splits the body of an incoming message into separate sections (represented by a top-level `section` element) and then sends each section to the target in a separate message:

```
<camelContext id="splitterRoute" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <splitter>
      <xpath>/section</xpath>
      <to uri="seda:b"/>
    </splitter>
  </route>
</camelContext>
```

Aggregator

The `aggregator` processor is used to aggregate related incoming messages into a single message. In order to distinguish which messages are eligible to be aggregated together, you need to define a *correlation key* for the aggregator.

The correlation key is normally derived from a field in the message (for example, a header field). Messages that have the same correlation key value are eligible to be aggregated together. You can also optionally specify an aggregation algorithm to the `aggregator` processor (the default algorithm is to pick the latest message with a given value of the correlation key and to discard the older messages with that correlation key value).

For example, if you are monitoring a data stream that reports stock prices in real time, you might only be interested in the *latest* price of each stock symbol. In this case, you could configure an aggregator to transmit only the latest price for a given stock and discard the older (out-of-date) price notifications. The following rule implements this functionality, where the correlation key is read from the `stockSymbol` header and the default aggregator algorithm is used:

```
<camelContext id="aggregatorRoute" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <aggregator>
      <simple>header.stockSymbol</simple>
      <to uri="TargetURL"/>
    </aggregator>
  </route>
</camelContext>
```

Resequencer

A `resequencer` processor is used to re-arrange the order in which incoming messages are transmitted. The `resequencer` element needs to be provided with a sequence number (where the sequence number is calculated from the contents of a field in the incoming message). Naturally, before you can start re-ordering messages, you need to wait until a certain number of messages have been received from the source. There are a couple of different ways to specify how long the `resequencer` processor should wait before attempting to re-order the accumulated messages and forward them to the target, as follows:

- *Batch resequencing*—(the default) wait until a specified number of messages have accumulated before starting to re-order and forward messages. For example, the following resequencing rule would re-order messages based on the `timeOfDay` header, waiting until at least 300 messages have accumulated or 4000 ms have elapsed since the last message received.

```

<camelContext id="batchResequencer" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL" />
    <resequencer>
      <!-- Sequence ordering based on timeOfDay header -->
      <simple>header.timeOfDay</simple>
      <to uri="TargetURL" />
      <!--
        batch-config can be omitted for default (batch)
        resequencer settings
      -->
      <batch-config batchSize="300" batchTimeout="4000" />
    </resequencer>
  </route>
</camelContext>

```

- *Stream resequencing*—transmit messages as soon as they arrive *unless* the resequencer detects a gap in the incoming message stream (missing sequence numbers), in which case the resequencer waits until the missing messages arrive and then forwards the messages in the correct order. To avoid the resequencer blocking forever, you can specify a timeout (default is 1000 ms), after which time the message sequence is transmitted with unresolved gaps. For example, the following resequencing rule detects gaps in the message stream by monitoring the value of the `sequenceNumber` header, where the maximum buffer size is limited to 5000 and the timeout is specified to be 4000 ms:

```

<camelContext id="streamResequencer" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <resequencer>
      <simple>header.sequenceNumber</simple>
      <to uri="TargetURL" />
      <stream-config capacity="5000" timeout="4000"/>
    </resequencer>
  </route>
</camelContext>

```

Throttler

The `throttler` processor is used to ensure that a target endpoint does not get overloaded. The throttler works by limiting the number of messages that can pass through per second. If the incoming messages exceed the specified rate, the throttler accumulates excess messages in a buffer and transmits

them more slowly to the target endpoint. For example, to limit the rate of throughput to 100 messages per second, you can define the following rule:

```
<camelContext id="throttlerRoute" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <throttler maximumRequestsPerPeriod="100" timePeriodMillis="1000">
      <to uri="TargetURL"/>
    </throttler>
  </route>
</camelContext>
```

Delayer

The `delayer` processor is used to hold up messages for a specified length of time. The delay can either be relative (wait a specified length of time after receipt of the incoming message) or absolute (wait until a specific time). For example, to add a delay of 2 seconds before transmitting received messages, you can use the following rule:

```
<camelContext id="delayerRelative" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <delayer>
      <delay>2000</delay>
      <to uri="TargetURL"/>
    </delayer>
  </route>
</camelContext>
```

To wait until the absolute time specified in the `processAfter` header, you can use the following rule:

```
<camelContext id="delayerRelative" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <delayer>
      <simple>header.processAfter</simple>
      <to uri="TargetURL"/>
    </delayer>
  </route>
</camelContext>
```

```
</route>  
</camelContext>
```

Load balancer

The `loadBalance` processor is used to load balance message exchanges over a list of target endpoints. For example, to load balance incoming messages exchanges using a round robin algorithm (each endpoint in the target list is tried in sequence), you can use the following rule:

```
<camelContext id="loadBalancer" xmlns="http://camel.apache.org/schema/spring">  
  <route>  
    <from uri="SourceURL"/>  
    <loadBalance>  
      <to uri="TargetURL_01"/>  
      <to uri="TargetURL_02"/>  
      <roundRobin/>  
    </loadBalance>  
  </route>  
</camelContext>
```

Currently, it is not possible to customize the load balancing algorithm in XML.

Languages for Expressions and Predicates

Overview

In the definition of a route, it is frequently necessary to evaluate expressions and predicates. For example, if a route includes a filter processor, you need to evaluate a predicate to determine whether or not a message is to be allowed through the filter. To facilitate the evaluation of expressions and predicates, Java Router supports multiple language plug-ins, which can be accessed through XML elements.

Elements for expressions and predicates

Table 5 on page 52 lists the elements that you can insert whenever the context demands an expression or a predicate. The content of the element must be a script written in the relevant language. At runtime, the return value of the script is read by the parent element.

Table 5. Elements for Expression and Predicate Languages

Element	Language	Description
simple	N/A	A simple expression language, native to Java Router (see Simple on page 30).
xpath	XPath	The XPath language, which is used to select element, attribute, and text nodes from XML documents (see http://www.w3schools.com/xpath/default.asp). The XPath expression is applied to the current message.
xquery	XQuery	The XQuery language, which is an extension of XPath (see http://www.w3schools.com/xquery/default.asp). The XQuery expression is applied to the current message.
sql	JoSQL	The JoSQL language, which is a language for extracting and manipulating data from collections of Java objects, using a SQL-like syntax (see http://josql.sourceforge.net/).
ognl	OGNL	The OGNL (Object Graph Navigation Language) language (see http://www.ognl.org/).

Element	Language	Description
el	EL	The Unified Expression Language (EL), originally developed as part of the JSP standard (see http://juel.sourceforge.net/).
groovy	Groovy	The Groovy scripting language (see http://groovy.codehaus.org/).
javaScript	JavaScript	The JavaScript scripting language (see http://developer.mozilla.org/en/docs/JavaScript), also known as ECMAScript (see http://www.ecmascript.org/).
php	PHP	The PHP scripting language (see http://www.php.net/).
python	Python	The Python scripting language (see http://www.python.org/).
ruby	Ruby	The Ruby scripting language (see http://www.ruby-lang.org/).
bean	Bean	Not really a language. The <code>bean</code> element is actually a mechanism for integrating with Java beans. You use the <code>bean</code> element to obtain an expression or predicate by invoking a method on a Java bean.

Transforming Message Content

Overview

This section describes how you can transform messages using the features provided in XML configuration.

Marshalling and unmarshalling

You can convert between low-level and high-level message formats using the following elements:

- `marshal`—convert a high-level data format to a low-level data format.
- `unmarshal`—convert a low-level data format to a high-level data format.

Java Router supports marshalling and unmarshalling of the following data formats:

- *Java serialization*—enables you to convert a Java object to a blob of binary data. For this data format, unmarshalling converts a binary blob to a Java object and marshalling converts a Java object to a binary blob. For example, to read a serialized Java object from an endpoint, `SourceURL`, and convert it to a Java object, you could use the following rule:

```
<camelContext id="serialization" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <unmarshal>
      <serialization/>
    </unmarshal>
    <to uri="TargetURL"/>
  </route>
</camelContext>
```

- *JAXB*—provides a mapping between XML schema types and Java types (see <https://jaxb.dev.java.net/>). For JAXB, unmarshalling converts an XML data type to a Java object and marshalling converts a Java object to an XML data type. Before you can use JAXB data formats, you must compile your XML schema using a JAXB compiler in order to generate the Java classes that represent the XML data types in the schema. This is called *binding* the schema. After you have bound the schema, you can define a rule to unmarshal XML data to a Java object, as follows:

```

<camelContext id="jaxb" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <unmarshal>
      <jaxb prettyPrint="true" contextPath="GeneratedPackageName"/>
    </unmarshal>
    <to uri="TargetURL"/>
  </route>
</camelContext>

```

Where *GeneratedPackagename* is the name of the Java package generated by the JAXB compiler, which contains the Java classes representing your XML schema.

- *XMLBeans*—provides an alternative mapping between XML schema types and Java types (see <http://xmlbeans.apache.org/>). For XMLBeans, unmarshalling converts an XML data type to a Java object and marshalling converts a Java object to an XML data type. For example, to unmarshal XML data to a Java object using XMLBeans, define a rule like the following:

```

<camelContext id="xmlBeans" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <unmarshal>
      <xmlBeans prettyPrint="true"/>
    </unmarshal>
    <to uri="TargetURL"/>
  </route>
</camelContext>

```

- *XStream*—is currently *not* supported in XML configuration.

Basic Principles of Route Building

Java Router provides a great number of different processors and components, which you can link together to build a route. This chapter aims to give you some basic orientation by explaining the principles of building a route using the provided building blocks.

Pipeline Processing	58
Multiple Inputs	63
Exception Handling	68
Bean Integration	70

Pipeline Processing

Overview

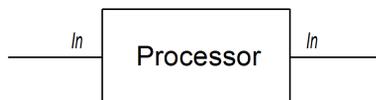
In Java Router, pipelining is the dominant paradigm for connecting nodes in a route definition. The pipeline concept is probably most familiar from the UNIX operating system, where it is used to join together operating system commands. For example, `ls | more` is an example of a command that pipes a directory listing, `ls`, to a page scrolling utility, `more`. The basic idea of a pipeline is that the *output* of one command is fed into the *input* of the next. The natural analogy in the case of a route is for the *Out* message from one processor to be copied to the *In* message of the next processor.

Processor nodes

Every node in a route, except for the initial endpoint, is a *processor* (in the sense that they inherit from the `org.apache.camel.Processor` interface). In other words, processors make up the basic building blocks of a DSL route. For example, DSL commands such as `filter()`, `delayer()`, `setBody()`, `setHeader()`, `to()`, and so on, all represent processors. When considering how processors connect together to build up a route, it is important to distinguish two different processing approaches.

The first approach is where the processor simply modifies the exchange's *In* message, as shown in [Figure 2 on page 58](#). The exchange's *Out* message remains `null` in this case.

Figure 2. Processor Modifying an In Message

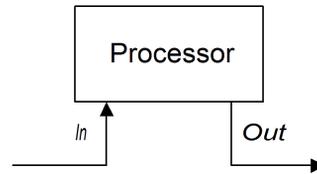


For example, the following route shows a `setHeader()` command that modifies the current *In* message by adding (or modifying) the `BillingSystem` heading:

```

from("activemq:orderQueue")
  .setHeader("BillingSystem", xpath("/order/billingSystem"))
  .to("activemq:billingQueue");
  
```

The second approach is where the processor creates an *Out* message to represent the result of the processing, as shown in [Figure 3 on page 59](#).

Figure 3. Processor Creating an Out Message

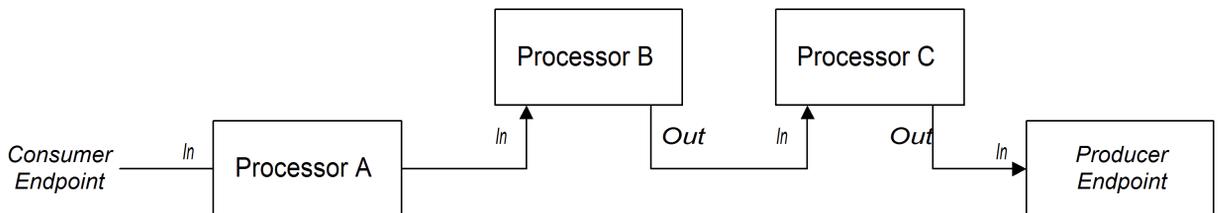
For example, the following route shows a `transform()` command that creates an `Out` message with a message body containing the string, `DummyBody`:

```
from("activemq:orderQueue")
  .transform(constant("DummyBody"))
  .to("activemq:billingQueue");
```

Where `constant("DummyBody")` represents a constant expression (you cannot pass the string, `DummyBody`, directly, because the argument to `transform()` must be an expression type).

Pipeline for InOnly exchanges

Figure 4 on page 59 shows an example of a processor pipeline for `InOnly` exchanges. Processor A acts by modifying the `In` message, while processors B and C create an `Out` message. The route builder links the processors together as shown. In particular, processors B and C are linked together in the form of a *pipeline*: that is, processor B's `Out` message is moved to the `In` message before feeding the exchange into processor C, and processor C's `Out` message is moved to the `In` message before feeding the exchange into the producer endpoint. Thus the processors' outputs and inputs are joined into a continuous pipeline, as shown in Figure 4 on page 59.

Figure 4. Sample Pipeline for InOnly Exchanges

Java Router employs the pipeline pattern by default. Hence, you do not need to use any special syntax to create a pipeline in your routes. For example, the following route pulls messages from a `userdataQueue` queue, pipes the

message through a Velocity template (to produce a customer address in text format), and then sends the resulting text address to the queue, `envelopeAddressQueue`:

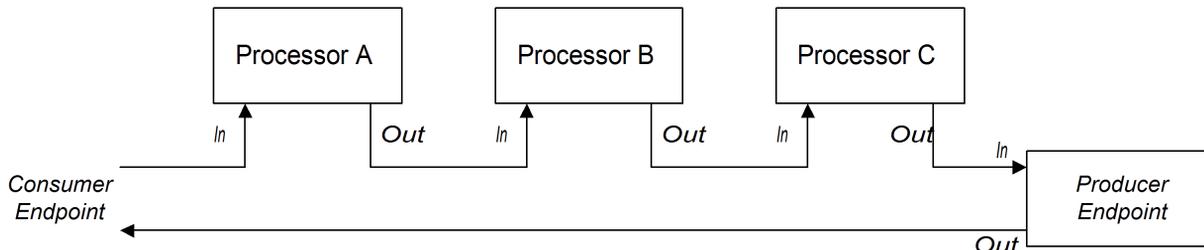
```
from("activemq:userdataQueue")
    .to("velocity:file:AdressTemplate.vm")
    .to("activemq:envelopeAddresses");
```

Where the Velocity endpoint, `velocity:file:AdressTemplate.vm`, specifies the location of a Velocity template file, `file:AdressTemplate.vm`, in the file system.

Pipeline for InOut exchanges

Figure 5 on page 60 shows an example of a processor pipeline for *InOut* exchanges, which you typically use to support remote procedure call (RPC) semantics. Processors A, B, and C are linked together in the form of a pipeline, with the output of each processor being fed into the input of the next. The final *Out* message produced by the producer endpoint is sent all the way back to the consumer endpoint, where it provides the reply to the original request.

Figure 5. Sample Pipeline for InOut Exchanges



Note that in order to support the *InOut* exchange pattern, it is *essential* that the last node in the route (whether it is a producer endpoint or some other kind of processor) creates an *Out* message. Otherwise, any client that connects to the consumer endpoint would hang, waiting indefinitely for a reply message. In particular, you should be aware that not all producer endpoints create *Out* messages.

For example, consider the following route that processes payment requests, by processing incoming HTTP requests:

```
from("jetty:http://localhost:8080/foo")
    .to("cxf:bean:addAccountDetails")
    .to("cxf:bean:getCreditRating")
    .to("cxf:bean:processTransaction");
```

Where the incoming payment request is processed by passing it through a pipeline of Web services, `cx:bean:addAccountDetails`, `cx:bean:getCreditRating`, and `cx:bean:processTransaction`. The final Web service, `processTransaction`, generates a response (*Out* message) that is sent back through the JETTY endpoint.

When the pipeline consists of just a sequence of endpoints, it is also possible to use the following alternative syntax:

```
from("jetty:http://localhost:8080/foo")
  .pipeline("cx:bean:addAccountDetails",
    "cx:bean:getCreditRating", "cx:bean:processTransaction");
```

Comparison of pipelining and interceptor chaining

An alternative paradigm for linking together the nodes of a route is *interceptor chaining*, where a processor in the route processes the exchange both before and after dispatching the exchange to the next processor in the chain. This style of processing is also supported by Java Router, but it is not the usual approach to use. [Figure 6 on page 61](#) shows an example of an interceptor processor that implements a custom encryption algorithm.

Figure 6. Example of Interceptor Chaining



In this example, incoming messages are encrypted in a custom format. The interceptor first decrypts the *In* message, then dispatches it to the Web services endpoint, `cx:bean:processTxn`, and finally, the reply (*Out* message) is encrypted using the custom format, before being sent back through the consumer endpoint. Using the interceptor chaining approach, therefore, a single interceptor instance can modify both the request and the response.

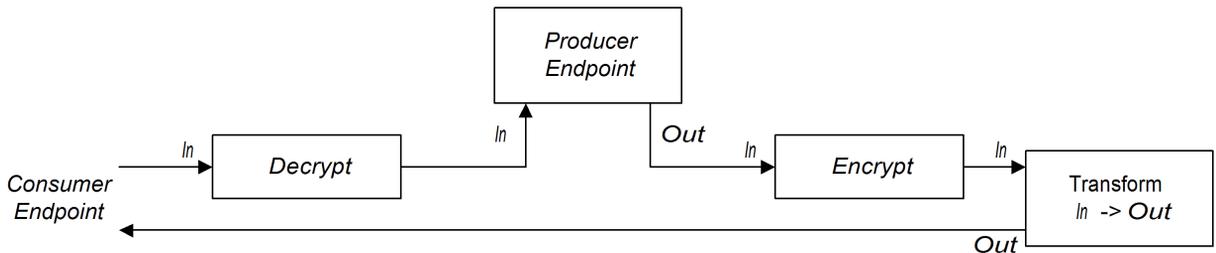
For example, if you want to define a route with a HTTP port that services incoming requests encoded using custom encryption, you could define a route like the following:

```
from("jetty:http://localhost:8080/foo")
  .intercept(new MyDecryptEncryptInterceptor())
  .to("cx:bean:processTxn");
```

Where the class, `MyDecryptEncryptInterceptor`, is implemented by inheriting from the class, `org.apache.camel.processor.DelegateProcessor`.

Although it is possible to implement this kind of functionality using an interceptor processor, this is not a very idiomatic way of programming in Java Router. A more typical approach is shown in [Figure 7 on page 62](#).

Figure 7. Pipeline Alternative to Interceptor Chaining



In this example, the encrypt functionality is implemented in a separate processor from the decrypt functionality. The resulting processor pipeline is semantically equivalent to the original interceptor chain shown in [Figure 6 on page 61](#). One slight complication of this route, however, is that it turns out to be necessary to add a `transform` processor at the end of the route, in order to copy the `In` message to the `Out` message. This processor ensures that the reply message is available to the HTTP consumer endpoint (an alternative solution to this problem would be to implement the encrypt processor such that it creates an `Out` message directly).

For example, to implement the pipeline approach shown in [Figure 7 on page 62](#), you could define a route like the following:

```

from("jetty:http://localhost:8080/foo")
  .process(new MyDecryptProcessor())
  .to("cxf:bean:processTxn")
  .process(new MyEncryptProcessor())
  .transform(body());
  
```

Where the final processor node, `transform(body())`, has the effect of copying the `In` message to the `Out` message (the `In` message body is copied explicitly and the `In` message headers are copied implicitly).

Multiple Inputs

Overview

A standard route takes its input from just a single endpoint, using the `from(EndpointURL)` syntax in the Java DSL. But what if you need to define multiple inputs for your route? For example, you might want to merge incoming messages from two different messaging systems and process them using the same route. In most cases, you can deal with multiple inputs by dividing your route into segments, as shown in [Figure 8 on page 63](#).

Figure 8. Processing Multiple Inputs with Segmented Routes

```

from("activemq:Nyse").to(InternalUrl)
                                     ↘
                                     from(InternalUrl).to("activemq:USTxn")
                                     ↗
from("activemq:Nasdaq").to(InternalUrl)

```

The initial segments of the route take their inputs from some external queues—for example, `activemq:Nyse` and `activemq:Nasdaq`—and send the incoming exchanges to an internal endpoint, `InternalUrl`. The second route segment merges the incoming exchanges, taking them from the internal endpoint and sending them to the destination queue, `activemq:USTxn`. The `InternalUrl` is the URL for a kind of endpoint that is intended only for use *within* a router application. The following types of endpoint are suitable for internal use:

- [Direct endpoints on page 64](#).
- [SEDA endpoints on page 65](#).
- [VM endpoints on page 66](#).

The main purpose of these endpoints is to enable you to glue together different segments of a route. In particular, they all provide an effective way of merging multiple inputs into a single route.



Note

The direct, SEDA, and VM components work only for the *InOnly* exchange pattern. If one of your inputs requires an *InOut* exchange

pattern, you should take a look at the [Content enricher pattern on page 66](#) instead.

Direct endpoints

The direct component provides the simplest mechanism for linking together routes. The event model for the direct component is *synchronous*, so that subsequent segments of the route run in the same thread as the first segment. The general format of a direct URL is `direct:EndpointID`, where the endpoint ID, `EndpointID`, is simply a unique alphanumeric string that identifies the endpoint instance.

For example, say you want to take the input from two message queues, `activemq:Nyse` and `activemq:Nasdaq`, and merge them into a single message queue, `activemq:USTxn`, you could do this by defining the following set of routes:

```
from("activemq:Nyse").to("direct:mergeTxns");
from("activemq:Nasdaq").to("direct:mergeTxns");

from("direct:mergeTxns").to("activemq:USTxn");
```

Where the first two routes take the input from the message queues, `Nyse` and `Nasdaq`, and send them to the endpoint, `direct:mergeTxns`. The last queue combines the inputs from the previous two queues and sends the combined message stream into the `activemq:USTxn` queue.

The implementation of the direct endpoint is very simple: whenever an exchange arrives at a producer endpoint (for example, `to("direct:mergeTxns")`), the direct endpoint passes the exchange directly to all of the consumers endpoints that have the same endpoint ID (for example, `from("direct:mergeTxns")`). Direct endpoints can only be used to communicate between routes that belong to the same `CamelContext` in the same Java virtual machine (JVM) instance.



Note

If you connect multiple consumers to a direct endpoint, every exchange that passes through the endpoint will be processed by *all* of the attached consumers (multicast). Hence, each exchange would

be processed *more than once*. If you don't want this to happen, consider using a SEDA endpoint, which behaves differently.

SEDA endpoints

The SEDA component provides an alternative mechanism for linking together routes. You can use it in a similar way to the direct component, but it has a different underlying event and threading model, as follows:

- Processing of a SEDA endpoint is *not* synchronous. That is, when you send an exchange to a SEDA producer endpoint, control immediately returns to the preceding processor in the route.
- SEDA endpoints contain a queue buffer (of `java.util.concurrent.BlockingQueue` type), which stores all of the incoming exchanges prior to processing by the next route segment.
- Each SEDA consumer endpoint creates a thread pool (default size is 5) to process exchange objects from the blocking queue.
- The SEDA component supports the *competing consumers* pattern, which guarantees that each incoming exchange is processed only once, even if there are multiple consumers attached to a specific endpoint.

One of the main advantages of using a SEDA endpoint is that the routes can be more responsive, owing to the built-in consumer thread pool. For example, the stock transactions example can be re-written to use SEDA endpoints instead of direct endpoints, as follows:

```
from("activemq:Nyse").to("seda:mergeTxns");
from("activemq:Nasdaq").to("seda:mergeTxns");

from("seda:mergeTxns").to("activemq:USTxn");
```

The main difference between this example and the direct example is that when using SEDA, the second route segment (from `seda:mergeTxns` to `activemq:USTxn`) is processed by a pool of five threads.



Note

There is more to SEDA than simply pasting together route segments. The staged event-driven architecture (SEDA) encompasses a design philosophy for building more manageable multi-threaded applications. The purpose of the SEDA component in Java Router is simply to enable you to apply this design philosophy to your applications. For

more details about SEDA, see <http://www.eecs.harvard.edu/~mdw/proj/seda/>.

VM endpoints

The VM component is very similar to the SEDA endpoint. The only difference is that, whereas the SEDA component is limited to linking together route segments from within the same `CamelContext`, the VM component enables you to link together routes from distinct Java Router applications, as long as they are running within the same Java virtual machine.

For example, the stock transactions example can be re-written to use VM endpoints instead of SEDA endpoints, as follows:

```
from("activemq:Nyse").to("vm:mergeTxns");
from("activemq:Nasdaq").to("vm:mergeTxns");
```

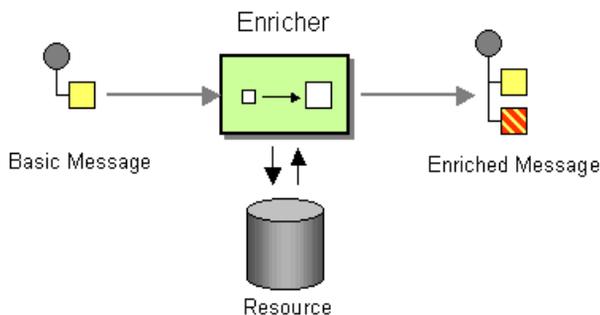
And in a separate router application (running in the same Java VM), you can define the second segment of the route:

```
from("vm:mergeTxns").to("activemq:USTxn");
```

Content enricher pattern

The content enricher pattern defines a fundamentally different way of dealing with multiple inputs to a route. A content enricher is a kind of processor that you can insert into a route, as shown in [Figure 9 on page 66](#). When an exchange enters the enricher processor, the enricher contacts an external resource to retrieve information, which is then added to the original message. In this pattern, the external resource effectively represents a second input to the message.

Figure 9. Processing Multiple Inputs with a Content Enricher



The key difference between the content enricher approach and the segmented route approach is that the content enricher is based on a radically different event model. In the segmented route approach, each of the input sources independently generate new message events. Whereas, in the content enricher approach, only one input source generates new message events, while a second resource (accessed by the enricher) is effectively a *slave* of the first stream of events. That is, the enricher's resource is only polled for input whenever a message arrives on the main route.

Exception Handling

Overview

Two exception handling methods are provided for catching exceptions in Java DSL routes, as follows:

- `errorHandler()` clause—is a simple catch-all mechanism that re-routes the current exchange to a dead letter channel, if an error occurs. This mechanism is *not* able to discriminate between different exception types.
- `onException()` clause—enables you to discriminate between different exception types, performing different kinds of processing for different exception types.

Dead letter channel pattern

This section provides a brief introduction to exception handling in Java Router.

`errorHandler` clause

The `errorHandler()` clause is defined in a `RouteBuilder` class and applies to all of the routes in that `RouteBuilder` class. It is triggered whenever an exception of *any kind* occurs in one of the applicable routes. For example, to define an error handler that routes all failed exchanges to the ActiveMQ `deadLetter` queue, you could define a `RouteBuilder` as follows:

```
public class MyRouteBuilder extends RouteBuilder {  
    public void configure() {  
        errorHandler(deadLetterChannel("activemq:deadLetter"));  
  
        // The preceding error handler applies to all of the  
        // following routes:  
        from("activemq:orderQueue").to("pop3://fulfill  
ment@acme.com");  
        from("file:src/data?noop=true").to("file:target/mes  
sages");  
        // ...  
    }  
}
```

Redirection to the dead letter channel does not occur, however, until attempts at redelivery have been exhausted (see [Redelivery policy on page 69](#)).

onException clause

The `onException(Class exceptionType)` clause is defined in a `RouteBuilder` class and applies to all of the routes in that `RouteBuilder` class. It is triggered whenever an exception of the specified type, `exceptionType`, occurs in one of the applicable routes. For example, you can define `onException` clauses for catching `NullPointerException`, `IOException`, and `Exception` exceptions as follows:

```
public class MyRouteBuilder extends RouteBuilder {

    public void configure() {
        onException(NullPointerException.class).to("activemq:ex.npex");
        onException(IOException.class).to("activemq:ex.ioex");

        onException(Exception.class).to("activemq:ex");

        // The preceding onException() clauses apply to all
        // of the following routes:
        from("activemq:orderQueue").to("pop3://fulfillment@acme.com");
        from("file:src/data?noop=true").to("file:target/messages");
        // ...
    }
}
```

When an exception occurs, Java Router selects the `onException` clause that matches the given exception type most closely. If no other clause matches the raised exception, the `onException(Exception.class)` clause (if present) matches by default, because `java.lang.Exception` is the base class of all Java exceptions. The applicable `onException` clause does not initiate processing, however, until attempts at redelivery have been exhausted (see [Redelivery policy on page 69](#)).

Redelivery policy

Both the `errorHandler` clause and the `onException` clause support a *redelivery policy* that specifies how often Java Router attempts to redeliver the failed exchange before giving up and triggering the actions defined by the relevant error handler. The most important redelivery policy setting is the *maximum redeliveries* value, which specifies how many times redelivery is attempted. The default value is 6.

Bean Integration

Overview

Bean integration provides a general purpose mechanism for processing messages using arbitrary Java objects. By inserting a bean reference into a route, you can call an arbitrary method on a Java object, which can then access and modify the incoming exchange. The mechanism that maps an exchange's contents to the parameters and return values of a bean method is known as *bean binding*. Bean binding can use either or both of the following approaches in order to initialize a method's parameters:

- *Conventional method signatures*—if the method signature conforms to certain conventions, the bean binding can use Java reflection to determine what parameters to pass.
- *Annotations and dependency injection*—for a more flexible binding mechanism, employ Java annotations to specify what to inject into the method's arguments. This dependency injection mechanism relies on Spring component scanning. Normally, if you are deploying your Java Router application into a Spring container, the dependency injection mechanism will work automatically.

Accessing a bean created in Java

To process exchange objects using a Java bean (which is just a plain old Java object or POJO), use the `bean()` processor, which binds the inbound exchange to a method on the Java object. For example, to process inbound exchanges using the class, `MyBeanProcessor`, define a route like the following:

```
from("file:data/inbound")
  .bean(MyBeanProcessor.class, "processBody")
  .to("file:data/outbound");
```

Where the `bean()` processor creates an instance of `MyBeanProcessor` type and invokes the `processBody()` method to process inbound exchanges. This approach is adequate, if you only want to access the `MyBeanProcessor` instance from a single route. If you want to access the same `MyBeanProcessor` instance from multiple routes, however, use the variant of `bean()` that takes the `Object` type as its first argument, as follows:

```
MyBeanProcessor myBean = new MyBeanProcessor();

from("file:data/inbound")
  .bean(myBean, "processBody")
  .to("file:data/outbound");
```

```
from("activemq:inboundData")
    .bean(myBean, "processBody")
    .to("activemq:outboundData");
```

Basic method signatures

In order to bind exchanges to a bean method, you can define a method signature that conforms to certain conventions. In particular, there are two basic conventions for method signatures:

- [Method signature for processing message bodies on page 71.](#)
- [Method signature for processing exchanges on page 71.](#)

Method signature for processing message bodies

If you want to implement a bean method that access or modifies the incoming message body, define a method signature that takes a single `String` argument and returns a `String` value. For example:

```
// Java
package com.acme;

public class MyBeanProcessor {
    public String processBody(String body) {
        // Do whatever you like to 'body'...
        return newBody;
    }
}
```

Method signature for processing exchanges

For greater flexibility, you can implement a bean method that accesses the incoming exchange. This enables you to access or modify all headers, bodies, and exchange properties. For processing exchanges, the method signature takes a single `org.apache.camel.Exchange` parameter and returns `void`. For example:

```
// Java
package com.acme;

public class MyBeanProcessor {
    public void processExchange(Exchange exchange) {
        // Do whatever you like to 'exchange'...
        exchange.getIn().setBody("Here is a new message body!");
    }
}
```

```
}
}
```

Accessing a bean created in Spring XML

Instead of creating a bean instance in Java, you can create an instance using Spring XML. This is the only feasible approach, in fact, if you are defining your routes in XML. To define a bean in XML, use the standard Spring `bean` element. For example, to create an instance of `MyBeanProcessor`:

```
<beans ...>
  ...
  <bean id="myBeanId" class="com.acme.MyBeanProcessor"/>
</beans>
```

It is also possible to pass data to the bean's constructor arguments using Spring syntax. For full details of how to use the Spring `bean` element, see [The IoC Container](#)¹ from the Spring reference guide.

When you create an object instance using the `bean` element, you can reference it later using the bean's ID (the value of the `bean` element's `id` attribute). For example, given the `bean` element with ID equal to `myBeanId`, you can reference the bean in a Java DSL route using the `beanRef()` processor, as follows:

```
from("file:data/inbound").beanRef("myBeanId", "process
Body").to("file:data/outbound");
```

Where the `beanRef()` processor invokes the `MyBeanProcessor.processBody()` method on the specified bean instance. You can also invoke the bean from within a Spring XML route, using the Camel schema's `bean` element. For example:

```
<camelContext id="CamelContextID" xmlns="ht
tp://camel.apache.org/schema/spring">
  <route>
    <from uri="file:data/inbound"/>
    <bean ref="myBeanId" method="processBody"/>
    <to uri="file:data/outbound"/>
  </route>
</camelContext>
```

Bean binding annotations

The basic bean bindings described in [Basic method signatures on page 71](#) might not always be convenient to use. For example, if you have a legacy

¹ <http://static.springsource.org/spring/docs/3.0.x/reference/beans.html>

Java class that performs some data manipulation, you might want to extract data from an inbound exchange and map it to the arguments of an existing method signature. For this kind of bean binding, Java Router provides the following kinds of Java annotation:

- [Basic annotations on page 73.](#)
- [Expression language annotations on page 74.](#)

Basic annotations

[Table 6 on page 73](#) shows the annotations from the `org.apache.camel` Java package that you can use to inject message data into the arguments of a bean method.

Table 6. Basic Bean Annotations

Annotation	Meaning	Parameter?
<code>@Body</code>	Binds to an inbound message body.	
<code>@Header</code>	Binds to an inbound message header.	String name of the header.
<code>@Headers</code>	Binds to a <code>java.util.Map</code> of the inbound message headers.	
<code>@OutHeaders</code>	Binds to a <code>java.util.Map</code> of the outbound message headers.	
<code>@Property</code>	Binds to a named exchange property.	String name of the property.
<code>@Properties</code>	Binds to a <code>java.util.Map</code> of the exchange properties.	

For example, the following class shows you how to use basic annotations to inject message data into the `processExchange()` method arguments.

```
// Java
import org.apache.camel.*;

public class MyBeanProcessor {
    public void processExchange(
        @Header(name="user") String user,
        @Body String body,
        Exchange exchange
    ) {
```

```

// Do whatever you like to 'exchange'...
exchange.getIn().setBody(body + "UserName = " + user);

}
}

```

Notice how you are able to mix the annotations with the default conventions: as well as injecting the annotated arguments, the bean binding also automatically injects the exchange object into the `org.apache.camel.Exchange` argument.

Expression language annotations

The expression language annotations provide a powerful mechanism for injecting message data into a bean method's arguments. Using these annotations, you can invoke an arbitrary script, written in the scripting language of your choice, in order to extract data from an inbound exchange and inject this data into a method argument. [Table 7 on page 74](#) shows the annotations from the `org.apache.camel.language` package (and sub-packages, for the non-core annotations) that you can use to inject message data into the arguments of a bean method.

Table 7. Expression Language Annotations

Annotation	Description
@Bean	Inject a Bean expression.
@BeanShell	Inject a BeanShell expression.
@Constant	Inject a Constant expression
@EL	Inject an EL expression.
@Groovy	Inject a Groovy expression.
@Header	Inject a Header expression.
@JavaScript	Inject a JavaScript expression.
@OGNL	Inject an OGNL expression.
@PHP	Inject a PHP expression.
@Python	Inject a Python expression.
@Ruby	Inject a Ruby expression.
@Simple	Inject a Simple expression.

Annotation	Description
@XPath	Inject an XPath expression.
@XQuery	Inject an XQuery expression.

For example, the following class shows you how to use the @XPath annotation to extract a username and a password from the body of an incoming message in XML format:

```
// Java
import org.apache.camel.language.*;

public class MyBeanProcessor {
    public void checkCredentials(
        @XPath("/credentials/username") String user,
        @XPath("/credentials/password") String pass
    ) {
        // Check the user/pass credentials...
        ...
    }
}
```

The @Bean annotation is a special case. It enables you to inject the result of invoking a registered bean. For example, to inject a correlation ID into a method argument, you could use the @Bean annotation to invoke an ID generator class, as follows:

```
// Java
import org.apache.camel.language.*;

public class MyBeanProcessor {
    public void processCorrelatedMsg(
        @Bean("myCorrIdGenerator") String corrId,
        @Body String body
    ) {
        // Check the user/pass credentials...
        ...
    }
}
```

Where the string, `myCorrIdGenerator`, is the bean ID of the ID generator instance. The ID generator class can be instantiated using the spring `bean` element, as follows:

```
<beans ...>
    ...
    <bean id="myCorrIdGenerator" class="com.acme.MyIdGenerat
```

```
or"/>
</beans>
```

Where the `MySimpleIdGenerator` class could be defined as follows:

```
// Java
package com.acme;

public class MyIdGenerator {

    private UserManager userManager;

    public String generate(
        @Header(name = "user") String user,
        @Body String payload
    ) throws Exception {
        User user = userManager.lookupUser(user);
        String userId = user.getPrimaryId();
        String id = userId + generateHashCodeForPayload(payload);
        return id;
    }
}
```

Notice that you can also use annotations in the referenced bean class, `MyIdGenerator`. The only restriction on the `generate()` method signature is that it must return the correct type to inject into the argument annotated by `@Bean`. Because the `@Bean` annotation does not let you specify a method name, the injection mechanism simply invokes the first method in the referenced bean that has the matching return type.



Note

Some of the language annotations are available in the core component (`@Bean`, `@Constant`, `@Simple`, and `@XPath`). For non-core components, however, you will have to make sure that you load the relevant component. For example, to use the OGNL script, you would have to load the `camel-ognl` component.