



Red Hat JBoss Fuse 6.1

Apache Camel Development Guide

Develop applications with Apache Camel

Red Hat JBoss Fuse 6.1 Apache Camel Development Guide

Develop applications with Apache Camel

JBoss A-MQ Docs Team
Content Services
fuse-docs-support@redhat.com

Legal Notice

Copyright © 2013 Red Hat.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

Guide to developing routes with Apache Camel.

Table of Contents

PART I. IMPLEMENTING ENTERPRISE INTEGRATION PATTERNS	9
CHAPTER 1. BUILDING BLOCKS FOR ROUTE DEFINITIONS	10
1.1. IMPLEMENTING A ROUTEBUILDER CLASS	10
1.2. BASIC JAVA DSL SYNTAX	11
1.3. ROUTER SCHEMA IN A SPRING XML FILE	14
1.4. ENDPOINTS	16
1.5. PROCESSORS	20
CHAPTER 2. BASIC PRINCIPLES OF ROUTE BUILDING	30
2.1. PIPELINE PROCESSING	30
2.2. MULTIPLE INPUTS	33
2.3. EXCEPTION HANDLING	36
2.4. BEAN INTEGRATION	52
2.5. CREATING EXCHANGE INSTANCES	62
2.6. TRANSFORMING MESSAGE CONTENT	63
2.7. PROPERTY PLACEHOLDERS	74
2.8. ASPECT ORIENTED PROGRAMMING	84
2.9. THREADING MODEL	85
2.10. CONTROLLING START-UP AND SHUTDOWN OF ROUTES	93
2.11. SCHEDULED ROUTE POLICY	98
2.12. JMX NAMING	106
2.13. PERFORMANCE AND OPTIMIZATION	108
CHAPTER 3. INTRODUCING ENTERPRISE INTEGRATION PATTERNS	110
3.1. OVERVIEW OF THE PATTERNS	110
CHAPTER 4. MESSAGING SYSTEMS	117
4.1. MESSAGE	117
4.2. MESSAGE CHANNEL	118
4.3. MESSAGE ENDPOINT	120
4.4. PIPES AND FILTERS	121
4.5. MESSAGE ROUTER	123
4.6. MESSAGE TRANSLATOR	125
CHAPTER 5. MESSAGING CHANNELS	127
5.1. POINT-TO-POINT CHANNEL	127
5.2. PUBLISH-SUBSCRIBE CHANNEL	128
5.3. DEAD LETTER CHANNEL	130
5.4. GUARANTEED DELIVERY	139
5.5. MESSAGE BUS	141
CHAPTER 6. MESSAGE CONSTRUCTION	143
6.1. CORRELATION IDENTIFIER	143
6.2. EVENT MESSAGE	143
6.3. RETURN ADDRESS	145
CHAPTER 7. MESSAGE ROUTING	147
7.1. CONTENT-BASED ROUTER	147
7.2. MESSAGE FILTER	148
7.3. RECIPIENT LIST	150
7.4. SPLITTER	159
7.5. AGGREGATOR	168

7.6. RESEQUENCER	187
7.7. ROUTING SLIP	191
7.8. THROTTLER	193
7.9. DELAYER	196
7.10. LOAD BALANCER	198
7.11. MULTICAST	207
7.12. COMPOSED MESSAGE PROCESSOR	214
7.13. SCATTER-GATHER	216
7.14. LOOP	219
7.15. SAMPLING	221
7.16. DYNAMIC ROUTER	223
CHAPTER 8. MESSAGE TRANSFORMATION	227
8.1. CONTENT ENRICHER	227
8.2. CONTENT FILTER	232
8.3. NORMALIZER	233
8.4. CLAIM CHECK	234
8.5. SORT	236
8.6. VALIDATE	238
CHAPTER 9. MESSAGING ENDPOINTS	240
9.1. MESSAGING MAPPER	240
9.2. EVENT DRIVEN CONSUMER	241
9.3. POLLING CONSUMER	242
9.4. COMPETING CONSUMERS	242
9.5. MESSAGE DISPATCHER	244
9.6. SELECTIVE CONSUMER	246
9.7. DURABLE SUBSCRIBER	248
9.8. IDEMPOTENT CONSUMER	251
9.9. TRANSACTIONAL CLIENT	257
9.10. MESSAGING GATEWAY	257
9.11. SERVICE ACTIVATOR	258
CHAPTER 10. SYSTEM MANAGEMENT	261
10.1. DETOUR	261
10.2. LOGEIP	262
10.3. WIRE TAP	263
APPENDIX A. MIGRATING FROM SERVICEMIX EIP	269
A.1. MIGRATING ENDPOINTS	269
A.2. COMMON ELEMENTS	271
A.3. SERVICEMIX EIP PATTERNS	272
A.4. CONTENT-BASED ROUTER	273
A.5. CONTENT ENRICHER	275
A.6. MESSAGE FILTER	277
A.7. PIPELINE	278
A.8. RESEQUENCER	279
A.9. STATIC RECIPIENT LIST	281
A.10. STATIC ROUTING SLIP	282
A.11. WIRE TAP	283
A.12. XPATH SPLITTER	285
PART II. ROUTING EXPRESSION AND PREDICATE LANGUAGES	287
CHAPTER 11. INTRODUCTION	288

11.1. OVERVIEW OF THE LANGUAGES	288
11.2. HOW TO INVOKE AN EXPRESSION LANGUAGE	289
CHAPTER 12. CONSTANT	294
OVERVIEW	294
XML EXAMPLE	294
JAVA EXAMPLE	294
CHAPTER 13. EL	295
OVERVIEW	295
ADDING JUEL PACKAGE	295
STATIC IMPORT	295
VARIABLES	295
EXAMPLE	296
CHAPTER 14. THE FILE LANGUAGE	297
14.1. WHEN TO USE THE FILE LANGUAGE	297
14.2. FILE VARIABLES	298
14.3. EXAMPLES	300
CHAPTER 15. GROOVY	303
OVERVIEW	303
ADDING THE SCRIPT MODULE	303
STATIC IMPORT	303
BUILT-IN ATTRIBUTES	303
EXAMPLE	304
USING THE PROPERTIES COMPONENT	304
CHAPTER 16. HEADER	306
OVERVIEW	306
XML EXAMPLE	306
JAVA EXAMPLE	306
CHAPTER 17. JAVASCRIPT	307
OVERVIEW	307
ADDING THE SCRIPT MODULE	307
STATIC IMPORT	307
BUILT-IN ATTRIBUTES	307
EXAMPLE	308
USING THE PROPERTIES COMPONENT	308
CHAPTER 18. JOSQL	310
OVERVIEW	310
ADDING THE JOSQL MODULE	310
STATIC IMPORT	310
VARIABLES	310
EXAMPLE	311
CHAPTER 19. JXPath	312
OVERVIEW	312
ADDING JXPath PACKAGE	312
VARIABLES	312
EXAMPLE	313
CHAPTER 20. MVEL	314
OVERVIEW	314

SYNTAX	314
ADDING THE MVEL MODULE	314
BUILT-IN VARIABLES	314
EXAMPLE	315
CHAPTER 21. THE OBJECT-GRAPH NAVIGATION LANGUAGE(OGNL)	316
OVERVIEW	316
ADDING THE OGNL MODULE	316
STATIC IMPORT	316
BUILT-IN VARIABLES	316
EXAMPLE	317
CHAPTER 22. PHP	318
OVERVIEW	318
ADDING THE SCRIPT MODULE	318
STATIC IMPORT	318
BUILT-IN ATTRIBUTES	318
EXAMPLE	319
USING THE PROPERTIES COMPONENT	319
CHAPTER 23. PROPERTY	320
OVERVIEW	320
XML EXAMPLE	320
JAVA EXAMPLE	320
CHAPTER 24. PYTHON	321
OVERVIEW	321
ADDING THE SCRIPT MODULE	321
STATIC IMPORT	321
BUILT-IN ATTRIBUTES	321
EXAMPLE	322
USING THE PROPERTIES COMPONENT	322
CHAPTER 25. REF	323
OVERVIEW	323
STATIC IMPORT	323
XML EXAMPLE	323
JAVA EXAMPLE	323
CHAPTER 26. RUBY	324
OVERVIEW	324
ADDING THE SCRIPT MODULE	324
STATIC IMPORT	324
BUILT-IN ATTRIBUTES	324
EXAMPLE	325
USING THE PROPERTIES COMPONENT	325
CHAPTER 27. THE SIMPLE LANGUAGE	326
27.1. JAVA DSL	326
27.2. XML DSL	327
27.3. INVOKING AN EXTERNAL SCRIPT	328
27.4. EXPRESSIONS	329
27.5. PREDICATES	331
27.6. VARIABLE REFERENCE	333
27.7. OPERATOR REFERENCE	337

CHAPTER 28. SPEL	340
OVERVIEW	340
SYNTAX	340
ADDING SPEL PACKAGE	340
VARIABLES	340
XML EXAMPLE	341
JAVA EXAMPLE	341
CHAPTER 29. THE XPATH LANGUAGE	343
29.1. JAVA DSL	343
29.2. XML DSL	344
29.3. XPATH INJECTION	346
29.4. XPATH BUILDER	347
29.5. ENABLING SAXON	348
29.6. EXPRESSIONS	350
29.7. PREDICATES	353
29.8. USING VARIABLES AND FUNCTIONS	354
29.9. VARIABLE NAMESPACES	355
29.10. FUNCTION REFERENCE	356
CHAPTER 30. XQUERY	358
OVERVIEW	358
JAVA SYNTAX	358
ADDING THE SAXON MODULE	358
STATIC IMPORT	358
VARIABLES	358
EXAMPLE	359
PART III. WEB SERVICES AND ROUTING WITH CAMEL CXF	360
CHAPTER 31. DEMONSTRATION CODE FOR CAMEL/CXF	361
31.1. DOWNLOADING AND INSTALLING THE DEMONSTRATIONS	361
31.2. RUNNING THE DEMONSTRATIONS	361
CHAPTER 32. JAVA-FIRST SERVICE IMPLEMENTATION	365
32.1. JAVA-FIRST OVERVIEW	365
32.2. DEFINE SEI AND RELATED CLASSES	366
32.3. ANNOTATE SEI FOR JAX-WS	369
32.4. INSTANTIATE THE WS ENDPOINT	372
32.5. JAVA-TO-WSDL MAVEN PLUG-IN	374
CHAPTER 33. WSDL-FIRST SERVICE IMPLEMENTATION	377
33.1. WSDL-FIRST OVERVIEW	377
33.2. CUSTOMERSERVICE WSDL CONTRACT	378
33.3. WSDL-TO-JAVA MAVEN PLUG-IN	381
33.4. INSTANTIATE THE WS ENDPOINT	383
33.5. DEPLOY TO AN OSGI CONTAINER	384
CHAPTER 34. IMPLEMENTING A WS CLIENT	388
34.1. WS CLIENT OVERVIEW	388
34.2. WSDL-TO-JAVA MAVEN PLUG-IN	389
34.3. INSTANTIATE THE WS CLIENT PROXY	391
34.4. INVOKE WS OPERATIONS	393
34.5. DEPLOY TO AN OSGI CONTAINER	393

CHAPTER 35. POJO-BASED ROUTE	397
35.1. PROCESSING MESSAGES IN POJO FORMAT	397
35.2. WSDL-TO-JAVA MAVEN PLUG-IN	398
35.3. INSTANTIATE THE WS ENDPOINT	400
35.4. SORT MESSAGES BY OPERATION NAME	403
35.5. PROCESS OPERATION PARAMETERS	404
35.6. DEPLOY TO OSGI	406
CHAPTER 36. PAYLOAD-BASED ROUTE	409
36.1. PROCESSING MESSAGES IN PAYLOAD FORMAT	409
36.2. INSTANTIATE THE WS ENDPOINT	410
36.3. SORT MESSAGES BY OPERATION NAME	412
36.4. SOAP/HTTP-TO-JMS BRIDGE USE CASE	413
36.5. GENERATING RESPONSES USING TEMPLATES	416
36.6. TYPECONVERTER FOR CXFPAYLOAD	419
36.7. DEPLOY TO OSGI	420
CHAPTER 37. PROVIDER-BASED ROUTE	423
37.1. PROVIDER-BASED JAX-WS ENDPOINT	423
37.2. CREATE A PROVIDER<?> IMPLEMENTATION CLASS	425
37.3. INSTANTIATE THE WS ENDPOINT	425
37.4. SORT MESSAGES BY OPERATION NAME	427
37.5. SOAP/HTTP-TO-JMS BRIDGE USE CASE	427
37.6. GENERATING RESPONSES USING TEMPLATES	430
37.7. TYPECONVERTER FOR SAXSOURCE	433
37.8. DEPLOY TO OSGI	433
CHAPTER 38. PROXYING A WEB SERVICE	436
38.1. PROXYING WITH HTTP	436
38.2. PROXYING WITH POJO FORMAT	438
38.3. PROXYING WITH PAYLOAD FORMAT	440
38.4. HANDLING HTTP HEADERS	441
CHAPTER 39. FILTERING SOAP MESSAGE HEADERS	444
39.1. BASIC CONFIGURATION	444
39.2. HEADER FILTERING	446
39.3. IMPLEMENTING A CUSTOM FILTER	447
39.4. INSTALLING FILTERS	450
PART IV. PROGRAMMING EIP COMPONENTS	452
CHAPTER 40. UNDERSTANDING MESSAGE FORMATS	453
40.1. EXCHANGES	453
40.2. MESSAGES	454
40.3. BUILT-IN TYPE CONVERTERS	458
40.4. BUILT-IN UUID GENERATORS	460
CHAPTER 41. IMPLEMENTING A PROCESSOR	463
41.1. PROCESSING MODEL	463
41.2. IMPLEMENTING A SIMPLE PROCESSOR	463
41.3. ACCESSING MESSAGE CONTENT	464
41.4. THE EXCHANGEHELPER CLASS	466
CHAPTER 42. TYPE CONVERTERS	468
42.1. TYPE CONVERTER ARCHITECTURE	468

42.2. IMPLEMENTING TYPE CONVERTER USING ANNOTATIONS	470
42.3. IMPLEMENTING A TYPE CONVERTER DIRECTLY	473
CHAPTER 43. PRODUCER AND CONSUMER TEMPLATES	475
43.1. USING THE PRODUCER TEMPLATE	475
43.2. USING THE CONSUMER TEMPLATE	489
CHAPTER 44. IMPLEMENTING A COMPONENT	492
44.1. COMPONENT ARCHITECTURE	492
44.2. HOW TO IMPLEMENT A COMPONENT	500
44.3. AUTO-DISCOVERY AND CONFIGURATION	501
CHAPTER 45. COMPONENT INTERFACE	505
45.1. THE COMPONENT INTERFACE	505
45.2. IMPLEMENTING THE COMPONENT INTERFACE	506
CHAPTER 46. ENDPOINT INTERFACE	511
46.1. THE ENDPOINT INTERFACE	511
46.2. IMPLEMENTING THE ENDPOINT INTERFACE	514
CHAPTER 47. CONSUMER INTERFACE	521
47.1. THE CONSUMER INTERFACE	521
47.2. IMPLEMENTING THE CONSUMER INTERFACE	525
CHAPTER 48. PRODUCER INTERFACE	533
48.1. THE PRODUCER INTERFACE	533
48.2. IMPLEMENTING THE PRODUCER INTERFACE	535
CHAPTER 49. EXCHANGE INTERFACE	538
49.1. THE EXCHANGE INTERFACE	538
CHAPTER 50. MESSAGE INTERFACE	542
50.1. THE MESSAGE INTERFACE	542
50.2. IMPLEMENTING THE MESSAGE INTERFACE	544
INDEX	545

PART I. IMPLEMENTING ENTERPRISE INTEGRATION PATTERNS

Abstract

This part describes how to build routes using Apache Camel. It covers the basic building blocks and EIP components.

CHAPTER 1. BUILDING BLOCKS FOR ROUTE DEFINITIONS

Abstract

Apache Camel supports two alternative *Domain Specific Languages* (DSL) for defining routes: a Java DSL and a Spring XML DSL. The basic building blocks for defining routes are *endpoints* and *processors*, where the behavior of a processor is typically modified by *expressions* or logical *predicates*. Apache Camel enables you to define expressions and predicates using a variety of different languages.

1.1. IMPLEMENTING A ROUTEBUILDER CLASS

Overview

To use the *Domain Specific Language* (DSL), you extend the **RouteBuilder** class and override its **configure()** method (where you define your routing rules).

You can define as many **RouteBuilder** classes as necessary. Each class is instantiated once and is registered with the **CamelContext** object. Normally, the lifecycle of each **RouteBuilder** object is managed automatically by the container in which you deploy the router.

RouteBuilder classes

As a router developer, your core task is to implement one or more **RouteBuilder** classes. There are two alternative **RouteBuilder** classes that you can inherit from:

- **org.apache.camel.builder.RouteBuilder**—this is the generic **RouteBuilder** base class that is suitable for deploying into *any* container type. It is provided in the **camel-core** artifact.
- **org.apache.camel.spring.SpringRouteBuilder**—this base class is specially adapted to the Spring container. In particular, it provides extra support for the following Spring specific features: looking up beans in the Spring registry (using the **beanRef()** Java DSL command) and transactions (see the *Transactions Guide* for details). It is provided in the **camel-spring** artifact.

The **RouteBuilder** class defines methods used to initiate your routing rules (for example, **from()**, **intercept()**, and **exception()**).

Implementing a RouteBuilder

[Example 1.1, “Implementation of a RouteBuilder Class”](#) shows a minimal **RouteBuilder** implementation. The **configure()** method body contains a routing rule; each rule is a single Java statement.

Example 1.1. 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.
    // ...
}
}

```

The form of the rule `from(URL1).to(URL2)` instructs the router to read files from the directory `src/data` and send them to the directory `target/messages`. The option `?noop=true` instructs the router to retain (not delete) the source files in the `src/data` directory.

1.2. BASIC JAVA DSL SYNTAX

What is a DSL?

A Domain Specific Language (DSL) is a mini-language designed for a special purpose. A DSL does not have to be logically complete but needs enough expressive power to describe problems adequately in the chosen domain. Typically, a DSL does not require a dedicated parser, interpreter, or compiler. A DSL can piggyback on top of an existing object-oriented host language, provided DSL constructs map cleanly to constructs in the host language API.

Consider the following sequence of commands in a hypothetical DSL:

```

command01;
command02;
command03;

```

You can map these commands to Java method invocations, as follows:

```

command01().command02().command03()

```

You can even map blocks to Java method invocations. For example:

```


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

```


The DSL syntax is implicitly defined by the data types of the host language API. For example, the return type of a Java method determines which methods you can legally invoke next (equivalent to the next command in the DSL).

Router rule syntax

Apache Camel 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 1.1, “Local Routing Rules”](#) shows an overview of the basic syntax for defining local routing rules.

Figure 1.1. Local Routing Rules

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

**NOTE**

You can also define a global routing rule, by starting the rule with a special processor type (such as `intercept()`, `exception()`, or `errorHandler()`). Global rules are outside the scope of this guide.

Consumers and producers

A local rule always starts by defining a consumer endpoint, using `from("EndpointURL")`, and typically (but not always) ends by defining a producer 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`, an Apache CXF endpoint, `cxf:MyServiceName`, or an Apache ActiveMQ endpoint, `activemq:queue:MyQName`. For a complete list of component types, see

Exchanges

An *exchange* object consists of a message, augmented by metadata. Exchanges are of central importance in Apache Camel, because the exchange is the standard form in which messages are propagated through routing rules. The main constituents of an exchange are, as follows:

- *In* message—is the current message encapsulated by the exchange. As the exchange progresses through a route, this message may be modified. So the *In* message at the start of a route is typically *not* the same as the *In* message at the end of the route. The `org.apache.camel.Message` type provides a generic model of a message, with the following parts:
 - Body.
 - Headers.
 - Attachments.

It is important to realize that this is a *generic* model of a message. Apache Camel supports a large variety of protocols and endpoint types. Hence, it is *not* possible to standardize the format of the message body or the message headers. For example, the body of a JMS message would have a completely different format to the body of a HTTP message or a Web services message. For this reason, the body and the headers are declared to be of **Object** type. The original content of the body and the headers is then determined by the endpoint that created the exchange instance (that is, the endpoint appearing in the **from()** command).

- *Out* message—is a temporary holding area for a reply message or for a transformed message. Certain processing nodes (in particular, the **to()** command) can modify the current message by treating the *In* message as a request, sending it to a producer endpoint, and then receiving a reply from that endpoint. The reply message is then inserted into the *Out* message slot in the exchange.

Normally, if an *Out* message has been set by the current node, Apache Camel modifies the exchange as follows before passing it to the next node in the route: the old *In* message is discarded and the *Out* message is moved to the *In* message slot. Thus, the reply becomes the new current message. For a more detailed discussion of how Apache Camel connects nodes together in a route, see [Section 2.1, “Pipeline Processing”](#).

There is one special case where an *Out* message is treated differently, however. If the consumer endpoint at the start of a route is expecting a reply message, the *Out* message at the very end of the route is taken to be the consumer endpoint's reply message (and, what is more, in this case the final node *must* create an *Out* message or the consumer endpoint would hang) .

- Message exchange pattern (MEP)—affects the interaction between the exchange and endpoints in the route, as follows:
 - *Consumer endpoint*—the consumer endpoint that creates the original exchange sets the initial value of the MEP. The initial value indicates whether the consumer endpoint expects to receive a reply (for example, the *InOut* MEP) or not (for example, the *InOnly* MEP).
 - *Producer endpoints*—the MEP affects the producer endpoints that the exchange encounters along the route (for example, when an exchange passes through a **to()** node). For example, if the current MEP is *InOnly*, a **to()** node would not expect to receive a reply from the endpoint. Sometimes you need to change the current MEP in order to customize the exchange's interaction with a producer endpoint. For more details, see [Section 1.4, “Endpoints”](#).
- Exchange properties—a list of named properties containing metadata for the current message.

Message exchange patterns

Using an **Exchange** object makes it easy to generalize message processing to different *message exchange patterns*. For example, an asynchronous protocol might define an MEP that consists of a single message that flows from the consumer endpoint to the producer endpoint (an *InOnly* MEP). An RPC protocol, on the other hand, might define an MEP that consists of a request message and a reply message (an *InOut* MEP). Currently, Apache Camel supports the following MEPs:

- **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`.

Grouped exchanges

Sometimes it is useful to have a single exchange that encapsulates multiple exchange instances. For this purpose, you can use a *grouped exchange*. A grouped exchange is essentially an exchange instance that contains a `java.util.List` of **Exchange** objects stored in the `Exchange.GROUPED_EXCHANGE` exchange property. For an example of how to use grouped exchanges, see [Section 7.5, “Aggregator”](#).

Processors

A *processor* is a node in a route that can access and modify the stream of exchanges passing through the route. Processors can take *expression* or *predicate* arguments, that modify their behavior. For example, the rule shown in [Figure 1.1, “Local Routing Rules”](#) 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. 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 [Expression and Predicate Languages](#)).

1.3. ROUTER SCHEMA IN A SPRING XML FILE

Namespace

The router schema—which defines the XML DSL—belongs to the following XML schema namespace:

```
http://camel.apache.org/schema/spring
```

Specifying the schema location

The location of the router schema is normally specified to be <http://camel.apache.org/schema/spring/camel-spring.xsd>, which references the latest version of the schema on the Apache Web site. For example, the root **beans** element of an Apache Camel Spring file is normally configured as shown in [Example 1.2, “Specifying the Router Schema Location”](#).

Example 1.2. Specifying the Router Schema Location

```
<beans xmlns="http://www.springframework.org/schema/beans"
      xmlns:camel="http://camel.apache.org/schema/spring"
      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:schemaLocation="
        http://www.springframework.org/schema/beans
        http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
        http://camel.apache.org/schema/spring
        http://camel.apache.org/schema/spring/camel-spring.xsd">

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

Runtime schema location

At run time, Apache Camel does *not* download the router schema from schema location specified in the Spring file. Instead, Apache Camel automatically picks up a copy of the schema from the root directory of the **camel-spring** JAR file. This ensures that the version of the schema used to parse the Spring file always matches the current runtime version. This is important, because the latest version of the schema posted up on the Apache Web site might not match the version of the runtime you are currently using.

Using an XML editor

Generally, it is recommended that you edit your Spring files using a full-feature XML editor. An XML editor's auto-completion features make it much easier to author XML that complies with the router schema and the editor can warn you instantly, if the XML is badly-formed.

XML editors generally *do* rely on downloading the schema from the location that you specify in the **xsi:schemaLocation** attribute. In order to be sure you are using the correct schema version whilst editing, it is usually a good idea to select a specific version of the **camel-spring.xsd** file. For example, to edit a Spring file for the 2.3 version of Apache Camel, you could modify the beans element as follows:

```
<beans xmlns="http://www.springframework.org/schema/beans"
      xmlns:camel="http://camel.apache.org/schema/spring"
      xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
      xsi:schemaLocation="
        http://www.springframework.org/schema/beans
        http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
```

```
http://camel.apache.org/schema/spring
http://camel.apache.org/schema/spring/camel-spring-2.3.0.xsd">
...
```

Change back to the default, **camel-spring.xsd**, when you are finished editing. To see which schema versions are currently available for download, navigate to the Web page, <http://camel.apache.org/schema/spring>.

1.4. ENDPOINTS

Overview

Apache Camel endpoints are the sources and sinks of messages in a route. An endpoint is a very general sort of building block: the only requirement it must satisfy is that it acts either as a source of messages (a consumer endpoint) or as a sink of messages (a producer endpoint). Hence, there are a great variety of different endpoint types supported in Apache Camel, ranging from protocol supporting endpoints, such as HTTP, to simple timer endpoints, such as Quartz, that generate dummy messages at regular time intervals. One of the major strengths of Apache Camel is that it is relatively easy to add a custom component that implements a new endpoint type.

Endpoint URIs

Endpoints are identified by *endpoint URIs*, which have the following general form:

```
scheme:contextPath[?queryOptions]
```

The URI *scheme* identifies a protocol, such as **http**, and the *contextPath* provides URI details that are interpreted by the protocol. In addition, most schemes allow you to define query options, *queryOptions*, which are specified in the following format:

```
?option01=value01&option02=value02&...
```

For example, the following HTTP URI can be used to connect to the Google search engine page:

```
http://www.google.com
```

The following File URI can be used to read all of the files appearing under the **C:\temp\src\data** directory:

```
file://C:/temp/src/data
```

Not every *scheme* represents a protocol. Sometimes a *scheme* just provides access to a useful utility, such as a timer. For example, the following Timer endpoint URI generates an exchange every second (=1000 milliseconds). You could use this to schedule activity in a route.

```
timer://tickTock?period=1000
```

Specifying time periods in a URI

Many of the Apache Camel components have options whose value is a time period (for example, for specifying timeout values and so on). By default, such time period options are normally specified as a pure number, which is interpreted as a millisecond time period. But Apache Camel also supports a more readable syntax for time periods, which enables you to express the period in hours, minutes, and seconds. Formally, the human-readable time period is a string that conforms to the following syntax:

```
[NHour(h|hour)][NMin(m|minute)][NSec(s|second)]
```

Where each term in square brackets, `[]`, is optional and the notation, `(A|B)`, indicates that **A** and **B** are alternatives.

For example, you can configure **timer** endpoint with a 45 minute period as follows:

```
from("timer:foo?period=45m")
  .to("log:foo");
```

You can also use arbitrary combinations of the hour, minute, and second units, as follows:

```
from("timer:foo?period=1h15m")
  .to("log:foo");
from("timer:bar?period=2h30s")
  .to("log:bar");
from("timer:bar?period=3h45m58s")
  .to("log:bar");
```

Specifying raw values in URI options

By default, the option values that you specify in a URI are automatically URI-encoded. In some cases this is undesirable behavior. For example, when setting a password option, it is preferable to transmit the raw character string *without* URI encoding.

It is possible to switch off URI encoding by specifying an option value with the syntax, **RAW(RawValue)**. For example,

```
from("SourceURI")
  .to("ftp:joe@myftpserver.com?password=RAW(se+re?t&23)&binary=true");
```

In this example, the password value is transmitted as the literal value, **se+re?t&23**.

Apache Camel components

Each URI *scheme* maps to a *Apache Camel component*, where a Apache Camel component is essentially an endpoint factory. In other words, to use a particular type of endpoint, you must deploy the corresponding Apache Camel component in your runtime container. For example, to use JMS endpoints, you would deploy the JMS component in your container.

Apache Camel provides a large variety of different components that enable you to integrate your application with various transport protocols and third-party products. For example, some of the more commonly used components are: File, JMS, CXF (Web services), HTTP, Jetty, Direct, and Mock. For the full list of supported components, see the [Apache Camel component documentation](#).

Most of the Apache Camel components are packaged separately to the Camel core. If you

use Maven to build your application, you can easily add a component (and its third-party dependencies) to your application simply by adding a dependency on the relevant component artifact. For example, to include the HTTP component, you would add the following Maven dependency to your project POM file:

```
<!-- Maven POM File -->
<properties>
  <camel-version>2.12.0.redhat-610379</camel-version>
  ...
</properties>

<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-http</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

The following components are built-in to the Camel core (in the **camel-core** artifact), so they are always available:

- Bean
- Browse
- Dataset
- Direct
- File
- Log
- Mock
- Properties
- Ref
- SEDA
- Timer
- VM

Consumer endpoints

A *consumer endpoint* is an endpoint that appears at the *start* of a route (that is, in a **from()** DSL command). In other words, the consumer endpoint is responsible for initiating processing in a route: it creates a new exchange instance (typically, based on some message that it has received or obtained), and provides a thread to process the exchange in the rest of the route.

For example, the following JMS consumer endpoint pulls messages off the **payments** queue and processes them in the route:

```
from("jms:queue:payments")
    .process(SomeProcessor)
    .to("TargetURI");
```

Or equivalently, in Spring XML:

```
<camelContext id="CamelContextID"
    xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="jms:queue:payments"/>
        <process ref="someProcessorId"/>
        <to uri="TargetURI"/>
    </route>
</camelContext>
```

Some components are *consumer only*—that is, they can only be used to define consumer endpoints. For example, the Quartz component is used exclusively to define consumer endpoints. The following Quartz endpoint generates an event every second (1000 milliseconds):

```
from("quartz://secondTimer?trigger.repeatInterval=1000")
    .process(SomeProcessor)
    .to("TargetURI");
```

If you like, you can specify the endpoint URI as a formatted string, using the **fromF()** Java DSL command. For example, to substitute the username and password into the URI for an FTP endpoint, you could write the route in Java, as follows:

```
fromF("ftp:%s@fusesource.com?password=%s", username, password)
    .process(SomeProcessor)
    .to("TargetURI");
```

Where the first occurrence of **%s** is replaced by the value of the **username** string and the second occurrence of **%s** is replaced by the **password** string. This string formatting mechanism is implemented by **String.format()** and is similar to the formatting provided by the C **printf()** function. For details, see [java.util.Formatter](#).

Producer endpoints

A *producer endpoint* is an endpoint that appears in the *middle* or at the *end* of a route (for example, in a **to()** DSL command). In other words, the producer endpoint receives an existing exchange object and sends the contents of the exchange to the specified endpoint.

For example, the following JMS producer endpoint pushes the contents of the current exchange onto the specified JMS queue:

```
from("SourceURI")
    .process(SomeProcessor)
    .to("jms:queue:orderForms");
```

Or equivalently in Spring XML:

```
<camelContext id="CamelContextID"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURI"/>
    <process ref="someProcessorId"/>
    <to uri="jms:queue:orderForms"/>
  </route>
</camelContext>
```

Some components are *producer only*—that is, they can only be used to define producer endpoints. For example, the HTTP endpoint is used exclusively to define producer endpoints.

```
from("SourceURI")
  .process(SomeProcessor)
  .to("http://www.google.com/search?hl=en&q=camel+router");
```

If you like, you can specify the endpoint URI as a formatted string, using the **toF()** Java DSL command. For example, to substitute a custom Google query into the HTTP URI, you could write the route in Java, as follows:

```
from("SourceURI")
  .process(SomeProcessor)
  .toF("http://www.google.com/search?hl=en&q=%s", myGoogleQuery);
```

Where the occurrence of **%s** is replaced by your custom query string, **myGoogleQuery**. For details, see [java.util.Formatter](#).

1.5. PROCESSORS

Overview

To enable the router to do something more interesting than simply connecting a consumer endpoint to a producer endpoint, you can add *processors* to your route. A processor is a command you can insert into a routing rule to perform arbitrary processing of messages that flow through the rule. Apache Camel provides a wide variety of different processors, as shown in [Table 1.1, “Apache Camel Processors”](#).

Table 1.1. Apache Camel Processors

Java DSL	XML DSL	Description
aggregate()	aggregate	Aggregator EIP : Creates an aggregator, which combines multiple incoming exchanges into a single exchange.
aop()	aop	Use Aspect Oriented Programming (AOP) to do work before and after a specified sub-route. See Section 2.8, “Aspect Oriented Programming” .

Java DSL	XML DSL	Description
<code>bean()</code> , <code>beanRef()</code>	<code>bean</code>	Process the current exchange by invoking a method on a Java object (or bean). See Section 2.4, “Bean Integration” .
<code>choice()</code>	<code>choice</code>	Content Based Router EIP : Selects a particular sub-route based on the exchange content, using when and otherwise clauses.
<code>convertBodyTo()</code>	<code>convertBodyTo</code>	Converts the <i>In</i> message body to the specified type.
<code>delay()</code>	<code>delay</code>	Delayer EIP : Delays the propagation of the exchange to the latter part of the route.
<code>doTry()</code>	<code>doTry</code>	Creates a try/catch block for handling exceptions, using doCatch , doFinally , and end clauses.
<code>end()</code>	N/A	Ends the current command block.
<code>enrich()</code> , <code>enrichRef()</code>	<code>enrich</code>	Content Enricher EIP : Combines the current exchange with data requested from a specified <i>producer</i> endpoint URI.
<code>filter()</code>	<code>filter</code>	Message Filter EIP : Uses a predicate expression to filter incoming exchanges.
<code>idempotentConsumer()</code>	<code>idempotentConsumer</code>	Idempotent Consumer EIP : Implements a strategy to suppress duplicate messages.
<code>inheritErrorHandler()</code>	<code>@inheritErrorHandler</code>	Boolean option that can be used to disable the inherited error handler on a particular route node (defined as a sub-clause in the Java DSL and as an attribute in the XML DSL).

Java DSL	XML DSL	Description
<code>inOnly()</code>	<code>inOnly</code>	Either sets the current exchange's MEP to <i>InOnly</i> (if no arguments) or sends the exchange as an <i>InOnly</i> to the specified endpoint(s).
<code>inOut()</code>	<code>inOut</code>	Either sets the current exchange's MEP to <i>InOut</i> (if no arguments) or sends the exchange as an <i>InOut</i> to the specified endpoint(s).
<code>loadBalance()</code>	<code>loadBalance</code>	Load Balancer EIP : Implements load balancing over a collection of endpoints.
<code>log()</code>	<code>log</code>	Logs a message to the console.
<code>loop()</code>	<code>loop</code>	Loop EIP : Repeatedly resends each exchange to the latter part of the route.
<code>markRollbackOnly()</code>	<code>@markRollbackOnly</code>	<i>(Transactions)</i> Marks the current transaction for rollback only (no exception is raised). In the XML DSL, this option is set as a boolean attribute on the rollback element. See " Transaction Guide ".
<code>markRollbackOnlyLast()</code>	<code>@markRollbackOnlyLast</code>	<i>(Transactions)</i> If one or more transactions have previously been associated with this thread and then suspended, this command marks the latest transaction for rollback only (no exception is raised). In the XML DSL, this option is set as a boolean attribute on the rollback element. See " Transaction Guide ".
<code>marshal()</code>	<code>marshal</code>	Transforms into a low-level or binary format using the specified data format, in preparation for sending over a particular transport protocol.

Java DSL	XML DSL	Description
<code>multicast()</code>	<code>multicast</code>	Multicast EIP : Multicasts the current exchange to multiple destinations, where each destination gets its own copy of the exchange.
<code>onCompletion()</code>	<code>onCompletion</code>	Defines a sub-route (terminated by <code>end()</code> in the Java DSL) that gets executed after the main route has completed. For conditional execution, use the <code>onWhen</code> sub-clause. Can also be defined on its own line (not in a route).
<code>onException()</code>	<code>onException</code>	Defines a sub-route (terminated by <code>end()</code> in the Java DSL) that gets executed whenever the specified exception occurs. Usually defined on its own line (not in a route).
<code>pipeline()</code>	<code>pipeline</code>	Pipes and Filters EIP : Sends the exchange to a series of endpoints, where the output of one endpoint becomes the input of the next endpoint. See also Section 2.1, "Pipeline Processing" .
<code>policy()</code>	<code>policy</code>	Apply a policy to the current route (currently only used for transactional policies—see "Transaction Guide").
<code>pollEnrich().pollEnrichRef()</code>	<code>pollEnrich</code>	Content Enricher EIP : Combines the current exchange with data polled from a specified <i>consumer</i> endpoint URI.
<code>process().processRef</code>	<code>process</code>	Execute a custom processor on the current exchange. See the section called "Custom processor" and Part IV, "Programming EIP Components" .
<code>recipientList()</code>	<code>recipientList</code>	Recipient List EIP : Sends the exchange to a list of recipients that is calculated at runtime (for example, based on the contents of a header).

Java DSL	XML DSL	Description
<code>removeHeader()</code>	<code>removeHeader</code>	Removes the specified header from the exchange's <i>In</i> message.
<code>removeHeaders()</code>	<code>removeHeaders</code>	Removes the headers matching the specified pattern from the exchange's <i>In</i> message. The pattern can have the form, prefix* —in which case it matches every name starting with prefix—otherwise, it is interpreted as a regular expression.
<code>removeProperty()</code>	<code>removeProperty</code>	Removes the specified exchange property from the exchange.
<code>resequence()</code>	<code>resequence</code>	Resequencer EIP : Re-orders incoming exchanges on the basis of a specified comparator operation. Supports a <i>batch</i> mode and a <i>stream</i> mode.
<code>rollback()</code>	<code>rollback</code>	<i>(Transactions)</i> Marks the current transaction for rollback only (also raising an exception, by default). See "Transaction Guide" .
<code>routingSlip()</code>	<code>routingSlip</code>	Routing Slip EIP : Routes the exchange through a pipeline that is constructed dynamically, based on the list of endpoint URIs extracted from a slip header.
<code>sample()</code>	<code>sample</code>	Creates a sampling throttler, allowing you to extract a sample of exchanges from the traffic on a route.
<code>setBody()</code>	<code>setBody</code>	Sets the message body of the exchange's <i>In</i> message.
<code>setExchangePattern()</code>	<code>setExchangePattern</code>	Sets the current exchange's MEP to the specified value. See the section called "Message exchange patterns" .

Java DSL	XML DSL	Description
<code>setHeader()</code>	<code>setHeader</code>	Sets the specified header in the exchange's <i>In</i> message.
<code>setOutHeader()</code>	<code>setOutHeader</code>	Sets the specified header in the exchange's <i>Out</i> message.
<code>setProperty()</code>	<code>setProperty()</code>	Sets the specified exchange property.
<code>sort()</code>	<code>sort</code>	Sorts the contents of the <i>In</i> message body (where a custom comparator can optionally be specified).
<code>split()</code>	<code>split</code>	Splitter EIP : Splits the current exchange into a sequence of exchanges, where each split exchange contains a fragment of the original message body.
<code>stop()</code>	<code>stop</code>	Stops routing the current exchange and marks it as completed.
<code>threads()</code>	<code>threads</code>	Creates a thread pool for concurrent processing of the latter part of the route.
<code>throttle()</code>	<code>throttle</code>	Throttler EIP : Limit the flow rate to the specified level (exchanges per second).
<code>throwException()</code>	<code>throwException</code>	Throw the specified Java exception.
<code>to()</code>	<code>to</code>	Send the exchange to one or more endpoints. See Section 2.1, "Pipeline Processing" .
<code>toF()</code>	<i>N/A</i>	Send the exchange to an endpoint, using string formatting. That is, the endpoint URI string can embed substitutions in the style of the C <code>printf()</code> function.
<code>transacted()</code>	<code>transacted</code>	Create a Spring transaction scope that encloses the latter part of the route. See "Transaction Guide" .

Java DSL	XML DSL	Description
<code>transform()</code>	<code>transform</code>	Message Translator EIP : Copy the <i>In</i> message headers to the <i>Out</i> message headers and set the <i>Out</i> message body to the specified value.
<code>unmarshal()</code>	<code>unmarshal</code>	Transforms the <i>In</i> message body from a low-level or binary format to a high-level format, using the specified data format.
<code>validate()</code>	<code>validate</code>	Takes a predicate expression to test whether the current message is valid. If the predicate returns false , throws a PredicateValidationException exception.
<code>wireTap()</code>	<code>wireTap</code>	Wire Tap EIP : Sends a copy of the current exchange to the specified wire tap URI, using the ExchangePattern.InOnly MEP.

Some sample processors

To get some idea of how to use processors in a route, see the following examples:

- [the section called “Choice”](#).
- [the section called “Filter”](#).
- [the section called “Throttler”](#).
- [the section called “Custom processor”](#).

Choice

The `choice()` processor is a conditional statement that is used to route incoming messages to alternative producer endpoints. Each alternative producer endpoint is 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");
```

Or equivalently in Spring XML:

```
<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>
```

In the Java DSL, there is a special case where you might need to use the **endChoice()** command. Some of the standard Apache Camel processors enable you to specify extra parameters using special sub-clauses, effectively opening an extra level of nesting which is usually terminated by the **end()** command. For example, you could specify a load balancer clause as **loadBalance().roundRobin().to("mock:foo").to("mock:bar").end()**, which load balances messages between the **mock:foo** and **mock:bar** endpoints. If the load balancer clause is embedded in a choice condition, however, it is necessary to terminate the clause using the **endChoice()** command, as follows:

```
from("direct:start")
  .choice()
    .when(bodyAs(String.class).contains("Camel"))

  .loadBalance().roundRobin().to("mock:foo").to("mock:bar").endChoice()
    .otherwise()
      .to("mock:result");
```

Filter

The **filter()** processor can be used to prevent uninteresting messages from reaching the producer endpoint. It takes a single predicate argument: if the predicate is true, the message exchange is allowed through to the producer; 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");
```

Or equivalently in Spring XML:

```
<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>

```

Throttler

The `throttle()` processor ensures that a producer 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 producer endpoint. For example, to limit the rate of throughput to 100 messages per second, you can define the following rule:

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

Or equivalently in Spring XML:

```

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

```

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 overrides the `process()` method. The following custom processor, `MyProcessor`, removes the header named `foo` from incoming messages:

Example 1.3. 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 1.3, “Implementing a Custom Processor Class”](#):

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

CHAPTER 2. BASIC PRINCIPLES OF ROUTE BUILDING

Abstract

Apache Camel provides several processors and components that you can link together in a route. This chapter provides a basic orientation by explaining the principles of building a route using the provided building blocks.

2.1. PIPELINE PROCESSING

Overview


In Apache Camel, pipelining is the dominant paradigm for connecting nodes in a route definition. The pipeline concept is probably most familiar to users of the UNIX operating system, where it is used to join operating system commands. For example, `ls | more` is an example of a command that pipes a directory listing, `ls`, to the 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()`, and `to()` 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 2.1, "Processor Modifying an In Message"](#). The exchange's *Out* message remains `null` in this case.

Figure 2.1. Processor Modifying an In Message




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 2.2, "Processor Creating an Out Message"](#).

Figure 2.2. Processor Creating an Out Message


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 2.3, “Sample Pipeline for InOnly Exchanges” 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 2.3, “Sample Pipeline for InOnly Exchanges”.

Figure 2.3. Sample Pipeline for InOnly Exchanges

Apache Camel employs the pipeline pattern by default, so 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(ExchangePattern.InOut, "velocity:file:AdressTemplate.vm")
  .to("activemq:envelopeAddresses");
```

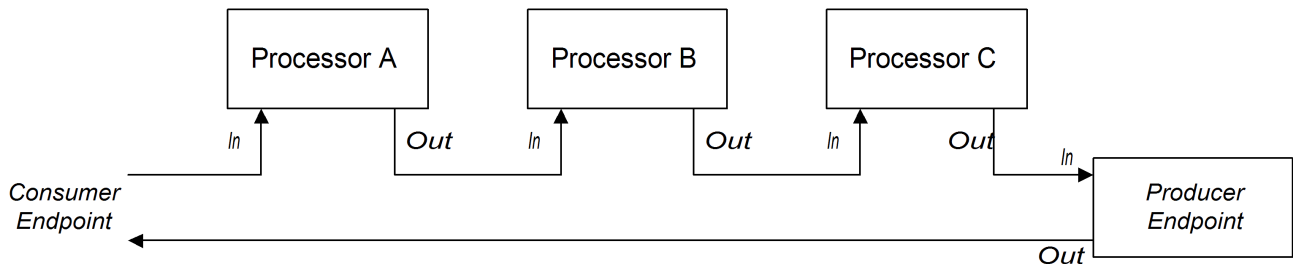
Where the Velocity endpoint, **velocity:file:AdressTemplate.vm**, specifies the location of a Velocity template file, **file:AdressTemplate.vm**, in the file system. The **to()** command changes the exchange pattern to *InOut* before sending the exchange to the Velocity

endpoint and then changes it back to *InOnly* afterwards. For more details of the Velocity endpoint, see .

Pipeline for InOut exchanges

Figure 2.4, “Sample Pipeline for InOut Exchanges” 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 2.4. 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 and wait indefinitely for a reply message. You should be aware that not all producer endpoints create *Out* messages.

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, **`cxf:bean:addAccountDetails`**, **`cxf:bean:getCreditRating`**, and **`cxf: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("cxf:bean:addAccountDetails", "cxf:bean:getCreditRating",
    "cxf:bean:processTransaction");
  
```

Pipeline for InOptionalOut exchanges

The pipeline for *InOptionalOut* exchanges is essentially the same as the pipeline in Figure 2.4, “Sample Pipeline for InOut Exchanges”. The difference between *InOut* and *InOptionalOut* is that an exchange with the *InOptionalOut* exchange pattern is allowed to have a null *Out* message as a reply. That is, in the case of an *InOptionalOut* exchange, a

`null Out` message is copied to the `In` message of the next node in the pipeline. By contrast, in the case of an `InOut` exchange, a `null Out` message is discarded and the original `In` message from the current node would be copied to the `In` message of the next node instead.

2.2. 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? Apache Camel provides several alternatives for specifying multiple inputs to a route. The approach to take depends on whether you want the exchanges to be processed independently of each other or whether you want the exchanges from different inputs to be combined in some way (in which case, you should use the [the section called “Content enricher pattern”](#)).

Multiple independent inputs

The simplest way to specify multiple inputs is using the multi-argument form of the `from()` DSL command, for example:

```
from("URI1", "URI2", "URI3").to("DestinationUri");
```

Or you can use the following equivalent syntax:

```
from("URI1").from("URI2").from("URI3").to("DestinationUri");
```

In both of these examples, exchanges from each of the input endpoints, `URI1`, `URI2`, and `URI3`, are processed independently of each other and in separate threads. In fact, you can think of the preceding route as being equivalent to the following three separate routes:

```
from("URI1").to("DestinationUri");
from("URI2").to("DestinationUri");
from("URI3").to("DestinationUri");
```

Segmented routes

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 2.5, “Processing Multiple Inputs with Segmented Routes”](#).

Figure 2.5. 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 an endpoint that is intended only for use *within* a router application. The following types of endpoints are suitable for internal use:

- the section called “Direct endpoints”.
- the section called “SEDA endpoints”.
- the section called “VM endpoints”.

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

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, if 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 can 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 to the **activemq:USTxn** queue.

The implementation of the direct endpoint behaves as follows: 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.

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 (the 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. 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 Apache Camel 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 Apache Camel applications, as long as they are running within the same Java virtual machine.

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 as follows:

```
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. 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.

For example, suppose you are writing an application that processes credit requests. Before processing a credit request, you need to augment it with the data that assigns a credit rating to the customer, where the ratings data is stored in a file in the directory, **src/data/ratings**. You can combine the incoming credit request with data from the ratings file using the **pollEnrich()** pattern and a **GroupedExchangeAggregationStrategy** aggregation strategy, as follows:

```
from("jms:queue:creditRequests")
    .pollEnrich("file:src/data/ratings?noop=true", new
GroupedExchangeAggregationStrategy())
    .bean(new MergeCreditRequestAndRatings(), "merge")
    .to("jms:queue:reformattedRequests");
```

Where the **GroupedExchangeAggregationStrategy** class is a standard aggregation strategy from the **org.apache.camel.processor.aggregate** package that adds each new exchange to a **java.util.List** instance and stores the resulting list in the **Exchange.GROUPED_EXCHANGE** exchange property. In this case, the list contains two elements: the original exchange (from the **creditRequests** JMS queue); and the enricher exchange (from the file endpoint).

To access the grouped exchange, you can use code like the following:

```
public class MergeCreditRequestAndRatings {
    public void merge(Exchange ex) {
        // Obtain the grouped exchange
        List<Exchange> list = ex.getProperty(Exchange.GROUPED_EXCHANGE,
List.class);

        // Get the exchanges from the grouped exchange
        Exchange originalEx = list.get(0);
        Exchange ratingsEx = list.get(1);

        // Merge the exchanges
        ...
    }
}
```

An alternative approach to this application would be to put the merge code directly into the implementation of the custom aggregation strategy class.

For more details about the content enricher pattern, see [Section 8.1, “Content Enricher”](#).

2.3. EXCEPTION HANDLING

Abstract

Apache Camel provides several different mechanisms, which let you handle exceptions at different levels of granularity: you can handle exceptions within a route using **doTry**, **doCatch**, and **doFinally**; or you can specify what action to take for each exception type

and apply this rule to all routes in a **RouteBuilder** using **onException**; or you can specify what action to take for *all* exception types and apply this rule to all routes in a **RouteBuilder** using **errorHandler**.

For more details about exception handling, see [Section 5.3, “Dead Letter Channel”](#).

2.3.1. onException Clause

Overview

The **onException** clause is a powerful mechanism for trapping exceptions that occur in one or more routes: it is type-specific, enabling you to define distinct actions to handle different exception types; it allows you to define actions using essentially the same (actually, slightly extended) syntax as a route, giving you considerable flexibility in the way you handle exceptions; and it is based on a trapping model, which enables a single **onException** clause to deal with exceptions occurring at any node in any route.

Trapping exceptions using onException

The **onException** clause is a mechanism for *trapping*, rather than catching exceptions. That is, once you define an **onException** clause, it traps exceptions that occur at any point in a route. This contrasts with the Java try/catch mechanism, where an exception is caught, only if a particular code fragment is *explicitly* enclosed in a try block.

What really happens when you define an **onException** clause is that the Apache Camel runtime implicitly encloses each route node in a try block. This is why the **onException** clause is able to trap exceptions at any point in the route. But this wrapping is done for you automatically; it is not visible in the route definitions.

Java DSL example

In the following Java DSL example, the **onException** clause applies to all of the routes defined in the **RouteBuilder** class. If a **ValidationException** exception occurs while processing either of the routes (**from("seda:inputA")** or **from("seda:inputB")**), the **onException** clause traps the exception and redirects the current exchange to the **validationFailed** JMS queue (which serves as a deadletter queue).

```
// Java
public class MyRouteBuilder extends RouteBuilder {

    public void configure() {
        onException(ValidationException.class)
            .to("activemq:validationFailed");

        from("seda:inputA")
            .to("validation:foo/bar.xsd", "activemq:someQueue");

        from("seda:inputB").to("direct:foo")
            .to("rnc:mySchema.rnc", "activemq:anotherQueue");
    }
}
```

XML DSL example

The preceding example can also be expressed in the XML DSL, using the **onException** element to define the exception clause, as follows:

```
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:camel="http://camel.apache.org/schema/spring"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
    http://camel.apache.org/schema/spring
    http://camel.apache.org/schema/spring/camel-spring.xsd">

  <camelContext xmlns="http://camel.apache.org/schema/spring">
    <onException>
      <exception>com.mycompany.ValidationException</exception>
      <to uri="activemq:validationFailed"/>
    </onException>
    <route>
      <from uri="seda:inputA"/>
      <to uri="validation:foo/bar.xsd"/>
      <to uri="activemq:someQueue"/>
    </route>
    <route>
      <from uri="seda:inputB"/>
      <to uri="rnc:mySchema.rnc"/>
      <to uri="activemq:anotherQueue"/>
    </route>
  </camelContext>

</beans>
```

Trapping multiple exceptions

You can define multiple **onException** clauses to trap exceptions in a **RouteBuilder** scope. This enables you to take different actions in response to different exceptions. For example, the following series of **onException** clauses defined in the Java DSL define different deadletter destinations for **ValidationException**, **IOException**, and **Exception**:

```
onException(ValidationException.class).to("activemq:validationFailed");
onException(java.io.IOException.class).to("activemq:ioExceptions");
onException(Exception.class).to("activemq:exceptions");
```

You can define the same series of **onException** clauses in the XML DSL as follows:

```
<onException>
  <exception>com.mycompany.ValidationException</exception>
  <to uri="activemq:validationFailed"/>
</onException>
<onException>
  <exception>java.io.IOException</exception>
  <to uri="activemq:ioExceptions"/>
</onException>
<onException>
```

```

    <exception>java.lang.Exception</exception>
    <to uri="activemq:exceptions"/>
</onException>

```

You can also group multiple exceptions together to be trapped by the same **onException** clause. In the Java DSL, you can group multiple exceptions as follows:

```

onException(ValidationException.class, BuesinessException.class)
    .to("activemq:validationFailed");

```

In the XML DSL, you can group multiple exceptions together by defining more than one **exception** element inside the **onException** element, as follows:

```

<onException>
    <exception>com.mycompany.ValidationException</exception>
    <exception>com.mycompany.BuesinessException</exception>
    <to uri="activemq:validationFailed"/>
</onException>

```

When trapping multiple exceptions, the order of the **onException** clauses is significant. Apache Camel initially attempts to match the thrown exception against the *first* clause. If the first clause fails to match, the next **onException** clause is tried, and so on until a match is found. Each matching attempt is governed by the following algorithm:

1. If the thrown exception is a [chained exception](#) (that is, where an exception has been caught and rethrown as a different exception), the most nested exception type serves initially as the basis for matching. This exception is tested as follows:
 - a. If the exception-to-test has exactly the type specified in the **onException** clause (tested using **instanceof**), a match is triggered.
 - b. If the exception-to-test is a sub-type of the type specified in the **onException** clause, a match is triggered.
2. If the most nested exception fails to yield a match, the next exception in the chain (the wrapping exception) is tested instead. The testing continues up the chain until either a match is triggered or the chain is exhausted.

Deadletter channel

The basic examples of **onException** usage have so far all exploited the *deadletter channel* pattern. That is, when an **onException** clause traps an exception, the current exchange is routed to a special destination (the deadletter channel). The deadletter channel serves as a holding area for failed messages that have *not* been processed. An administrator can inspect the messages at a later time and decide what action needs to be taken.

For more details about the deadletter channel pattern, see [Section 5.3, “Dead Letter Channel”](#).

Use original message

By the time an exception is raised in the middle of a route, the message in the exchange could have been modified considerably (and might not even be readable by a human). Often, it is easier for an administrator to decide what corrective actions to take, if the

messages visible in the deadletter queue are the *original* messages, as received at the start of the route.

In the Java DSL, you can replace the message in the exchange by the original message, using the `useOriginalMessage()` DSL command, as follows:

```
onException(ValidationException.class)
    .useOriginalMessage()
    .to("activemq:validationFailed");
```

In the XML DSL, you can retrieve the original message by setting the `useOriginalMessage` attribute on the `onException` element, as follows:

```
<onException useOriginalMessage="true">
  <exception>com.mycompany.ValidationException</exception>
  <to uri="activemq:validationFailed"/>
</onException>
```

Redelivery policy

Instead of interrupting the processing of a message and giving up as soon as an exception is raised, Apache Camel gives you the option of attempting to *redeliver* the message at the point where the exception occurred. In networked systems, where timeouts can occur and temporary faults arise, it is often possible for failed messages to be processed successfully, if they are redelivered shortly after the original exception was raised.

The Apache Camel redelivery supports various strategies for redelivering messages after an exception occurs. Some of the most important options for configuring redelivery are as follows:

`maximumRedeliveries()`

Specifies the maximum number of times redelivery can be attempted (default is **0**). A negative value means redelivery is always attempted (equivalent to an infinite value).

`retryWhile()`

Specifies a predicate (of **Predicate** type), which determines whether Apache Camel ought to continue redelivering. If the predicate evaluates to **true** on the current exchange, redelivery is attempted; otherwise, redelivery is stopped and no further redelivery attempts are made.

This option takes precedence over the `maximumRedeliveries()` option.

In the Java DSL, redelivery policy options are specified using DSL commands in the `onException` clause. For example, you can specify a maximum of six redeliveries, after which the exchange is sent to the `validationFailed` deadletter queue, as follows:

```
onException(ValidationException.class)
    .maximumRedeliveries(6)
    .retryAttemptedLogLevel(org.apache.camel.LogginLevel.WARN)
    .to("activemq:validationFailed");
```

In the XML DSL, redelivery policy options are specified by setting attributes on the **redeliveryPolicy** element. For example, the preceding route can be expressed in XML DSL as follows:

```
<onException useOriginalMessage="true">
  <exception>com.mycompany.ValidationException</exception>
  <redeliveryPolicy maximumRedeliveries="6"/>
  <to uri="activemq:validationFailed"/>
</onException>
```

The latter part of the route—after the redelivery options are set—is not processed until after the last redelivery attempt has failed. For detailed descriptions of all the redelivery options, see [Section 5.3, “Dead Letter Channel”](#).

Alternatively, you can specify redelivery policy options in a **redeliveryPolicyProfile** instance. You can then reference the **redeliveryPolicyProfile** instance using the **onException** element's **redeliveryPolicyRef** attribute. For example, the preceding route can be expressed as follows:

```
<redeliveryPolicyProfile id="redelivPolicy" maximumRedeliveries="6"
  retryAttemptedLogLevel="WARN"/>

<onException useOriginalMessage="true"
  redeliveryPolicyRef="redelivPolicy">
  <exception>com.mycompany.ValidationException</exception>
  <to uri="activemq:validationFailed"/>
</onException>
```



NOTE

The approach using **redeliveryPolicyProfile** is useful, if you want to re-use the same redelivery policy in multiple **onException** clauses.

Conditional trapping

Exception trapping with **onException** can be made conditional by specifying the **onWhen** option. If you specify the **onWhen** option in an **onException** clause, a match is triggered only when the thrown exception matches the clause *and* the **onWhen** predicate evaluates to **true** on the current exchange.

For example, in the following Java DSL fragment, the first **onException** clause triggers, only if the thrown exception matches **MyUserException** and the **user** header is non-null in the current exchange:

```
// Java

// Here we define onException() to catch MyUserException when
// there is a header[user] on the exchange that is not null
onException(MyUserException.class)
  .onWhen(header("user").isNotNull())
  .maximumRedeliveries(2)
  .to(ERROR_USER_QUEUE);

// Here we define onException to catch MyUserException as a kind
```

```
// of fallback when the above did not match.
// Noitce: The order how we have defined these onException is
// important as Camel will resolve in the same order as they
// have been defined
onException(MyUserException.class)
    .maximumRedeliveries(2)
    .to(ERROR_QUEUE);
```

The preceding **onException** clauses can be expressed in the XML DSL as follows:

```
<redeliveryPolicyProfile id="twoRedeliveries" maximumRedeliveries="2"/>
<onException redeliveryPolicyRef="twoRedeliveries">
  <exception>com.mycompany.MyUserException</exception>
  <onWhen>
    <simple>${header.user} != null</simple>
  </onWhen>
  <to uri="activemq:error_user_queue"/>
</onException>

<onException redeliveryPolicyRef="twoRedeliveries">
  <exception>com.mycompany.MyUserException</exception>
  <to uri="activemq:error_queue"/>
</onException>
```

Handling exceptions

By default, when an exception is raised in the middle of a route, processing of the current exchange is interrupted and the thrown exception is propagated back to the consumer endpoint at the start of the route. When an **onException** clause is triggered, the behavior is essentially the same, except that the **onException** clause performs some processing before the thrown exception is propagated back.

But this default behavior is *not* the only way to handle an exception. The **onException** provides various options to modify the exception handling behavior, as follows:

- [the section called “Suppressing exception rethrow”](#)—you have the option of suppressing the rethrown exception after the **onException** clause has completed. In other words, in this case the exception does *not* propagate back to the consumer endpoint at the start of the route.
- [the section called “Continuing processing”](#)—you have the option of resuming normal processing of the exchange from the point where the exception originally occurred. Implicitly, this approach also suppresses the rethrown exception.
- [the section called “Sending a response”](#)—in the special case where the consumer endpoint at the start of the route expects a reply (that is, having an *InOut* MEP), you might prefer to construct a custom fault reply message, rather than propagating the exception back to the consumer endpoint.

Suppressing exception rethrow

To prevent the current exception from being rethrown and propagated back to the consumer endpoint, you can set the **handled()** option to **true** in the Java DSL, as follows:

```
onException(ValidationException.class)
    .handled(true)
    .to("activemq:validationFailed");
```

In the Java DSL, the argument to the **handled()** option can be of boolean type, of **Predicate** type, or of **Expression** type (where any non-boolean expression is interpreted as **true**, if it evaluates to a non-null value).

The same route can be configured to suppress the rethrown exception in the XML DSL, using the **handled** element, as follows:

```
<onException>
  <exception>com.mycompany.ValidationException</exception>
  <handled>
    <constant>true</constant>
  </handled>
  <to uri="activemq:validationFailed"/>
</onException>
```

Continuing processing

To continue processing the current message from the point in the route where the exception was originally thrown, you can set the **continued** option to **true** in the Java DSL, as follows:

```
onException(ValidationException.class)
    .continued(true);
```

In the Java DSL, the argument to the **continued()** option can be of boolean type, of **Predicate** type, or of **Expression** type (where any non-boolean expression is interpreted as **true**, if it evaluates to a non-null value).

The same route can be configured in the XML DSL, using the **continued** element, as follows:

```
<onException>
  <exception>com.mycompany.ValidationException</exception>
  <continued>
    <constant>true</constant>
  </continued>
</onException>
```

Sending a response

When the consumer endpoint that starts a route expects a reply, you might prefer to construct a custom fault reply message, instead of simply letting the thrown exception propagate back to the consumer. There are two essential steps you need to follow in this case: suppress the rethrown exception using the **handled** option; and populate the exchange's *Out* message slot with a custom fault message.

For example, the following Java DSL fragment shows how to send a reply message containing the text string, **Sorry**, whenever the **MyFunctionalException** exception occurs:

```
// we catch MyFunctionalException and want to mark it as handled (= no
failure returned to client)
// but we want to return a fixed text response, so we transform OUT body
as Sorry.
onException(MyFunctionalException.class)
    .handled(true)
    .transform().constant("Sorry");
```

If you are sending a fault response to the client, you will often want to incorporate the text of the exception message in the response. You can access the text of the current exception message using the **exceptionMessage()** builder method. For example, you can send a reply containing just the text of the exception message whenever the **MyFunctionalException** exception occurs, as follows:

```
// we catch MyFunctionalException and want to mark it as handled (= no
failure returned to client)
// but we want to return a fixed text response, so we transform OUT body
and return the exception message
onException(MyFunctionalException.class)
    .handled(true)
    .transform(exceptionMessage());
```

The exception message text is also accessible from the Simple language, through the **exception.message** variable. For example, you could embed the current exception text in a reply message, as follows:

```
// we catch MyFunctionalException and want to mark it as handled (= no
failure returned to client)
// but we want to return a fixed text response, so we transform OUT body
and return a nice message
// using the simple language where we want insert the exception message
onException(MyFunctionalException.class)
    .handled(true)
    .transform().simple("Error reported: ${exception.message} - cannot
process this message.");
```

The preceding **onException** clause can be expressed in XML DSL as follows:

```
<onException>
  <exception>com.mycompany.MyFunctionalException</exception>
  <handled>
    <constant>>true</constant>
  </handled>
  <transform>
    <simple>Error reported: ${exception.message} - cannot process this
message.</simple>
  </transform>
</onException>
```

Exception thrown while handling an exception

An exception that gets thrown while handling an existing exception (in other words, one that gets thrown in the middle of processing an **onException** clause) is handled in a special way. Such an exception is handled by the special fallback exception handler, which handles

the exception as follows:

- All existing exception handlers are ignored and processing fails immediately.
- The new exception is logged.
- The new exception is set on the exchange object.

The simple strategy avoids complex failure scenarios which could otherwise end up with an **onException** clause getting locked into an infinite loop.

Scopes

The **onException** clauses can be effective in either of the following scopes:

- *RouteBuilder scope*—**onException** clauses defined as standalone statements inside a **RouteBuilder.configure()** method affect all of the routes defined in that **RouteBuilder** instance. On the other hand, these **onException** clauses *have no effect whatsoever* on routes defined inside any other **RouteBuilder** instance. The **onException** clauses *must* appear before the route definitions.

All of the examples up to this point are defined using the **RouteBuilder** scope.

- *Route scope*—**onException** clauses can also be embedded directly within a route. These **onException** clauses affect *only* the route in which they are defined.

Route scope

You can embed an **onException** clause anywhere inside a route definition, but you must terminate the embedded **onException** clause using the **end()** DSL command.

For example, you can define an embedded **onException** clause in the Java DSL, as follows:

```
// Java
from("direct:start")
  .onException(OrderFailedException.class)
    .maximumRedeliveries(1)
    .handled(true)
    .beanRef("orderService", "orderFailed")
    .to("mock:error")
  .end()
  .beanRef("orderService", "handleOrder")
  .to("mock:result");
```

You can define an embedded **onException** clause in the XML DSL, as follows:

```
<route errorHandlerRef="deadLetter">
  <from uri="direct:start"/>
  <onException>
    <exception>com.mycompany.OrderFailedException</exception>
    <redeliveryPolicy maximumRedeliveries="1"/>
    <handled>
      <constant>true</constant>
    </handled>
    <bean ref="orderService" method="orderFailed"/>
    <to uri="mock:error"/>
  </onException>
  <bean ref="orderService" method="handleOrder"/>
  <to uri="mock:result"/>
</route>
```

```

    </onException>
    <bean ref="orderService" method="handleOrder"/>
    <to uri="mock:result"/>
</route>

```

2.3.2. Error Handler

Overview

The `errorHandler()` clause provides similar features to the `onException` clause, except that this mechanism is *not* able to discriminate between different exception types. The `errorHandler()` clause is the original exception handling mechanism provided by Apache Camel and was available before the `onException` clause was implemented.

Java DSL example

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 can 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://fulfillment@acme.com");
        from("file:src/data?noop=true")
            .to("file:target/messages");
        // ...
    }
}

```

Redirection to the dead letter channel will not occur, however, until all attempts at redelivery have been exhausted.

XML DSL example

In the XML DSL, you define an error handler within a `camelContext` scope using the `errorHandler` element. For example, to define an error handler that routes all failed exchanges to the ActiveMQ `deadLetter` queue, you can define an `errorHandler` element as follows:

```

<beans xmlns="http://www.springframework.org/schema/beans"
    xmlns:camel="http://camel.apache.org/schema/spring"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="
        http://www.springframework.org/schema/beans
        http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
        http://camel.apache.org/schema/spring
        http://camel.apache.org/schema/spring/camel-spring.xsd">

```

```

<camelContext xmlns="http://camel.apache.org/schema/spring">
  <errorHandler type="DeadLetterChannel"
    deadLetterUri="activemq:deadLetter"/>
  <route>
    <from uri="activemq:orderQueue"/>
    <to uri="pop3://fulfillment@acme.com"/>
  </route>
  <route>
    <from uri="file:src/data?noop=true"/>
    <to uri="file:target/messages"/>
  </route>
</camelContext>
</beans>

```

Types of error handler

Table 2.1, “Error Handler Types” provides an overview of the different types of error handler you can define.

Table 2.1. Error Handler Types

Java DSL Builder	XML DSL Type Attribute	Description
<code>defaultErrorHandler()</code>	<code>DefaultErrorHandler</code>	Propagates exceptions back to the caller and supports the redelivery policy, but it does not support a dead letter queue.
<code>deadLetterChannel()</code>	<code>DeadLetterChannel</code>	Supports the same features as the default error handler and, in addition, supports a dead letter queue.
<code>loggingErrorChannel()</code>	<code>LoggingErrorChannel</code>	Logs the exception text whenever an exception occurs.
<code>noErrorHandler()</code>	<code>NoErrorHandler</code>	Dummy handler implementation that can be used to disable the error handler.
	<code>TransactionErrorHandler</code>	An error handler for transacted routes. A default transaction error handler instance is automatically used for a route that is marked as transacted.

2.3.3. doTry, doCatch, and doFinally


Overview

To handle exceptions within a route, you can use a combination of the **doTry**, **doCatch**, and **doFinally** clauses, which handle exceptions in a similar way to Java's **try**, **catch**, and **finally** blocks.

Similarities between doCatch and Java catch

In general, the **doCatch()** clause in a route definition behaves in an analogous way to the **catch()** statement in Java code. In particular, the following features are supported by the **doCatch()** clause:

- *Multiple doCatch clauses*—you can have multiple **doCatch** clauses within a single **doTry** block. The **doCatch** clauses are tested in the order they appear, just like Java **catch()** statements. Apache Camel executes the first **doCatch** clause that matches the thrown exception.



NOTE

This algorithm is different from the exception matching algorithm used by the **onException** clause—see [Section 2.3.1, “onException Clause”](#) for details.

- *Rethrowing exceptions*—you can rethrow the current exception from within a **doCatch** clause using the **handled** sub-clause (see [the section called “Rethrowing exceptions in doCatch”](#)).

Special features of doCatch

There are some special features of the **doCatch()** clause, however, that have no analogue in the Java **catch()** statement. The following features are specific to **doCatch()**:

- *Catching multiple exceptions*—the **doCatch** clause allows you to specify a list of exceptions to catch, in contrast to the Java **catch()** statement, which catches only one exception (see [the section called “Example”](#)).
- *Conditional catching*—you can catch an exception conditionally, by appending an **onWhen** sub-clause to the **doCatch** clause (see [the section called “Conditional exception catching using onWhen”](#)).

Example

The following example shows how to write a **doTry** block in the Java DSL, where the **doCatch()** clause will be executed, if either the **IOException** exception or the **IllegalStateException** exception are raised, and the **doFinally()** clause is *always* executed, irrespective of whether an exception is raised or not.

```
from("direct:start")
  .doTry()
    .process(new ProcessorFail())
    .to("mock:result")
  .doCatch(IOException.class, IllegalStateException.class)
```

```

        .to("mock:catch")
    .doFinally()
        .to("mock:finally")
    .end();

```

Or equivalently, in Spring XML:

```

<route>
  <from uri="direct:start"/>
  <!-- here the try starts. its a try .. catch .. finally just as
regular java code -->
  <doTry>
    <process ref="processorFail"/>
    <to uri="mock:result"/>
    <doCatch>
      <!-- catch multiple exceptions -->
      <exception>java.io.IOException</exception>
      <exception>java.lang.IllegalStateException</exception>
      <to uri="mock:catch"/>
    </doCatch>
    <doFinally>
      <to uri="mock:finally"/>
    </doFinally>
  </doTry>
</route>

```

Rethrowing exceptions in doCatch

It is possible to rethrow an exception in a `doCatch()` clause by calling the `handled()` sub-clause with its argument set to `false`, as follows:

```

from("direct:start")
  .doTry()
    .process(new ProcessorFail())
    .to("mock:result")
  .doCatch(IOException.class)
    // mark this as NOT handled, eg the caller will also get the
exception
    .handled(false)
    .to("mock:io")
  .doCatch(Exception.class)
    // and catch all other exceptions
    .to("mock:error")
  .end();

```

In the preceding example, if the `IOException` is caught by `doCatch()`, the current exchange is sent to the `mock:io` endpoint, and then the `IOException` is rethrown. This gives the consumer endpoint at the start of the route (in the `from()` command) an opportunity to handle the exception as well.

The following example shows how to define the same route in Spring XML:

```

<route>
  <from uri="direct:start"/>
  <doTry>

```

```

    <process ref="processorFail"/>
    <to uri="mock:result"/>
    <doCatch>
        <exception>java.io.IOException</exception>
        <!-- mark this as NOT handled, eg the caller will also get the
exception -->
        <handled>
            <constant>false</constant>
        </handled>
        <to uri="mock:io"/>
    </doCatch>
    <doCatch>
        <!-- and catch all other exceptions they are handled by
default (ie handled = true) -->
        <exception>java.lang.Exception</exception>
        <to uri="mock:error"/>
    </doCatch>
    </doTry>
</route>

```

Conditional exception catching using onWhen

A special feature of the Apache Camel `doCatch()` clause is that you can conditionalize the catching of exceptions based on an expression that is evaluated at run time. In other words, if you catch an exception using a clause of the form, `doCatch(ExceptionList).doWhen(Expression)`, an exception will only be caught, if the predicate expression, *Expression*, evaluates to **true** at run time.

For example, the following `doTry` block will catch the exceptions, `IOException` and `IllegalStateException`, only if the exception message contains the word, `Severe`:

```

from("direct:start")
    .doTry()
        .process(new ProcessorFail())
        .to("mock:result")
    .doCatch(IOException.class, IllegalStateException.class)
        .onWhen(exceptionMessage().contains("Severe"))
        .to("mock:catch")
    .doCatch(CamelExchangeException.class)
        .to("mock:catchCamel")
    .doFinally()
        .to("mock:finally")
    .end();

```

Or equivalently, in Spring XML:

```

<route>
    <from uri="direct:start"/>
    <doTry>
        <process ref="processorFail"/>
        <to uri="mock:result"/>
        <doCatch>
            <exception>java.io.IOException</exception>
            <exception>java.lang.IllegalStateException</exception>
            <onWhen>

```

```

        <simple>${exception.message} contains 'Severe'</simple>
    </onWhen>
    <to uri="mock:catch"/>
</doCatch>
<doCatch>
    <exception>org.apache.camel.CamelExchangeException</exception>
    <to uri="mock:catchCamel"/>
</doCatch>
<doFinally>
    <to uri="mock:finally"/>
</doFinally>
</doTry>
</route>

```

2.3.4. Propagating SOAP Exceptions

Overview

The Camel CXF component provides an integration with Apache CXF, enabling you to send and receive SOAP messages from Apache Camel endpoints. You can easily define Apache Camel endpoints in XML, which can then be referenced in a route using the endpoint's bean ID. For more details, see [CXF](#).

How to propagate stack trace information

It is possible to configure a CXF endpoint so that, when a Java exception is thrown on the server side, the stack trace for the exception is marshalled into a fault message and returned to the client. To enable this feature, set the **dataFormat** to **PAYLOAD** and set the **faultStackTraceEnabled** property to **true** in the **cxfEndpoint** element, as follows:

```

<cxf:cxfEndpoint id="router" address="http://localhost:9002/TestMessage"
    wsdlURL="ship.wsdl"
    endpointName="s:TestSoapEndpoint"
    serviceName="s:TestService"
    xmlns:s="http://test">
    <cxf:properties>
        <!-- enable sending the stack trace back to client; the default value
is false-->
        <entry key="faultStackTraceEnabled" value="true" />
        <entry key="dataFormat" value="PAYLOAD" />
    </cxf:properties>
</cxf:cxfEndpoint>

```

For security reasons, the stack trace does not include the causing exception (that is, the part of a stack trace that follows **Caused by**). If you want to include the causing exception in the stack trace, set the **exceptionMessageCauseEnabled** property to **true** in the **cxfEndpoint** element, as follows:

```

<cxf:cxfEndpoint id="router" address="http://localhost:9002/TestMessage"
    wsdlURL="ship.wsdl"
    endpointName="s:TestSoapEndpoint"
    serviceName="s:TestService"
    xmlns:s="http://test">
    <cxf:properties>

```

```

    <!-- enable to show the cause exception message and the default value
    is false -->
    <entry key="exceptionMessageCauseEnabled" value="true" />
    <!-- enable to send the stack trace back to client, the default value
    is false-->
    <entry key="faultStackTraceEnabled" value="true" />
    <entry key="dataFormat" value="PAYLOAD" />
  </cdf:properties>
</cdf:cxfEndpoint>

```



WARNING

You should only enable the **exceptionMessageCauseEnabled** flag for testing and diagnostic purposes. It is normal practice for servers to conceal the original cause of an exception to make it harder for hostile users to probe the server.

2.4. 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 *parameter binding*. Parameter binding can use any combination of the following approaches in order to initialize a method's parameters:

- *Conventional method signatures* — If the method signature conforms to certain conventions, the parameter 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 2.5 component scanning. Normally, if you are deploying your Apache Camel application into a Spring container, the dependency injection mechanism will work automatically.
- *Explicitly specified parameters* — You can specify parameters explicitly (either as constants or using the Simple language), at the point where the bean is invoked.

Bean registry

Beans are made accessible through a *bean registry*, which is a service that enables you to look up beans using either the class name or the bean ID as a key. The way that you create an entry in the bean registry depends on the underlying framework—for example, plain Java, Spring, Guice, or Blueprint. Registry entries are usually created implicitly (for example, when you instantiate a Spring bean in a Spring XML file).

Registry plug-in strategy

Apache Camel implements a plug-in strategy for the bean registry, defining an integration layer for accessing beans which makes the underlying registry implementation transparent. Hence, it is possible to integrate Apache Camel applications with a variety of different bean registries, as shown in [Table 2.2, “Registry Plug-Ins”](#).

Table 2.2. Registry Plug-Ins

Registry Implementation	Camel Component with Registry Plug-In
Spring bean registry	camel-spring
Guice bean registry	camel-guice
Blueprint bean registry	camel-blueprint
OSGi service registry	deployed in <i>OSGi container</i>

Normally, you do not have to worry about configuring bean registries, because the relevant bean registry is automatically installed for you. For example, if you are using the Spring framework to define your routes, the Spring **ApplicationContextRegistry** plug-in is automatically installed in the current **CamelContext** instance.

Deployment in an OSGi container is a special case. When an Apache Camel route is deployed into the OSGi container, the **CamelContext** automatically sets up a registry chain for resolving bean instances: the registry chain consists of the OSGi registry, followed by the Blueprint (or Spring) registry.

Accessing a bean created in Java

To process exchange objects using a Java bean (which is 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. However, if you want to access the same **MyBeanProcessor** instance from multiple routes, use the variant of **bean()** that takes the **Object** type as its first argument. For example:

```
MyBeanProcessor myBean = new MyBeanProcessor();

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

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

Accessing overloaded bean methods

If a bean defines overloaded methods, you can choose which of the overloaded methods to invoke by specifying the method name along with its parameter types. For example, if the **MyBeanProcessor** class has two overloaded methods, **processBody(String)** and **processBody(String, String)**, you can invoke the latter overloaded method as follows:

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

Alternatively, if you want to identify a method by the number of parameters it takes, rather than specifying the type of each parameter explicitly, you can use the wildcard character, *****. For example, to invoke a method named **processBody** that takes two parameters, irrespective of the exact type of the parameters, invoke the **bean()** processor as follows:

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

When specifying the method, you can use either a simple unqualified type name—for example, **processBody(Exchange)**—or a fully qualified type name—for example, **processBody(org.apache.camel.Exchange)**.



NOTE

In the current implementation, the specified type name must be an exact match of the parameter type. Type inheritance is not taken into account.

Specify parameters explicitly

You can specify parameter values explicitly, when you call the bean method. The following simple type values can be passed:

- Boolean: **true** or **false**.
- Numeric: **123**, **7**, and so on.
- String: **'In single quotes'** or **"In double quotes"**.
- Null object: **null**.

The following example shows how you can mix explicit parameter values with type specifiers in the same method invocation:

```
from("file:data/inbound")
    .bean(MyBeanProcessor.class, "processBody(String, 'Sample string value',
true, 7)")
    .to("file:data/outbound");
```

In the preceding example, the value of the first parameter would presumably be determined by a parameter binding annotation (see [the section called “Basic annotations”](#)).

In addition to the simple type values, you can also specify parameter values using the Simple language ([Chapter 27, The Simple Language](#)). This means that the *full power of the Simple language is available* when specifying parameter values. For example, to pass the message body and the value of the **title** header to a bean method:

```
from("file:data/inbound")
    .bean(MyBeanProcessor.class,
"processBodyAndHeader(${body},${header.title}")
    .to("file:data/outbound");
```

You can also pass the entire header hash map as a parameter. For example, in the following example, the second method parameter must be declared to be of type **java.util.Map**:

```
from("file:data/inbound")
    .bean(MyBeanProcessor.class,
"processBodyAndAllHeaders(${body},${header}")
    .to("file:data/outbound");
```

Basic method signatures

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:

- [the section called “Method signature for processing message bodies”](#)
- [the section called “Method signature for processing exchanges”](#).

Method signature for processing message bodies

If you want to implement a bean method that accesses or modifies the incoming message body, you must 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. In fact, this is the only feasible approach if you are defining your routes in XML. To define a bean in XML, use the standard Spring **bean** element. The following example shows how 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](#) 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",
    "processBody").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="http://camel.apache.org/schema/spring">
    <route>
        <from uri="file:data/inbound"/>
        <bean ref="myBeanId" method="processBody"/>
        <to uri="file:data/outbound"/>
    </route>
</camelContext>
```

Parameter binding annotations

The basic parameter bindings described in [the section called “Basic method signatures”](#) might not always be convenient to use. For example, if you have a legacy 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 parameter binding, Apache Camel provides the following kinds of Java annotation:

- the section called “Basic annotations”.
- the section called “Expression language annotations”.
- the section called “Inherited annotations”.

Basic annotations

Table 2.3, “Basic Bean Annotations” 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 2.3. Basic Bean Annotations

Annotation	Meaning	Parameter?
<code>@Attachments</code>	Binds to a list of attachments.	
<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 parameter 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, to extract data from an inbound exchange and inject the data into a method argument. [Table 2.4, "Expression Language Annotations"](#) 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 2.4. Expression Language Annotations

Annotation	Description
<code>@Bean</code>	Injects a Bean expression.
<code>@Constant</code>	Injects a Constant expression
<code>@EL</code>	Injects an EL expression.
<code>@Groovy</code>	Injects a Groovy expression.
<code>@Header</code>	Injects a Header expression.
<code>@JavaScript</code>	Injects a JavaScript expression.
<code>@OGNL</code>	Injects an OGNL expression.
<code>@PHP</code>	Injects a PHP expression.
<code>@Python</code>	Injects a Python expression.
<code>@Ruby</code>	Injects a Ruby expression.
<code>@Simple</code>	Injects a Simple expression.
<code>@XPath</code>	Injects an XPath expression.
<code>@XQuery</code>	Injects 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/text()") String user,
        @XPath("/credentials/password/text()") String pass
    ) {
        // Check the user/pass credentials...
        ...
    }
}
```

The **@Bean** annotation is a special case, because it enables you to inject the result of invoking a registered bean. For example, to inject a correlation ID into a method argument, you can 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.MyIdGenerator"/>
</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 must load the **camel-ognl** component.

Inherited annotations

Parameter binding annotations can be inherited from an interface or from a superclass. For example, if you define a Java interface with a **Header** annotation and a **Body** annotation, as follows:

```

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

public interface MyBeanProcessorIntf {
    void processExchange(
        @Header(name="user") String user,
        @Body String body,
        Exchange exchange
    );
}

```

The overloaded methods defined in the implementation class, **MyBeanProcessor**, now inherit the annotations defined in the base interface, as follows:

```

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

public class MyBeanProcessor implements MyBeanProcessorIntf {
    public void processExchange(
        String user, // Inherits Header annotation
        String body, // Inherits Body annotation
        Exchange exchange
    ) {
        ...
    }
}

```

Interface implementations

The class that implements a Java interface is often **protected**, **private** or in **package-only** scope. If you try to invoke a method on an implementation class that is restricted in this

way, the bean binding falls back to invoking the corresponding interface method, which is publicly accessible.

For example, consider the following public **BeanIntf** interface:

```
// Java
public interface BeanIntf {
    void processBodyAndHeader(String body, String title);
}
```

Where the **BeanIntf** interface is implemented by the following protected **BeanIntfImpl** class:

```
// Java
protected class BeanIntfImpl implements BeanIntf {
    void processBodyAndHeader(String body, String title) {
        ...
    }
}
```

The following bean invocation would fall back to invoking the public **BeanIntf.processBodyAndHeader** method:

```
from("file:data/inbound")
    .bean(BeanIntfImpl.class, "processBodyAndHeader(${body},
    ${header.title}")
    .to("file:data/outbound");
```

Invoking static methods

Bean integration has the capability to invoke static methods *without* creating an instance of the associated class. For example, consider the following Java class that defines the static method, **changeSomething()**:

```
// Java
...
public final class MyStaticClass {
    private MyStaticClass() {
    }

    public static String changeSomething(String s) {
        if ("Hello World".equals(s)) {
            return "Bye World";
        }
        return null;
    }

    public void doSomething() {
        // noop
    }
}
```

You can use bean integration to invoke the static **changeSomething** method, as follows:

```
from("direct:a")
  .bean(MyStaticClass.class, "changeSomething")
  .to("mock:a");
```

Note that, although this syntax looks identical to the invocation of an ordinary function, bean integration exploits Java reflection to identify the method as static and proceeds to invoke the method *without* instantiating **MyStaticClass**.

Invoking an OSGi service

In the special case where a route is deployed into a Red Hat JBoss Fuse container, it is possible to invoke an OSGi service directly using bean integration. For example, assuming that one of the bundles in the OSGi container has exported the service, **org.fusesource.example>HelloWorldOsgiService**, you could invoke the **sayHello** method using the following bean integration code:

```
from("file:data/inbound")
  .bean(org.fusesource.example>HelloWorldOsgiService.class, "sayHello")
  .to("file:data/outbound");
```

You could also invoke the OSGi service from within a Spring or Blueprint XML file, using the bean component, as follows:

```
<to uri="bean:org.fusesource.example>HelloWorldOsgiService?
method=sayHello"/>
```

The way this works is that Apache Camel sets up a chain of registries when it is deployed in the OSGi container. First of all, it looks up the specified class name in the OSGi service registry; if this lookup fails, it then falls back to the local Spring DM or Blueprint registry.

2.5. CREATING EXCHANGE INSTANCES

Overview

When processing messages with Java code (for example, in a bean class or in a processor class), it is often necessary to create a fresh exchange instance. If you need to create an **Exchange** object, the easiest approach is to invoke the methods of the **ExchangeBuilder** class, as described here.

ExchangeBuilder class

The fully qualified name of the **ExchangeBuilder** class is as follows:

```
org.apache.camel.builder.ExchangeBuilder
```

The **ExchangeBuilder** exposes the static method, **anExchange**, which you can use to start building an exchange object.

Example

For example, the following code creates a new exchange object containing the message body string, **Hello World!**, and with headers containing username and password credentials:

```
// Java
import org.apache.camel.Exchange;
import org.apache.camel.builder.ExchangeBuilder;
...
Exchange exch = ExchangeBuilder.anExchange(camelCtx)
    .withBody("Hello World!")
    .withHeader("username", "jdoe")
    .withHeader("password", "pass")
    .build();
```

ExchangeBuilder methods

The **ExchangeBuilder** class supports the following methods:

ExchangeBuilder anExchange(CamelContext context)

(static method) Initiate building an exchange object.

Exchange build()

Build the exchange.

ExchangeBuilder withBody(Object body)

Set the message body on the exchange (that is, sets the exchange's *In* message body).

ExchangeBuilder withHeader(String key, Object value)

Set a header on the exchange (that is, sets a header on the exchange's *In* message).

ExchangeBuilder withPattern(ExchangePattern pattern)

Sets the exchange pattern on the exchange.

ExchangeBuilder withProperty(String key, Object value)

Sets a property on the exchange.

2.6. TRANSFORMING MESSAGE CONTENT

Abstract

Apache Camel supports a variety of approaches to transforming message content. In addition to a simple native API for modifying message content, Apache Camel supports integration with several different third-party libraries and transformation standards.

2.6.1. Simple Message Transformations

Overview

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 2.1, “Simple Transformation of Incoming Messages”](#) appends the text, **World!**, to the end of the incoming message body.

Example 2.1. 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.

API for simple transformations

You can use the following API classes to perform simple transformations of the message content in a router rule:

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

ProcessorDefinition 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 2.1, “Simple Transformation of Incoming Messages”](#). [Table 2.5, “Transformation Methods from the ProcessorDefinition Class”](#) shows the `ProcessorDefinition` methods that are relevant to transforming message content:

Table 2.5. Transformation Methods from the ProcessorDefinition Class

Method	Description
Type <code>convertBodyTo(Class type)</code>	Converts the IN message body to the specified type.
Type <code>removeFaultHeader(String name)</code>	Adds a processor which removes the header on the FAULT message.
Type <code>removeHeader(String name)</code>	Adds a processor which removes the header on the IN message.
Type <code>removeProperty(String name)</code>	Adds a processor which removes the exchange property.
ExpressionClause<ProcessorDefinition<Type>> <code>setBody()</code>	Adds a processor which sets the body on the IN message.
Type <code>setFaultBody(Expression expression)</code>	Adds a processor which sets the body on the FAULT message.

Method	Description
Type <code>setFaultHeader(String name, Expression expression)</code>	Adds a processor which sets the header on the FAULT message.
ExpressionClause<ProcessorDefinition<Type>> <code>setHeader(String name)</code>	Adds a processor which sets the header on the IN message.
Type <code>setHeader(String name, Expression expression)</code>	Adds a processor which sets the header on the IN message.
ExpressionClause<ProcessorDefinition<Type>> <code>setOutHeader(String name)</code>	Adds a processor which sets the header on the OUT message.
Type <code>setOutHeader(String name, Expression expression)</code>	Adds a processor which sets the header on the OUT message.
ExpressionClause<ProcessorDefinition<Type>> <code>setProperty(String name)</code>	Adds a processor which sets the exchange property.
Type <code>setProperty(String name, Expression expression)</code>	Adds a processor which sets the exchange property.
ExpressionClause<ProcessorDefinition<Type>> <code>transform()</code>	Adds a processor which sets the body on the OUT message.
Type <code>transform(Expression expression)</code>	Adds a processor which sets the body on the OUT message.

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 2.1, “Simple Transformation of Incoming Messages”](#). [Table 2.6, “Methods from the Builder Class”](#) summarizes the static methods available in the **Builder** class.

Table 2.6. Methods from the Builder Class

Method	Description
static <E extends Exchange> ValueBuilder<E> <code>body()</code>	Returns a predicate and value builder for the inbound body on an exchange.
static <E extends Exchange, T> ValueBuilder<E> <code>bodyAs(Class<T> type)</code>	Returns a predicate and value builder for the inbound message body as a specific type.

Method	Description
static <E extends Exchange> ValueBuilder<E> constant(Object value)	Returns a constant expression.
static <E extends Exchange> ValueBuilder<E> faultBody()	Returns a predicate and value builder for the fault body on an exchange.
static <E extends Exchange, T> ValueBuilder<E> faultBodyAs(Class<T> type)	Returns a predicate and value builder for the fault message body as a specific type.
static <E extends Exchange> ValueBuilder<E> header(String name)	Returns a predicate and value builder for headers on an exchange.
static <E extends Exchange> ValueBuilder<E> outBody()	Returns a predicate and value builder for the outbound body on an exchange.
static <E extends Exchange> ValueBuilder<E> outBodyAs(Class<T> type)	Returns a predicate and value builder for the outbound message body as a specific type.
static ValueBuilder property(String name)	Returns a predicate and value builder for properties on an exchange.
static ValueBuilder regexReplaceAll(Expression content, String regex, Expression replacement)	Returns an expression that replaces all occurrences of the regular expression with the given replacement.
static ValueBuilder regexReplaceAll(Expression content, String regex, String replacement)	Returns an expression that replaces all occurrences of the regular expression with the given replacement.
static ValueBuilder sendTo(String uri)	Returns an expression processing the exchange to the given endpoint uri.
static <E extends Exchange> ValueBuilder<E> systemProperty(String name)	Returns an expression for the given system property.
static <E extends Exchange> ValueBuilder<E> systemProperty(String name, String defaultValue)	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 2.7, “Modifier Methods from the ValueBuilder Class”](#) 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 2.7. Modifier Methods from the ValueBuilder Class

Method	Description
ValueBuilder<E> append(Object value)	Appends the string evaluation of this expression with the given value.
Predicate contains(Object value)	Create a predicate that the left hand expression contains the value of the right hand expression.
ValueBuilder<E> convertTo(Class type)	Converts the current value to the given type using the registered type converters.
ValueBuilder<E> convertToString()	Converts the current value a String using the registered type converters.
Predicate endsWith(Object value)	
<T> T evaluate(Exchange exchange, Class<T> type)	
Predicate in(Object... values)	
Predicate in(Predicate... predicates)	
Predicate isEqualTo(Object value)	Returns true, if the current value is equal to the given value argument.
Predicate isGreaterThan(Object value)	Returns true, if the current value is greater than the given value argument.
Predicate isGreaterThanOrEqualTo(Object value)	Returns true, if the current value is greater than or equal to the given value argument.
Predicate isInstanceOf(Class type)	Returns true, if the current value is an instance of the given type.
Predicate isLessThan(Object value)	Returns true, if the current value is less than the given value argument.

Method	Description
Predicate isLessThanOrEqualTo(Object value)	Returns true, if the current value is less than or equal to the given value argument.
Predicate isNotEqualTo(Object value)	Returns true, if the current value is not equal to the given value argument.
Predicate isNotNull()	Returns true, if the current value is not null .
Predicate isNull()	Returns true, if the current value is null .
Predicate matches(Expression expression)	
Predicate not(Predicate predicate)	Negates the predicate argument.
ValueBuilder prepend(Object value)	Prepends the string evaluation of this expression to the given value.
Predicate regex(String regex)	
ValueBuilder<E> regexReplaceAll(String regex, Expression<E> replacement)	Replaces all occurrences of the regular expression with the given replacement.
ValueBuilder<E> regexReplaceAll(String regex, String replacement)	Replaces all occurrences of the regular expression with the given replacement.
ValueBuilder<E> regexTokenize(String regex)	Tokenizes the string conversion of this expression using the given regular expression.
ValueBuilder sort(Comparator comparator)	Sorts the current value using the given comparator.
Predicate startsWith(Object value)	Returns true, if the current value matches the string value of the value argument.
ValueBuilder<E> tokenize()	Tokenizes the string conversion of this expression using the comma token separator.
ValueBuilder<E> tokenize(String token)	Tokenizes the string conversion of this expression using the given token separator.

2.6.2. Marshalling and Unmarshalling

Java DSL commands

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

- **marshal()**— Converts a high-level data format to a low-level data format.
- **unmarshal()** — Converts a low-level data format to a high-level data format.

Data formats

Apache Camel supports marshalling and unmarshalling of the following data formats:

- Java serialization
- JAXB
- XMLBeans
- XStream

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 use a rule like the following:

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

Or alternatively, in Spring XML:

```
<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 to generate the Java classes that represent the XML data types in the schema. This is called *binding* the schema. After the schema is bound, you 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("GeneratedPackageName");
```

```
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.

Or alternatively, in Spring XML:

```
<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>
```

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 use code like the following:

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


Or alternatively, in Spring XML:

```
<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

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.

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



NOTE

The XStream data format is currently *not* supported in Spring XML.

2.6.3. Endpoint Bindings

What is a binding?

In Apache Camel, a binding is a way of wrapping an endpoint in a contract—for example, by applying a Data Format, a Content Enricher or a validation step. A condition or transformation is applied to the messages coming in, and a complementary condition or transformation is applied to the messages going out.

DataFormatBinding

The `DataFormatBinding` class is useful for the specific case where you want to define a binding that marshals and unmarshals a particular data format (see [Section 2.6.2, “Marshalling and Unmarshalling”](#)). In this case, all that you need to do to create a binding is to create a `DataFormatBinding` instance, passing a reference to the relevant data format in the constructor.

For example, the XML DSL snippet in [Example 2.2, “JAXB Binding”](#) shows a binding (with ID, `jaxb`) that is capable of marshalling and unmarshalling the JAXB data format when it is associated with an Apache Camel endpoint:

Example 2.2. JAXB Binding

```
<beans ... >
  ...
  <bean id="jaxb"
class="org.apache.camel.processor.binding.DataFormatBinding">
    <constructor-arg ref="jaxbformat"/>
  </bean>

  <bean id="jaxbformat"
class="org.apache.camel.model.dataformat.JaxbDataFormat">
    <property name="prettyPrint" value="true"/>
    <property name="contextPath" value="org.apache.camel.example"/>
  </bean>
</beans>
```

Associating a binding with an endpoint

The following alternatives are available for associating a binding with an endpoint:

- [the section called “Binding URI”](#)
- [the section called “BindingComponent”](#)

Binding URI

To associate a binding with an endpoint, you can prefix the endpoint URI with `binding:NameOfBinding`, where `NameOfBinding` is the bean ID of the binding (for example, the ID of a binding bean created in Spring XML).

For example, the following example shows how to associate ActiveMQ endpoints with the JAXB binding defined in [Example 2.2, "JAXB Binding"](#).

```
<beans ...>
  ...
  <camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="binding:jaxb:activemq:orderQueue"/>
      <to uri="binding:jaxb:activemq:otherQueue"/>
    </route>
  </camelContext>
  ...
</beans>
```

BindingComponent

Instead of using a prefix to associate a binding with an endpoint, you can make the association implicit, so that the binding does not need to appear in the URI. For existing endpoints that do not have an implicit binding, the easiest way to achieve this is to wrap the endpoint using the **BindingComponent** class.

For example, to associate the **jaxb** binding with **activemq** endpoints, you could define a new **BindingComponent** instance as follows:

```
<beans ... >
  ...
  <bean id="jaxbmq"
class="org.apache.camel.component.binding.BindingComponent">
    <constructor-arg ref="jaxb"/>
    <constructor-arg value="activemq:foo."/>
  </bean>

  <bean id="jaxb"
class="org.apache.camel.processor.binding.DataFormatBinding">
    <constructor-arg ref="jaxbformat"/>
  </bean>

  <bean id="jaxbformat"
class="org.apache.camel.model.dataformat.JaxbDataFormat">
    <property name="prettyPrint" value="true"/>
    <property name="contextPath" value="org.apache.camel.example"/>
  </bean>

</beans>
```

Where the (optional) second constructor argument to **jaxbmq** defines a URI prefix. You can now use the **jaxbmq** ID as the scheme for an endpoint URI. For example, you can define the following route using this binding component:

```
<beans ...>
  ...
  <camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="jaxbmq:firstQueue"/>
      <to uri="jaxbmq:otherQueue"/>
    </route>
  </camelContext>
  ...
</beans>
```

```

        </route>
    </camelContext>
    ...
</beans>

```

The preceding route is equivalent to the following route, which uses the binding URI approach:

```

<beans ...>
    ...
    <camelContext xmlns="http://camel.apache.org/schema/spring">
        <route>
            <from uri="binding:jaxb:activemq:foo.firstQueue"/>
            <to uri="binding:jaxb:activemq:foo.otherQueue"/>
        </route>
    </camelContext>
    ...
</beans>

```



NOTE

For developers that implement a custom Apache Camel component, it is possible to achieve this by implementing an endpoint class that inherits from the `org.apache.camel.spi.HasBinding` interface.

BindingComponent constructors

The `BindingComponent` class supports the following constructors:

`public BindingComponent()`

No arguments form. Use property injection to configure the binding component instance.

`public BindingComponent(Binding binding)`

Associate this binding component with the specified `Binding` object, `binding`.

`public BindingComponent(Binding binding, String uriPrefix)`

Associate this binding component with the specified `Binding` object, `binding`, and URI prefix, `uriPrefix`. This is the most commonly used constructor.

`public BindingComponent(Binding binding, String uriPrefix, String uriPostfix)`

This constructor supports the additional URI post-fix, `uriPostfix`, argument, which is automatically appended to any URIs defined using this binding component.

Implementing a custom binding

In addition to the `DataFormatBinding`, which is used for marshalling and unmarshalling data formats, you can implement your own custom bindings. Define a custom binding as follows:

1. Implement an `org.apache.camel.Processor` class to perform a transformation on messages incoming to a consumer endpoint (appearing in a `from` element).

2. Implement a complementary `org.apache.camel.Processor` class to perform the reverse transformation on messages outgoing from a producer endpoint (appearing in a `to` element).
3. Implement the `org.apache.camel.spi.Binding` interface, which acts as a factory for the processor instances.

Binding interface

[Example 2.3, “The `org.apache.camel.spi.Binding` Interface”](#) shows the definition of the `org.apache.camel.spi.Binding` interface, which you must implement to define a custom binding.

Example 2.3. The `org.apache.camel.spi.Binding` Interface

```
// Java
package org.apache.camel.spi;

import org.apache.camel.Processor;

/**
 * Represents a Binding or contract
 * which can be applied to an Endpoint; such as ensuring that a
 * particular
 * Data Format is
 * used on messages in and out of an endpoint.
 */
public interface Binding {

    /**
     * Returns a new {@link Processor} which is used by a producer on an
     * endpoint to implement
     * the producer side binding before the message is sent to the
     * underlying endpoint.
     */
    Processor createProduceProcessor();

    /**
     * Returns a new {@link Processor} which is used by a consumer on an
     * endpoint to process the
     * message with the binding before its passed to the endpoint
     * consumer producer.
     */
    Processor createConsumeProcessor();
}
```

When to use bindings

Bindings are useful when you need to apply the same kind of transformation to many different kinds of endpoint.

2.7. PROPERTY PLACEHOLDERS

Overview

The property placeholders feature can be used to substitute strings into various contexts (such as endpoint URIs and attributes in XML DSL elements), where the placeholder settings are stored in Java properties files. This feature can be useful, if you want to share settings between different Apache Camel applications or if you want to centralize certain configuration settings.

For example, the following route sends requests to a Web server, whose host and port are substituted by the placeholders, `{{remote.host}}` and `{{remote.port}}`:

```
from("direct:start").to("http://{{remote.host}}:{{remote.port}}");
```

The placeholder values are defined in a Java properties file, as follows:

```
# Java properties file
remote.host=myserver.com
remote.port=8080
```

Property files

Property settings are stored in one or more Java properties files and must conform to the standard Java properties file format. Each property setting appears on its own line, in the format **Key=Value**. Lines with `#` or `!` as the first non-blank character are treated as comments.

For example, a property file could have content as shown in [Example 2.4, “Sample Property File”](#).

Example 2.4. Sample Property File

```
# Property placeholder settings
# (in Java properties file format)
cool.end=mock:result
cool.result=result
cool.concat=mock:{{cool.result}}
cool.start=direct:cool
cool.showid=true

cheese.end=mock:cheese
cheese.quote=Camel rocks
cheese.type=Gouda

bean.foo=foo
bean.bar=bar
```

Resolving properties

The properties component must be configured with the locations of one or more property files before you can start using it in route definitions. You must provide the property values using one of the following resolvers:

```
classpath:PathName,PathName,...
```

(Default) Specifies locations on the classpath, where *PathName* is a file pathname delimited using forward slashes.

file:*PathName*,*PathName*,...

Specifies locations on the file system, where *PathName* is a file pathname delimited using forward slashes.

ref:*BeanID*

Specifies the ID of a `java.util.Properties` object in the registry.

blueprint:*BeanID*

Specifies the ID of a `cm:property-placeholder` bean, which is used in the context of an OSGi Blueprint file to access properties defined in the *OSGi Configuration Admin* service. For details, see [the section called “Integration with OSGi Blueprint property placeholders”](#).

For example, to specify the `com/fusesource/cheese.properties` property file and the `com/fusesource/bar.properties` property file, both located on the classpath, you would use the following location string:

```
classpath:com/fusesource/cheese.properties,com/fusesource/bar.properties
```



NOTE

You can omit the `classpath:` prefix in this example, because the classpath resolver is used by default.

Specifying locations using system properties and environment variables

You can embed Java system properties and O/S environment variables in a location *PathName*.

Java system properties can be embedded in a location resolver using the syntax, `${PropertyName}`. For example, if the root directory of Red Hat JBoss Fuse is stored in the Java system property, `karaf.home`, you could embed that directory value in a file location, as follows:

```
file:${karaf.home}/etc/foo.properties
```

O/S environment variables can be embedded in a location resolver using the syntax, `${env:VarName}`. For example, if the root directory of JBoss Fuse is stored in the environment variable, `SMX_HOME`, you could embed that directory value in a file location, as follows:

```
file:${env:SMX_HOME}/etc/foo.properties
```

Configuring the properties component

Before you can start using property placeholders, you must configure the properties component, specifying the locations of one or more property files.

In the Java DSL, you can configure the properties component with the property file locations, as follows:

```
// Java
import org.apache.camel.component.properties.PropertiesComponent;
...
PropertiesComponent pc = new PropertiesComponent();
pc.setLocation("com/fusesource/cheese.properties,com/fusesource/bar.properties");
context.addComponent("properties", pc);
```

As shown in the `addComponent()` call, the name of the properties component *must* be set to **properties**.

In the XML DSL, you can configure the properties component using the dedicated **propertyPlaceholder** element, as follows:

```
<camelContext ...>
  <propertyPlaceholder
    id="properties"

    location="com/fusesource/cheese.properties,com/fusesource/bar.properties"
  />
</camelContext>
```

If you want the properties component to ignore any missing **.properties** files when it is being initialized, you can set the **ignoreMissingLocation** option to **true** (normally, a missing **.properties** file would result in an error being raised).

Placeholder syntax

After it is configured, the property component automatically substitutes placeholders (in the appropriate contexts). The syntax of a placeholder depends on the context, as follows:

- *In endpoint URIs and in Spring XML files*—the placeholder is specified as **{{Key}}**.
- *When setting XML DSL attributes*—**xs:string** attributes are set using the following syntax:

```
AttributeName="{{Key}}"
```

Other attribute types (for example, **xs:int** or **xs:boolean**) must be set using the following syntax:

```
prop:AttributeName="Key"
```

Where **prop** is associated with the **http://camel.apache.org/schema/placeholder** namespace.

- *When setting Java DSL EIP options*—to set an option on an Enterprise Integration Pattern (EIP) command in the Java DSL, add a **placeholder()** clause like the following to the fluent DSL:

```
.placeholder("OptionName", "Key")
```

- In *Simple language expressions*—the placeholder is specified as `${properties:Key}`.

Substitution in endpoint URIs

Wherever an endpoint URI string appears in a route, the first step in parsing the endpoint URI is to apply the property placeholder parser. The placeholder parser automatically substitutes any property names appearing between double braces, `{{Key}}`. For example, given the property settings shown in [Example 2.4, “Sample Property File”](#), you could define a route as follows:

```
from("${cool.start}")
    .to("log:${cool.start}?showBodyType=false&showExchangeId=${cool.showid}")
    .to("mock:${cool.result}");
```

By default, the placeholder parser looks up the **properties** bean ID in the registry to find the property component. If you prefer, you can explicitly specify the scheme in the endpoint URIs. For example, by prefixing **properties:** to each of the endpoint URIs, you can define the following equivalent route:

```
from("properties:${cool.start}")
    .to("properties:log:${cool.start}?showBodyType=false&showExchangeId=${cool.showid}")
    .to("properties:mock:${cool.result}");
```

When specifying the scheme explicitly, you also have the option of specifying options to the properties component. For example, to override the property file location, you could set the **location** option as follows:

```
from("direct:start").to("properties:${bar.end}?
location=com/mycompany/bar.properties");
```

Substitution in Spring XML files

You can also use property placeholders in the XML DSL, for setting various attributes of the DSL elements. In this context, the placeholder syntax also uses double braces, `{{Key}}`. For example, you could define a **jmxAgent** element using property placeholders, as follows:

```
<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <propertyPlaceholder id="properties"
location="org/apache/camel/spring/jmx.properties"/>

  <!-- we can use property placeholders when we define the JMX agent -->
  <jmxAgent id="agent" registryPort="{{myjmx.port}}"
    usePlatformMBeanServer="{{myjmx.usePlatform}}"
    createConnector="true"
    statisticsLevel="RoutesOnly"
  />

  <route>
    <from uri="seda:start"/>
```

```

    <to uri="mock:result"/>
  </route>
</camelContext>

```

Substitution of XML DSL attribute values

You can use the regular placeholder syntax for specifying attribute values of **xs:string** type—for example, `<jmxAgent registryPort="{myjmx.port}" ...>`. But for attributes of any other type (for example, **xs:int** or **xs:boolean**), you must use the special syntax, **prop:AttributeName="Key"**.

For example, given that a property file defines the **stop.flag** property to have the value, **true**, you can use this property to set the **stopOnException** boolean attribute, as follows:

```

<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:prop="http://camel.apache.org/schema/placeholder"
       ... >

  <bean id="illegal" class="java.lang.IllegalArgumentException">
    <constructor-arg index="0" value="Good grief!"/>
  </bean>

  <camelContext xmlns="http://camel.apache.org/schema/spring">

    <propertyPlaceholder id="properties"


location="classpath:org/apache/camel/component/properties/myprop.properties"
      xmlns="http://camel.apache.org/schema/spring"/>

      <route>
        <from uri="direct:start"/>
        <multicast prop:stopOnException="stop.flag">
          <to uri="mock:a"/>
          <throwException ref="damn"/>
          <to uri="mock:b"/>
        </multicast>
      </route>

    </camelContext>

  </beans>

```



IMPORTANT

The **prop** prefix must be explicitly assigned to the **http://camel.apache.org/schema/placeholder** namespace in your Spring file, as shown in the **beans** element of the preceding example.

Substitution of Java DSL EIP options

When invoking an EIP command in the Java DSL, you can set any EIP option using the value of a property placeholder, by adding a sub-clause of the form, **placeholder("OptionName", "Key")**.

For example, given that a property file defines the **stop.flag** property to have the value, **true**, you can use this property to set the **stopOnException** option of the multicast EIP, as follows:

```
from("direct:start")
    .multicast().placeholder("stopOnException", "stop.flag")
        .to("mock:a").throwException(new
IllegalAccessException("Damn")).to("mock:b");
```

Substitution in Simple language expressions

You can also substitute property placeholders in Simple language expressions, but in this case the syntax of the placeholder is **#{properties:Key}**. For example, you can substitute the **cheese.quote** placeholder inside a Simple expression, as follows:

```
from("direct:start")
    .transform().simple("Hi ${body} do you think
${properties:cheese.quote}?");
```

It is also possible to override the location of the property file using the syntax, **#{properties:Location:Key}**. For example, to substitute the **bar.quote** placeholder using the settings from the **com/mycompany/bar.properties** property file, you can define a Simple expression as follows:

```
from("direct:start")
    .transform().simple("Hi ${body}.
${properties:com/mycompany/bar.properties:bar.quote}.");
```

Integration with OSGi Blueprint property placeholders

If you deploy your route into the Red Hat JBoss Fuse OSGi container, you can integrate the Apache Camel property placeholder mechanism with JBoss Fuse's Blueprint property placeholder mechanism (in fact, the integration is enabled by default). There are two basic approaches to setting up the integration, as follows:

- [the section called "Implicit Blueprint integration"](#).
- [the section called "Explicit Blueprint integration"](#).

Implicit Blueprint integration

If you define a **camelContext** element inside an OSGi Blueprint file, the Apache Camel property placeholder mechanism automatically integrates with the Blueprint property placeholder mechanism. That is, placeholders obeying the Apache Camel syntax (for example, **#{cool.end}**) that appear within the scope of **camelContext** are implicitly resolved by looking up the *Blueprint property placeholder* mechanism.

For example, consider the following route defined in an OSGi Blueprint file, where the last endpoint in the route is defined by the property placeholder, **#{result}**:

■

```

<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
           xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
           xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-
cm/v1.0.0"
           xsi:schemaLocation="
           http://www.osgi.org/xmlns/blueprint/v1.0.0
http://www.osgi.org/xmlns/blueprint/v1.0.0/blueprint.xsd">

    <!-- OSGi blueprint property placeholder -->
    <cm:property-placeholder id="myblueprint.placeholder" persistent-
id="camel.blueprint">
        <!-- list some properties for this test -->
        <cm:default-properties>
            <cm:property name="result" value="mock:result"/>
        </cm:default-properties>
    </cm:property-placeholder>

    <camelContext xmlns="http://camel.apache.org/schema/blueprint">
        <!-- in the route we can use {{ }} placeholders which will look up
in blueprint,
           as Camel will auto detect the OSGi blueprint property
placeholder and use it -->
        <route>
            <from uri="direct:start"/>
            <to uri="mock:foo"/>
            <to uri="{{result}}"/>
        </route>
    </camelContext>

</blueprint>

```

The Blueprint property placeholder mechanism is initialized by creating a **cm:property-placeholder** bean. In the preceding example, the **cm:property-placeholder** bean is associated with the **camel.blueprint** persistent ID, where a persistent ID is the standard way of referencing a group of related properties from the *OSGi Configuration Admin* service. In other words, the **cm:property-placeholder** bean provides access to all of the properties defined under the **camel.blueprint** persistent ID. It is also possible to specify default values for some of the properties (using the nested **cm:property** elements).

In the context of Blueprint, the Apache Camel placeholder mechanism searches for an instance of **cm:property-placeholder** in the bean registry. If it finds such an instance, it automatically integrates the Apache Camel placeholder mechanism, so that placeholders like, **{{result}}**, are resolved by looking up the key in the Blueprint property placeholder mechanism (in this example, through the **myblueprint.placeholder** bean).



NOTE

The default Blueprint placeholder syntax (accessing the Blueprint properties directly) is **\${Key}**. Hence, *outside* the scope of a **camelContext** element, the placeholder syntax you must use is **\${Key}**. Whereas, *inside* the scope of a **camelContext** element, the placeholder syntax you must use is **{{Key}}**.

Explicit Blueprint integration

If you want to have more control over where the Apache Camel property placeholder mechanism finds its properties, you can define a **propertyPlaceholder** element and specify the resolver locations explicitly.

For example, consider the following Blueprint configuration, which differs from the previous example in that it creates an explicit **propertyPlaceholder** instance:

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:cm="http://aries.apache.org/blueprint/xmlns/blueprint-
cm/v1.0.0"
  xsi:schemaLocation="
  http://www.osgi.org/xmlns/blueprint/v1.0.0
http://www.osgi.org/xmlns/blueprint/v1.0.0/blueprint.xsd">

  <!-- OSGI blueprint property placeholder -->
  <cm:property-placeholder id="myblueprint.placeholder" persistent-
id="camel.blueprint">
    <!-- list some properties for this test -->
    <cm:default-properties>
      <cm:property name="result" value="mock:result"/>
    </cm:default-properties>
  </cm:property-placeholder>

  <camelContext xmlns="http://camel.apache.org/schema/blueprint">

    <!-- using Camel properties component and refer to the blueprint
property placeholder by its id -->
    <propertyPlaceholder id="properties"
location="blueprint:myblueprint.placeholder"/>

    <!-- in the route we can use {{ }} placeholders which will lookup
in blueprint -->
    <route>
      <from uri="direct:start"/>
      <to uri="mock:foo"/>
      <to uri="{{result}}"/>
    </route>

  </camelContext>

</blueprint>
```

In the preceding example, the **propertyPlaceholder** element specifies explicitly which **cm:property-placeholder** bean to use by setting the location to **blueprint:myblueprint.placeholder**. That is, the **blueprint:** resolver explicitly references the ID, **myblueprint.placeholder**, of the **cm:property-placeholder** bean.

This style of configuration is useful, if there is more than one **cm:property-placeholder** bean defined in the Blueprint file and you need to specify which one to use. It also makes it possible to source properties from multiple locations, by specifying a comma-separated list of locations. For example, if you wanted to look up properties both from the **cm:property-placeholder** bean and from the properties file, **myproperties.properties**, on the classpath, you could define the **propertyPlaceholder** element as follows:

```
<propertyPlaceholder id="properties"
```

```
location="blueprint:myblueprint.placeholder,classpath:myproperties.properties"/>
```

Integration with Spring property placeholders

If you define your Apache Camel application using XML DSL in a Spring XML file, you can integrate the Apache Camel property placeholder mechanism with Spring property placeholder mechanism by declaring a Spring bean of type, **org.apache.camel.spring.spi.BridgePropertyPlaceholderConfigurer**.

Define a **BridgePropertyPlaceholderConfigurer**, which replaces both Apache Camel's **propertyPlaceholder** element and Spring's **ctx:property-placeholder** element in the Spring XML file. You can then refer to the configured properties using either the Spring **#{PropName}** syntax or the Apache Camel **{{PropName}}** syntax.

For example, defining a bridge property placeholder that reads its property settings from the **cheese.properties** file:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:ctx="http://www.springframework.org/schema/context"
  xsi:schemaLocation="
    http://www.springframework.org/schema/beans
    http://www.springframework.org/schema/beans/spring-beans.xsd
    http://www.springframework.org/schema/context
    http://www.springframework.org/schema/context/spring-context.xsd"
  >

  <!-- Bridge Spring property placeholder with Camel -->
  <!-- Do not use <ctx:property-placeholder ... > at the same time -->
  <bean id="bridgePropertyPlaceholder"

class="org.apache.camel.spring.spi.BridgePropertyPlaceholderConfigurer">
  <property name="location"

value="classpath:org/apache/camel/component/properties/cheese.properties"/
  >
  </bean>

  <!-- A bean that uses Spring property placeholder -->
  <!-- The ${hi} is a spring property placeholder -->
  <bean id="hello"
class="org.apache.camel.component.properties.HelloBean">
  <property name="greeting" value="${hi}"/>
  </bean>

  <camelContext xmlns="http://camel.apache.org/schema/spring">
  <!-- Use Camel's property placeholder {{ }} style -->
  <route>
    <from uri="direct:{{cool.bar}}"/>
    <bean ref="hello"/>
    <to uri="{{cool.end}}"/>
  </route>
```

```
</camelContext>
</beans>
```



NOTE

Alternatively, you can set the **location** attribute of the **BridgePropertyPlaceholderConfigurer** to point at a Spring properties file. The Spring properties file syntax is fully supported.

2.8. ASPECT ORIENTED PROGRAMMING

Overview

The aspect oriented programming (AOP) feature in Apache Camel enables you to apply *before* and *after* processing to a specified portion of a route. As a matter of fact, AOP does *not* provide anything that you could not do with the regular route syntax. The advantage of the AOP syntax, however, is that it enables you to specify before and after processing at a *single point* in the route. In some cases, this gives a more readable syntax. The typical use case for AOP is the application of a symmetrical pair of operations before and after a route fragment is processed. For example, typical pairs of operations that you might want to apply using AOP are: encrypt and decrypt; begin transaction and commit transaction; allocate resources and deallocate resources; and so on.

Java DSL example

In Java DSL, the route fragment to which you apply before and after processing is bracketed between **aop()** and **end()**. For example, the following route performs AOP processing around the route fragment that calls the bean methods:

```
from("jms:queue:inbox")
    .aop().around("log:before", "log:after")
        .to("bean:order?method=validate")
        .to("bean:order?method=handle")
    .end()
    .to("jms:queue:order");
```

Where the **around()** subclause specifies an endpoint, **log:before**, where the exchange is routed *before* processing the route fragment and an endpoint, **log:after**, where the exchange is routed *after* processing the route fragment.

AOP options in the Java DSL

Starting an AOP block with **aop().around()** is probably the most common use case, but the AOP block supports other subclauses, as follows:

- **around()**—specifies *before* and *after* endpoints.
- **begin()**—specifies *before* endpoint only.
- **after()**—specifies *after* endpoint only.
- **aroundFinally()**—specifies a *before* endpoint, and an *after* endpoint that is always called, even when an exception occurs in the enclosed route fragment.

- **afterFinally()**—specifies an *after* endpoint that is always called, even when an exception occurs in the enclosed route fragment.

Spring XML example

In the XML DSL, the route fragment to which you apply *before* and *after* processing is enclosed in the **aop** element. For example, the following Spring XML route performs AOP processing around the route fragment that calls the bean methods:

```
<route>
  <from uri="jms:queue:inbox"/>
  <aop beforeUri="log:before" afterUri="log:after">
    <to uri="bean:order?method=validate"/>
    <to uri="bean:order?method=handle"/>
  </aop>
  <to uri="jms:queue:order"/>
</route>
```

Where the **beforeUri** attribute specifies the endpoint where the exchange is routed *before* processing the route fragment, and the **afterUri** attribute specifies the endpoint where the exchange is routed *after* processing the route fragment.

AOP options in the Spring XML

The **aop** element supports the following optional attributes:

- **beforeUri**
- **afterUri**
- **afterFinallyUri**

The various use cases described for the Java DSL can be obtained in Spring XML using the appropriate combinations of these attributes. For example, the **aroundFinally()** Java DSL subclass is equivalent to the combination of **beforeUri** and **afterFinallyUri** in Spring XML.

2.9. THREADING MODEL

Java thread pool API

The Apache Camel threading model is based on the powerful Java concurrency API, java.util.concurrent, that first became available in Sun's JDK 1.5. The key interface in this API is the **ExecutorService** interface, which represents a thread pool. Using the concurrency API, you can create many different kinds of thread pool, covering a wide range of scenarios.

Apache Camel thread pool API

The Apache Camel thread pool API builds on the Java concurrency API by providing a central factory (of **org.apache.camel.spi.ExecutorServiceManager** type) for all of the thread pools in your Apache Camel application. Centralising the creation of thread pools in this way provides several advantages, including:

- Simplified creation of thread pools, using utility classes.
- Integrating thread pools with graceful shutdown.
- Threads automatically given informative names, which is beneficial for logging and management.

Component threading model

Some Apache Camel components—such as SEDA, JMS, and Jetty—are inherently multi-threaded. These components have all been implemented using the Apache Camel threading model and thread pool API.

If you are planning to implement your own Apache Camel component, it is recommended that you integrate your threading code with the Apache Camel threading model. For example, if your component needs a thread pool, it is recommended that you create it using the CamelContext's **ExecutorServiceManager** object.

Processor threading model

Some of the standard processors in Apache Camel create their own thread pool by default. These threading-aware processors are also integrated with the Apache Camel threading model and they provide various options that enable you to customize customize the thread pools that they use.

[Table 2.8, “Processor Threading Options”](#) shows the various options for controlling and setting thread pools on the threading-aware processors built-in to Apache Camel.

Table 2.8. Processor Threading Options

Processor	Java DSL	XML DSL
aggregate	<pre>parallelProcessing() executorService() executorServiceRef()</pre>	<pre>@parallelProcessing @executorServiceRef</pre>
multicast	<pre>parallelProcessing() executorService() executorServiceRef()</pre>	<pre>@parallelProcessing @executorServiceRef</pre>
recipientList	<pre>parallelProcessing() executorService() executorServiceRef()</pre>	<pre>@parallelProcessing @executorServiceRef</pre>
split	<pre>parallelProcessing() executorService() executorServiceRef()</pre>	<pre>@parallelProcessing @executorServiceRef</pre>

Processor	Java DSL	XML DSL
threads	<pre> executorService() executorServiceRef() poolSize() maxPoolSize() keepAliveTime() timeUnit() maxQueueSize() rejectedPolicy() </pre>	<pre> @executorServiceRef @poolSize @maxPoolSize @keepAliveTime @timeUnit @maxQueueSize @rejectedPolicy </pre>
wireTap	<pre> wireTap(String uri, ExecutorService executorService) wireTap(String uri, String executorServiceRef) </pre>	<pre> @executorServiceRef </pre>

Creating a default thread pool

To create a default thread pool for one of the threading-aware processors, enable the **parallelProcessing** option, using the **parallelProcessing()** sub-clause, in the Java DSL, or the **parallelProcessing** attribute, in the XML DSL.

For example, in the Java DSL, you can invoke the multicast processor with a default thread pool (where the thread pool is used to process the multicast destinations concurrently) as follows:

```

from("direct:start")
  .multicast().parallelProcessing()
  .to("mock:first")
  .to("mock:second")
  .to("mock:third");

```

You can define the same route in XML DSL as follows

```

<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <multicast parallelProcessing="true">
      <to uri="mock:first"/>
      <to uri="mock:second"/>
      <to uri="mock:third"/>
    </multicast>
  </route>
</camelContext>

```

Default thread pool profile settings

The default thread pools are automatically created by a thread factory that takes its settings from the *default thread pool profile*. The default thread pool profile has the settings shown in [Table 2.9, “Default Thread Pool Profile Settings”](#) (assuming that these settings have not been modified by the application code).

Table 2.9. Default Thread Pool Profile Settings

Thread Option	Default Value
<code>maxQueueSize</code>	1000
<code>poolSize</code>	10
<code>maxPoolSize</code>	20
<code>keepAliveTime</code>	60 (seconds)
<code>rejectedPolicy</code>	CallerRuns

Changing the default thread pool profile

It is possible to change the default thread pool profile settings, so that all subsequent default thread pools will be created with the custom settings. You can change the profile either in Java or in Spring XML.

For example, in the Java DSL, you can customize the `poolSize` option and the `maxQueueSize` option in the default thread pool profile, as follows:

```
// Java
import org.apache.camel.spi.ExecutorServiceManager;
import org.apache.camel.spi.ThreadPoolProfile;
...
ExecutorServiceManager manager = context.getExecutorServiceManager();
ThreadPoolProfile defaultProfile = manager.getDefaultThreadPoolProfile();

// Now, customize the profile settings.
defaultProfile.setPoolSize(3);
defaultProfile.setMaxQueueSize(100);
...
```

In the XML DSL, you can customize the default thread pool profile, as follows:

```
<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <threadPoolProfile
    id="changedProfile"
    defaultProfile="true"
    poolSize="3"
    maxQueueSize="100"/>
    ...
</camelContext>
```

Note that it is essential to set the `defaultProfile` attribute to `true` in the preceding XML DSL example, otherwise the thread pool profile would be treated like a custom thread pool

profile (see [the section called “Creating a custom thread pool profile”](#)), instead of replacing the default thread pool profile.

Customizing a processor's thread pool

It is also possible to specify the thread pool for a threading-aware processor more directly, using either the `executorService` or `executorServiceRef` options (where these options are used instead of the `parallelProcessing` option). There are two approaches you can use to customize a processor's thread pool, as follows:

- *Specify a custom thread pool*—explicitly create an `ExecutorService` (thread pool) instance and pass it to the `executorService` option.
- *Specify a custom thread pool profile*—create and register a custom thread pool factory. When you reference this factory using the `executorServiceRef` option, the processor automatically uses the factory to create a custom thread pool instance.

When you pass a bean ID to the `executorServiceRef` option, the threading-aware processor first tries to find a custom thread pool with that ID in the registry. If no thread pool is registered with that ID, the processor then attempts to look up a custom thread pool profile in the registry and uses the custom thread pool profile to instantiate a custom thread pool.

Creating a custom thread pool

A custom thread pool can be any thread pool of `java.util.concurrent.ExecutorService` type. The following approaches to creating a thread pool instance are recommended in Apache Camel:

- Use the `org.apache.camel.builder.ThreadPoolBuilder` utility to build the thread pool class.
- Use the `org.apache.camel.spi.ExecutorServiceManager` instance from the current `CamelContext` to create the thread pool class.

Ultimately, there is not much difference between the two approaches, because the `ThreadPoolBuilder` is actually defined using the `ExecutorServiceManager` instance. Normally, the `ThreadPoolBuilder` is preferred, because it offers a simpler approach. But there is at least one kind of thread (the `ScheduledExecutorService`) that can only be created by accessing the `ExecutorServiceManager` instance directory.

[Table 2.10, “Thread Pool Builder Options”](#) shows the options supported by the `ThreadPoolBuilder` class, which you can set when defining a new custom thread pool.

Table 2.10. Thread Pool Builder Options

Builder Option	Description
<code>maxQueueSize()</code>	Sets the maximum number of pending tasks that this thread pool can store in its incoming task queue. A value of <code>-1</code> specifies an unbounded queue. Default value is taken from default thread pool profile.

Builder Option	Description
poolSize()	Sets the minimum number of threads in the pool (this is also the initial pool size). Default value is taken from default thread pool profile.
maxPoolSize()	Sets the maximum number of threads that can be in the pool. Default value is taken from default thread pool profile.
keepAliveTime()	If any threads are idle for longer than this period of time (specified in seconds), they are terminated. This allows the thread pool to shrink when the load is light. Default value is taken from default thread pool profile.
rejectedPolicy()	<p>Specifies what course of action to take, if the incoming task queue is full. You can specify four possible values:</p> <p>CallerRuns <i>(Default value)</i> Gets the caller thread to run the latest incoming task. As a side effect, this option prevents the caller thread from receiving any more tasks until it has finished processing the latest incoming task.</p> <p>Abort Aborts the latest incoming task by throwing an exception.</p> <p>Discard Quietly discards the latest incoming task.</p> <p>DiscardOldest Discards the oldest unhandled task and then attempts to enqueue the latest incoming task in the task queue.</p>
build()	Finishes building the custom thread pool and registers the new thread pool under the ID specified as the argument to build() .

In Java DSL, you can define a custom thread pool using the **ThreadPoolBuilder**, as follows:

```
// Java
import org.apache.camel.builder.ThreadPoolBuilder;
import java.util.concurrent.ExecutorService;
...
ThreadPoolBuilder poolBuilder = new ThreadPoolBuilder(context);
ExecutorService customPool =
```

```

poolBuilder.poolSize(5).maxPoolSize(5).maxQueueSize(100).build("customPool");
...

from("direct:start")
  .multicast().executorService(customPool)
    .to("mock:first")
    .to("mock:second")
    .to("mock:third");

```

Instead of passing the object reference, **customPool**, directly to the **executorService()** option, you can look up the thread pool in the registry, by passing its bean ID to the **executorServiceRef()** option, as follows:

```

// Java
from("direct:start")
  .multicast().executorServiceRef("customPool")
    .to("mock:first")
    .to("mock:second")
    .to("mock:third");

```

In XML DSL, you access the **ThreadPoolBuilder** using the **threadPool** element. You can then reference the custom thread pool using the **executorServiceRef** attribute to look up the thread pool by ID in the Spring registry, as follows:

```

<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <threadPool id="customPool"
    poolSize="5"
    maxPoolSize="5"
    maxQueueSize="100" />

  <route>
    <from uri="direct:start"/>
    <multicast executorServiceRef="customPool">
      <to uri="mock:first"/>
      <to uri="mock:second"/>
      <to uri="mock:third"/>
    </multicast>
  </route>
</camelContext>

```

Creating a custom thread pool profile

If you have many custom thread pool instances to create, you might find it more convenient to define a custom thread pool profile, which acts as a factory for thread pools. Whenever you reference a thread pool profile from a threading-aware processor, the processor automatically uses the profile to create a new thread pool instance. You can define a custom thread pool profile either in Java DSL or in XML DSL.

For example, in Java DSL you can create a custom thread pool profile with the bean ID, **customProfile**, and reference it from within a route, as follows:

```

// Java
import org.apache.camel.spi.ThreadPoolProfile;

```

```

import org.apache.camel.impl.ThreadPoolProfileSupport;
...
// Create the custom thread pool profile
ThreadPoolProfile customProfile = new
ThreadPoolProfileSupport("customProfile");
customProfile.setPoolSize(5);
customProfile.setMaxPoolSize(5);
customProfile.setMaxQueueSize(100);
context.getExecutorServiceManager().registerThreadPoolProfile(customProfile);
...
// Reference the custom thread pool profile in a route
from("direct:start")
    .multicast().executorServiceRef("customProfile")
        .to("mock:first")
        .to("mock:second")
        .to("mock:third");

```

In XML DSL, use the **threadPoolProfile** element to create a custom pool profile (where you let the **defaultProfile** option default to **false**, because this is *not* a default thread pool profile). You can create a custom thread pool profile with the bean ID, **customProfile**, and reference it from within a route, as follows:

```

<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <threadPoolProfile
    id="customProfile"
    poolSize="5"
    maxPoolSize="5"
    maxQueueSize="100" />

  <route>
    <from uri="direct:start"/>
    <multicast executorServiceRef="customProfile">
      <to uri="mock:first"/>
      <to uri="mock:second"/>
      <to uri="mock:third"/>
    </multicast>
  </route>
</camelContext>

```

Sharing a thread pool between components

Some of the standard poll-based components—such as File and FTP—allow you to specify the thread pool to use. This makes it possible for different components to share the same thread pool, reducing the overall number of threads in the JVM.

For example, the [File component](#) and the [FTP component](#) both expose the **scheduledExecutorService** property, which you can use to specify the component's **ExecutorService** object.

Customizing thread names

To make the application logs more readable, it is often a good idea to customize the thread names (which are used to identify threads in the log). To customize thread names, you can configure the *thread name pattern* by calling the **setThreadNamePattern** method on the

ExecutorServiceStrategy class or the **ExecutorServiceManager** class. Alternatively, an easier way to set the thread name pattern is to set the **threadNamePattern** property on the **CamelContext** object.

The following placeholders can be used in a thread name pattern:

#camelId#

The name of the current **CamelContext**.

#counter#

A unique thread identifier, implemented as an incrementing counter.

#name#

The regular Camel thread name.

#longName#

The long thread name—which can include endpoint parameters and so on.

The following is a typical example of a thread name pattern:

```
Camel (#camelId#) thread #counter# - #name#
```

The following example shows how to set the **threadNamePattern** attribute on a Camel context using XML DSL:

```
<camelContext xmlns="http://camel.apache.org/schema/spring"
  threadNamePattern="Riding the thread #counter#" >
  <route>
    <from uri="seda:start"/>
    <to uri="log:result"/>
    <to uri="mock:result"/>
  </route>
</camelContext>
```

2.10. CONTROLLING START-UP AND SHUTDOWN OF ROUTES

Overview

By default, routes are automatically started when your Apache Camel application (as represented by the **CamelContext** instance) starts up and routes are automatically shut down when your Apache Camel application shuts down. For non-critical deployments, the details of the shutdown sequence are usually not very important. But in a production environment, it is often crucial that existing tasks should run to completion during shutdown, in order to avoid data loss. You typically also want to control the order in which routes shut down, so that dependencies are not violated (which would prevent existing tasks from running to completion).

For this reason, Apache Camel provides a set of features to support *graceful shutdown* of applications. Graceful shutdown gives you full control over the stopping and starting of routes, enabling you to control the shutdown order of routes and enabling current tasks to run to completion.

Setting the route ID

It is good practice to assign a route ID to each of your routes. As well as making logging messages and management features more informative, the use of route IDs enables you to apply greater control over the stopping and starting of routes.

For example, in the Java DSL, you can assign the route ID, **myCustomerRouteId**, to a route by invoking the **routeId()** command as follows:

```
from("SourceURI").routeId("myCustomRouteId").process(...).to(TargetURI);
```

In the XML DSL, set the **route** element's **id** attribute, as follows:

```
<camelContext id="CamelContextID"
xmlns="http://camel.apache.org/schema/spring">
  <route id="myCustomRouteId" >
    <from uri="SourceURI"/>
    <process ref="someProcessorId"/>
    <to uri="TargetURI"/>
  </route>
</camelContext>
```

Disabling automatic start-up of routes

By default, all of the routes that the CamelContext knows about at start time will be started automatically. If you want to control the start-up of a particular route manually, however, you might prefer to disable automatic start-up for that route.

To control whether a Java DSL route starts up automatically, invoke the **autoStartup** command, either with a **boolean** argument (**true** or **false**) or a **String** argument (**true** or **false**). For example, you can disable automatic start-up of a route in the Java DSL, as follows:

```
from("SourceURI")
  .routeId("nonAuto")
  .autoStartup(false)
  .to(TargetURI);
```

You can disable automatic start-up of a route in the XML DSL by setting the **autoStartup** attribute to **false** on the **route** element, as follows:

```
<camelContext id="CamelContextID"
xmlns="http://camel.apache.org/schema/spring">
  <route id="nonAuto" autoStartup="false">
    <from uri="SourceURI"/>
    <to uri="TargetURI"/>
  </route>
</camelContext>
```

Manually starting and stopping routes

You can manually start or stop a route at any time in Java by invoking the **startRoute()** and **stopRoute()** methods on the **CamelContext** instance. For example, to start the route having the route ID, **nonAuto**, invoke the **startRoute()** method on the **CamelContext**

instance, **context**, as follows:

```
// Java
context.startRoute("nonAuto");
```

To stop the route having the route ID, **nonAuto**, invoke the **stopRoute()** method on the **CamelContext** instance, **context**, as follows:

```
// Java
context.stopRoute("nonAuto");
```

Startup order of routes

By default, Apache Camel starts up routes in a non-deterministic order. In some applications, however, it can be important to control the startup order. To control the startup order in the Java DSL, use the **startupOrder()** command, which takes a positive integer value as its argument. The route with the lowest integer value starts first, followed by the routes with successively higher startup order values.

For example, the first two routes in the following example are linked together through the **seda:buffer** endpoint. You can ensure that the first route segment starts *after* the second route segment by assigning startup orders (2 and 1 respectively), as follows:

Example 2.5. Startup Order in Java DSL

```
from("jetty:http://fooserver:8080")
    .routeId("first")
    .startupOrder(2)
    .to("seda:buffer");

from("seda:buffer")
    .routeId("second")
    .startupOrder(1)
    .to("mock:result");

// This route's startup order is unspecified
from("jms:queue:foo").to("jms:queue:bar");
```

Or in Spring XML, you can achieve the same effect by setting the **route** element's **startupOrder** attribute, as follows:

Example 2.6. Startup Order in XML DSL

```
<route id="first" startupOrder="2">
  <from uri="jetty:http://fooserver:8080"/>
  <to uri="seda:buffer"/>
</route>

<route id="second" startupOrder="1">
  <from uri="seda:buffer"/>
  <to uri="mock:result"/>
</route>
```

```

<!-- This route's startup order is unspecified -->
<route>
  <from uri="jms:queue:foo"/>
  <to uri="jms:queue:bar"/>
</route>

```

Each route must be assigned a *unique* startup order value. You can choose any positive integer value that is less than 1000. Values of 1000 and over are reserved for Apache Camel, which automatically assigns these values to routes without an explicit startup value. For example, the last route in the preceding example would automatically be assigned the startup value, 1000 (so it starts up after the first two routes).

Shutdown sequence


When a **CamelContext** instance is shutting down, Apache Camel controls the shutdown sequence using a pluggable *shutdown strategy*. The default shutdown strategy implements the following shutdown sequence:

1. Routes are shut down in the *reverse* of the start-up order.
2. Normally, the shutdown strategy waits until the currently active exchanges have finished processing. The treatment of running tasks is configurable, however.
3. Overall, the shutdown sequence is bound by a timeout (default, 300 seconds). If the shutdown sequence exceeds this timeout, the shutdown strategy will force shutdown to occur, even if some tasks are still running.

Shutdown order of routes

Routes are shut down in the reverse of the start-up order. That is, when a start-up order is defined using the **startupOrder()** command (in Java DSL) or **startupOrder** attribute (in XML DSL), the first route to shut down is the route with the *highest* integer value assigned by the start-up order and the last route to shut down is the route with the *lowest* integer value assigned by the start-up order.

For example, in [Example 2.5, “Startup Order in Java DSL”](#), the first route segment to be shut down is the route with the ID, **first**, and the second route segment to be shut down is the route with the ID, **second**. This example illustrates a general rule, which you should observe when shutting down routes: *the routes that expose externally-accessible consumer endpoints should be shut down first*, because this helps to throttle the flow of messages through the rest of the route graph.



NOTE

Apache Camel also provides the option **shutdownRoute(Defer)**, which enables you to specify that a route must be amongst the last routes to shut down (overriding the start-up order value). But you should rarely ever need this option. This option was mainly needed as a workaround for earlier versions of Apache Camel (prior to 2.3), for which routes would shut down in the *same* order as the start-up order.

Shutting down running tasks in a route

If a route is still processing messages when the shutdown starts, the shutdown strategy

normally waits until the currently active exchange has finished processing before shutting down the route. This behavior can be configured on each route using the **shutdownRunningTask** option, which can take either of the following values:

ShutdownRunningTask.CompleteCurrentTaskOnly

(Default) Usually, a route operates on just a single message at a time, so you can safely shut down the route after the current task has completed.

ShutdownRunningTask.CompleteAllTasks

Specify this option in order to shut down *batch consumers* gracefully. Some consumer endpoints (for example, File, FTP, Mail, iBATIS, and JPA) operate on a batch of messages at a time. For these endpoints, it is more appropriate to wait until all of the messages in the current batch have completed.

For example, to shut down a File consumer endpoint gracefully, you should specify the **CompleteAllTasks** option, as shown in the following Java DSL fragment:

```
// Java
public void configure() throws Exception {
    from("file:target/pending")
        .routeId("first").startupOrder(2)
        .shutdownRunningTask(ShutdownRunningTask.CompleteAllTasks)
        .delay(1000).to("seda:foo");

    from("seda:foo")
        .routeId("second").startupOrder(1)
        .to("mock:bar");
}
```

The same route can be defined in the XML DSL as follows:

```
<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <!-- let this route complete all its pending messages when asked to
  shut down -->
  <route id="first"
    startupOrder="2"
    shutdownRunningTask="CompleteAllTasks">
    <from uri="file:target/pending"/>
    <delay><constant>1000</constant></delay>
    <to uri="seda:foo"/>
  </route>

  <route id="second" startupOrder="1">
    <from uri="seda:foo"/>
    <to uri="mock:bar"/>
  </route>
</camelContext>
```

Shutdown timeout

The shutdown timeout has a default value of 300 seconds. You can change the value of the timeout by invoking the **setTimeout()** method on the shutdown strategy. For example, you can change the timeout value to 600 seconds, as follows:

```
// Java
// context = CamelContext instance
context.getShutdownStrategy().setTimeout(600);
```

Integration with custom components

If you are implementing a custom Apache Camel component (which also inherits from the `org.apache.camel.Service` interface), you can ensure that your custom code receives a shutdown notification by implementing the `org.apache.camel.spi.ShutdownPrepared` interface. This gives the component an opportunity to execute custom code in preparation for shutdown.

2.11. SCHEDULED ROUTE POLICY

2.11.1. Overview of Scheduled Route Policies

Overview

A scheduled route policy can be used to trigger events that affect a route at runtime. In particular, the implementations that are currently available enable you to start, stop, suspend, or resume a route at any time (or times) specified by the policy.

Scheduling tasks

The scheduled route policies are capable of triggering the following kinds of event:

- *Start a route*—start the route at the time (or times) specified. This event only has an effect, if the route is currently in a stopped state, awaiting activation.
- *Stop a route*—stop the route at the time (or times) specified. This event only has an effect, if the route is currently active.
- *Suspend a route*—temporarily de-activate the consumer endpoint at the start of the route (as specified in `from()`). The rest of the route is still active, but clients will not be able to send new messages into the route.
- *Resume a route*—re-activate the consumer endpoint at the start of the route, returning the route to a fully active state.

Quartz component

The Quartz component is a timer component based on Terracotta's [Quartz](#), which is an open source implementation of a job scheduler. The Quartz component provides the underlying implementation for both the simple scheduled route policy and the cron scheduled route policy.

2.11.2. Simple Scheduled Route Policy

Overview

The simple scheduled route policy is a route policy that enables you to start, stop, suspend, and resume routes, where the timing of these events is defined by providing the time and

date of an initial event and (optionally) by specifying a certain number of subsequent repetitions. To define a simple scheduled route policy, create an instance of the following class:

```
org.apache.camel.routePolicy.quartz.SimpleScheduledRoutePolicy
```

Dependency

The simple scheduled route policy depends on the Quartz component, **camel-quartz**. For example, if you are using Maven as your build system, you would need to add a dependency on the **camel-quartz** artifact.

Java DSL example


[Example 2.7, “Java DSL Example of Simple Scheduled Route”](#) shows how to schedule a route to start up using the Java DSL. The initial start time, **startTime**, is defined to be 3 seconds after the current time. The policy is also configured to start the route a *second* time, 3 seconds after the initial start time, which is configured by setting **routeStartRepeatCount** to 1 and **routeStartRepeatInterval** to 3000 milliseconds.

In Java DSL, you attach the route policy to the route by calling the **routePolicy()** DSL command in the route.

Example 2.7. Java DSL Example of Simple Scheduled Route

```
// Java
SimpleScheduledRoutePolicy policy = new SimpleScheduledRoutePolicy();
long startTime = System.currentTimeMillis() + 3000L;
policy.setRouteStartDate(new Date(startTime));
policy.setRouteStartRepeatCount(1);
policy.setRouteStartRepeatInterval(3000);

from("direct:start")
    .routeId("test")
    .routePolicy(policy)
    .to("mock:success");
```



NOTE

You can specify multiple policies on the route by calling **routePolicy()** with multiple arguments.

XML DSL example

[Example 2.8, “XML DSL Example of Simple Scheduled Route”](#) shows how to schedule a route to start up using the XML DSL.

In XML DSL, you attach the route policy to the route by setting the **routePolicyRef** attribute on the **route** element.

Example 2.8. XML DSL Example of Simple Scheduled Route

```

<bean id="date" class="java.util.Date"/>

<bean id="startPolicy"
class="org.apache.camel.routepolicy.quartz.SimpleScheduledRoutePolicy">
  <property name="routeStartDate" ref="date"/>
  <property name="routeStartRepeatCount" value="1"/>
  <property name="routeStartRepeatInterval" value="3000"/>
</bean>

<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route id="myroute" routePolicyRef="startPolicy">
    <from uri="direct:start"/>
    <to uri="mock:success"/>
  </route>
</camelContext>

```



NOTE

You can specify multiple policies on the route by setting the value of **routePolicyRef** as a comma-separated list of bean IDs.

Defining dates and times

The initial times of the triggers used in the simple scheduled route policy are specified using the **java.util.Date** type. The most flexible way to define a **Date** instance is through the **java.util.GregorianCalendar** class. Use the convenient constructors and methods of the **GregorianCalendar** class to define a date and then obtain a **Date** instance by calling **GregorianCalendar.getTime()**.

For example, to define the time and date for January 1, 2011 at noon, call a **GregorianCalendar** constructor as follows:

```

// Java
import java.util.GregorianCalendar;
import java.util.Calendar;
...
GregorianCalendar gc = new GregorianCalendar(
    2011,
    Calendar.JANUARY,
    1,
    12, // hourOfDay
    0, // minutes
    0 // seconds
);

java.util.Date triggerDate = gc.getTime();

```

The **GregorianCalendar** class also supports the definition of times in different time zones. By default, it uses the local time zone on your computer.

Graceful shutdown

When you configure a simple scheduled route policy to stop a route, the route stopping

algorithm is automatically integrated with the graceful shutdown procedure (see [Section 2.10, “Controlling Start-Up and Shutdown of Routes”](#)). This means that the task waits until the current exchange has finished processing before shutting down the route. You can set a timeout, however, that forces the route to stop after the specified time, irrespective of whether or not the route has finished processing the exchange.

Scheduling tasks

You can use a simple scheduled route policy to define one or more of the following scheduling tasks:

- [the section called “Starting a route”](#).
- [the section called “Stopping a route”](#).
- [the section called “Suspending a route”](#).
- [the section called “Resuming a route”](#).

Starting a route

The following table lists the parameters for scheduling one or more route starts.

Parameter	Type	Default	Description
<code>routeStartDate</code>	<code>java.util.Date</code>	<i>None</i>	Specifies the date and time when the route is started for the first time.
<code>routeStartRepeat Count</code>	<code>int</code>	<code>0</code>	When set to a non-zero value, specifies how many times the route should be started.
<code>routeStartRepeat Interval</code>	<code>long</code>	<code>0</code>	Specifies the time interval between starts, in units of milliseconds.

Stopping a route

The following table lists the parameters for scheduling one or more route stops.

Parameter	Type	Default	Description
<code>routeStopDate</code>	<code>java.util.Date</code>	<i>None</i>	Specifies the date and time when the route is stopped for the first time.

Parameter	Type	Default	Description
<code>routeStopRepeatCount</code>	<code>int</code>	<code>0</code>	When set to a non-zero value, specifies how many times the route should be stopped.
<code>routeStopRepeatInterval</code>	<code>long</code>	<code>0</code>	Specifies the time interval between stops, in units of milliseconds.
<code>routeStopGracePeriod</code>	<code>int</code>	<code>10000</code>	Specifies how long to wait for the current exchange to finish processing (grace period) before forcibly stopping the route. Set to 0 for an infinite grace period.
<code>routeStopTimeUnit</code>	<code>long</code>	<code>TimeUnit.MILLISECONDS</code>	Specifies the time unit of the grace period.

Suspending a route

The following table lists the parameters for scheduling the suspension of a route one or more times.

Parameter	Type	Default	Description
<code>routeSuspendDate</code>	<code>java.util.Date</code>	<i>None</i>	Specifies the date and time when the route is suspended for the first time.
<code>routeSuspendRepeatCount</code>	<code>int</code>	<code>0</code>	When set to a non-zero value, specifies how many times the route should be suspended.
<code>routeSuspendRepeatInterval</code>	<code>long</code>	<code>0</code>	Specifies the time interval between suspends, in units of milliseconds.

Resuming a route

The following table lists the parameters for scheduling the resumption of a route one or more times.

Parameter	Type	Default	Description
<code>routeResumeDate</code>	<code>java.util.Date</code>	<i>None</i>	Specifies the date and time when the route is resumed for the first time.
<code>routeResumeRepeatCount</code>	<code>int</code>	0	When set to a non-zero value, specifies how many times the route should be resumed.
<code>routeResumeRepeatInterval</code>	<code>long</code>	0	Specifies the time interval between resumes, in units of milliseconds.

2.11.3. Cron Scheduled Route Policy

Overview

The cron scheduled route policy is a route policy that enables you to start, stop, suspend, and resume routes, where the timing of these events is specified using cron expressions. To define a cron scheduled route policy, create an instance of the following class:

```
org.apache.camel.routepolicy.quartz.CronScheduledRoutePolicy
```

Dependency

The simple scheduled route policy depends on the Quartz component, **camel-quartz**. For example, if you are using Maven as your build system, you would need to add a dependency on the **camel-quartz** artifact.

Java DSL example


[Example 2.9, “Java DSL Example of a Cron Scheduled Route”](#) shows how to schedule a route to start up using the Java DSL. The policy is configured with the cron expression, `*/*/*/*/*/?`, which triggers a start event every 3 seconds.

In Java DSL, you attach the route policy to the route by calling the `routePolicy()` DSL command in the route.

Example 2.9. Java DSL Example of a Cron Scheduled Route

```
// Java
CronScheduledRoutePolicy policy = new CronScheduledRoutePolicy();
policy.setRouteStartTime("*/3 * * * * ?");
```

```
from("direct:start")
  .routeId("test")
  .routePolicy(policy)
  .to("mock:success");;
```

**NOTE**

You can specify multiple policies on the route by calling **routePolicy()** with multiple arguments.

XML DSL example

[Example 2.10, “XML DSL Example of a Cron Scheduled Route”](#) shows how to schedule a route to start up using the XML DSL.

In XML DSL, you attach the route policy to the route by setting the **routePolicyRef** attribute on the **route** element.

Example 2.10. XML DSL Example of a Cron Scheduled Route

```
<bean id="date" class="org.apache.camel.routepolicy.quartz.SimpleDate"/>
<bean id="startPolicy"
class="org.apache.camel.routepolicy.quartz.CronScheduledRoutePolicy">
  <property name="routeStartTime" value="*/3 * * * * ?"/>
</bean>

<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route id="testRoute" routePolicyRef="startPolicy">
    <from uri="direct:start"/>
    <to uri="mock:success"/>
  </route>
</camelContext>
```

**NOTE**

You can specify multiple policies on the route by setting the value of **routePolicyRef** as a comma-separated list of bean IDs.

Defining cron expressions

The *cron expression* syntax has its origins in the UNIX **cron** utility, which schedules jobs to run in the background on a UNIX system. A cron expression is effectively a syntax for wildcarding dates and times that enables you to specify either a single event or multiple events that recur periodically.

A cron expression consists of 6 or 7 fields in the following order:

```
Seconds Minutes Hours DayOfMonth Month DayOfWeek [Year]
```

The **Year** field is optional and usually omitted, unless you want to define an event that occurs once and once only. Each field consists of a mixture of literals and special characters. For example, the following cron expression specifies an event that fires once every day at midnight:

```
0 0 24 * * ?
```

The ***** character is a wildcard that matches every value of a field. Hence, the preceding expression matches every day of every month. The **?** character is a dummy placeholder that means *ignore this field*. It always appears either in the **DayOfMonth** field or in the **DayOfWeek** field, because it is not logically consistent to specify both of these fields at the same time. For example, if you want to schedule an event that fires once a day, but only from Monday to Friday, use the following cron expression:

```
0 0 24 ? * MON-FRI
```

Where the hyphen character specifies a range, **MON-FRI**. You can also use the forward slash character, **/**, to specify increments. For example, to specify that an event fires every 5 minutes, use the following cron expression:

```
0 0/5 * * * ?
```

For a full explanation of the cron expression syntax, see the Wikipedia article on [CRON expressions](#).

Scheduling tasks

You can use a cron scheduled route policy to define one or more of the following scheduling tasks:

- [the section called “Starting a route”](#).
- [the section called “Stopping a route”](#).
- [the section called “Suspending a route”](#).
- [the section called “Resuming a route”](#).

Starting a route

The following table lists the parameters for scheduling one or more route starts.

Parameter	Type	Default	Description
<code>routeStartString</code>	String	None	Specifies a cron expression that triggers one or more route start events.

Stopping a route

The following table lists the parameters for scheduling one or more route stops.

Parameter	Type	Default	Description
<code>routeStopTime</code>	<code>String</code>	<i>None</i>	Specifies a cron expression that triggers one or more route stop events.
<code>routeStopGracePeriod</code>	<code>int</code>	10000	Specifies how long to wait for the current exchange to finish processing (grace period) before forcibly stopping the route. Set to 0 for an infinite grace period.
<code>routeStopTimeUnit</code>	<code>long</code>	TimeUnit.MILLISECONDS	Specifies the time unit of the grace period.

Suspending a route

The following table lists the parameters for scheduling the suspension of a route one or more times.

Parameter	Type	Default	Description
<code>routeSuspendTime</code>	<code>String</code>	<i>None</i>	Specifies a cron expression that triggers one or more route suspend events.

Resuming a route

The following table lists the parameters for scheduling the resumption of a route one or more times.

Parameter	Type	Default	Description
<code>routeResumeTime</code>	<code>String</code>	<i>None</i>	Specifies a cron expression that triggers one or more route resume events.

2.12. JMX NAMING

Overview

Apache Camel allows you to customise the name of a **CamelContext** bean as it appears in JMX, by defining a *management name pattern* for it. For example, you can customise the name pattern of an XML **CamelContext** instance, as follows:

```
<camelContext id="myCamel" managementNamePattern="#name#">
    ...
</camelContext>
```

If you do not explicitly set a name pattern for the **CamelContext** bean, Apache Camel reverts to a default naming strategy.

Default naming strategy

By default, the JMX name of a **CamelContext** bean is equal to the value of the bean's **id** attribute, prefixed by the current bundle ID. For example, if the **id** attribute on a **camelContext** element is **myCamel** and the current bundle ID is 250, the JMX name would be **250-myCamel**. In cases where there is more than one **CamelContext** instance with the same **id** in the bundle, the JMX name is disambiguated by adding a counter value as a suffix. For example, if there are multiple instances of **myCamel** in the bundle, the corresponding JMX MBeans are named as follows:

```
250-myCamel-1
250-myCamel-2
250-myCamel-3
...
```

Customising the JMX naming strategy

One drawback of the default naming strategy is that you cannot guarantee that a given **CamelContext** bean will have the same JMX name between runs. If you want to have greater consistency between runs, you can control the JMX name more precisely by defining a *JMX name pattern* for the **CamelContext** instances.

Specifying a name pattern in Java

To specify a name pattern on a **CamelContext** in Java, call the **setNamePattern** method, as follows:

```
// Java
context.getManagementNameStrategy().setNamePattern("#name#");
```

Specifying a name pattern in XML

To specify a name pattern on a **CamelContext** in XML, set the **managementNamePattern** attribute on the **camelContext** element, as follows:

```
<camelContext id="myCamel" managementNamePattern="#name#">
```

Name pattern tokens

You can construct a JMX name pattern by mixing literal text with any of the following tokens:

Table 2.11. JMX Name Pattern Tokens

Token	Description
#camelId#	Value of the id attribute on the CamelContext bean.
#name#	Same as #camelId# .
#counter#	An incrementing counter (starting at 1).
#bundleId#	The OSGi bundle ID of the deployed bundle (<i>OSGi only</i>).
#symbolicName#	The OSGi symbolic name (<i>OSGi only</i>).
#version#	The OSGi bundle version (<i>OSGi only</i>).

Examples

Here are some examples of JMX name patterns you could define using the supported tokens:

```
<camelContext id="fooContext" managementNamePattern="FooApplication-
#name#">
    ...
</camelContext>
<camelContext id="myCamel" managementNamePattern="#bundleID#-
#symbolicName#-#name#">
    ...
</camelContext>
```

Ambiguous names

Because the customised naming pattern overrides the default naming strategy, it is possible to define ambiguous JMX MBean names using this approach. For example:

```
<camelContext id="foo" managementNamePattern="Same0ldSame0ld"> ...
</camelContext>
...
<camelContext id="bar" managementNamePattern="Same0ldSame0ld"> ...
</camelContext>
```

In this case, Apache Camel would fail on start-up and report an *MBean already exists* exception. You should, therefore, take extra care to ensure that you do not define ambiguous name patterns.

2.13. PERFORMANCE AND OPTIMIZATION

Avoid unnecessary message copying

You can avoid making an unnecessary copy of the original message, by setting the **allowUseOriginalMessage** option to **false** on the **CamelContext** object. For example, in Blueprint XML you can set this option as follows:

```
<camelContext xmlns="http://camel.apache.org/schema/blueprint"
              allowUseOriginalMessage="false">
    ...
</camelContext>
```

You can safely set **allowUseOriginalMessage** to **false**, if the following conditions are satisfied:

- You do not set **useOriginalMessage=true** on any of the error handlers or on the **onException** element.
- You do not use the **getOriginalMessage** method anywhere in your Java application code.

CHAPTER 3. INTRODUCING ENTERPRISE INTEGRATION PATTERNS

Abstract

The Apache Camel's *Enterprise Integration Patterns* are inspired by a book of the same name written by Gregor Hohpe and Bobby Woolf. The patterns described by these authors provide an excellent toolbox for developing enterprise integration projects. In addition to providing a common language for discussing integration architectures, many of the patterns can be implemented directly using Apache Camel's programming interfaces and XML configuration.

3.1. OVERVIEW OF THE PATTERNS



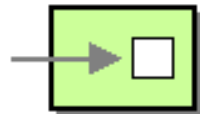

Enterprise Integration Patterns book

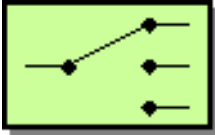
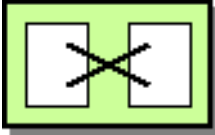
Apache Camel supports most of the patterns from the book, [Enterprise Integration Patterns](#) by Gregor Hohpe and Bobby Woolf.

Messaging systems

The messaging systems patterns, shown in [Table 3.1, "Messaging Systems"](#), introduce the fundamental concepts and components that make up a messaging system.

Table 3.1. Messaging Systems



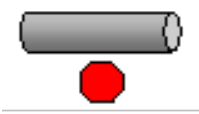

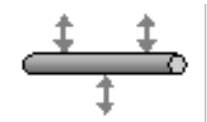
Icon	Name	Use Case
	Message	How can two applications connected by a message channel exchange a piece of information?
	Message Channel	How does one application communicate with another application using messaging?
	Message Endpoint	How does an application connect to a messaging channel to send and receive messages?
	Pipes and Filters	How can we perform complex processing on a message while still maintaining independence and flexibility?

Icon	Name	Use Case
	Message Router	How can you decouple individual processing steps so that messages can be passed to different filters depending on a set of defined conditions?
	Message Translator	How do systems using different data formats communicate with each other using messaging?

Messaging channels

A messaging channel is the basic component used for connecting the participants in a messaging system. The patterns in [Table 3.2, “Messaging Channels”](#) describe the different kinds of messaging channels available.

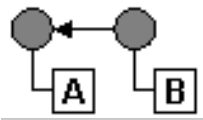

Table 3.2. Messaging Channels

Icon	Name	Use Case
	Point to Point Channel	How can the caller be sure that exactly one receiver will receive the document or will perform the call?
	Publish Subscribe Channel	How can the sender broadcast an event to all interested receivers?
	Dead Letter Channel	What will the messaging system do with a message it cannot deliver?
	Guaranteed Delivery	How does the sender make sure that a message will be delivered, even if the messaging system fails?
	Message Bus	What is an architecture that enables separate, decoupled applications to work together, such that one or more of the applications can be added or removed without affecting the others?

Message construction

The message construction patterns, shown in [Table 3.3, “Message Construction”](#), describe the various forms and functions of the messages that pass through the system.

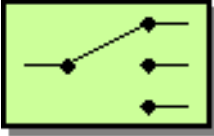
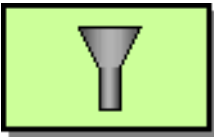
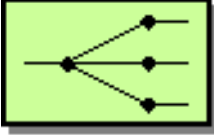
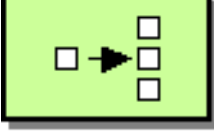
Table 3.3. Message Construction

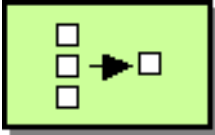
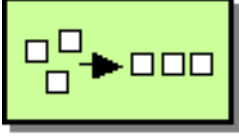
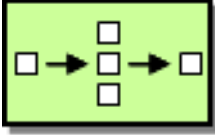
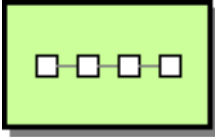
Icon	Name	Use Case
	Correlation Identifier	How does a requestor identify the request that generated the received reply?
	Return Address	How does a replier know where to send the reply?

Message routing

The message routing patterns, shown in [Table 3.4, “Message Routing”](#), describe various ways of linking message channels together, including various algorithms that can be applied to the message stream (without modifying the body of the message).

Table 3.4. Message Routing

Icon	Name	Use Case
	Content Based Router	How do we handle a situation where the implementation of a single logical function (e.g., inventory check) is spread across multiple physical systems?
	Message Filter	How does a component avoid receiving uninteresting messages?
	Recipient List	How do we route a message to a list of dynamically specified recipients?
	Splitter	How can we process a message if it contains multiple elements, each of which might have to be processed in a different way?

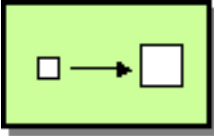
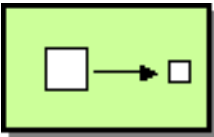
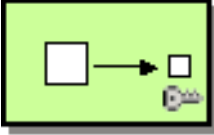

Icon	Name	Use Case
	Aggregator	How do we combine the results of individual, but related messages so that they can be processed as a whole?
	Resequencer	How can we get a stream of related, but out-of-sequence, messages back into the correct order?
	Composed Message Processor	How can you maintain the overall message flow when processing a message consisting of multiple elements, each of which may require different processing?
	Scatter-Gather	How do you maintain the overall message flow when a message needs to be sent to multiple recipients, each of which may send a reply?
	Routing Slip	How do we route a message consecutively through a series of processing steps when the sequence of steps is not known at design-time, and might vary for each message?
	Throttler	How can I throttle messages to ensure that a specific endpoint does not get overloaded, or that we don't exceed an agreed SLA with some external service?
	Delayer	How can I delay the sending of a message?
	Load Balancer	How can I balance load across a number of endpoints?
	Multicast	How can I route a message to a number of endpoints at the same time?
	Loop	How can I repeat processing a message in a loop?

Icon	Name	Use Case
	Sampling	How can I sample one message out of many in a given period to avoid downstream route does not get overloaded?

Message transformation

The message transformation patterns, shown in [Table 3.5, “Message Transformation”](#), describe how to modify the contents of messages for various purposes.

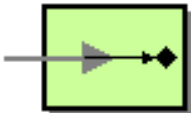
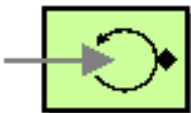
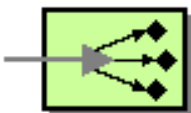
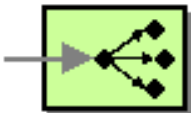
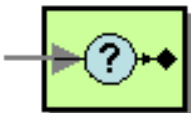
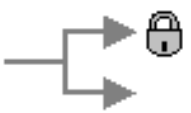
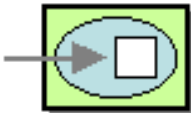
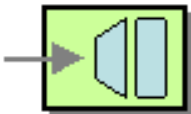
Table 3.5. Message Transformation

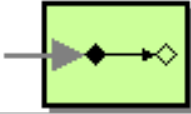
Icon	Name	Use Case
	Content Enricher	How do we communicate with another system if the message originator does not have all the required data items available?
	Content Filter	How do you simplify dealing with a large message, when you are interested in only a few data items?
	Claim Check	How can we reduce the data volume of message sent across the system without sacrificing information content?
	Normalizer	How do you process messages that are semantically equivalent, but arrive in a different format?
	Sort	How can I sort the body of a message?

Messaging endpoints

A messaging endpoint denotes the point of contact between a messaging channel and an application. The messaging endpoint patterns, shown in [Table 3.6, “Messaging Endpoints”](#), describe various features and qualities of service that can be configured on an endpoint.

Table 3.6. Messaging Endpoints

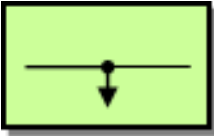
Icon	Name	Use Case
	Messaging Mapper	How do you move data between domain objects and the messaging infrastructure while keeping the two independent of each other?
	Event Driven Consumer	How can an application automatically consume messages as they become available?
	Polling Consumer	How can an application consume a message when the application is ready?
	Competing Consumers	How can a messaging client process multiple messages concurrently?
	Message Dispatcher	How can multiple consumers on a single channel coordinate their message processing?
	Selective Consumer	How can a message consumer select which messages it wants to receive?
	Durable Subscriber	How can a subscriber avoid missing messages when it's not listening for them?
	Idempotent Consumer	How can a message receiver deal with duplicate messages?
	Transactional Client	How can a client control its transactions with the messaging system?
	Messaging Gateway	How do you encapsulate access to the messaging system from the rest of the application?

Icon	Name	Use Case
	Service Activator	How can an application design a service to be invoked both via various messaging technologies and via non-messaging techniques?

System management

The system management patterns, shown in [Table 3.7, “System Management”](#), describe how to monitor, test, and administer a messaging system.

Table 3.7. System Management

Icon	Name	Use Case
	Wire Tap	How do you inspect messages that travel on a point-to-point channel?

CHAPTER 4. MESSAGING SYSTEMS

Abstract


This chapter introduces the fundamental building blocks of a messaging system, such as endpoints, messaging channels, and message routers.

4.1. MESSAGE

Overview

A *message* is the smallest unit for transmitting data in a messaging system (represented by the grey dot in the figure below). The message itself might have some internal structure—for example, a message containing multiple parts—which is represented by geometrical figures attached to the grey dot in [Figure 4.1, “Message Pattern”](#).

Figure 4.1. Message Pattern



Types of message

Apache Camel defines the following distinct message types:

- *In* message — A message that travels through a route from a consumer endpoint to a producer endpoint (typically, initiating a message exchange).
- *Out* message — A message that travels through a route from a producer endpoint back to a consumer endpoint (usually, in response to an *In* message).

All of these message types are represented internally by the `org.apache.camel.Message` interface.

Message structure

By default, Apache Camel applies the following structure to all message types:

- *Headers* — Contains metadata or header data extracted from the message.
- *Body* — Usually contains the entire message in its original form.
- *Attachments* — Message attachments (required for integrating with certain messaging systems, such as [JBI](#)).

It is important to remember that this division into headers, body, and attachments is an abstract model of the message. Apache Camel supports many different components, that generate a wide variety of message formats. Ultimately, it is the underlying component implementation that decides what gets placed into the headers and body of a message.

Correlating messages

Internally, Apache Camel remembers the message IDs, which are used to correlate individual messages. In practice, however, the most important way that Apache Camel correlates messages is through *exchange* objects.

Exchange objects

An exchange object is an entity that encapsulates related messages, where the collection of related messages is referred to as a *message exchange* and the rules governing the sequence of messages are referred to as an *exchange pattern*. For example, two common exchange patterns are: one-way event messages (consisting of an *In* message), and request-reply exchanges (consisting of an *In* message, followed by an *Out* message).

Accessing messages

When defining a routing rule in the Java DSL, you can access the headers and body of a message using the following DSL builder methods:

- **header(String name), body()** — Returns the named header and the body of the current *In* message.
- **outBody()** — Returns the body of the current *Out* message.

For example, to populate the *In* message's **username** header, you can use the following Java DSL route:


```
from(SourceURL).setHeader("username", "John.Doe").to(TargetURL);
```

4.2. MESSAGE CHANNEL

Overview

A *message channel* is a logical channel in a messaging system. That is, sending messages to different message channels provides an elementary way of sorting messages into different message types. Message queues and message topics are examples of message channels. You should remember that a logical channel is *not* the same as a physical channel. There can be several different ways of physically realizing a logical channel.

In Apache Camel, a message channel is represented by an endpoint URI of a message-oriented component as shown in [Figure 4.2, “Message Channel Pattern”](#).

Figure 4.2. Message Channel Pattern

Message-oriented components

The following message-oriented components in Apache Camel support the notion of a message channel:

- [ActiveMQ](#)
- [JMS](#)
- [AMQP](#)

ActiveMQ

In ActiveMQ, message channels are represented by *queues* or *topics*. The endpoint URI for a specific queue, *QueueName*, has the following format:

```
activemq:QueueName
```

The endpoint URI for a specific topic, *TopicName*, has the following format:

```
activemq:topic:TopicName
```

For example, to send messages to the queue, **Foo.Bar**, use the following endpoint URI:

```
activemq:Foo.Bar
```

See for more details and instructions on setting up the ActiveMQ component.

JMS

The Java Messaging Service (JMS) is a generic wrapper layer that is used to access many different kinds of message systems (for example, you can use it to wrap ActiveMQ, MQSeries, Tibco, BEA, Sonic, and others). In JMS, message channels are represented by queues, or topics. The endpoint URI for a specific queue, *QueueName*, has the following format:

```
jms:QueueName
```

The endpoint URI for a specific topic, *TopicName*, has the following format:

```
jms:topic:TopicName
```

See for more details and instructions on setting up the JMS component.

AMQP

In AMQP, message channels are represented by queues, or topics. The endpoint URI for a specific queue, *QueueName*, has the following format:

```
amqp:QueueName
```

The endpoint URI for a specific topic, *TopicName*, has the following format:

```
amqp:topic:TopicName
```


See for more details and instructions on setting up the AMQP component.

4.3. MESSAGE ENDPOINT

Overview

A *message endpoint* is the interface between an application and a messaging system. As shown in [Figure 4.3, “Message Endpoint Pattern”](#), you can have a sender endpoint, sometimes called a proxy or a service consumer, which is responsible for sending *In* messages, and a receiver endpoint, sometimes called an endpoint or a service, which is responsible for receiving *In* messages.

Figure 4.3. Message Endpoint Pattern



Types of endpoint

Apache Camel defines two basic types of endpoint:

- *Consumer endpoint* — Appears at the start of a Apache Camel route and reads *n* messages from an incoming channel (equivalent to a *receiver* endpoint).
- *Producer endpoint* — Appears at the end of a Apache Camel route and writes *n* messages to an outgoing channel (equivalent to a *sender* endpoint). It is possible to define a route with multiple producer endpoints.

Endpoint URIs

In Apache Camel, an endpoint is represented by an *endpoint URI*, which typically encapsulates the following kinds of data:

- *Endpoint URI for a consumer endpoint*— Advertises a specific location (for example, to expose a service to which senders can connect). Alternatively, the URI can specify a message source, such as a message queue. The endpoint URI can include settings to configure the endpoint.
- *Endpoint URI for a producer endpoint*— Contains details of where to send messages and includes the settings to configure the endpoint. In some cases, the URI specifies the location of a remote receiver endpoint; in other cases, the destination can have an abstract form, such as a queue name.

An endpoint URI in Apache Camel has the following general form:

```
ComponentPrefix:ComponentSpecificURI
```

Where *ComponentPrefix* is a URI prefix that identifies a particular Apache Camel component (see for details of all the supported components). The remaining part of the URI, *ComponentSpecificURI*, has a syntax defined by the particular component. For example, to connect to the JMS queue, **Foo.Bar**, you can define an endpoint URI like the following:

```
jms:Foo.Bar
```

To define a route that connects the consumer endpoint, **file://local/router/messages/foo**, directly to the producer endpoint, **jms:Foo.Bar**, you can use the following Java DSL fragment:

```
from("file://local/router/messages/foo").to("jms:Foo.Bar");
```


Alternatively, you can define the same route in XML, as follows:

```
<camelContext id="CamelContextID"
  xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="file://local/router/messages/foo"/>
    <to uri="jms:Foo.Bar"/>
  </route>
</camelContext>
```

4.4. PIPES AND FILTERS


Overview

The *pipes and filters* pattern, shown in [Figure 4.4, “Pipes and Filters Pattern”](#), describes a way of constructing a route by creating a chain of filters, where the output of one filter is fed into the input of the next filter in the pipeline (analogous to the UNIX **pipe** command). The advantage of the pipeline approach is that it enables you to compose services (some of which can be external to the Apache Camel application) to create more complex forms of message processing.

Figure 4.4. Pipes and Filters Pattern

Pipeline for the InOut exchange pattern

Normally, all of the endpoints in a pipeline have an input (*In* message) and an output (*Out* message), which implies that they are compatible with the *InOut* message exchange pattern. A typical message flow through an *InOut* pipeline is shown in [Figure 4.5, “Pipeline for InOut Exchanges”](#).

Figure 4.5. Pipeline for InOut Exchanges

Where the pipeline connects the output of each endpoint to the input of the next one. The *Out* message from the final endpoint gets sent back to the original caller. You can define a route for this pipeline, as follows:

```
from("jms:RawOrders").pipeline("cx:bean:decrypt",
"cx:bean:authenticate", "cx:bean:dedup", "jms:CleanOrders");
```

The same route can be configured in XML, as follows:

```
<camelContext id="buildPipeline"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="jms:RawOrders"/>
    <to uri="cx:bean:decrypt"/>
    <to uri="cx:bean:authenticate"/>
    <to uri="cx:bean:dedup"/>
    <to uri="jms:CleanOrders"/>
  </route>
</camelContext>
```


There is no dedicated pipeline element in XML. The preceding combination of **from** and **to** elements is semantically equivalent to a pipeline. See [the section called “Comparison of pipeline\(\) and to\(\) DSL commands”](#).

Pipeline for the InOnly and RobustInOnly exchange patterns

When there are no *Out* messages available from the endpoints in the pipeline (as is the case for the **InOnly** and **RobustInOnly** exchange patterns), a pipeline cannot be connected in the normal way. In this special case, the pipeline is constructed by passing a copy of the

original *In* message to each of the endpoints in the pipeline, as shown in [Figure 4.6, “Pipeline for InOnly Exchanges”](#). This type of pipeline is equivalent to a recipient list with fixed destinations (see [Section 7.3, “Recipient List”](#)).

Figure 4.6. Pipeline for InOnly Exchanges



The route for this pipeline is defined using the same syntax as an *InOut* pipeline (either in Java DSL or in XML).

Comparison of `pipeline()` and `to()` DSL commands

In the Java DSL, you can define a pipeline route using either of the following syntaxes:

- *Using the `pipeline()` processor command* — Use the pipeline processor to construct a pipeline route as follows:

```
from(SourceURI).pipeline(FilterA, FilterB, TargetURI);
```

- *Using the `to()` command* — Use the `to()` command to construct a pipeline route as follows:

```
from(SourceURI).to(FilterA, FilterB, TargetURI);
```

Alternatively, you can use the equivalent syntax:


```
from(SourceURI).to(FilterA).to(FilterB).to(TargetURI);
```

Exercise caution when using the `to()` command syntax, because it *is not* always equivalent to a pipeline processor. In Java DSL, the meaning of `to()` can be modified by the preceding command in the route. For example, when the `multicast()` command precedes the `to()` command, it binds the listed endpoints into a multicast pattern, instead of a pipeline pattern (see [Section 7.11, “Multicast”](#)).

4.5. MESSAGE ROUTER

Overview

A *message router*, shown in [Figure 4.7, “Message Router Pattern”](#), is a type of filter that consumes messages from a single consumer endpoint and redirects them to the appropriate target endpoint, based on a particular decision criterion. A message router is concerned only with redirecting messages; it does not modify the message content.

Figure 4.7. Message Router Pattern

A message router can easily be implemented in Apache Camel using the **choice()** processor, where each of the alternative target endpoints can be selected using a **when()** subclause (for details of the choice processor, see [Section 1.5, "Processors"](#)).

Java DSL example

The following Java DSL example shows how to route messages to three alternative destinations (either **seda:a**, **seda:b**, or **seda:c**) depending on the contents of the **foo** header:

```
from("seda:a").choice()
    .when(header("foo").isEqualTo("bar")).to("seda:b")
    .when(header("foo").isEqualTo("cheese")).to("seda:c")
    .otherwise().to("seda:d");
```

XML configuration example

The following example shows how to configure the same route in XML:

```
<camelContext id="buildSimpleRouteWithChoice"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <choice>
      <when>
        <xpath>$foo = 'bar'</xpath>
        <to uri="seda:b"/>
      </when>
      <when>
        <xpath>$foo = 'cheese'</xpath>
        <to uri="seda:c"/>
      </when>
      <otherwise>
        <to uri="seda:d"/>
      </otherwise>
    </choice>
  </route>
</camelContext>
```

Choice without otherwise


If you use `choice()` without an `otherwise()` clause, any unmatched exchanges are dropped by default.

4.6. MESSAGE TRANSLATOR

Overview

The *message translator* pattern, shown in [Figure 4.8, “Message Translator Pattern”](#) describes a component that modifies the contents of a message, translating it to a different format. You can use Apache Camel's bean integration feature to perform the message translation.

Figure 4.8. Message Translator Pattern



Bean integration

You can transform a message using bean integration, which enables you to call a method on any registered bean. For example, to call the method, `myMethodName()`, on the bean with ID, `myTransformerBean`:

```
from("activemq:SomeQueue")
  .beanRef("myTransformerBean", "myMethodName")
  .to("mqseries:AnotherQueue");
```

Where the `myTransformerBean` bean is defined in either a Spring XML file or in JNDI. If, you omit the method name parameter from `beanRef()`, the bean integration will try to deduce the method name to invoke by examining the message exchange.

You can also add your own explicit **Processor** instance to perform the transformation, as follows:

```
from("direct:start").process(new Processor() {
    public void process(Exchange exchange) {
        Message in = exchange.getIn();
        in.setBody(in.getBody(String.class) + " World!");
    }
}).to("mock:result");
```

Or, you can use the DSL to explicitly configure the transformation, as follows:

```
from("direct:start").setBody(body().append(" World!")).to("mock:result");
```

You can also use *templating* to consume a message from one destination, transform it with something like [Velocity](#) or XQuery and then send it on to another destination. For example, using the *InOnly* exchange pattern (one-way messaging) :

```
from("activemq:My.Queue").
  to("velocity:com/acme/MyResponse.vm").
  to("activemq:Another.Queue");
```

If you want to use *InOut* (request-reply) semantics to process requests on the **My.Queue** queue on [ActiveMQ](#) with a template generated response, then you could use a route like the following to send responses back to the **JMSReplyTo** destination:

```
from("activemq:My.Queue").
  to("velocity:com/acme/MyResponse.vm");
```

CHAPTER 5. MESSAGING CHANNELS

Abstract


Messaging channels provide the plumbing for a messaging application. This chapter describes the different kinds of messaging channels available in a messaging system, and the roles that they play.

5.1. POINT-TO-POINT CHANNEL

Overview

A *point-to-point channel*, shown in [Figure 5.1, “Point to Point Channel Pattern”](#) is a [message channel](#) that guarantees that only one receiver consumes any given message. This is in contrast with a [publish-subscribe channel](#), which allows multiple receivers to consume the same message. In particular, with a point-to-point channel, it is possible for multiple receivers to subscribe to the same channel. If more than one receiver competes to consume a message, it is up to the message channel to ensure that only one receiver actually consumes the message.

Figure 5.1. Point to Point Channel Pattern



Components that support point-to-point channel

The following Apache Camel components support the point-to-point channel pattern:

- [JMS](#)
- [ActiveMQ](#)
- [SEDA](#)
- [JPA](#)
- [XMPP](#)

JMS

In JMS, a point-to-point channel is represented by a *queue*. For example, you can specify the endpoint URI for a JMS queue called **Foo.Bar** as follows:

```
jms:queue:Foo.Bar
```

The qualifier, **queue:**, is optional, because the JMS component creates a queue endpoint by default. Therefore, you can also specify the following equivalent endpoint URI:

■

```
jms:Foo.Bar
```

See for more details.

ActiveMQ

In ActiveMQ, a point-to-point channel is represented by a queue. For example, you can specify the endpoint URI for an ActiveMQ queue called **Foo.Bar** as follows:

```
activemq:queue:Foo.Bar
```

See for more details.

SEDA

The Apache Camel Staged Event-Driven Architecture (SEDA) component is implemented using a blocking queue. Use the SEDA component if you want to create a lightweight point-to-point channel that is *internal* to the Apache Camel application. For example, you can specify an endpoint URI for a SEDA queue called **SedaQueue** as follows:

```
seda:SedaQueue
```

JPA

The Java Persistence API (JPA) component is an EJB 3 persistence standard that is used to write entity beans out to a database. See for more details.


XMPP

The XMPP (Jabber) component supports the point-to-point channel pattern when it is used in the person-to-person mode of communication. See for more details.

5.2. PUBLISH-SUBSCRIBE CHANNEL

Overview

A *publish-subscribe channel*, shown in [Figure 5.2, “Publish Subscribe Channel Pattern”](#), is a [message channel](#) that enables multiple subscribers to consume any given message. This is in contrast with a [point-to-point channel](#). Publish-subscribe channels are frequently used as a means of broadcasting events or notifications to multiple subscribers.

Figure 5.2. Publish Subscribe Channel Pattern

Components that support publish-subscribe channel

The following Apache Camel components support the publish-subscribe channel pattern:

- [JMS](#)
- [ActiveMQ](#)
- [XMPP](#)
- [SEDA](#) for working with SEDA in the same [CamelContext](#) which can work in pub-sub, but allowing multiple consumers.
- [VM](#) as SEDA, but for use within the same JVM.

JMS

In JMS, a publish-subscribe channel is represented by a *topic*. For example, you can specify the endpoint URI for a JMS topic called **StockQuotes** as follows:

```
jms:topic:StockQuotes
```

See for more details.

ActiveMQ

In ActiveMQ, a publish-subscribe channel is represented by a topic. For example, you can specify the endpoint URI for an ActiveMQ topic called **StockQuotes**, as follows:

```
activemq:topic:StockQuotes
```

See for more details.

XMPP

The XMPP (Jabber) component supports the publish-subscribe channel pattern when it is used in the group communication mode. See for more details.


Static subscription lists

If you prefer, you can also implement publish-subscribe logic within the Apache Camel application itself. A simple approach is to define a *static subscription list*, where the target endpoints are all explicitly listed at the end of the route. However, this approach is not as flexible as a JMS or ActiveMQ topic.

Java DSL example

The following Java DSL example shows how to simulate a publish-subscribe channel with a single publisher, **seda:a**, and three subscribers, **seda:b**, **seda:c**, and **seda:d**:

```
from("seda:a").to("seda:b", "seda:c", "seda:d");
```



NOTE

This only works for the *InOnly* message exchange pattern.

XML configuration example

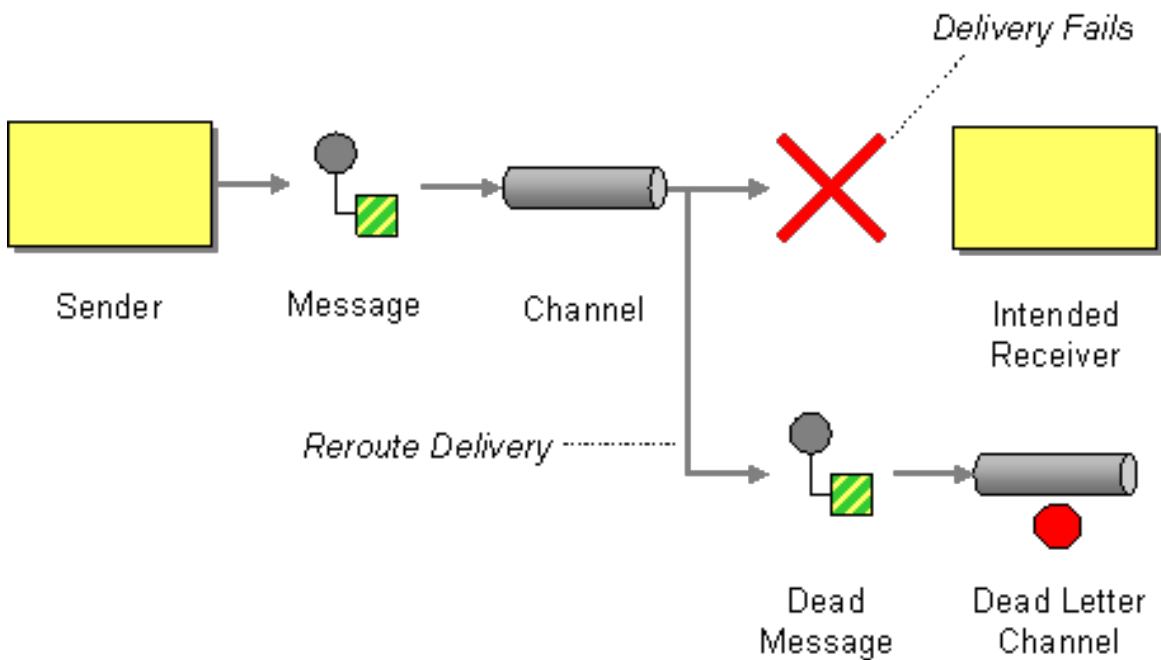
The following example shows how to configure the same route in XML:

```
<camelContext id="buildStaticRecipientList"
  xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <to uri="seda:b"/>
    <to uri="seda:c"/>
    <to uri="seda:d"/>
  </route>
</camelContext>
```

5.3. DEAD LETTER CHANNEL

Overview

The *dead letter channel* pattern, shown in [Figure 5.3, “Dead Letter Channel Pattern”](#), describes the actions to take when the messaging system fails to deliver a message to the intended recipient. This includes such features as retrying delivery and, if delivery ultimately fails, sending the message to a dead letter channel, which archives the undelivered messages.

Figure 5.3. Dead Letter Channel Pattern

Creating a dead letter channel in Java DSL

The following example shows how to create a dead letter channel using Java DSL:

```
errorHandler(deadLetterChannel("seda:errors"));
from("seda:a").to("seda:b");
```

Where the `errorHandler()` method is a Java DSL interceptor, which implies that *all* of the routes defined in the current route builder are affected by this setting. The `deadLetterChannel()` method is a Java DSL command that creates a new dead letter channel with the specified destination endpoint, `seda:errors`.

The `errorHandler()` interceptor provides a catch-all mechanism for handling *all* error types. If you want to apply a more fine-grained approach to exception handling, you can use the `onException` clauses instead (see [the section called "onException clause"](#)).

XML DSL example

You can define a dead letter channel in the XML DSL, as follows:

```
<route errorHandlerRef="myDeadLetterErrorHandler">
  ...
</route>

<bean id="myDeadLetterErrorHandler"
class="org.apache.camel.builder.DeadLetterChannelBuilder">
  <property name="deadLetterUri" value="jms:queue:dead"/>
  <property name="redeliveryPolicy" ref="myRedeliveryPolicyConfig"/>
</bean>

<bean id="myRedeliveryPolicyConfig"
class="org.apache.camel.processor.RedeliveryPolicy">
```

```

    <property name="maximumRedeliveries" value="3"/>
    <property name="redeliveryDelay" value="5000"/>
</bean>

```

Redelivery policy

Normally, you do not send a message straight to the dead letter channel, if a delivery attempt fails. Instead, you re-attempt delivery up to some maximum limit, and after all redelivery attempts fail you would send the message to the dead letter channel. To customize message redelivery, you can configure the dead letter channel to have a *redelivery policy*. For example, to specify a maximum of two redelivery attempts, and to apply an exponential backoff algorithm to the time delay between delivery attempts, you can configure the dead letter channel as follows:

```

errorHandler(deadLetterChannel("seda:errors").maximumRedeliveries(2).useExponentialBackOff());
from("seda:a").to("seda:b");

```

Where you set the redelivery options on the dead letter channel by invoking the relevant methods in a chain (each method in the chain returns a reference to the current **RedeliveryPolicy** object). [Table 5.1, “Redelivery Policy Settings”](#) summarizes the methods that you can use to set redelivery policies.

Table 5.1. Redelivery Policy Settings

Method Signature	Default	Description
backOffMultiplier(double multiplier)	2	If exponential backoff is enabled, let m be the backoff multiplier and let d be the initial delay. The sequence of redelivery attempts are then timed as follows: <pre>d, m*d, m*m*d, m*m*m*d, ...</pre>
collisionAvoidancePercent(double collisionAvoidancePercent)	15	If collision avoidance is enabled, let p be the collision avoidance percent. The collision avoidance policy then tweaks the next delay by a random amount, up to plus/minus p% of its current value.
delayPattern(String delayPattern)	None	Apache Camel 2.0:

Method Signature	Default	Description
<code>disableRedelivery()</code>	<code>true</code>	Apache Camel 2.0: Disables the redelivery feature. To enable redelivery, set <code>maximumRedeliveries()</code> to a positive integer value.
<code>handled(boolean handled)</code>	<code>true</code>	Apache Camel 2.0: If <code>true</code> , the current exception is cleared when the message is moved to the dead letter channel; if <code>false</code> , the exception is propagated back to the client.
<code>initialRedeliveryDelay(long initialRedeliveryDelay)</code>	<code>1000</code>	Specifies the delay (in milliseconds) before attempting the first redelivery.
<code>logStackTrace(boolean logStackTrace)</code>	<code>false</code>	Apache Camel 2.0: If <code>true</code> , the JVM stack trace is included in the error logs.
<code>maximumRedeliveries(int maximumRedeliveries)</code>	<code>0</code>	Apache Camel 2.0: Maximum number of delivery attempts.
<code>maximumRedeliveryDelay(long maxDelay)</code>	<code>60000</code>	Apache Camel 2.0: When using an exponential backoff strategy (see <code>useExponentialBackOff()</code>), it is theoretically possible for the redelivery delay to increase without limit. This property imposes an upper limit on the redelivery delay (in milliseconds)
<code>onRedelivery(Processor processor)</code>	<code>None</code>	Apache Camel 2.0: Configures a processor that gets called before every redelivery attempt.
<code>redeliveryDelay(long int)</code>	<code>0</code>	Apache Camel 2.0: Specifies the delay (in milliseconds) between redelivery attempts.

Method Signature	Default	Description
<code>retriesExhaustedLogLevel(LoggingLevel logLevel)</code>	<code>LoggingLevel.ERROR</code>	Apache Camel 2.0: Specifies the logging level at which to log delivery failure (specified as an <code>org.apache.camel.LoggingLevel</code> constant).
<code>retryAttemptedLogLevel(LoggingLevel logLevel)</code>	<code>LoggingLevel.DEBUG</code>	Apache Camel 2.0: Specifies the logging level at which to redelivery attempts (specified as an <code>org.apache.camel.LoggingLevel</code> constant).
<code>useCollisionAvoidance()</code>	<code>false</code>	Enables collision avoidance, which adds some randomization to the backoff timings to reduce contention probability.
<code>useOriginalMessage()</code>	<code>false</code>	Apache Camel 2.0: If this feature is enabled, the message sent to the dead letter channel is a copy of the <i>original</i> message exchange, as it existed at the beginning of the route (in the <code>from()</code> node).
<code>useExponentialBackOff()</code>	<code>false</code>	Enables exponential backoff.
<code>allowRedeliveryWhileStopping()</code>	<code>true</code>	Controls whether redelivery is attempted during graceful shutdown or while a route is stopping. A delivery that is already in progress when stopping is initiated will <i>not</i> be interrupted.

Redelivery headers

If Apache Camel attempts to redeliver a message, it automatically sets the headers described in [Table 5.2, “Dead Letter Redelivery Headers”](#) on the *In* message.

Table 5.2. Dead Letter Redelivery Headers

Header Name	Type	Description
-------------	------	-------------

Header Name	Type	Description
CamelRedeliveryCounter	Integer	Apache Camel 2.0: Counts the number of unsuccessful delivery attempts. This value is also set in Exchange.REDELIVERY_COUNTER .
CamelRedelivered	Boolean	Apache Camel 2.0: True, if one or more redelivery attempts have been made. This value is also set in Exchange.REDELIVERED .
CamelRedeliveryMaxCounter	Integer	Apache Camel 2.6: Holds the maximum redelivery setting (also set in the Exchange.REDELIVERY_MAX_COUNTER exchange property). This header is absent if you use retryWhile or have unlimited maximum redelivery configured.

Redelivery exchange properties

If Apache Camel attempts to redeliver a message, it automatically sets the exchange properties described in [Table 5.3, “Redelivery Exchange Properties”](#).

Table 5.3. Redelivery Exchange Properties

Exchange Property Name	Type	Description
Exchange.FAILURE_ROUTE_ID	String	Provides the route ID of the route that failed. The literal name of this property is CamelFailureRouteId .

Using the original message

Available as of Apache Camel 2.0 Because an exchange object is subject to modification as it passes through the route, the exchange that is current when an exception is raised is not necessarily the copy that you would want to store in the dead letter channel. In many cases, it is preferable to log the message that arrived at the start of the route, before it was subject to any kind of transformation by the route. For example, consider the following route:

```
from("jms:queue:order:input")
    .to("bean:validateOrder");
    .to("bean:transformOrder")
```

```
.to("bean:handleOrder");
```

The preceding route listens for incoming JMS messages and then processes the messages using the sequence of beans: **validateOrder**, **transformOrder**, and **handleOrder**. But when an error occurs, we do not know in which state the message is in. Did the error happen before the **transformOrder** bean or after? We can ensure that the original message from **jms:queue:order:input** is logged to the dead letter channel by enabling the **useOriginalMessage** option as follows:

```
// will use original body
errorHandler(deadLetterChannel("jms:queue:dead")
    .useOriginalMessage().maximumRedeliveries(5).redeliveryDelay(5000);
```

Redeliver delay pattern

Available as of Apache Camel 2.0 The **delayPattern** option is used to specify delays for particular ranges of the redelivery count. The delay pattern has the following syntax: **limit1:delay1;limit2:delay2;limit3:delay3;...**, where each *delayN* is applied to redeliveries in the range **limitN ≤ redeliveryCount < limitN+1**

For example, consider the pattern, **5:1000;10:5000;20:20000**, which defines three groups and results in the following redelivery delays:

- Attempt number 1..4 = 0 milliseconds (as the first group starts with 5).
- Attempt number 5..9 = 1000 milliseconds (the first group).
- Attempt number 10..19 = 5000 milliseconds (the second group).
- Attempt number 20.. = 20000 milliseconds (the last group).

You can start a group with limit 1 to define a starting delay. For example, **1:1000;5:5000** results in the following redelivery delays:

- Attempt number 1..4 = 1000 millis (the first group)
- Attempt number 5.. = 5000 millis (the last group)

There is no requirement that the next delay should be higher than the previous and you can use any delay value you like. For example, the delay pattern, **1:5000;3:1000**, starts with a 5 second delay and then reduces the delay to 1 second.

Which endpoint failed?

When Apache Camel routes messages, it updates an **Exchange** property that contains the *last* endpoint the **Exchange** was sent to. Hence, you can obtain the URI for the current exchange's most recent destination using the following code:

```
// Java
String lastEndpointUri = exchange.getProperty(Exchange.TO_ENDPOINT,
String.class);
```

Where **Exchange.TO_ENDPOINT** is a string constant equal to **CamelToEndpoint**. This property is updated whenever Camel sends a message to *any* endpoint.

If an error occurs during routing and the exchange is moved into the dead letter queue, Apache Camel will additionally set a property named **CamelFailureEndpoint**, which identifies the last destination the exchange was sent to before the error occurred. Hence, you can access the failure endpoint from within a dead letter queue using the following code:

```
// Java
String failedEndpointUri = exchange.getProperty(Exchange.FAILURE_ENDPOINT,
String.class);
```

Where **Exchange.FAILURE_ENDPOINT** is a string constant equal to **CamelFailureEndpoint**.

NOTE

These properties remain set in the current exchange, even if the failure occurs *after* the given destination endpoint has finished processing. For example, consider the following route:

```
from("activemq:queue:foo")
.to("http://someserver/somepath")
.beanRef("foo");
```

Now suppose that a failure happens in the **foo** bean. In this case the **Exchange.TO_ENDPOINT** property and the **Exchange.FAILURE_ENDPOINT** property still contain the value, <http://someserver/somepath>.

onRedelivery processor

When a dead letter channel is performing redeliveries, it is possible to configure a **Processor** that is executed *just before* every redelivery attempt. This can be used for situations where you need to alter the message before it is redelivered.

For example, the following dead letter channel is configured to call the **MyRedeliverProcessor** before redelivering exchanges:

```
// we configure our Dead Letter Channel to invoke
// MyRedeliveryProcessor before a redelivery is
// attempted. This allows us to alter the message before
errorHandler(deadLetterChannel("mock:error").maximumRedeliveries(5)
.onRedelivery(new MyRedeliverProcessor())
// setting delay to zero is just to make unit testing faster
.redeliveryDelay(0L));
```

Where the **MyRedeliverProcessor** process is implemented as follows:

```
// This is our processor that is executed before every redelivery attempt
// here we can do what we want in the java code, such as altering the
message
public class MyRedeliverProcessor implements Processor {

    public void process(Exchange exchange) throws Exception {
        // the message is being redelivered so we can alter it

        // we just append the redelivery counter to the body
```

```

        // you can of course do all kind of stuff instead
        String body = exchange.getIn().getBody(String.class);
        int count =
exchange.getIn().getHeader(Exchange.REDELIVERY_COUNTER, Integer.class);

        exchange.getIn().setBody(body + count);

        // the maximum redelivery was set to 5
        int max =
exchange.getIn().getHeader(Exchange.REDELIVERY_MAX_COUNTER,
Integer.class);
        assertEquals(5, max);
    }
}

```

Control redelivery during shutdown or stopping

If you stop a route or initiate graceful shutdown, the default behavior of the error handler is to continue attempting redelivery. Because this is typically not the desired behavior, you have the option of disabling redelivery during shutdown or stopping, by setting the **allowRedeliveryWhileStopping** option to **false**, as shown in the following example:

```

errorHandler(deadLetterChannel("jms:queue:dead")
    .allowRedeliveryWhileStopping(false)
    .maximumRedeliveries(20)
    .redeliveryDelay(1000)
    .retryAttemptedLogLevel(LoggingLevel.INFO));

```



NOTE

The **allowRedeliveryWhileStopping** option is **true** by default, for backwards compatibility reasons. During aggressive shutdown, however, redelivery is always suppressed, irrespective of this option setting (for example, after graceful shutdown has timed out).

onException clause

Instead of using the **errorHandler()** interceptor in your route builder, you can define a series of **onException()** clauses that define different redelivery policies and different dead letter channels for various exception types. For example, to define distinct behavior for each of the **NullPointerException**, **IOException**, and **Exception** types, you can define the following rules in your route builder using Java DSL:

```

onException(NullPointerException.class)
    .maximumRedeliveries(1)
    .setHeader("messageInfo", "Oh dear! An NPE.")
    .to("mock:npe_error");

onException(IOException.class)
    .initialRedeliveryDelay(5000L)
    .maximumRedeliveries(3)
    .backOffMultiplier(1.0)
    .useExponentialBackOff()
    .setHeader("messageInfo", "Oh dear! Some kind of I/O exception.")

```

```

        .to("mock:io_error");

onException(Exception.class)
    .initialRedeliveryDelay(1000L)
    .maximumRedeliveries(2)
    .setHeader("messageInfo", "Oh dear! An exception.")
    .to("mock:error");

from("seda:a").to("seda:b");

```

Where the redelivery options are specified by chaining the redelivery policy methods (as listed in [Table 5.1, “Redelivery Policy Settings”](#)), and you specify the dead letter channel's endpoint using the `to()` DSL command. You can also call other Java DSL commands in the `onException()` clauses. For example, the preceding example calls `setHeader()` to record some error details in a message header named, `messageInfo`.


In this example, the `NullPointerException` and the `IOException` exception types are configured specially. All other exception types are handled by the generic `Exception` exception interceptor. By default, Apache Camel applies the exception interceptor that most closely matches the thrown exception. If it fails to find an exact match, it tries to match the closest base type, and so on. Finally, if no other interceptor matches, the interceptor for the `Exception` type matches all remaining exceptions.

5.4. GUARANTEED DELIVERY

Overview

Guaranteed delivery means that once a message is placed into a message channel, the messaging system guarantees that the message will reach its destination, even if parts of the application should fail. In general, messaging systems implement the guaranteed delivery pattern, shown in [Figure 5.4, “Guaranteed Delivery Pattern”](#), by writing messages to persistent storage before attempting to deliver them to their destination.

Figure 5.4. Guaranteed Delivery Pattern



Components that support guaranteed delivery

The following Apache Camel components support the guaranteed delivery pattern:

- [JMS](#)
- [ActiveMQ](#)

- [ActiveMQ Journal](#)
- [File Component](#)

JMS

In JMS, the **deliveryPersistent** query option indicates whether or not persistent storage of messages is enabled. Usually it is unnecessary to set this option, because the default behavior is to enable persistent delivery. To configure all the details of guaranteed delivery, it is necessary to set configuration options on the JMS provider. These details vary, depending on what JMS provider you are using. For example, MQSeries, TibCo, BEA, Sonic, and others, all provide various qualities of service to support guaranteed delivery.

See for more details.

ActiveMQ

In ActiveMQ, message persistence is enabled by default. From version 5 onwards, ActiveMQ uses the AMQ message store as the default persistence mechanism. There are several different approaches you can use to enable message persistence in ActiveMQ.

The simplest option (different from [Figure 5.4, “Guaranteed Delivery Pattern”](#)) is to enable persistence in a central broker and then connect to that broker using a reliable protocol. After a message is sent to the central broker, delivery to consumers is guaranteed. For example, in the Apache Camel configuration file, **META-INF/spring/camel-context.xml**, you can configure the ActiveMQ component to connect to the central broker using the OpenWire/TCP protocol as follows:

```
<beans ... >
  ...
  <bean id="activemq"
class="org.apache.activemq.camel.component.ActiveMQComponent">
    <property name="brokerURL" value="tcp://somehost:61616"/>
  </bean>
  ...
</beans>
```

If you prefer to implement an architecture where messages are stored locally before being sent to a remote endpoint (similar to [Figure 5.4, “Guaranteed Delivery Pattern”](#)), you do this by instantiating an embedded broker in your Apache Camel application. A simple way to achieve this is to use the ActiveMQ Peer-to-Peer protocol, which implicitly creates an embedded broker to communicate with other peer endpoints. For example, in the **camel-context.xml** configuration file, you can configure the ActiveMQ component to connect to all of the peers in group, **GroupA**, as follows:

```
<beans ... >
  ...
  <bean id="activemq"
class="org.apache.activemq.camel.component.ActiveMQComponent">
    <property name="brokerURL" value="peer://GroupA/broker1"/>
  </bean>
  ...
</beans>
```

Where **broker1** is the broker name of the embedded broker (other peers in the group

should use different broker names). One limiting feature of the Peer-to-Peer protocol is that it relies on IP multicast to locate the other peers in its group. This makes it unsuitable for use in wide area networks (and in some local area networks that do not have IP multicast enabled).

A more flexible way to create an embedded broker in the ActiveMQ component is to exploit ActiveMQ's VM protocol, which connects to an embedded broker instance. If a broker of the required name does not already exist, the VM protocol automatically creates one. You can use this mechanism to create an embedded broker with custom configuration. For example:

```
<beans ... >
  ...
  <bean id="activemq"
class="org.apache.activemq.camel.component.ActiveMQComponent">
    <property name="brokerURL" value="vm://broker1?
brokerConfig=xbean:activemq.xml"/>
  </bean>
  ...
</beans>
```

Where `activemq.xml` is an ActiveMQ file which configures the embedded broker instance. Within the ActiveMQ configuration file, you can choose to enable one of the following persistence mechanisms:

- *AMQ persistence (the default)* — A fast and reliable message store that is native to ActiveMQ. For details, see [amqPersistenceAdapter](#) and [AMQ Message Store](#).
- *JDBC persistence* — Uses JDBC to store messages in any JDBC-compatible database. For details, see [jdbcPersistenceAdapter](#) and [ActiveMQ Persistence](#).
- *Journal persistence* — A fast persistence mechanism that stores messages in a rolling log file. For details, see [journalPersistenceAdapter](#) and [ActiveMQ Persistence](#).
- *Kaha persistence* — A persistence mechanism developed specifically for ActiveMQ. For details, see [kahaPersistenceAdapter](#) and [ActiveMQ Persistence](#).

See for more details.

ActiveMQ Journal


The ActiveMQ Journal component is optimized for a special use case where multiple, concurrent producers write messages to queues, but there is only one active consumer. Messages are stored in rolling log files and concurrent writes are aggregated to boost efficiency.

See for more details.

5.5. MESSAGE BUS

Overview

Message bus refers to a messaging architecture, shown in [Figure 5.5, “Message Bus Pattern”](#), that enables you to connect diverse applications running on diverse computing platforms. In effect, the Apache Camel and its components constitute a message bus.

Figure 5.5. Message Bus Pattern

The following features of the message bus pattern are reflected in Apache Camel:

- *Common communication infrastructure* — The router itself provides the core of the common communication infrastructure in Apache Camel. However, in contrast to some message bus architectures, Apache Camel provides a heterogeneous infrastructure: messages can be sent into the bus using a wide variety of different transports and using a wide variety of different message formats.
- *Adapters* — Where necessary, Apache Camel can translate message formats and propagate messages using different transports. In effect, Apache Camel is capable of behaving like an adapter, so that external applications can hook into the message bus without refactoring their messaging protocols.

In some cases, it is also possible to integrate an adapter directly into an external application. For example, if you develop an application using Apache CXF, where the service is implemented using JAX-WS and JAXB mappings, it is possible to bind a variety of different transports to the service. These transport bindings function as adapters.

CHAPTER 6. MESSAGE CONSTRUCTION

Abstract

The message construction patterns describe the various forms and functions of the messages that pass through the system.

6.1. CORRELATION IDENTIFIER

Overview

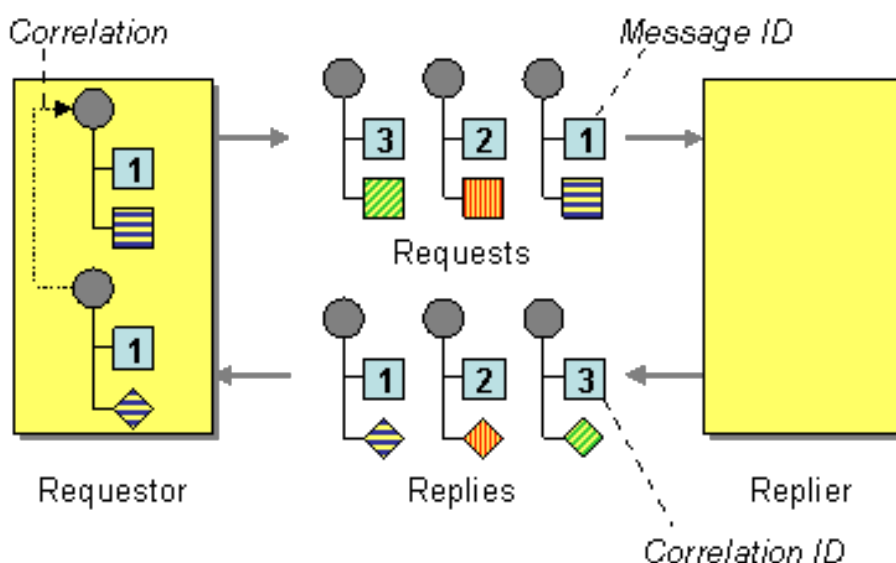
The *correlation identifier* pattern, shown in [Figure 6.1, “Correlation Identifier Pattern”](#), describes how to match reply messages with request messages, given that an asynchronous messaging system is used to implement a request-reply protocol. The essence of this idea is that request messages should be generated with a unique token, the *request ID*, that identifies the request message and reply messages should include a token, the *correlation ID*, that contains the matching request ID.

Apache Camel supports the Correlation Identifier from the EIP patterns by getting or setting a header on a Message.

When working with the [ActiveMQ](#) or [JMS](#) components, the correlation identifier header is called `JMSCorrelationID`. You can add your own correlation identifier to any message exchange to help correlate messages together in a single conversation (or business process). A correlation identifier is usually stored in a Apache Camel message header.

Some [EIP](#) patterns spin off a sub message and, in those cases, Apache Camel adds a correlation ID to the [Exchange](#) as a property with the key, `Exchange.CORRELATION_ID`, which links back to the source [Exchange](#). For example, the [Splitter](#), [Multicast](#), [Recipient List](#), and [Wire Tap](#) EIPs do this.

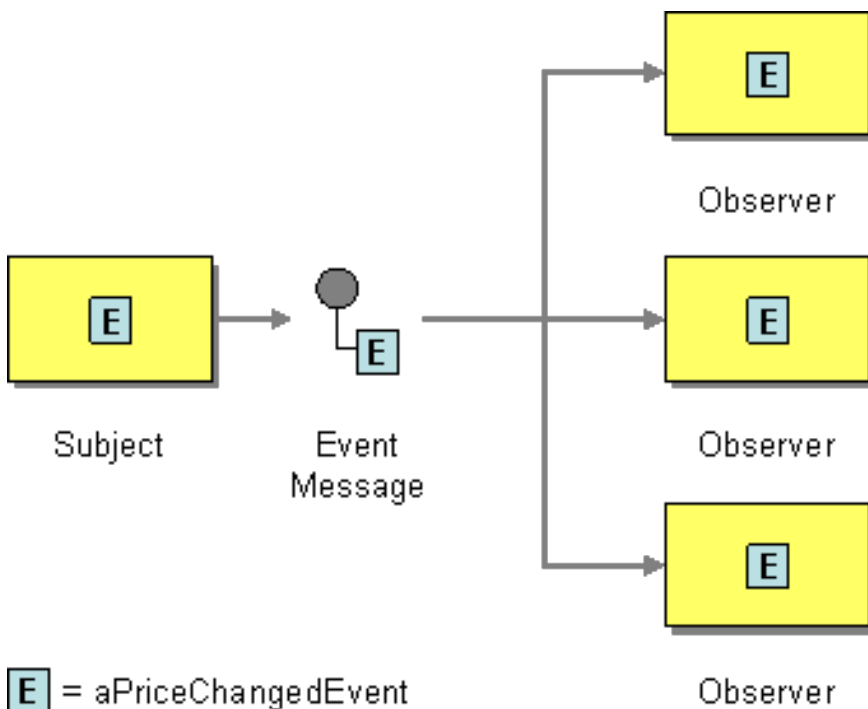
Figure 6.1. Correlation Identifier Pattern



6.2. EVENT MESSAGE

Event Message

Camel supports the [Event Message](#) from the [Introducing Enterprise Integration Patterns](#) by supporting the [Exchange Pattern](#) on a [Message](#) which can be set to **InOnly** to indicate a oneway event message. Camel Components then implement this pattern using the underlying transport or protocols.



The default behaviour of many Components is InOnly such as for [JMS](#), [File](#) or [SEDA](#)

Explicitly specifying InOnly

If you are using a component which defaults to InOut you can override the [Exchange Pattern](#) for an endpoint using the pattern property.

```
foo:bar?exchangePattern=InOnly
```

From 2.0 onwards on Camel you can specify the [Exchange Pattern](#) using the dsl.

Using the [Fluent Builders](#)

```
from("mq:someQueue").
  inOnly().
  bean(Foo.class);
```

or you can invoke an endpoint with an explicit pattern

```
from("mq:someQueue").
  inOnly("mq:anotherQueue");
```

Using the [Spring XML Extensions](#)

```
<route>
  <from uri="mq:someQueue"/>
```

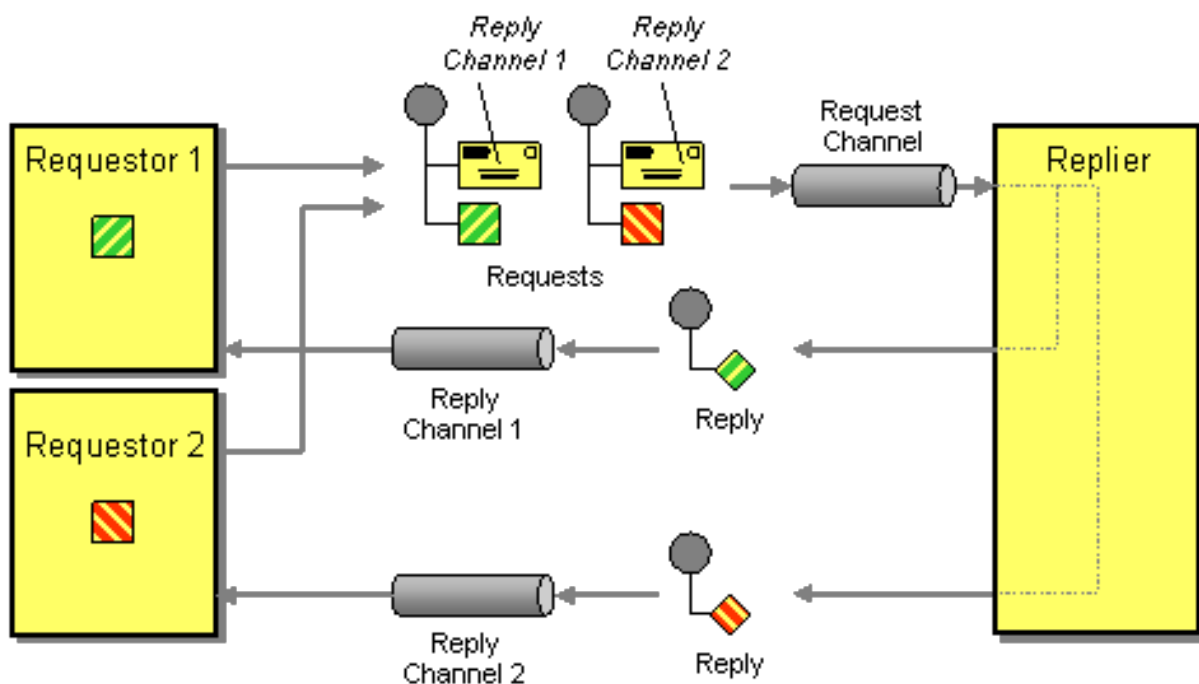
```
<inOnly uri="bean:foo"/>
</route>
```

```
<route>
  <from uri="mq:someQueue"/>
  <inOnly uri="mq:anotherQueue"/>
</route>
```

6.3. RETURN ADDRESS

Return Address

Apache Camel supports the [Return Address](#) from the [Introducing Enterprise Integration Patterns](#) using the `JMSReplyTo` header.



For example when using [JMS](#) with *InOut*, the component will by default be returned to the address given in `JMSReplyTo`.

Example

Requestor Code

```
getMockEndpoint("mock:bar").expectedBodiesReceived("Bye World");
template.sendBodyAndHeader("direct:start", "World", "JMSReplyTo",
"queue:bar");
```

Route Using the [Fluent Builders](#)

```
from("direct:start").to("activemq:queue:foo?preserveMessageQos=true");
from("activemq:queue:foo").transform(body().prepend("Bye "));
from("activemq:queue:bar?disableReplyTo=true").to("mock:bar");
```

Route Using the [Spring XML Extensions](#)

```
<route>
  <from uri="direct:start"/>
  <to uri="activemq:queue:foo?preserveMessageQos=true"/>
</route>

<route>
  <from uri="activemq:queue:foo"/>
  <transform>
    <simple>Bye ${in.body}</simple>
  </transform>
</route>

<route>
  <from uri="activemq:queue:bar?disableReplyTo=true"/>
  <to uri="mock:bar"/>
</route>
```

For a complete example of this pattern, see this [junit test case](#)

CHAPTER 7. MESSAGE ROUTING

Abstract


The message routing patterns describe various ways of linking message channels together. This includes various algorithms that can be applied to the message stream (without modifying the body of the message).

7.1. CONTENT-BASED ROUTER

Overview

A *content-based router*, shown in [Figure 7.1, “Content-Based Router Pattern”](#), enables you to route messages to the appropriate destination based on the message contents.

Figure 7.1. Content-Based Router Pattern



Java DSL example

The following example shows how to route a request from an input, **seda:a**, endpoint to either **seda:b**, **queue:c**, or **seda:d** depending on the evaluation of various predicate expressions:

```
RouteBuilder builder = new RouteBuilder() {
    public void configure() {
        from("seda:a").choice()
            .when(header("foo").isEqualTo("bar")).to("seda:b")
            .when(header("foo").isEqualTo("cheese")).to("seda:c")
            .otherwise().to("seda:d");
    }
};
```

XML configuration example

The following example shows how to configure the same route in XML:

```
<camelContext id="buildSimpleRouteWithChoice"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <choice>
      <when>
        <xpath>$foo = 'bar'</xpath>
```

```

        <to uri="seda:b"/>
    </when>
    <when>
        <xpath>$foo = 'cheese'</xpath>
        <to uri="seda:c"/>
    </when>
    <otherwise>
        <to uri="seda:d"/>
    </otherwise>
</choice>
</route>
</camelContext>


```

7.2. MESSAGE FILTER

Overview

A *message filter* is a processor that eliminates undesired messages based on specific criteria. In Apache Camel, the message filter pattern, shown in [Figure 7.2, “Message Filter Pattern”](#), is implemented by the `filter()` Java DSL command. The `filter()` command takes a single predicate argument, which controls the filter. When the predicate is **true**, the incoming message is allowed to proceed, and when the predicate is **false**, the incoming message is blocked.

Figure 7.2. Message Filter Pattern



Java DSL example

The following example shows how to create a route from endpoint, **seda:a**, to endpoint, **seda:b**, that blocks all messages except for those messages whose **foo** header have the value, **bar**:

```

RouteBuilder builder = new RouteBuilder() {
    public void configure() {

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

```


To evaluate more complex filter predicates, you can invoke one of the supported scripting languages, such as XPath, XQuery, or SQL (see [Expression and Predicate Languages](#)). The following example defines a route that blocks all messages except for those containing a **person** element whose **name** attribute is equal to **James**:


```
from("direct:start").
    filter().xpath("/person[@name='James']").
    to("mock:result");
```

XML configuration example

The following example shows how to configure the route with an XPath predicate in XML (see [Expression and Predicate Languages](#)):

```
<camelContext id="simpleFilterRoute"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <filter>
      <xpath>$foo = 'bar'</xpath>
      <to uri="seda:b"/>
    </filter>
  </route>
</camelContext>
```



FILTERED ENDPOINT REQUIRED INSIDE </FILTER> TAG

Make sure you put the endpoint you want to filter (for example, **<to uri="seda:b"/>**) before the closing **</filter>** tag or the filter will not be applied (in 2.8+, omitting this will result in an error).

Filtering with beans

Here is an example of using a bean to define the filter behavior:

```
from("direct:start")
    .filter().method(MyBean.class,
"isGoldCustomer").to("mock:result").end()
    .to("mock:end");

public static class MyBean {
    public boolean isGoldCustomer(@Header("level") String level) {
        return level.equals("gold");
    }
}
```

Using stop()

Available as of Camel 2.0

Stop is a special type of filter that filters out *all* messages. Stop is convenient to use in a [Content-Based Router](#) when you need to stop further processing in one of the predicates.

In the following example, we do not want messages with the word **Bye** in the message body to propagate any further in the route. We prevent this in the **when()** predicate using **.stop()**.

```
from("direct:start")
```

```

.choice()
  .when(bodyAs(String.class).contains("Hello")).to("mock:hello")
  .when(bodyAs(String.class).contains("Bye")).to("mock:bye").stop()
  .otherwise().to("mock:other")
.end()
.to("mock:result");

```

Knowing if Exchange was filtered or not

Available as of Camel 2.5

The [Message Filter](#) EIP will add a property on the Exchange which states if it was filtered or not.

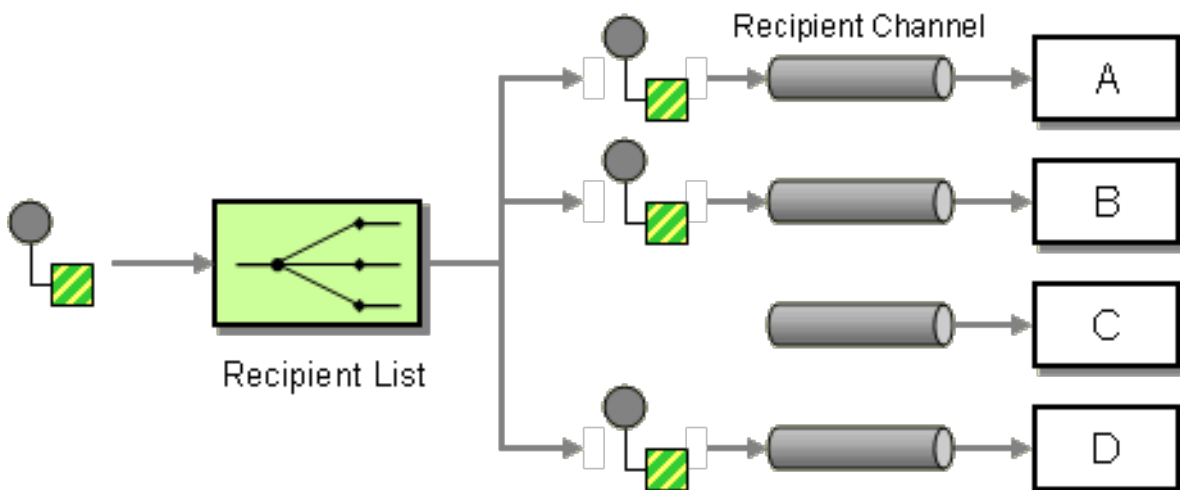
The property has the key **Exchange.FILTER_MATCHED** which has the String value of **CamelFilterMatched**. Its value is a boolean indicating **true** or **false**. If the value is **true** then the Exchange was routed in the filter block.

7.3. RECIPIENT LIST

Overview

A *recipient list*, shown in [Figure 7.3, "Recipient List Pattern"](#), is a type of router that sends each incoming message to multiple different destinations. In addition, a recipient list typically requires that the list of recipients be calculated at run time.

Figure 7.3. Recipient List Pattern



Recipient list with fixed destinations

The simplest kind of recipient list is where the list of destinations is fixed and known in advance, and the exchange pattern is *InOnly*. In this case, you can hardwire the list of destinations into the **to()** Java DSL command.



NOTE

The examples given here, for the recipient list with fixed destinations, work *only* with the *InOnly* exchange pattern (similar to a [pipeline](#)). If you want to create a recipient list for exchange patterns with *Out* messages, use the [multicast](#) pattern instead.

Java DSL example

The following example shows how to route an *InOnly* exchange from a consumer endpoint, **queue:a**, to a fixed list of destinations:

```
from("seda:a").to("seda:b", "seda:c", "seda:d");
```

XML configuration example

The following example shows how to configure the same route in XML:

```
<camelContext id="buildStaticRecipientList"
  xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <to uri="seda:b"/>
    <to uri="seda:c"/>
    <to uri="seda:d"/>
  </route>
</camelContext>
```

Recipient list calculated at run time

In most cases, when you use the recipient list pattern, the list of recipients should be calculated at runtime. To do this use the **recipientList()** processor, which takes a list of destinations as its sole argument. Because Apache Camel applies a type converter to the list argument, it should be possible to use most standard Java list types (for example, a collection, a list, or an array). For more details about type converters, see [Section 40.3, “Built-In Type Converters”](#).

The recipients receive a copy of the *same* exchange instance and Apache Camel executes them sequentially.

Java DSL example

The following example shows how to extract the list of destinations from a message header called **recipientListHeader**, where the header value is a comma-separated list of endpoint URIs:

```
from("direct:a").recipientList(header("recipientListHeader").tokenize(",")
);
```

In some cases, if the header value is a list type, you might be able to use it directly as the argument to **recipientList()**. For example:

```
from("seda:a").recipientList(header("recipientListHeader"));
```

However, this example is entirely dependent on how the underlying component parses this particular header. If the component parses the header as a simple string, this example will *not* work. The header must be parsed into some type of Java list.

XML configuration example

The following example shows how to configure the preceding route in XML, where the header value is a comma-separated list of endpoint URIs:

```
<camelContext id="buildDynamicRecipientList"
  xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <recipientList delimiter=",">
      <header>recipientListHeader</header>
    </recipientList>
  </route>
</camelContext>
```

Sending to multiple recipients in parallel

Available as of Camel 2.2

The [Recipient List](#) supports **parallelProcessing**, which is similar to the corresponding feature in [Splitter](#). Use the parallel processing feature to send the exchange to multiple recipients concurrently—for example:

```
from("direct:a").recipientList(header("myHeader")).parallelProcessing();
```

In Spring XML, the parallel processing feature is implemented as an attribute on the **recipientList** tag—for example:

```
<route>
  <from uri="direct:a"/>
  <recipientList parallelProcessing="true">
    <header>myHeader</header>
  </recipientList>
</route>
```

Stop on exception

Available as of Camel 2.2

The [Recipient List](#) supports the **stopOnException** feature, which you can use to stop sending to any further recipients, if any recipient fails.

```
from("direct:a").recipientList(header("myHeader")).stopOnException();
```

And in Spring XML its an attribute on the recipient list tag.

In Spring XML, the stop on exception feature is implemented as an attribute on the **recipientList** tag—for example:

```
<route>
  <from uri="direct:a"/>
  <recipientList stopOnException="true">
    <header>myHeader</header>
  </recipientList>
</route>
```

**NOTE**

You can combine **parallelProcessing** and **stopOnException** in the same route.

Ignore invalid endpoints**Available as of Camel 2.3**

The [Recipient List](#) supports the **ignoreInvalidEndpoints** option, which enables the recipient list to skip invalid endpoints ([Routing Slip](#) also supports this option). For example:

```
from("direct:a").recipientList(header("myHeader")).ignoreInvalidEndpoints(
);
```

And in Spring XML, you can enable this option by setting the **ignoreInvalidEndpoints** attribute on the **recipientList** tag, as follows

```
<route>
  <from uri="direct:a"/>
  <recipientList ignoreInvalidEndpoints="true">
    <header>myHeader</header>
  </recipientList>
</route>
```

Consider the case where **myHeader** contains the two endpoints, **direct:foo,xxx:bar**. The first endpoint is valid and works. The second is invalid and, therefore, ignored. Apache Camel logs at **INFO** level whenever an invalid endpoint is encountered.

Using custom AggregationStrategy**Available as of Camel 2.2**

You can use a custom **AggregationStrategy** with the [Recipient List](#), which is useful for aggregating replies from the recipients in the list. By default, Apache Camel uses the **UseLatestAggregationStrategy** aggregation strategy, which keeps just the last received reply. For a more sophisticated aggregation strategy, you can define your own implementation of the **AggregationStrategy** interface—see [Aggregator](#) EIP for details. For example, to apply the custom aggregation strategy, **MyOwnAggregationStrategy**, to the reply messages, you can define a Java DSL route as follows:

```
from("direct:a")
  .recipientList(header("myHeader")).aggregationStrategy(new
MyOwnAggregationStrategy())
  .to("direct:b");
```

In Spring XML, you can specify the custom aggregation strategy as an attribute on the **recipientList** tag, as follows:

```
<route>
  <from uri="direct:a"/>
  <recipientList strategyRef="myStrategy">
    <header>myHeader</header>
  </recipientList>
  <to uri="direct:b"/>
</route>

<bean id="myStrategy" class="com.mycompany.MyOwnAggregationStrategy"/>
```

Using custom thread pool

Available as of Camel 2.2

This is only needed when you use **parallelProcessing**. By default Camel uses a thread pool with 10 threads. Notice this is subject to change when we overhaul thread pool management and configuration later (hopefully in Camel 2.2).

You configure this just as you would with the custom aggregation strategy.

Using method call as recipient list

You can use a [Bean](#) to provide the recipients, for example:

```
from("activemq:queue:test").recipientList().method(MessageRouter.class,
"routeTo");
```

Where the **MessageRouter** bean is defined as follows:

```
public class MessageRouter {
    public String routeTo() {
        String queueName = "activemq:queue:test2";
        return queueName;
    }
}
```

Bean as recipient list

You can make a bean behave as a recipient list by adding the **@RecipientList** annotation to a methods that returns a list of recipients. For example:

```
public class MessageRouter {
    @RecipientList
    public String routeTo() {
        String queueList = "activemq:queue:test1,activemq:queue:test2";
        return queueList;
    }
}
```

In this case, do *not* include the **recipientList** DSL command in the route. Define the route as follows:

```
from("activemq:queue:test").bean(MessageRouter.class, "routeTo");
```

Using timeout

Available as of Camel 2.5

If you use **parallelProcessing**, you can configure a total **timeout** value in milliseconds. Camel will then process the messages in parallel until the timeout is hit. This allows you to continue processing if one message is slow.

In the example below, the **recipientList** header has the value, **direct:a, direct:b, direct:c**, so that the message is sent to three recipients. We have a timeout of 250 milliseconds, which means only the last two messages can be completed within the timeframe. The aggregation therefore yields the string result, **BC**.


```
from("direct:start")
    .recipientList(header("recipients"), ",")
    .aggregationStrategy(new AggregationStrategy() {
        public Exchange aggregate(Exchange oldExchange, Exchange
newExchange) {
            if (oldExchange == null) {
                return newExchange;
            }

            String body = oldExchange.getIn().getBody(String.class);
            oldExchange.getIn().setBody(body +
newExchange.getIn().getBody(String.class));
            return oldExchange;
        }
    })
    .parallelProcessing().timeout(250)
    // use end to indicate end of recipientList clause
    .end()
    .to("mock:result");

from("direct:a").delay(500).to("mock:A").setBody(constant("A"));

from("direct:b").to("mock:B").setBody(constant("B"));

from("direct:c").to("mock:C").setBody(constant("C"));
```



NOTE

This **timeout** feature is also supported by **splitter** and both **multicast** and **recipientList**.

By default if a timeout occurs the **AggregationStrategy** is not invoked. However you can implement a specialized version

```
// Java
public interface TimeoutAwareAggregationStrategy extends
```

```

AggregationStrategy {

    /**
     * A timeout occurred
     *
     * @param oldExchange the oldest exchange (is <tt>null</tt> on
first aggregation as we only have the new exchange)
     * @param index       the index
     * @param total       the total
     * @param timeout     the timeout value in millis
     */
    void timeout(Exchange oldExchange, int index, int total, long
timeout);
}

```

This allows you to deal with the timeout in the **AggregationStrategy** if you really need to.



TIMEOUT IS TOTAL

The timeout is total, which means that after X time, Camel will aggregate the messages which has completed within the timeframe. The remainders will be cancelled. Camel will also only invoke the **timeout** method in the **TimeoutAwareAggregationStrategy** once, for the first index which caused the timeout.

Apply custom processing to the outgoing messages

Before **recipientList** sends a message to one of the recipient endpoints, it creates a message replica, which is a shallow copy of the original message. If you want to perform some custom processing on each message replica before the replica is sent to its endpoint, you can invoke the **onPrepare** DSL command in the **recipientList** clause. The **onPrepare** command inserts a custom processor just *after* the message has been shallow-copied and just *before* the message is dispatched to its endpoint. For example, in the following route, the **CustomProc** processor is invoked on the message replica *foreach recipient endpoint*:

```

from("direct:start")
    .recipientList().onPrepare(new CustomProc());

```

A common use case for the **onPrepare** DSL command is to perform a deep copy of some or all elements of a message. This allows each message replica to be modified independently of the others. For example, the following **CustomProc** processor class performs a deep copy of the message body, where the message body is presumed to be of type, **BodyType**, and the deep copy is performed by the method, **BodyType.deepCopy()**.

```

// Java
import org.apache.camel.*;
...
public class CustomProc implements Processor {

    public void process(Exchange exchange) throws Exception {
        BodyType body = exchange.getIn().getBody(BodyType.class);

        // Make a _deep_ copy of of the body object
        BodyType clone = BodyType.deepCopy();
        exchange.getIn().setBody(clone);
    }
}

```



```

    // Headers and attachments have already been
    // shallow-copied. If you need deep copies,
    // add some more code here.
  }
}

```

Options

The `recipientList` DSL command supports the following options:

Name	Default Value	Description
<code>delimiter</code>	,	Delimiter used if the Expression returned multiple endpoints.
<code>strategyRef</code>		Refers to an AggregationStrategy to be used to assemble the replies from the recipients, into a single outgoing message from the Recipient List . By default Camel will use the last reply as the outgoing message.
<code>parallelProcessing</code>	<code>false</code>	Camel 2.2: If enables then sending messages to the recipients occurs concurrently. Note the caller thread will still wait until all messages has been fully processed, before it continues. Its only the sending and processing the replies from the recipients which happens concurrently.
<code>executorServiceRef</code>		Camel 2.2: Refers to a custom Thread Pool to be used for parallel processing. Notice if you set this option, then parallel processing is automatic implied, and you do not have to enable that option as well.

stopOnException	false	Camel 2.2: Whether or not to stop continue processing immediately when an exception occurred. If disable, then Camel will send the message to all recipients regardless if one of them failed. You can deal with exceptions in the AggregationStrategy class where you have full control how to handle that.
ignoreInvalidEndpoints	false	Camel 2.3: If an endpoint uri could not be resolved, should it be ignored. Otherwise Camel will throw an exception stating the endpoint uri is not valid.
streaming	false	Camel 2.5: If enabled then Camel will process replies out-of-order, eg in the order they come back. If disabled, Camel will process replies in the same order as the Expression specified.
timeout		Camel 2.5: Sets a total timeout specified in millis. If the Recipient List hasn't been able to send and process all replies within the given timeframe, then the timeout triggers and the Recipient List breaks out and continues. Notice if you provide a TimeoutAwareAggregationStrategy then the timeout method is invoked before breaking out.
onPrepareRef		Camel 2.8: Refers to a custom Processor to prepare the copy of the Exchange each recipient will receive. This allows you to do any custom logic, such as deep-cloning the message payload if that's needed etc.


<code>shareUnitOfWork</code>	<code>false</code>	Camel 2.8: Whether the unit of work should be shared. See the same option on Splitter for more details.
------------------------------	--------------------	--

7.4. SPLITTER

Overview

A *splitter* is a type of router that splits an incoming message into a series of outgoing messages. Each of the outgoing messages contains a piece of the original message. In Apache Camel, the splitter pattern, shown in [Figure 7.4, “Splitter Pattern”](#), is implemented by the `split()` Java DSL command.

Figure 7.4. Splitter Pattern



The Apache Camel splitter actually supports two patterns, as follows:

- *Simple splitter*—implements the splitter pattern on its own.
- *Splitter/aggregator*—combines the splitter pattern with the aggregator pattern, such that the pieces of the message are recombined after they have been processed.

Java DSL example

The following example defines a route from `seda:a` to `seda:b` that splits messages by converting each line of an incoming message into a separate outgoing message:

```
RouteBuilder builder = new RouteBuilder() {
    public void configure() {
        from("seda:a")
            .split(bodyAs(String.class).tokenize("\n"))
            .to("seda:b");
    }
};
```

The splitter can use any expression language, so you can split messages using any of the supported scripting languages, such as XPath, XQuery, or SQL (see [Part II, “Routing Expression and Predicate Languages”](#)). The following example extracts `bar` elements from an incoming message and insert them into separate outgoing messages:

```
from("activemq:my.queue")
  .split(xpath("//foo/bar"))
  .to("file://some/directory")
```

XML configuration example

The following example shows how to configure a splitter route in XML, using the XPath scripting language:

```
<camelContext id="buildSplitter"
  xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <split>
      <xpath>//foo/bar</xpath>
      <to uri="seda:b"/>
    </split>
  </route>
</camelContext>
```

You can use the `tokenize` expression in the XML DSL to split bodies or headers using a token, where the `tokenize` expression is defined using the `tokenize` element. In the following example, the message body is tokenized using the `\n` separator character. To use a regular expression pattern, set `regex=true` in the `tokenize` element.

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <split>
      <tokenize token="\n"/>
      <to uri="mock:result"/>
    </split>
  </route>
</camelContext>
```

Splitting into groups of lines

To split a big file into chunks of 1000 lines, you can define a splitter route as follows in the Java DSL:

```
from("file:inbox")
  .split().tokenize("\n", 1000).streaming()
  .to("activemq:queue:order");
```

The second argument to `tokenize` specifies the number of lines that should be grouped into a single chunk. The `streaming()` clause directs the splitter not to read the whole file at once (resulting in much better performance if the file is large).

The same route can be defined in XML DSL as follows:

```
<route>
  <from uri="file:inbox"/>
  <split streaming="true">
```

```

    <tokenize token="\n" group="1000"/>
    <to uri="activemq:queue:order"/>
  </split>
</route>

```

The output when using the **group** option is always of **java.lang.String** type.

Splitter reply

If the exchange that enters the splitter has the *InOut* message-exchange pattern (that is, a reply is expected), the splitter returns a copy of the original input message as the reply message in the *Out* message slot. You can override this default behavior by implementing your own [aggregation strategy](#).

Parallel execution

If you want to execute the resulting pieces of the message in parallel, you can enable the parallel processing option, which instantiates a thread pool to process the message pieces. For example:

```

XPathBuilder xpathBuilder = new XPathBuilder("//foo/bar");
from("activemq:my.queue").split(xpathBuilder).parallelProcessing().to("act
ivemq:my.parts");

```

You can customize the underlying **ThreadPoolExecutor** used in the parallel splitter. For example, you can specify a custom executor in the Java DSL as follows:

```

XPathBuilder xpathBuilder = new XPathBuilder("//foo/bar");
ThreadPoolExecutor threadPoolExecutor = new ThreadPoolExecutor(8, 16, 0L,
TimeUnit.MILLISECONDS, new LinkedBlockingQueue());
from("activemq:my.queue")
    .split(xpathBuilder)
    .parallelProcessing()
    .executorService(threadPoolExecutor)
    .to("activemq:my.parts");

```

You can specify a custom executor in the XML DSL as follows:

```

<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:parallel-custom-pool"/>
    <split executorServiceRef="threadPoolExecutor">
      <xpath>/invoice/lineItems</xpath>
      <to uri="mock:result"/>
    </split>
  </route>
</camelContext>

<bean id="threadPoolExecutor"
class="java.util.concurrent.ThreadPoolExecutor">
  <constructor-arg index="0" value="8"/>
  <constructor-arg index="1" value="16"/>
  <constructor-arg index="2" value="0"/>
  <constructor-arg index="3" value="MILLISECONDS"/>

```

```
<constructor-arg index="4"><bean
class="java.util.concurrent.LinkedBlockingQueue"/></constructor-arg>
</bean>
```

Using a bean to perform splitting

As the splitter can use *any* expression to do the splitting, we can use a bean to perform splitting, by invoking the `method()` expression. The bean should return an iterable value such as: `java.util.Collection`, `java.util.Iterator`, or an array.

The following route defines a `method()` expression that calls a method on the `mySplitterBean` bean instance:

```
from("direct:body")
    // here we use a POJO bean mySplitterBean to do the split of the
payload
    .split()
    .method("mySplitterBean", "splitBody")
    .to("mock:result");
from("direct:message")
    // here we use a POJO bean mySplitterBean to do the split of the
message
    // with a certain header value
    .split()
    .method("mySplitterBean", "splitMessage")
    .to("mock:result");
```

Where `mySplitterBean` is an instance of the `MySplitterBean` class, which is defined as follows:

```
public class MySplitterBean {

    /**
     * The split body method returns something that is iterable such as
a java.util.List.
     *
     * @param body the payload of the incoming message
     * @return a list containing each part split
     */
    public List<String> splitBody(String body) {
        // since this is based on an unit test you can of couse
        // use different logic for splitting as Apache Camel have out
        // of the box support for splitting a String based on comma
        // but this is for show and tell, since this is java code
        // you have the full power how you like to split your messages
        List<String> answer = new ArrayList<String>();
        String[] parts = body.split(",");
        for (String part : parts) {
            answer.add(part);
        }
        return answer;
    }

    /**
     * The split message method returns something that is iterable such
```

```

as a java.util.List.
 *
 * @param header the header of the incoming message with the name user
 * @param body the payload of the incoming message
 * @return a list containing each part split
 */
public List<Message> splitMessage(@Header(value = "user") String
header, @Body String body) {
    // we can leverage the Parameter Binding Annotations
    // http://camel.apache.org/parameter-binding-annotations.html
    // to access the message header and body at same time,
    // then create the message that we want, splitter will
    // take care rest of them.
    // *NOTE* this feature requires Apache Camel version >= 1.6.1
    List<Message> answer = new ArrayList<Message>();
    String[] parts = header.split(",");
    for (String part : parts) {
        DefaultMessage message = new DefaultMessage();
        message.setHeader("user", part);
        message.setBody(body);
        answer.add(message);
    }
    return answer;
}
}

```

Exchange properties

The following properties are set on each split exchange:

header	type	description
CamelSplitIndex	int	Apache Camel 2.0: A split counter that increases for each Exchange being split. The counter starts from 0.
CamelSplitSize	int	Apache Camel 2.0: The total number of Exchanges that was split. This header is not applied for stream based splitting.
CamelSplitComplete	boolean	Apache Camel 2.4: Whether or not this Exchange is the last.

Splitter/aggregator pattern

It is a common pattern for the message pieces to be aggregated back into a single exchange, after processing of the individual pieces has completed. To support this pattern, the **split()** DSL command lets you provide an **AggregationStrategy** object as the second argument.

Java DSL example

The following example shows how to use a custom aggregation strategy to recombine a split message after all of the message pieces have been processed:

```
from("direct:start")
    .split(body().tokenize("@"), new MyOrderStrategy())
    // each split message is then send to this bean where we can
process it
    .to("bean:MyOrderService?method=handleOrder")
    // this is important to end the splitter route as we do not want
to do more routing
    // on each split message
    .end()
    // after we have split and handled each message we want to send a
single combined
    // response back to the original caller, so we let this bean build it
for us
    // this bean will receive the result of the aggregate strategy:
MyOrderStrategy
    .to("bean:MyOrderService?method=buildCombinedResponse")
```

AggregationStrategy implementation

The custom aggregation strategy, **MyOrderStrategy**, used in the preceding route is implemented as follows:

```
/**
 * This is our own order aggregation strategy where we can control
 * how each split message should be combined. As we do not want to
 * lose any message, we copy from the new to the old to preserve the
 * order lines as long we process them
 */
public static class MyOrderStrategy implements AggregationStrategy {

    public Exchange aggregate(Exchange oldExchange, Exchange newExchange)
    {
        // put order together in old exchange by adding the order from new
exchange

        if (oldExchange == null) {
            // the first time we aggregate we only have the new exchange,
            // so we just return it
            return newExchange;
        }

        String orders = oldExchange.getIn().getBody(String.class);
        String newLine = newExchange.getIn().getBody(String.class);

        LOG.debug("Aggregate old orders: " + orders);
        LOG.debug("Aggregate new order: " + newLine);

        // put orders together separating by semi colon
        orders = orders + ";" + newLine;
        // put combined order back on old to preserve it
```



```

        oldExchange.getIn().setBody(orders);

        // return old as this is the one that has all the orders gathered
until now
        return oldExchange;
    }
}

```

Stream based processing

When parallel processing is enabled, it is theoretically possible for a later message piece to be ready for aggregation before an earlier piece. In other words, the message pieces might arrive at the aggregator out of order. By default, this does not happen, because the splitter implementation rearranges the message pieces back into their original order before passing them into the aggregator.

If you would prefer to aggregate the message pieces as soon as they are ready (and possibly out of order), you can enable the streaming option, as follows:

```

from("direct:streaming")
    .split(body().tokenize(","), new MyOrderStrategy())
    .parallelProcessing()
    .streaming()
    .to("activemq:my.parts")
    .end()
    .to("activemq:all.parts");

```

You can also supply a custom iterator to use with streaming, as follows:

```

// Java
import static org.apache.camel.builder.ExpressionBuilder.beanExpression;
...
from("direct:streaming")
    .split(beanExpression(new MyCustomIteratorFactory(), "iterator"))
    .streaming().to("activemq:my.parts")

```



STREAMING AND XPATH

You cannot use streaming mode in conjunction with XPath. XPath requires the complete DOM XML document in memory.

Stream based processing with XML

If an incoming messages is a very large XML file, you can process the message most efficiently using the **tokenizeXML** sub-command in streaming mode.

For example, given a large XML file that contains a sequence of **order** elements, you can split the file into **order** elements using a route like the following:

```

from("file:inbox")
    .split().tokenizeXML("order").streaming()
    .to("activemq:queue:order");

```

You can do the same thing in XML, by defining a route like the following:

```

<route>
  <from uri="file:inbox"/>
  <split streaming="true">
    <tokenize token="order" xml="true"/>
    <to uri="activemq:queue:order"/>
  </split>
</route>

```

It is often the case that you need access to namespaces that are defined in one of the enclosing (ancestor) elements of the token elements. You can copy namespace definitions from one of the ancestor elements into the token element, by specifying which element you want to inherit namespace definitions from.

In the Java DSL, you specify the ancestor element as the second argument of **tokenizeXML**. For example, to inherit namespace definitions from the enclosing **orders** element:

```

from("file:inbox")
  .split().tokenizeXML("order", "orders").streaming()
  .to("activemq:queue:order");

```

In the XML DSL, you specify the ancestor element using the **inheritNamespaceTagName** attribute. For example:

```

<route>
  <from uri="file:inbox"/>
  <split streaming="true">
    <tokenize token="order"
      xml="true"
      inheritNamespaceTagName="orders"/>
    <to uri="activemq:queue:order"/>
  </split>
</route>

```

Options

The **split** DSL command supports the following options:

Name	Default Value	Description
strategyRef		Refers to an AggregationStrategy to be used to assemble the replies from the sub-messages, into a single outgoing message from the Splitter . See the section titled <i>What does the splitter return</i> below for whats used by default.

parallelProcessing	false	If enables then processing the sub-messages occurs concurrently. Note the caller thread will still wait until all sub-messages has been fully processed, before it continues.
executorServiceRef		Refers to a custom Thread Pool to be used for parallel processing. Notice if you set this option, then parallel processing is automatic implied, and you do not have to enable that option as well.
stopOnException	false	Camel 2.2: Whether or not to stop continue processing immediately when an exception occurred. If disable, then Camel continue splitting and process the sub-messages regardless if one of them failed. You can deal with exceptions in the AggregationStrategy class where you have full control how to handle that.
streaming	false	If enabled then Camel will split in a streaming fashion, which means it will split the input message in chunks. This reduces the memory overhead. For example if you split big messages its recommended to enable streaming. If streaming is enabled then the sub-message replies will be aggregated out-of-order, eg in the order they come back. If disabled, Camel will process sub-message replies in the same order as they where splitted.


<code>timeout</code>		Camel 2.5: Sets a total timeout specified in millis. If the Recipient List hasn't been able to split and process all replies within the given timeframe, then the timeout triggers and the Splitter breaks out and continues. Notice if you provide a TimeoutAwareAggregationStrategy then the <code>timeout</code> method is invoked before breaking out.
<code>onPrepareRef</code>		Camel 2.8: Refers to a custom Processor to prepare the sub-message of the Exchange, before its processed. This allows you to do any custom logic, such as deep-cloning the message payload if that's needed etc.
<code>shareUnitOfWork</code>	<code>false</code>	Camel 2.8: Whether the unit of work should be shared. See further below for more details.

7.5. AGGREGATOR

Overview

The *aggregator* pattern, shown in [Figure 7.5, “Aggregator Pattern”](#), enables you to combine a batch of related messages into a single message.

Figure 7.5. Aggregator Pattern



To control the aggregator's behavior, Apache Camel allows you to specify the properties described in *Enterprise Integration Patterns*, as follows:

- *Correlation expression* — Determines which messages should be aggregated together. The correlation expression is evaluated on each incoming message to produce a *correlation key*. Incoming messages with the same correlation key are

then grouped into the same batch. For example, if you want to aggregate *all* incoming messages into a single message, you can use a constant expression.

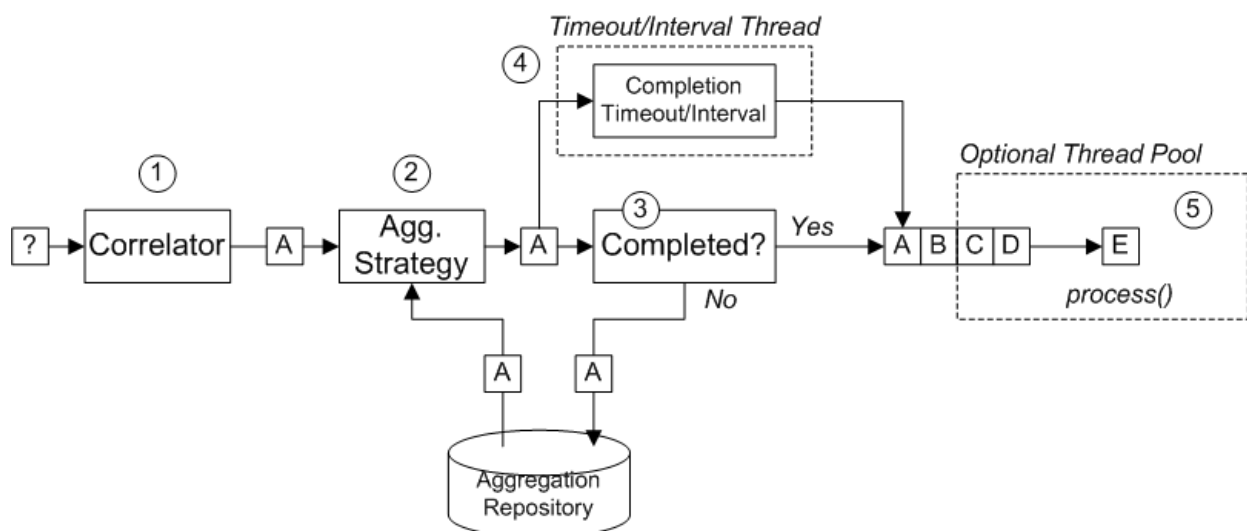
- *Completeness condition* — Determines when a batch of messages is complete. You can specify this either as a simple size limit or, more generally, you can specify a predicate condition that flags when the batch is complete.
- *Aggregation algorithm* — Combines the message exchanges for a single correlation key into a single message exchange.

For example, consider a stock market data system that receives 30,000 messages per second. You might want to throttle down the message flow if your GUI tool cannot cope with such a massive update rate. The incoming stock quotes can be aggregated together simply by choosing the latest quote and discarding the older prices. (You can apply a delta processing algorithm, if you prefer to capture some of the history.)

How the aggregator works

Figure 7.6, “Aggregator Implementation” shows an overview of how the aggregator works, assuming it is fed with a stream of exchanges that have correlation keys such as A, B, C, or D.

Figure 7.6. Aggregator Implementation



The incoming stream of exchanges shown in Figure 7.6, “Aggregator Implementation” is processed as follows:

1. The *correlator* is responsible for sorting exchanges based on the correlation key. For each incoming exchange, the correlation expression is evaluated, yielding the correlation key. For example, for the exchange shown in Figure 7.6, “Aggregator Implementation”, the correlation key evaluates to A.
2. The *aggregation strategy* is responsible for merging exchanges with the same correlation key. When a new exchange, A, comes in, the aggregator looks up the corresponding *aggregate exchange*, A', in the aggregation repository and combines it with the new exchange.

Until a particular aggregation cycle is completed, incoming exchanges are continuously aggregated with the corresponding aggregate exchange. An aggregation cycle lasts until terminated by one of the completion mechanisms.

3. If a completion predicate is specified on the aggregator, the aggregate exchange is tested to determine whether it is ready to be sent to the next processor in the route. Processing continues as follows:
 - If complete, the aggregate exchange is processed by the latter part of the route. There are two alternative models for this: *synchronous* (the default), which causes the calling thread to block, or *asynchronous* (if parallel processing is enabled), where the aggregate exchange is submitted to an executor thread pool (as shown in [Figure 7.6, “Aggregator Implementation”](#)).
 - If not complete, the aggregate exchange is saved back to the aggregation repository.
4. In parallel with the synchronous completion tests, it is possible to enable an asynchronous completion test by enabling *either* the **completionTimeout** option or the **completionInterval** option. These completion tests run in a separate thread and, whenever the completion test is satisfied, the corresponding exchange is marked as complete and starts to be processed by the latter part of the route (either synchronously or asynchronously, depending on whether parallel processing is enabled or not).
5. If parallel processing is enabled, a thread pool is responsible for processing exchanges in the latter part of the route. By default, this thread pool contains ten threads, but you have the option of customizing the pool ([the section called “Threading options”](#)).

Java DSL example

The following example aggregates exchanges with the same **StockSymbol** header value, using the **UseLatestAggregationStrategy** aggregation strategy. For a given **StockSymbol** value, if more than three seconds elapse since the last exchange with that correlation key was received, the aggregated exchange is deemed to be complete and is sent to the **mock** endpoint.

```
from("direct:start")
    .aggregate(header("id"), new UseLatestAggregationStrategy())
        .completionTimeout(3000)
        .to("mock:aggregated");
```

XML DSL example

The following example shows how to configure the same route in XML:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <aggregate strategyRef="aggregatorStrategy"
      completionTimeout="3000">
      <correlationExpression>
        <simple>header.StockSymbol</simple>
      </correlationExpression>
      <to uri="mock:aggregated"/>
    </aggregate>
  </route>
</camelContext>
```

```
<bean id="aggregatorStrategy"
class="org.apache.camel.processor.aggregate.UseLatestAggregationStrategy"/
>
```

Specifying the correlation expression

In the Java DSL, the correlation expression is always passed as the first argument to the **aggregate()** DSL command. You are not limited to using the Simple expression language here. You can specify a correlation expression using any of the expression languages or scripting languages, such as XPath, XQuery, SQL, and so on.

For example, to correlate exchanges using an XPath expression, you could use the following Java DSL route:

```
from("direct:start")
    .aggregate(xpath("/stockQuote/@symbol"), new
UseLatestAggregationStrategy())
    .completionTimeout(3000)
    .to("mock:aggregated");
```

If the correlation expression cannot be evaluated on a particular incoming exchange, the aggregator throws a **CamelExchangeException** by default. You can suppress this exception by setting the **ignoreInvalidCorrelationKeys** option. For example, in the Java DSL:

```
from(...).aggregate(...).ignoreInvalidCorrelationKeys()
```

In the XML DSL, you can set the **ignoreInvalidCorrelationKeys** option is set as an attribute, as follows:

```
<aggregate strategyRef="aggregatorStrategy"
            ignoreInvalidCorrelationKeys="true"
            ...>
    ...
</aggregate>
```

Specifying the aggregation strategy

In Java DSL, you can either pass the aggregation strategy as the second argument to the **aggregate()** DSL command or specify it using the **aggregationStrategy()** clause. For example, you can use the **aggregationStrategy()** clause as follows:

```
from("direct:start")
    .aggregate(header("id"))
    .aggregationStrategy(new UseLatestAggregationStrategy())
    .completionTimeout(3000)
    .to("mock:aggregated");
```

Apache Camel provides the following basic aggregation strategies (where the classes belong to the **org.apache.camel.processor.aggregate** Java package):

UseLatestAggregationStrategy

Return the last exchange for a given correlation key, discarding all earlier exchanges with this key. For example, this strategy could be useful for throttling the feed from a stock exchange, where you just want to know the latest price of a particular stock symbol.

UseOriginalAggregationStrategy

Return the first exchange for a given correlation key, discarding all later exchanges with this key. You must set the first exchange by calling

UseOriginalAggregationStrategy.setOriginal() before you can use this strategy.

GroupedExchangeAggregationStrategy

Concatenates *all* of the exchanges for a given correlation key into a list, which is stored in the **Exchange.GROUPED_EXCHANGE** exchange property. See [the section called “Grouped exchanges”](#).

Implementing a custom aggregation strategy

If you want to apply a different aggregation strategy, you can implement one of the following aggregation strategy base interfaces:

org.apache.camel.processor.aggregate.AggregationStrategy

The basic aggregation strategy interface.

org.apache.camel.processor.aggregate.TimeoutAwareAggregationStrategy

Implement this interface, if you want your implementation to receive a notification when an aggregation cycle times out. The **timeout** notification method has the following signature:

```
void timeout(Exchange oldExchange, int index, int total, long timeout)
```

org.apache.camel.processor.aggregate.CompletionAwareAggregationStrategy

Implement this interface, if you want your implementation to receive a notification when an aggregation cycle completes normally. The notification method has the following signature:

```
void onCompletion(Exchange exchange)
```

For example, the following code shows two different custom aggregation strategies, **StringAggregationStrategy** and **ArrayListAggregationStrategy**:

```
//simply combines Exchange String body values using '+' as a delimiter
class StringAggregationStrategy implements AggregationStrategy {
    public Exchange aggregate(Exchange oldExchange, Exchange newExchange)
    {
        if (oldExchange == null) {
            return newExchange;
        }

        String oldBody = oldExchange.getIn().getBody(String.class);
```




```

        String newBody = newExchange.getIn().getBody(String.class);
        oldExchange.getIn().setBody(oldBody + "+" + newBody);
        return oldExchange;
    }
}

//simply combines Exchange body values into an ArrayList<Object>
class ArrayListAggregationStrategy implements AggregationStrategy {

    public Exchange aggregate(Exchange oldExchange, Exchange newExchange)
    {
        Object newBody = newExchange.getIn().getBody();
        ArrayList<Object> list = null;
        if (oldExchange == null) {
            list = new ArrayList<Object>();
            list.add(newBody);
            newExchange.getIn().setBody(list);
            return newExchange;
        } else {
            list = oldExchange.getIn().getBody(ArrayList.class);
            list.add(newBody);
            return oldExchange;
        }
    }
}

```



NOTE

Since Apache Camel 2.0, the **AggregationStrategy.aggregate()** callback method is also invoked for the very first exchange. On the first invocation of the **aggregate** method, the **oldExchange** parameter is **null** and the **newExchange** parameter contains the first incoming exchange.

To aggregate messages using the custom strategy class, **ArrayListAggregationStrategy**, define a route like the following:

```

from("direct:start")
    .aggregate(header("StockSymbol"), new ArrayListAggregationStrategy())
    .completionTimeout(3000)
    .to("mock:result");

```

You can also configure a route with a custom aggregation strategy in XML, as follows:

```

<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <aggregate strategyRef="aggregatorStrategy"
      completionTimeout="3000">
      <correlationExpression>
        <simple>header.StockSymbol</simple>
      </correlationExpression>
      <to uri="mock:aggregated"/>
    </aggregate>
  </route>

```

```

</camelContext>

<bean id="aggregatorStrategy"
class="com.my_package_name.ArrayListAggregationStrategy"/>

```

Controlling the lifecycle of a custom aggregation strategy

You can implement a custom aggregation strategy so that its lifecycle is aligned with the lifecycle of the enterprise integration pattern that is controlling it. This can be useful for ensuring that the aggregation strategy can shut down gracefully.

To implement an aggregation strategy with lifecycle support, you must implement the **org.apache.camel.Service** interface (in addition to the **AggregationStrategy** interface) and provide implementations of the **start()** and **stop()** lifecycle methods. For example, the following code example shows an outline of an aggregation strategy with lifecycle support:

```

// Java
import org.apache.camel.processor.aggregate.AggregationStrategy;
import org.apache.camel.Service;
import java.lang.Exception;
...
class MyAggStrategyWithLifecycleControl
    implements AggregationStrategy, Service {

    public Exchange aggregate(Exchange oldExchange, Exchange newExchange)
    {
        // Implementation not shown...
        ...
    }

    public void start() throws Exception {
        // Actions to perform when the enclosing EIP starts up
        ...
    }

    public void stop() throws Exception {
        // Actions to perform when the enclosing EIP is stopping
        ...
    }
}

```

Exchange properties

The following properties are set on each aggregated exchange:

Table 7.1. Aggregated Exchange Properties

Header	Type	Description
Exchange.AGGREGATED_SIZE	int	The total number of exchanges aggregated into this exchange.

Header	Type	Description
<code>Exchange.AGGREGATED_COMPLETED_BY</code>	<code>String</code>	Indicates the mechanism responsible for completing the aggregate exchange. Possible values are: predicate , size , timeout , interval , or consumer .

The following properties are set on exchanges redelivered by the HawtDB aggregation repository (see [the section called “Persistent aggregation repository”](#)):

Table 7.2. Redelivered Exchange Properties

Header	Type	Description
<code>Exchange.REDELIVERY_COUNTER</code>	<code>int</code>	Sequence number of the current redelivery attempt (starting at 1).

Specifying a completion condition

It is mandatory to specify *at least one* completion condition, which determines when an aggregate exchange leaves the aggregator and proceeds to the next node on the route. The following completion conditions can be specified:

`completionPredicate`

Evaluates a predicate after each exchange is aggregated in order to determine completeness. A value of **true** indicates that the aggregate exchange is complete.

`completionSize`

Completes the aggregate exchange after the specified number of incoming exchanges are aggregated.

`completionTimeout`

(Incompatible with `completionInterval`) Completes the aggregate exchange, if no incoming exchanges are aggregated within the specified timeout.

In other words, the timeout mechanism keeps track of a timeout for *each* correlation key value. The clock starts ticking after the latest exchange with a particular key value is received. If another exchange with the same key value is *not* received within the specified timeout, the corresponding aggregate exchange is marked complete and sent to the next node on the route.

`completionInterval`

(Incompatible with `completionTimeout`) Completes *all* outstanding aggregate exchanges, after each time interval (of specified length) has elapsed.

The time interval is *not* tailored to each aggregate exchange. This mechanism forces simultaneous completion of all outstanding aggregate exchanges. Hence, in some cases, this mechanism could complete an aggregate exchange immediately after it started

aggregating.

completionFromBatchConsumer

When used in combination with a consumer endpoint that supports the *batch consumer* mechanism, this completion option automatically figures out when the current batch of exchanges is complete, based on information it receives from the consumer endpoint. See [the section called “Batch consumer”](#).

forceCompletionOnStop

When this option is enabled, it forces completion of all outstanding aggregate exchanges when the current route context is stopped.

The preceding completion conditions can be combined arbitrarily, except for the **completionTimeout** and **completionInterval** conditions, which cannot be simultaneously enabled. When conditions are used in combination, the general rule is that the first completion condition to trigger is the effective completion condition.

Specifying the completion predicate

You can specify an arbitrary predicate expression that determines when an aggregated exchange is complete. There are two possible ways of evaluating the predicate expression:

- *On the latest aggregate exchange*—this is the default behavior.
- *On the latest incoming exchange*—this behavior is selected when you enable the **eagerCheckCompletion** option.

For example, if you want to terminate a stream of stock quotes every time you receive an **ALERT** message (as indicated by the value of **aMsgType** header in the latest incoming exchange), you can define a route like the following:

```
from("direct:start")
  .aggregate(
    header("id"),
    new UseLatestAggregationStrategy()
  )
  .completionPredicate(
    header("MsgType").isEqualTo("ALERT")
  )
  .eagerCheckCompletion()
  .to("mock:result");
```

The following example shows how to configure the same route using XML:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <aggregate strategyRef="aggregatorStrategy"
      eagerCheckCompletion="true">
      <correlationExpression>
        <simple>header.StockSymbol</simple>
      </correlationExpression>
    </aggregate>
  </route>
</camelContext>
```

```

        <simple>$MsgType = 'ALERT'</simple>
    </completionPredicate>
    <to uri="mock:result"/>
</aggregate>
</route>
</camelContext>

<bean id="aggregatorStrategy"
class="org.apache.camel.processor.aggregate.UseLatestAggregationStrategy"/
>

```

Specifying a dynamic completion timeout

It is possible to specify a *dynamic completion timeout*, where the timeout value is recalculated for every incoming exchange. For example, to set the timeout value from the **timeout** header in each incoming exchange, you could define a route as follows:

```

from("direct:start")
    .aggregate(header("StockSymbol"), new UseLatestAggregationStrategy())
        .completionTimeout(header("timeout"))
    .to("mock:aggregated");

```

You can configure the same route in the XML DSL, as follows:

```

<camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="direct:start"/>
        <aggregate strategyRef="aggregatorStrategy">
            <correlationExpression>
                <simple>header.StockSymbol</simple>
            </correlationExpression>
            <completionTimeout>
                <header>timeout</header>
            </completionTimeout>
            <to uri="mock:aggregated"/>
        </aggregate>
    </route>
</camelContext>

<bean id="aggregatorStrategy"
class="org.apache.camel.processor.aggregate.UseLatestAggregationStrategy"/>

```



NOTE

You can also add a fixed timeout value and Apache Camel will fall back to use this value, if the dynamic value is **null** or **0**.

Specifying a dynamic completion size

It is possible to specify a *dynamic completion size*, where the completion size is recalculated for every incoming exchange. For example, to set the completion size from the **mySize** header in each incoming exchange, you could define a route as follows:

■

```

from("direct:start")
    .aggregate(header("StockSymbol"), new UseLatestAggregationStrategy())
        .completionSize(header("mySize"))
    .to("mock:aggregated");

```

And the same example using Spring XML:

```

<camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="direct:start"/>
        <aggregate strategyRef="aggregatorStrategy">
            <correlationExpression>
                <simple>header.StockSymbol</simple>
            </correlationExpression>
            <completionSize>
                <header>mySize</header>
            </completionSize>
            <to uri="mock:aggregated"/>
        </aggregate>
    </route>
</camelContext>

<bean id="aggregatorStrategy"
    class="org.apache.camel.processor.UseLatestAggregationStrategy"/>

```



NOTE

You can also add a fixed size value and Apache Camel will fall back to use this value, if the dynamic value is **null** or **0**.

Forcing completion with a special message

It is possible to force completion of all outstanding aggregate messages, by sending a message with a special header to the route. There are two alternative header settings you can use to force completion:

Exchange.AGGREGATION_COMPLETE_ALL_GROUPS

Set to **true**, to force completion of the current aggregation cycle. This message acts purely as a signal and is *not* included in any aggregation cycle. After processing this signal message, the content of the message is discarded.

Exchange.AGGREGATION_COMPLETE_ALL_GROUPS_INCLUSIVE

Set to **true**, to force completion of the current aggregation cycle. This message is *included* in the current aggregation cycle.

Enforcing unique correlation keys

In some aggregation scenarios, you might want to enforce the condition that the correlation key is unique for each batch of exchanges. In other words, when the aggregate exchange for a particular correlation key completes, you want to make sure that no further aggregate exchanges with that correlation key are allowed to proceed. For example, you might want to enforce this condition, if the latter part of the route expects to process exchanges with unique correlation key values.

Depending on how the completion conditions are configured, there might be a risk of more than one aggregate exchange being generated with a particular correlation key. For example, although you might define a completion predicate that is designed to wait until *all* the exchanges with a particular correlation key are received, you might also define a completion timeout, which could fire before all of the exchanges with that key have arrived. In this case, the late-arriving exchanges could give rise to a *second* aggregate exchange with the same correlation key value.

For such scenarios, you can configure the aggregator to suppress aggregate exchanges that duplicate previous correlation key values, by setting the **closeCorrelationKeyOnCompletion** option. In order to suppress duplicate correlation key values, it is necessary for the aggregator to record previous correlation key values in a cache. The size of this cache (the number of cached correlation keys) is specified as an argument to the **closeCorrelationKeyOnCompletion()** DSL command. To specify a cache of unlimited size, you can pass a value of zero or a negative integer. For example, to specify a cache size of **10000** key values:

```
from("direct:start")
    .aggregate(header("UniqueBatchID"), new MyConcatenateStrategy())
        .completionSize(header("mySize"))
        .closeCorrelationKeyOnCompletion(10000)
    .to("mock:aggregated");
```

If an aggregate exchange completes with a duplicate correlation key value, the aggregator throws a **ClosedCorrelationKeyException** exception.

Grouped exchanges

You can combine all of the aggregated exchanges in an outgoing batch into a single **org.apache.camel.impl.GroupedExchange** holder class. To enable grouped exchanges, specify the **groupExchanges()** option, as shown in the following Java DSL route:

```
from("direct:start")
    .aggregate(header("StockSymbol"))
        .completionTimeout(3000)
        .groupExchanges()
    .to("mock:result");
```

The grouped exchange that is sent to **mock:result** contains the list of aggregated exchanges stored in the exchange property, **Exchange.GROUPED_EXCHANGE**. The following line of code shows how a subsequent processor can access the contents of the grouped exchange in the form of a list:

```
// Java
List<Exchange> grouped = ex.getProperty(Exchange.GROUPED_EXCHANGE,
List.class);
```



NOTE

When you enable the grouped exchanges feature, you *must not* configure an aggregation strategy (the grouped exchanges feature is itself an aggregation strategy).

Batch consumer

The aggregator can work together with the *batch consumer* pattern to aggregate the total number of messages reported by the batch consumer (a batch consumer endpoint sets the **CamelBatchSize**, **CamelBatchIndex**, and **CamelBatchComplete** properties on the incoming exchange). For example, to aggregate all of the files found by a File consumer endpoint, you could use a route like the following:

```
from("file://inbox")
    .aggregate(xpath("//order/@customerId"), new
AggregateCustomerOrderStrategy())
    .completionFromBatchConsumer()
    .to("bean:processOrder");
```

Currently, the following endpoints support the batch consumer mechanism: File, FTP, Mail, iBatis, and JPA.

Persistent aggregation repository


If you want pending aggregated exchanges to be stored persistently, you can use either the [HawtDB](#) component or the [SQL Component](#) for persistence support as a persistent aggregation repository. For example, if using HawtDB, you need to include a dependency on the **camel-hawtodb** component in your Maven POM. You can then configure a route to use the HawtDB aggregation repository as follows:

```
public void configure() throws Exception {
    HawtDBAggregationRepository repo = new AggregationRepository("repo1",
"target/data/hawtodb.dat");

    from("direct:start")
        .aggregate(header("id"), new UseLatestAggregationStrategy())
        .completionTimeout(3000)
        .aggregationRepository(repo)
        .to("mock:aggregated");
}
```

The HawtDB aggregation repository has a feature that enables it to recover and retry any failed exchanges (that is, any exchange that raised an exception while it was being processed by the latter part of the route). [Figure 7.7, “Recoverable Aggregation Repository”](#) shows an overview of the recovery mechanism.

Figure 7.7. Recoverable Aggregation Repository



The recovery mechanism works as follows:

1. The aggregator creates a dedicated recovery thread, which runs in the background, scanning the aggregation repository to find any failed exchanges.
2. Each failed exchange is checked to see whether its current redelivery count exceeds the maximum redelivery limit. If it is under the limit, the recovery task resubmits the exchange for processing in the latter part of the route.
3. If the current redelivery count is over the limit, the failed exchange is passed to the dead letter queue.

For more details about the HawtDB component, see .

Threading options

As shown in Figure 7.6, “Aggregator Implementation”, the aggregator is decoupled from the latter part of the route, where the exchanges sent to the latter part of the route are processed by a dedicated thread pool. By default, this pool contains just a single thread. If you want to specify a pool with multiple threads, enable the `parallelProcessing` option, as follows:

```
from("direct:start")
    .aggregate(header("id"), new UseLatestAggregationStrategy())
      .completionTimeout(3000)
      .parallelProcessing()
    .to("mock:aggregated");
```

By default, this creates a pool with 10 worker threads.

If you want to exercise more control over the created thread pool, specify a custom `java.util.concurrent.ExecutorService` instance using the `executorService` option (in which case it is unnecessary to enable the `parallelProcessing` option).

Aggregating into a List

A common aggregation scenario involves aggregating a series of incoming message bodies into a `List` object. To facilitate this scenario, Apache Camel provides the `AbstractListAggregationStrategy` abstract class, which you can quickly extend to create an aggregation strategy for this case. Incoming message bodies of type, `T`, are aggregated into a completed exchange, with a message body of type `List<T>`.

For example, to aggregate a series of `Integer` message bodies into a `List<Integer>` object, you could use an aggregation strategy defined as follows:

```
import
org.apache.camel.processor.aggregate.AbstractListAggregationStrategy;
...
/**
 * Strategy to aggregate integers into a List<Integer>.
 */
public final class MyListOfNumbersStrategy extends
AbstractListAggregationStrategy<Integer> {

    @Override
    public Integer getValue(Exchange exchange) {
        // the message body contains a number, so just return that as-is
```

```

    return exchange.getIn().getBody(Integer.class);
  }
}

```

Aggregator options

The aggregator supports the following options:

Table 7.3. Aggregator Options

Option	Default	Description
correlationExpression		Mandatory Expression which evaluates the correlation key to use for aggregation. The Exchange which has the same correlation key is aggregated together. If the correlation key could not be evaluated an Exception is thrown. You can disable this by using the ignoreBadCorrelationKeys option.
aggregationStrategy		Mandatory AggregationStrategy which is used to <i>merge</i> the incoming Exchange with the existing already merged exchanges. At first call the oldExchange parameter is null . On subsequent invocations the oldExchange contains the merged exchanges and newExchange is of course the new incoming Exchange. From Camel 2.9.2 onwards, the strategy can optionally be a TimeoutAwareAggregationStrategy implementation, which supports a timeout callback
strategyRef		A reference to lookup the AggregationStrategy in the Registry .

Option	Default	Description
completionSize		Number of messages aggregated before the aggregation is complete. This option can be set as either a fixed value or using an Expression which allows you to evaluate a size dynamically - will use Integer as result. If both are set Camel will fallback to use the fixed value if the Expression result was null or 0 .
completionTimeout		Time in millis that an aggregated exchange should be inactive before its complete. This option can be set as either a fixed value or using an Expression which allows you to evaluate a timeout dynamically - will use Long as result. If both are set Camel will fallback to use the fixed value if the Expression result was null or 0 . You cannot use this option together with <code>completionInterval</code> , only one of the two can be used.
completionInterval		A repeating period in millis by which the aggregator will complete all current aggregated exchanges. Camel has a background task which is triggered every period. You cannot use this option together with <code>completionTimeout</code> , only one of them can be used.
completionPredicate		A Predicate to indicate when an aggregated exchange is complete.

Option	Default	Description
completionFromBatchConsumer	false	This option is if the exchanges are coming from a Batch Consumer . Then when enabled the Aggregator will use the batch size determined by the Batch Consumer in the message header CamelBatchSize . See more details at Batch Consumer . This can be used to aggregate all files consumed from a File endpoint in that given poll.
eagerCheckCompletion	false	Whether or not to eager check for completion when a new incoming Exchange has been received. This option influences the behavior of the completionPredicate option as the Exchange being passed in changes accordingly. When false the Exchange passed in the Predicate is the <i>aggregated</i> Exchange which means any information you may store on the aggregated Exchange from the AggregationStrategy is available for the Predicate . When true the Exchange passed in the Predicate is the <i>incoming</i> Exchange, which means you can access data from the incoming Exchange.
forceCompletionOnStop	false	If true , complete all aggregated exchanges when the current route context is stopped.

Option	Default	Description
groupExchanges	false	<p>If enabled then Camel will group all aggregated Exchanges into a single combined</p> <p>org.apache.camel.impl.GroupedExchange holder class that holds all the aggregated Exchanges. And as a result only one Exchange is being sent out from the aggregator. Can be used to combine many incoming Exchanges into a single output Exchange without coding a custom AggregationStrategy yourself.</p>
ignoreInvalidCorrelationKeys	false	<p>Whether or not to ignore correlation keys which could not be evaluated to a value. By default Camel will throw an Exception, but you can enable this option and ignore the situation instead.</p>
closeCorrelationKeyOnCompletion		<p>Whether or not <i>late</i> Exchanges should be accepted or not. You can enable this to indicate that if a correlation key has already been completed, then any new exchanges with the same correlation key be denied. Camel will then throw a closedCorrelationKeyException exception. When using this option you pass in a integer which is a number for a LRUcache which keeps that last X number of closed correlation keys. You can pass in 0 or a negative value to indicate a unbounded cache. By passing in a number you are ensured that cache wont grown too big if you use a log of different correlation keys.</p>

Option	Default	Description
discardOnCompletionTimeout	false	Camel 2.5: Whether or not exchanges which complete due to a timeout should be discarded. If enabled, then when a timeout occurs the aggregated message will not be sent out but dropped (discarded).
aggregationRepository		Allows you to plug in you own implementation of org.apache.camel.spi.AggregationRepository which keeps track of the current inflight aggregated exchanges. Camel uses by default a memory based implementation.
aggregationRepositoryRef		Reference to lookup a aggregationRepository in the Registry .
parallelProcessing	false	When aggregated are completed they are being send out of the aggregator. This option indicates whether or not Camel should use a thread pool with multiple threads for concurrency. If no custom thread pool has been specified then Camel creates a default pool with 10 concurrent threads.
executorService		If using parallelProcessing you can specify a custom thread pool to be used. In fact also if you are not using parallelProcessing this custom thread pool is used to send out aggregated exchanges as well.
executorServiceRef		Reference to lookup a executorService in the Registry


Option	Default	Description
<code>timeoutCheckerExecutorsService</code>		If using one of the <code>completionTimeout</code> , <code>completionTimeoutExpression</code> , or <code>completionInterval</code> options, a background thread is created to check for the completion for every aggregator. Set this option to provide a custom thread pool to be used rather than creating a new thread for every aggregator.
<code>timeoutCheckerExecutorsServiceRef</code>		Reference to look up a <code>timeoutCheckerExecutorsService</code> in the registry.
<code>optimisticLocking</code>	<code>false</code>	Turns on optimistic locking, which can be used in combination with an aggregation repository.
<code>optimisticLockRetryPolicy</code>		Configures the retry policy for optimistic locking.

7.6. RESEQUENCER

Overview

The *resequencer* pattern, shown in [Figure 7.8, “Resequencer Pattern”](#), enables you to resequence messages according to a sequencing expression. Messages that generate a low value for the sequencing expression are moved to the front of the batch and messages that generate a high value are moved to the back.

Figure 7.8. Resequencer Pattern



Apache Camel supports two resequencing algorithms:

- *Batch resequencing* — Collects messages into a batch, sorts the messages and sends them to their output.

- *Stream resequencing* — Re-orders (continuous) message streams based on the detection of gaps between messages.

By default the resequencer does not support duplicate messages and will only keep the last message, in cases where a message arrives with the same message expression. However, in batch mode you can enable the resequencer to allow duplicates.

Batch resequencing

The batch resequencing algorithm is enabled by default. For example, to resequence a batch of incoming messages based on the value of a timestamp contained in the **TimeStamp** header, you can define the following route in Java DSL:

```
from("direct:start").resequence(header("TimeStamp")).to("mock:result");
```

By default, the batch is obtained by collecting all of the incoming messages that arrive in a time interval of 1000 milliseconds (default *batch timeout*), up to a maximum of 100 messages (default *batch size*). You can customize the values of the batch timeout and the batch size by appending a **batch()** DSL command, which takes a **BatchResequencerConfig** instance as its sole argument. For example, to modify the preceding route so that the batch consists of messages collected in a 4000 millisecond time window, up to a maximum of 300 messages, you can define the Java DSL route as follows:

```
import org.apache.camel.model.config.BatchResequencerConfig;

RouteBuilder builder = new RouteBuilder() {
    public void configure() {
        from("direct:start").resequence(header("TimeStamp")).batch(new
BatchResequencerConfig(300,4000L)).to("mock:result");
    }
};
```

You can also specify a batch resequencer pattern using XML configuration. The following example defines a batch resequencer with a batch size of 300 and a batch timeout of 4000 milliseconds:

```
<camelContext id="resequencerBatch"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start" />
    <resequence>
      <!--
        batch-config can be omitted for default (batch) resequencer
settings
      -->
      <batch-config batchSize="300" batchTimeout="4000" />
      <simple>header.TimeStamp</simple>
      <to uri="mock:result" />
    </resequence>
  </route>
</camelContext>
```

Batch options

Table 7.4, “Batch Resequencer Options” shows the options that are available in batch mode only.

Table 7.4. Batch Resequencer Options

Java DSL	XML DSL	Default	Description
<code>allowDuplicates()</code>	<code>batch-config/@allowDuplicates</code>	<code>false</code>	If <code>true</code> , do not discard duplicate messages from the batch (where <i>duplicate</i> means that the message expression evaluates to the same value).
<code>reverse()</code>	<code>batch-config/@reverse</code>	<code>false</code>	If <code>true</code> , put the messages in reverse order (where the default ordering applied to a message expression is based on Java's string lexical ordering, as defined by <code>String.compareTo()</code>).

For example, if you want to resequence messages from JMS queues based on **JMSPriority**, you would need to combine the options, **allowDuplicates** and **reverse**, as follows:

```
from("jms:queue:foo")
    // sort by JMSPriority by allowing duplicates (message can have
    // same JMSPriority)
    // and use reverse ordering so 9 is first output (most important),
    // and 0 is last
    // use batch mode and fire every 3th second

    .resequence(header("JMSPriority")).batch().timeout(3000).allowDuplicates()
    .reverse()
    .to("mock:result");
```

Stream resequencing

To enable the stream resequencing algorithm, you must append **stream()** to the **resequence()** DSL command. For example, to resequence incoming messages based on the value of a sequence number in the **seqnum** header, you define a DSL route as follows:

```
from("direct:start").resequence(header("seqnum")).stream().to("mock:result");
```

The stream-processing resequencer algorithm is based on the detection of gaps in a message stream, rather than on a fixed batch size. Gap detection, in combination with timeouts, removes the constraint of needing to know the number of messages of a

sequence (that is, the batch size) in advance. Messages must contain a unique sequence number for which a predecessor and a successor is known. For example a message with the sequence number **3** has a predecessor message with the sequence number **2** and a successor message with the sequence number **4**. The message sequence **2, 3, 5** has a gap because the successor of **3** is missing. The resequencer therefore must retain message **5** until message **4** arrives (or a timeout occurs).

By default, the stream resequencer is configured with a timeout of 1000 milliseconds, and a maximum message capacity of 100. To customize the stream's timeout and message capacity, you can pass a **StreamResequencerConfig** object as an argument to **stream()**. For example, to configure a stream resequencer with a message capacity of 5000 and a timeout of 4000 milliseconds, you define a route as follows:

```
// Java
import org.apache.camel.model.config.StreamResequencerConfig;

RouteBuilder builder = new RouteBuilder() {
    public void configure() {
        from("direct:start").resequence(header("seqnum")).
            stream(new StreamResequencerConfig(5000, 4000L)).
            to("mock:result");
    }
};
```

If the maximum time delay between successive messages (that is, messages with adjacent sequence numbers) in a message stream is known, the resequencer's timeout parameter should be set to this value. In this case, you can guarantee that all messages in the stream are delivered in the correct order to the next processor. The lower the timeout value that is compared to the out-of-sequence time difference, the more likely it is that the resequencer will deliver messages out of sequence. Large timeout values should be supported by sufficiently high capacity values, where the capacity parameter is used to prevent the resequencer from running out of memory.

If you want to use sequence numbers of some type other than **long**, you would must define a custom comparator, as follows:

```
// Java
ExpressionResultComparator<Exchange> comparator = new MyComparator();
StreamResequencerConfig config = new StreamResequencerConfig(5000, 4000L,
    comparator);
from("direct:start").resequence(header("seqnum")).stream(config).to("mock:
result");
```

You can also specify a stream resequencer pattern using XML configuration. The following example defines a stream resequencer with a message capacity of 5000 and a timeout of 4000 milliseconds:

```
<camelContext id="resequencerStream"
xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="direct:start"/>
        <resequence>
            <stream-config capacity="5000" timeout="4000"/>
            <simple>header.seqnum</simple>
            <to uri="mock:result" />
        </resequence>
    </route>
</camelContext>
```

```

    </resequence>
  </route>
</camelContext>

```

Ignore invalid exchanges

The resequencer EIP throws a **CamelExchangeException** exception, if the incoming exchange is not valid—that is, if the sequencing expression cannot be evaluated for some reason (for example, due to a missing header). You can use the **ignoreInvalidExchanges** option to ignore these exceptions, which means the resequencer will skip any invalid exchanges.

```

from("direct:start")
  .resequence(header("seqno")).batch().timeout(1000)
  // ignore invalid exchanges (they are discarded)
  .ignoreInvalidExchanges()
  .to("mock:result");

```

Reject old messages

The **rejectOld** option can be used to prevent messages being sent out of order, regardless of the mechanism used to resequence messages. When the **rejectOld** option is enabled, the resequencer rejects an incoming message (by throwing a **MessageRejectedException** exception), if the incoming messages is *older* (as defined by the current comparator) than the last delivered message.

```

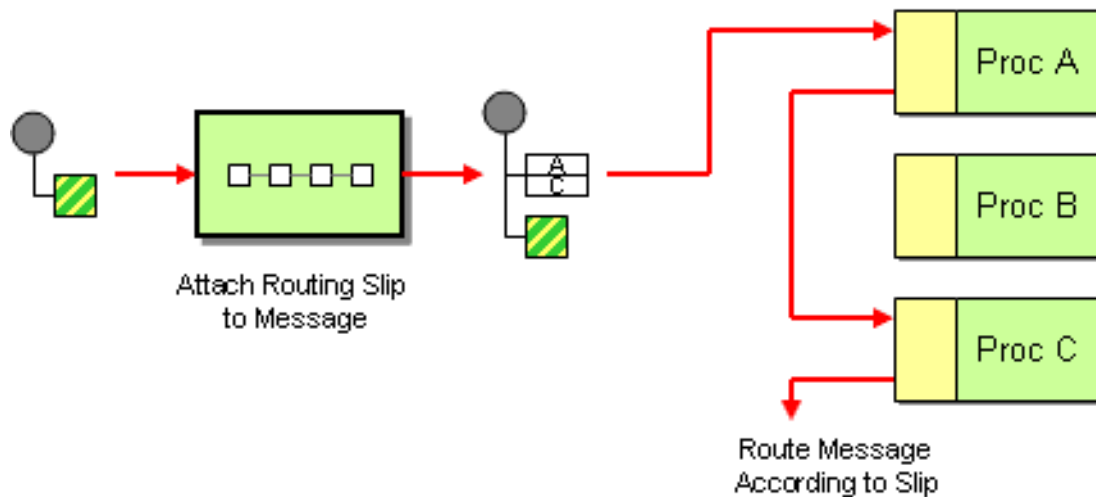
from("direct:start")
  .onException(MessageRejectedException.class).handled(true).to("mock:er
ror").end()
  .resequence(header("seqno")).stream().timeout(1000).rejectOld()
  .to("mock:result");

```

7.7. ROUTING SLIP

Overview

The *routing slip* pattern, shown in [Figure 7.9, “Routing Slip Pattern”](#), enables you to route a message consecutively through a series of processing steps, where the sequence of steps is not known at design time and can vary for each message. The list of endpoints through which the message should pass is stored in a header field (the *slip*), which Apache Camel reads at run time to construct a pipeline on the fly.

Figure 7.9. Routing Slip Pattern

The slip header

The routing slip appears in a user-defined header, where the header value is a comma-separated list of endpoint URIs. For example, a routing slip that specifies a sequence of security tasks—decrypting, authenticating, and de-duplicating a message—might look like the following:

```
cxf:bean:decrypt,cxf:bean:authenticate,cxf:bean:dedup
```

The current endpoint property

From Camel 2.5 the Routing Slip will set a property (**Exchange.SLIP_ENDPOINT**) on the exchange which contains the current endpoint as it advanced through the slip. This enables you to find out how far the exchange has progressed through the slip.

The [Routing Slip](#) will compute the slip *beforehand* which means, the slip is only computed once. If you need to compute the slip *on-the-fly* then use the [Dynamic Router](#) pattern instead.

Java DSL example

The following route takes messages from the **direct:a** endpoint and reads a routing slip from the **aRoutingSlipHeader** header:

```
from("direct:b").routingSlip("aRoutingSlipHeader");
```

You can specify the header name either as a string literal or as an expression.

You can also customize the URI delimiter using the two-argument form of **routingSlip()**. The following example defines a route that uses the **aRoutingSlipHeader** header key for the routing slip and uses the **#** character as the URI delimiter:

```
from("direct:c").routingSlip("aRoutingSlipHeader", "#");
```

XML configuration example

The following example shows how to configure the same route in XML:

```
<camelContext id="buildRoutingSlip"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:c"/>
    <routingSlip uriDelimiter="#">
      <headerName>aRoutingSlipHeader</headerName>
    </routingSlip>
  </route>
</camelContext>
```

Ignore invalid endpoints

The [Routing Slip](#) now supports **ignoreInvalidEndpoints**, which the [Recipient List](#) pattern also supports. You can use it to skip endpoints that are invalid. For example:

```
from("direct:a").routingSlip("myHeader").ignoreInvalidEndpoints();
```

In Spring XML, this feature is enabled by setting the **ignoreInvalidEndpoints** attribute on the **<routingSlip>** tag:

```
<route>
  <from uri="direct:a"/>
  <routingSlip ignoreInvalidEndpoints="true">
    <headerName>myHeader</headerName>
  </routingSlip>
</route>
```

Consider the case where **myHeader** contains the two endpoints, **direct:foo,xxx:bar**. The first endpoint is valid and works. The second is invalid and, therefore, ignored. Apache Camel logs at **INFO** level whenever an invalid endpoint is encountered.

Options

The **routingSlip** DSL command supports the following options:

Name	Default Value	Description
uriDelimiter	,	Delimiter used if the Expression returned multiple endpoints.
ignoreInvalidEndpoints	false	If an endpoint uri could not be resolved, should it be ignored. Otherwise Camel will throw an exception stating the endpoint uri is not valid.

7.8. THROTTLER

Overview

A *throttler* is a processor that limits the flow rate of incoming messages. You can use this pattern to protect a target endpoint from getting overloaded. In Apache Camel, you can implement the throttler pattern using the **throttle()** Java DSL command.

Video demo

There is a video demo of how to implement the throttler pattern at <http://vimeo.com/27592682>.

Java DSL example

To limit the flow rate to 100 messages per second, define a route as follows:

```
from("seda:a").throttle(100).to("seda:b");
```

If necessary, you can customize the time period that governs the flow rate using the **timePeriodMillis()** DSL command. For example, to limit the flow rate to 3 messages per 30000 milliseconds, define a route as follows:

```
from("seda:a").throttle(3).timePeriodMillis(30000).to("mock:result");
```

XML configuration example

The following example shows how to configure the preceding route in XML:

```
<camelContext id="throttleRoute"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <!-- throttle 3 messages per 30 sec -->
    <throttle timePeriodMillis="30000">
      <constant>3</constant>
      <to uri="mock:result"/>
    </throttle>
  </route>
</camelContext>
```

Dynamically changing maximum requests per period

Available as of Camel 2.8 Since we use an [Expression](#), you can adjust this value at runtime, for example you can provide a header with the value. At runtime Camel evaluates the expression and converts the result to a **java.lang.Long** type. In the example below we use a header from the message to determine the maximum requests per period. If the header is absent, then the [Throttler](#) uses the old value. So that allows you to only provide a header if the value is to be changed:

```
<camelContext id="throttleRoute"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:expressionHeader"/>
    <throttle timePeriodMillis="500">
```

```

        <!-- use a header to determine how many messages to throttle per 0.5
sec -->
        <header>throttleValue</header>
        <to uri="mock:result"/>
    </throttle>
</route>
</camelContext>

```

Asynchronous delaying

The throttler can enable *non-blocking asynchronous delaying*, which means that Apache Camel schedules a task to be executed in the future. The task is responsible for processing the latter part of the route (after the throttler). This allows the caller thread to unblock and service further incoming messages. For example:

```
from("seda:a").throttle(100).asyncDelayed().to("seda:b");
```

Options

The **throttle** DSL command supports the following options:

Name	Default Value	Description
maximumRequestsPerPeriod		Maximum number of requests per period to throttle. This option must be provided and a positive number. Notice, in the XML DSL, from Camel 2.8 onwards this option is configured using an Expression instead of an attribute.
timePeriodMillis	1000	The time period in millis, in which the throttler will allow at most maximumRequestsPerPeriod number of messages.
asyncDelayed	false	Camel 2.4: If enabled then any messages which is delayed happens asynchronously using a scheduled thread pool.
executorServiceRef		Camel 2.4: Refers to a custom Thread Pool to be used if asyncDelay has been enabled.

<code>callerRunsWhenRejected</code>	<code>true</code>	Camel 2.4: Is used if asyncDelayed was enabled. This controls if the caller thread should execute the task if the thread pool rejected the task.
-------------------------------------	-------------------	--

7.9. DELAYER

Overview

A *delayer* is a processor that enables you to apply a *relative* time delay to incoming messages.

Java DSL example

You can use the **delay()** command to add a *relative* time delay, in units of milliseconds, to incoming messages. For example, the following route delays all incoming messages by 2 seconds:

```
from("seda:a").delay(2000).to("mock:result");
```

Alternatively, you can specify the time delay using an expression:

```
from("seda:a").delay(header("MyDelay")).to("mock:result");
```

The DSL commands that follow **delay()** are interpreted as sub-clauses of **delay()**. Hence, in some contexts it is necessary to terminate the sub-clauses of **delay()** by inserting the **end()** command. For example, when **delay()** appears inside an **onException()** clause, you would terminate it as follows:

```
from("direct:start")
    .onException(Exception.class)
        .maximumRedeliveries(2)
        .backOffMultiplier(1.5)
        .handled(true)
        .delay(1000)
            .log("Halting for some time")
            .to("mock:halt")
        .end()
    .end()
.to("mock:result");
```

XML configuration example

The following example demonstrates the delay in XML DSL:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <delay>
```



```

        <header>MyDelay</header>
    </delay>
    <to uri="mock:result"/>
</route>
<route>
    <from uri="seda:b"/>
    <delay>
        <constant>1000</constant>
    </delay>
    <to uri="mock:result"/>
</route>
</camelContext>

```

Creating a custom delay

You can use an expression combined with a bean to determine the delay as follows:

```

from("activemq:foo").
    delay().expression().method("someBean", "computeDelay").
    to("activemq:bar");

```

Where the bean class could be defined as follows:

```

public class SomeBean {
    public long computeDelay() {
        long delay = 0;
        // use java code to compute a delay value in millis
        return delay;
    }
}

```

Asynchronous delaying

You can let the delayer use *non-blocking asynchronous delaying*, which means that Apache Camel schedules a task to be executed in the future. The task is responsible for processing the latter part of the route (after the delayer). This allows the caller thread to unblock and service further incoming messages. For example:

```

from("activemq:queue:foo")
    .delay(1000)
    .asyncDelayed()
    .to("activemq:aDelayedQueue");

```

The same route can be written in the XML DSL, as follows:

```

<route>
    <from uri="activemq:queue:foo"/>
    <delay asyncDelayed="true">
        <constant>1000</constant>
    </delay>
    <to uri="activemq:aDealyedQueue"/>
</route>

```

Options

The `delayer` pattern supports the following options:

Name	Default Value	Description
<code>asyncDelayed</code>	<code>false</code>	Camel 2.4: If enabled then delayed messages happens asynchronously using a scheduled thread pool.
<code>executorServiceRef</code>		Camel 2.4: Refers to a custom Thread Pool to be used if <code>asyncDelay</code> has been enabled.
<code>callerRunsWhenRejected</code>	<code>true</code>	Camel 2.4: Is used if <code>asyncDelayed</code> was enabled. This controls if the caller thread should execute the task if the thread pool rejected the task.

7.10. LOAD BALANCER

Overview

The *load balancer* pattern allows you to delegate message processing to one of several endpoints, using a variety of different load-balancing policies.

Java DSL example

The following route distributes incoming messages between the target endpoints, `mock:x`, `mock:y`, `mock:z`, using a round robin load-balancing policy:

```
from("direct:start").loadBalance().roundRobin().to("mock:x", "mock:y",
"mock:z");
```

XML configuration example

The following example shows how to configure the same route in XML:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <loadBalance>
      <roundRobin/>
      <to uri="mock:x"/>
      <to uri="mock:y"/>
      <to uri="mock:z"/>
    </loadBalance>
  </route>
</camelContext>
```

```

    </loadBalance>
  </route>
</camelContext>

```

Load-balancing policies

The Apache Camel load balancer supports the following load-balancing policies:

- [Round robin](#)
- [Random](#)
- [Sticky](#)
- [Topic](#)
- [the section called “Failover”](#)
- [the section called “Weighted round robin and weighted random”](#)
- [the section called “Custom Load Balancer”](#)

Round robin

The round robin load-balancing policy cycles through all of the target endpoints, sending each incoming message to the next endpoint in the cycle. For example, if the list of target endpoints is, **mock:x**, **mock:y**, **mock:z**, then the incoming messages are sent to the following sequence of endpoints: **mock:x**, **mock:y**, **mock:z**, **mock:x**, **mock:y**, **mock:z**, and so on.

You can specify the round robin load-balancing policy in Java DSL, as follows:

```

from("direct:start").loadBalance().roundRobin().to("mock:x", "mock:y",
"mock:z");

```

Alternatively, you can configure the same route in XML, as follows:

```

<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <loadBalance>
      <roundRobin/>
      <to uri="mock:x"/>
      <to uri="mock:y"/>
      <to uri="mock:z"/>
    </loadBalance>
  </route>
</camelContext>

```

Random

The random load-balancing policy chooses the target endpoint randomly from the specified list.

You can specify the random load-balancing policy in Java DSL, as follows:

```
from("direct:start").loadBalance().random().to("mock:x", "mock:y",
"mock:z");
```

Alternatively, you can configure the same route in XML, as follows:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <loadBalance>
      <random/>
      <to uri="mock:x"/>
      <to uri="mock:y"/>
      <to uri="mock:z"/>
    </loadBalance>
  </route>
</camelContext>
```

Sticky

The sticky load-balancing policy directs the *In* message to an endpoint that is chosen by calculating a hash value from a specified expression. The advantage of this load-balancing policy is that expressions of the same value are always sent to the same server. For example, by calculating the hash value from a header that contains a username, you ensure that messages from a particular user are always sent to the same target endpoint. Another useful approach is to specify an expression that extracts the session ID from an incoming message. This ensures that all messages belonging to the same session are sent to the same target endpoint.

You can specify the sticky load-balancing policy in Java DSL, as follows:

```
from("direct:start").loadBalance().sticky(header("username")).to("mock:x",
"mock:y", "mock:z");
```

Alternatively, you can configure the same route in XML, as follows:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <loadBalance>
      <sticky>
        <expression>
          <simple>header.username</simple>
        </expression>
      </sticky>
      <to uri="mock:x"/>
      <to uri="mock:y"/>
      <to uri="mock:z"/>
    </loadBalance>
  </route>
</camelContext>
```

Topic

The topic load-balancing policy sends a copy of each *In* message to *all* of the listed destination endpoints (effectively broadcasting the message to all of the destinations, like a JMS topic).

You can use the Java DSL to specify the topic load-balancing policy, as follows:

```
from("direct:start").loadBalance().topic().to("mock:x", "mock:y",
"mock:z");
```

Alternatively, you can configure the same route in XML, as follows:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <loadBalance>
      <topic/>
      <to uri="mock:x"/>
      <to uri="mock:y"/>
      <to uri="mock:z"/>
    </loadBalance>
  </route>
</camelContext>
```

Failover

Available as of Apache Camel 2.0 The **failover** load balancer is capable of trying the next processor in case an Exchange failed with an **exception** during processing. You can configure the **failover** with a list of specific exceptions that trigger failover. If you do not specify any exceptions, failover is triggered by any exception. The failover load balancer uses the same strategy for matching exceptions as the **onException** exception clause.



ENABLE STREAM CACHING IF USING STREAMS

If you use streaming, you should enable [Stream Caching](#) when using the failover load balancer. This is needed so the stream can be re-read when failing over.

The **failover** load balancer supports the following options:

Option	Type	Default	Description
--------	------	---------	-------------

inheritErrorHandler	boolean	true	<p>Camel 2.3: Specifies whether to use the errorHandler configured on the route. If you want to fail over immediately to the next endpoint, you should disable this option (value of false). If you enable this option, Apache Camel will first attempt to process the message using the errorHandler.</p> <p>For example, the errorHandler might be configured to redeliver messages and use delays between attempts. Apache Camel will initially try to redeliver to the <i>original</i> endpoint, and only fail over to the next endpoint when the errorHandler is exhausted.</p>
maximumFailoverAttempts	int	-1	<p>Camel 2.3: Specifies the maximum number of attempts to fail over to a new endpoint. The value, 0, implies that <i>no</i> failover attempts are made and the value, -1, implies an infinite number of failover attempts.</p>

<code>roundRobin</code>	<code>boolean</code>	<code>false</code>	Camel 2.3: Specifies whether the failover load balancer should operate in round robin mode or not. If not, it will <i>always</i> start from the first endpoint when a new message is to be processed. In other words it restarts from the top for every message. If round robin is enabled, it keeps state and continues with the next endpoint in a round robin fashion. When using round robin it will not <i>stick</i> to last known good endpoint, it will always pick the next endpoint to use.
-------------------------	----------------------	--------------------	--

The following example is configured to fail over, only if an **IOException** exception is thrown:

```
from("direct:start")
    // here we will load balance if IOException was thrown
    // any other kind of exception will result in the Exchange as failed
    // to failover over any kind of exception we can just omit the
exception
    // in the failOver DSL
    .loadBalance().failover(IOException.class)
    .to("direct:x", "direct:y", "direct:z");
```

You can optionally specify multiple exceptions to fail over, as follows:

```
// enable redelivery so failover can react
errorHandler(defaultErrorHandler().maximumRedeliveries(5));

from("direct:foo")
    .loadBalance()
    .failover(IOException.class, MyOtherException.class)
    .to("direct:a", "direct:b");
```

You can configure the same route in XML, as follows:

```
<route errorHandlerRef="myErrorHandler">
  <from uri="direct:foo"/>
  <loadBalance>
    <failover>
      <exception>java.io.IOException</exception>
      <exception>com.mycompany.MyOtherException</exception>
    </failover>
```

```

        <to uri="direct:a"/>
        <to uri="direct:b"/>
    </loadBalance>
</route>

```

The following example shows how to fail over in round robin mode:

```

from("direct:start")
    // Use failover load balancer in stateful round robin mode,
    // which means it will fail over immediately in case of an exception
    // as it does NOT inherit error handler. It will also keep retrying,
as
    // it is configured to retry indefinitely.
    .loadBalance().failover(-1, false, true)
    .to("direct:bad", "direct:bad2", "direct:good", "direct:good2");

```

You can configure the same route in XML, as follows:

```

<route>
  <from uri="direct:start"/>
  <loadBalance>
    <!-- failover using stateful round robin,
    which will keep retrying the 4 endpoints indefinitely.
    You can set the maximumFailoverAttempt to break out after X
attempts -->
    <failover roundRobin="true"/>
    <to uri="direct:bad"/>
    <to uri="direct:bad2"/>
    <to uri="direct:good"/>
    <to uri="direct:good2"/>
  </loadBalance>
</route>

```

Weighted round robin and weighted random

In many enterprise environments, where server nodes of unequal processing power are hosting services, it is usually preferable to distribute the load in accordance with the individual server processing capacities. A *weighted round robin* algorithm or a *weighted random* algorithm can be used to address this problem.

The weighted load balancing policy allows you to specify a processing load *distribution ratio* for each server with respect to the others. You can specify this value as a positive processing weight for each server. A larger number indicates that the server can handle a larger load. The processing weight is used to determine the payload distribution ratio of each processing endpoint with respect to the others.

The parameters that can be used are

Table 7.5. Weighted Options

Option	Type	Default	Description
--------	------	---------	-------------

roundRobin	boolean	false	The default value for round-robin is false . In the absence of this setting or parameter, the load-balancing algorithm used is random.
distributionRatioDelimiter	String	,	The distributionRatioDelimiter is the delimiter used to specify the distributionRatio . If this attribute is not specified, comma , is the default delimiter.

The following Java DSL examples show how to define a weighted round-robin route and a weighted random route:

```
// Java
// round-robin
from("direct:start")
  .loadBalance().weighted(true, "4:2:1" distributionRatioDelimiter=":")
  .to("mock:x", "mock:y", "mock:z");

//random
from("direct:start")
  .loadBalance().weighted(false, "4,2,1")
  .to("mock:x", "mock:y", "mock:z");
```

You can configure the round-robin route in XML, as follows:

```
<!-- round-robin -->
<route>
  <from uri="direct:start"/>
  <loadBalance>
    <weighted roundRobin="true" distributionRatio="4:2:1"
distributionRatioDelimiter=":" />
    <to uri="mock:x"/>
    <to uri="mock:y"/>
    <to uri="mock:z"/>
  </loadBalance>
</route>
```

Custom Load Balancer

You can use a custom load balancer (eg your own implementation) also.

An example using Java DSL:

```

from("direct:start")
  // using our custom load balancer
  .loadBalance(new MyLoadBalancer())
  .to("mock:x", "mock:y", "mock:z");

```

And the same example using XML DSL:

```

<!-- this is the implementation of our custom load balancer -->
<bean id="myBalancer"
class="org.apache.camel.processor.CustomLoadBalanceTest$MyLoadBalancer"/>

<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <loadBalance>
      <!-- refer to my custom load balancer -->
      <custom ref="myBalancer"/>
      <!-- these are the endpoints to balancer -->
      <to uri="mock:x"/>
      <to uri="mock:y"/>
      <to uri="mock:z"/>
    </loadBalance>
  </route>
</camelContext>

```

Notice in the XML DSL above we use `<custom>` which is only available in **Camel 2.8** onwards. In older releases you would have to do as follows instead:

```

<loadBalance ref="myBalancer">
  <!-- these are the endpoints to balancer -->
  <to uri="mock:x"/>
  <to uri="mock:y"/>
  <to uri="mock:z"/>
</loadBalance>

```

To implement a custom load balancer you can extend some support classes such as **LoadBalancerSupport** and **SimpleLoadBalancerSupport**. The former supports the asynchronous routing engine, and the latter does not. Here is an example:

```

public static class MyLoadBalancer extends LoadBalancerSupport {

  public boolean process(Exchange exchange, AsyncCallback callback) {
    String body = exchange.getIn().getBody(String.class);
    try {
      if ("x".equals(body)) {
        getProcessors().get(0).process(exchange);
      } else if ("y".equals(body)) {
        getProcessors().get(1).process(exchange);
      } else {
        getProcessors().get(2).process(exchange);
      }
    } catch (Throwable e) {
      exchange.setException(e);
    }
    callback.done(true);
  }
}

```

```

    return true;
}
}

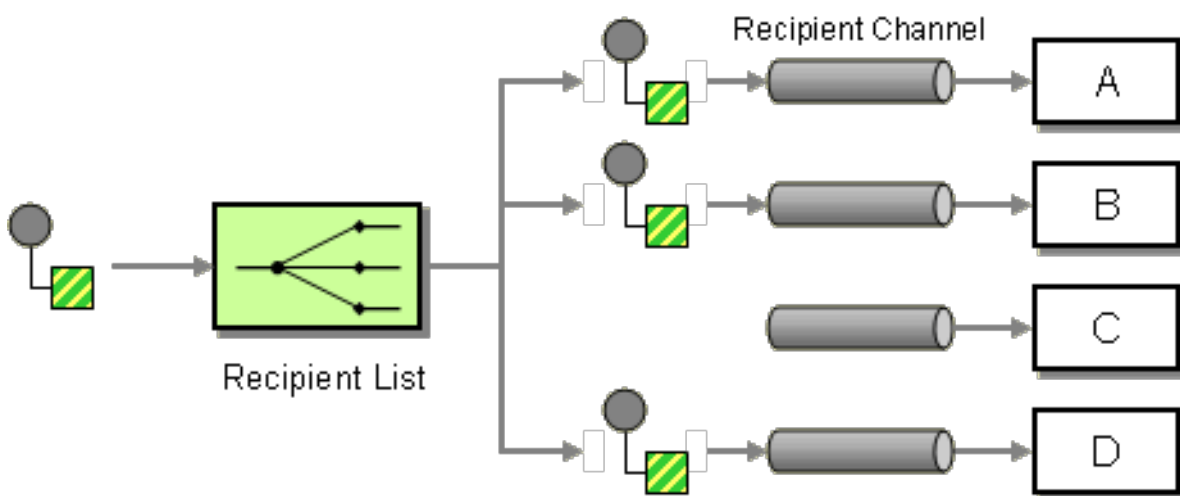
```

7.11. MULTICAST

Overview

The *multicast* pattern, shown in [Figure 7.10, “Multicast Pattern”](#), is a variation of the [recipient list](#) with a fixed destination pattern, which is compatible with the *Out* message exchange pattern. This is in contrast to recipient list, which is only compatible with the *InOnly* exchange pattern.

Figure 7.10. Multicast Pattern



Multicast with a custom aggregation strategy

Whereas the multicast processor receives multiple *Out* messages in response to the original request (one from each of the recipients), the original caller is only expecting to receive a *single* reply. Thus, there is an inherent mismatch on the reply leg of the message exchange, and to overcome this mismatch, you must provide a custom *aggregation strategy* to the multicast processor. The aggregation strategy class is responsible for aggregating all of the *Out* messages into a single reply message.

Consider the example of an electronic auction service, where a seller offers an item for sale to a list of buyers. The buyers each put in a bid for the item, and the seller automatically selects the bid with the highest price. You can implement the logic for distributing an offer to a fixed list of buyers using the `multicast()` DSL command, as follows:

```

from("cxf:bean:offer").multicast(new HighestBidAggregationStrategy()).
    to("cxf:bean:Buyer1", "cxf:bean:Buyer2", "cxf:bean:Buyer3");

```

Where the seller is represented by the endpoint, `cxf:bean:offer`, and the buyers are represented by the endpoints, `cxf:bean:Buyer1`, `cxf:bean:Buyer2`, `cxf:bean:Buyer3`. To consolidate the bids received from the various buyers, the multicast processor uses the aggregation strategy, `HighestBidAggregationStrategy`. You can implement the `HighestBidAggregationStrategy` in Java, as follows:

```

// Java

```

```
import org.apache.camel.processor.aggregate.AggregationStrategy;
import org.apache.camel.Exchange;

public class HighestBidAggregationStrategy implements AggregationStrategy
{
    public Exchange aggregate(Exchange oldExchange, Exchange newExchange)
    {
        float oldBid = oldExchange.getOut().getHeader("Bid", Float.class);
        float newBid = newExchange.getOut().getHeader("Bid", Float.class);
        return (newBid > oldBid) ? newExchange : oldExchange;
    }
}
```

Where it is assumed that the buyers insert the bid price into a header named, **Bid**. For more details about custom aggregation strategies, see [Section 7.5, “Aggregator”](#).

Parallel processing

By default, the multicast processor invokes each of the recipient endpoints one after another (in the order listed in the **to()** command). In some cases, this might cause unacceptably long latency. To avoid these long latency times, you have the option of enabling parallel processing by adding the **parallelProcessing()** clause. For example, to enable parallel processing in the electronic auction example, define the route as follows:

```
from("cxf:bean:offer")
    .multicast(new HighestBidAggregationStrategy())
    .parallelProcessing()
    .to("cxf:bean:Buyer1", "cxf:bean:Buyer2", "cxf:bean:Buyer3");
```

Where the multicast processor now invokes the buyer endpoints, using a thread pool that has one thread for each of the endpoints.

If you want to customize the size of the thread pool that invokes the buyer endpoints, you can invoke the **executorService()** method to specify your own custom executor service. For example:

```
from("cxf:bean:offer")
    .multicast(new HighestBidAggregationStrategy())
    .executorService(MyExecutor)
    .to("cxf:bean:Buyer1", "cxf:bean:Buyer2", "cxf:bean:Buyer3");
```

Where *MyExecutor* is an instance of [java.util.concurrent.ExecutorService](#) type.

When the exchange has an *InOut* pattern, an aggregation strategy is used to aggregate reply messages. The default aggregation strategy takes the latest reply message and discards earlier replies. For example, in the following route, the custom strategy, **MyAggregationStrategy**, is used to aggregate the replies from the endpoints, **direct:a**, **direct:b**, and **direct:c**:

```
from("direct:start")
    .multicast(new MyAggregationStrategy())
    .parallelProcessing()
    .timeout(500)
```

```

        .to("direct:a", "direct:b", "direct:c")
    .end()
    .to("mock:result");

```

XML configuration example

The following example shows how to configure a similar route in XML, where the route uses a custom aggregation strategy and a custom thread executor:

```

<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xsi:schemaLocation="
           http://www.springframework.org/schema/beans
           http://www.springframework.org/schema/beans/spring-beans-3.0.xsd
           http://camel.apache.org/schema/spring
           http://camel.apache.org/schema/spring/camel-spring.xsd"
       ">

    <camelContext xmlns="http://camel.apache.org/schema/spring">
        <route>
            <from uri="cxf:bean:offer"/>
            <multicast strategyRef="highestBidAggregationStrategy"
                    parallelProcessing="true"
                    threadPoolRef="myThreadExcutor">
                <to uri="cxf:bean:Buyer1"/>
                <to uri="cxf:bean:Buyer2"/>
                <to uri="cxf:bean:Buyer3"/>
            </multicast>
        </route>
    </camelContext>

    <bean id="highestBidAggregationStrategy"
          class="com.acme.example.HighestBidAggregationStrategy"/>
    <bean id="myThreadExcutor" class="com.acme.example.MyThreadExcutor"/>

</beans>

```

Where both the **parallelProcessing** attribute and the **threadPoolRef** attribute are optional. It is only necessary to set them if you want to customize the threading behavior of the multicast processor.

Apply custom processing to the outgoing messages

Before multicast sends a message to one of the recipient endpoints, it creates a message replica, which is a shallow copy of the original message. If you want to perform some custom processing on each message replica before the replica is sent to its endpoint, you can invoke the **onPrepare** DSL command in the **multicast** clause. The **onPrepare** command inserts a custom processor just *after* the message has been shallow-copied and just *before* the message is dispatched to its endpoint. For example, in the following route, the **CustomProc** processor is invoked on the message sent to **direct:a** and the **CustomProc** processor is also invoked on the message sent to **direct:b**.

```

from("direct:start")
    .multicast().onPrepare(new CustomProc())

```

```
.to("direct:a").to("direct:b");
```

A common use case for the **onPrepare** DSL command is to perform a deep copy of some or all elements of a message. For example, the following **CustomProc** processor class performs a deep copy of the message body, where the message body is presumed to be of type, **BodyType**, and the deep copy is performed by the method **BodyType.deepCopy()**.

```
// Java
import org.apache.camel.*;
...
public class CustomProc implements Processor {

    public void process(Exchange exchange) throws Exception {
        BodyType body = exchange.getIn().getBody(BodyType.class);

        // Make a _deep_ copy of of the body object
        BodyType clone = BodyType.deepCopy();
        exchange.getIn().setBody(clone);

        // Headers and attachments have already been
        // shallow-copied. If you need deep copies,
        // add some more code here.
    }
}
```



NOTE

Although the **multicast** syntax allows you to invoke the **process** DSL command in the **multicast** clause, this does not make sense semantically and it does *not* have the same effect as **onPrepare** (in fact, in this context, the **process** DSL command has no effect).

Using onPrepare to execute custom logic when preparing messages

The **Multicast** will copy the source **Exchange** and multicast each copy. However the copy is a shallow copy, so in case you have mutable message bodies, then any changes will be visible by the other copied messages. If you want to use a deep clone copy then you need to use a custom **onPrepare** which allows you to do this using the **Processor** interface.

Notice the **onPrepare** can be used for any kind of custom logic which you would like to execute before the **Exchange** is being multicasted.



NOTE

Its best practice to design for immutable objects.

For example if you have a mutable message body as this **Animal** class:

```
public class Animal implements Serializable {

    private int id;
    private String name;

    public Animal() {
```

```

    }

    public Animal(int id, String name) {
        this.id = id;
        this.name = name;
    }

    public Animal deepClone() {
        Animal clone = new Animal();
        clone.setId(getId());
        clone.setName(getName());
        return clone;
    }

    public int getId() {
        return id;
    }

    public void setId(int id) {
        this.id = id;
    }

    public String getName() {
        return name;
    }

    public void setName(String name) {
        this.name = name;
    }

    @Override
    public String toString() {
        return id + " " + name;
    }
}

```

Then we can create a deep clone processor which clones the message body:

```

public class AnimalDeepClonePrepare implements Processor {

    public void process(Exchange exchange) throws Exception {
        Animal body = exchange.getIn().getBody(Animal.class);

        // do a deep clone of the body which wont affect when doing
multicasting
        Animal clone = body.deepClone();
        exchange.getIn().setBody(clone);
    }
}

```

Then we can use the `AnimalDeepClonePrepare` class in the `Multicast` route using the `onPrepare` option as shown:

```
from("direct:start")
    .multicast().onPrepare(new
AnimalDeepClonePrepare()).to("direct:a").to("direct:b");
```

And the same example in XML DSL

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <!-- use on prepare with multicast -->
    <multicast onPrepareRef="animalDeepClonePrepare">
      <to uri="direct:a"/>
      <to uri="direct:b"/>
    </multicast>
  </route>

  <route>
    <from uri="direct:a"/>
    <process ref="processorA"/>
    <to uri="mock:a"/>
  </route>
  <route>
    <from uri="direct:b"/>
    <process ref="processorB"/>
    <to uri="mock:b"/>
  </route>
</camelContext>

<!-- the on prepare Processor which performs the deep cloning -->
<bean id="animalDeepClonePrepare"
class="org.apache.camel.processor.AnimalDeepClonePrepare"/>

<!-- processors used for the last two routes, as part of unit test -->
<bean id="processorA"
class="org.apache.camel.processor.MulticastOnPrepareTest$ProcessorA"/>
<bean id="processorB"
class="org.apache.camel.processor.MulticastOnPrepareTest$ProcessorB"/>
```

Options

The **multicast** DSL command supports the following options:

Name	Default Value	Description
strategyRef		Refers to an AggregationStrategy to be used to assemble the replies from the multicasts, into a single outgoing message from the Multicast . By default Camel will use the last reply as the outgoing message.

parallelProcessing	false	If enables then sending messages to the multicasts occurs concurrently. Note the caller thread will still wait until all messages has been fully processed, before it continues. Its only the sending and processing the replies from the multicasts which happens concurrently.
executorServiceRef		Refers to a custom Thread Pool to be used for parallel processing. Notice if you set this option, then parallel processing is automatic implied, and you do not have to enable that option as well.
stopOnException	false	Camel 2.2: Whether or not to stop continue processing immediately when an exception occurred. If disable, then Camel will send the message to all multicasts regardless if one of them failed. You can deal with exceptions in the AggregationStrategy class where you have full control how to handle that.
streaming	false	If enabled then Camel will process replies out-of-order, eg in the order they come back. If disabled, Camel will process replies in the same order as multicasted.
timeout		Camel 2.5: Sets a total timeout specified in millis. If the Multicast hasn't been able to send and process all replies within the given timeframe, then the timeout triggers and the Multicast breaks out and continues. Notice if you provide a TimeoutAwareAggregationStrategy then the timeout method is invoked before breaking out.


onPrepareRef		Camel 2.8: Refers to a custom Processor to prepare the copy of the Exchange each multicast will receive. This allows you to do any custom logic, such as deep-cloning the message payload if that's needed etc.
shareUnitOfWork	false	Camel 2.8: Whether the unit of work should be shared. See the same option on Splitter for more details.

7.12. COMPOSED MESSAGE PROCESSOR

Composed Message Processor

The *composed message processor* pattern, as shown in [Figure 7.11, “Composed Message Processor Pattern”](#), allows you to process a composite message by splitting it up, routing the sub-messages to appropriate destinations, and then re-aggregating the responses back into a single message.

Figure 7.11. Composed Message Processor Pattern



Java DSL example

The following example checks that a multipart order can be filled, where each part of the order requires a check to be made at a different inventory:

```
// split up the order so individual OrderItems can be validated by the
// appropriate bean
from("direct:start")
    .split().body()
    .choice()
        .when().method("orderItemHelper", "isWidget")
            .to("bean:widgetInventory")
        .otherwise()
            .to("bean:gadgetInventory")
```

```

        .end()
        .to("seda:aggregate");

// collect and re-assemble the validated OrderItems into an order again
from("seda:aggregate")
    .aggregate(new MyOrderAggregationStrategy())
    .header("orderId")
    .completionTimeout(1000L)
    .to("mock:result");

```

XML DSL example

The preceding route can also be written in XML DSL, as follows:

```

<route>
  <from uri="direct:start"/>
  <split>
    <simple>body</simple>
    <choice>
      <when>
        <method bean="orderItemHelper" method="isWidget"/>
      <to uri="bean:widgetInventory"/>
      </when>
      <otherwise>
      <to uri="bean:gadgetInventory"/>
      </otherwise>
    </choice>
    <to uri="seda:aggregate"/>
  </split>
</route>

<route>
  <from uri="seda:aggregate"/>
  <aggregate strategyRef="myOrderAggregatorStrategy"
completionTimeout="1000">
    <correlationExpression>
      <simple>header.orderId</simple>
    </correlationExpression>
    <to uri="mock:result"/>
  </aggregate>
</route>

```

Processing steps

Processing starts by splitting the order, using a [Splitter](#). The [Splitter](#) then sends individual **OrderItems** to a [Content Based Router](#), which routes messages based on the item type. *Widget* items get sent for checking in the **widgetInventory** bean and *gadget* items get sent to the **gadgetInventory** bean. Once these **OrderItems** have been validated by the appropriate bean, they are sent on to the [Aggregator](#) which collects and re-assembles the validated **OrderItems** into an order again.

Each received order has a header containing an *order ID*. We make use of the order ID during the aggregation step: the `.header("orderId")` qualifier on the `aggregate()` DSL command instructs the aggregator to use the header with the key, **orderId**, as the correlation expression.

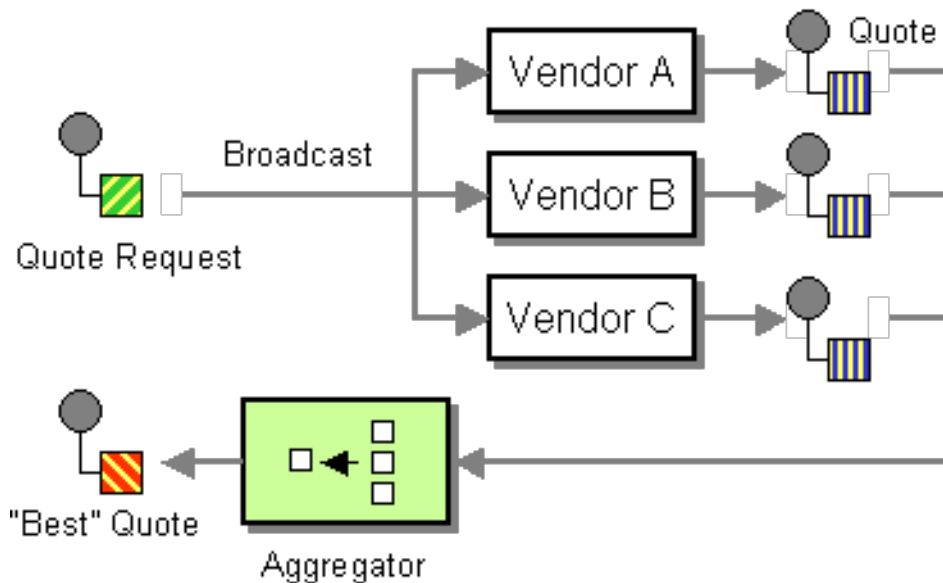
For full details, check the example source here:

7.13. SCATTER-GATHER

Scatter-Gather

The *scatter-gather pattern*, as shown in [Figure 7.12, “Scatter-Gather Pattern”](#), enables you to route messages to a number of dynamically specified recipients and re-aggregate the responses back into a single message.

Figure 7.12. Scatter-Gather Pattern



Dynamic scatter-gather example

The following example outlines an application that gets the best quote for beer from several different vendors. The examples uses a dynamic [Recipient List](#) to request a quote from all vendors and an [Aggregator](#) to pick the best quote out of all the responses. The routes for this application are defined as follows:

```
<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <recipientList>
      <header>listOfVendors</header>
    </recipientList>
  </route>
  <route>
    <from uri="seda:quoteAggregator"/>
    <aggregate strategyRef="aggregatorStrategy" completionTimeout="1000">
      <correlationExpression>
        <header>quoteRequestId</header>
      </correlationExpression>
      <to uri="mock:result"/>
    </aggregate>
  </route>
</camelContext>
```

In the first route, the [Recipient List](#) looks at the `listOfVendors` header to obtain the list of

recipients. Hence, the client that sends messages to this application needs to add a **listOfVendors** header to the message. [Example 7.1, “Messaging Client Sample”](#) shows some sample code from a messaging client that adds the relevant header data to outgoing messages.

Example 7.1. Messaging Client Sample

```
Map<String, Object> headers = new HashMap<String, Object>();
headers.put("listOfVendors", "bean:vendor1, bean:vendor2,
bean:vendor3");
headers.put("quoteRequestId", "quoteRequest-1");
template.sendBodyAndHeaders("direct:start", "<quote_request
item=\"beer\"/>", headers);
```

The message would be distributed to the following endpoints: **bean:vendor1**, **bean:vendor2**, and **bean:vendor3**. These beans are all implemented by the following class:

```
public class MyVendor {
    private int beerPrice;

    @Produce(uri = "seda:quoteAggregator")
    private ProducerTemplate quoteAggregator;

    public MyVendor(int beerPrice) {
        this.beerPrice = beerPrice;
    }

    public void getQuote(@XPath("/quote_request/@item") String item,
Exchange exchange) throws Exception {
        if ("beer".equals(item)) {
            exchange.getIn().setBody(beerPrice);
            quoteAggregator.send(exchange);
        } else {
            throw new Exception("No quote available for " + item);
        }
    }
}
```

The bean instances, **vendor1**, **vendor2**, and **vendor3**, are instantiated using Spring XML syntax, as follows:

```
<bean id="aggregatorStrategy"
class="org.apache.camel.spring.processor.scattergather.LowestQuoteAggregat
ionStrategy"/>

<bean id="vendor1"
class="org.apache.camel.spring.processor.scattergather.MyVendor">
    <constructor-arg>
        <value>1</value>
    </constructor-arg>
</bean>

<bean id="vendor2"
class="org.apache.camel.spring.processor.scattergather.MyVendor">
```

```

    <constructor-arg>
      <value>2</value>
    </constructor-arg>
  </bean>

  <bean id="vendor3"
    class="org.apache.camel.spring.processor.scattergather.MyVendor">
    <constructor-arg>
      <value>3</value>
    </constructor-arg>
  </bean>

```

Each bean is initialized with a different price for beer (passed to the constructor argument). When a message is sent to each bean endpoint, it arrives at the **MyVendor.getQuote** method. This method does a simple check to see whether this quote request is for beer and then sets the price of beer on the exchange for retrieval at a later step. The message is forwarded to the next step using [POJO Producing](#) (see the `@Produce` annotation).

At the next step, we want to take the beer quotes from all vendors and find out which one was the best (that is, the lowest). For this, we use an [Aggregator](#) with a custom aggregation strategy. The [Aggregator](#) needs to identify which messages are relevant to the current quote, which is done by correlating messages based on the value of the **quoteRequestId** header (passed to the **correlationExpression**). As shown in [Example 7.1, "Messaging Client Sample"](#), the correlation ID is set to **quoteRequest-1** (the correlation ID should be unique). To pick the lowest quote out of the set, you can use a custom aggregation strategy like the following:

```

public class LowestQuoteAggregationStrategy implements AggregationStrategy
{
    public Exchange aggregate(Exchange oldExchange, Exchange newExchange)
    {
        // the first time we only have the new exchange
        if (oldExchange == null) {
            return newExchange;
        }

        if (oldExchange.getIn().getBody(int.class) <
            newExchange.getIn().getBody(int.class)) {
            return oldExchange;
        } else {
            return newExchange;
        }
    }
}

```

Static scatter-gather example

You can specify the recipients explicitly in the scatter-gather application by employing a static [Recipient List](#). The following example shows the routes you would use to implement a static scatter-gather scenario:

```

from("direct:start").multicast().to("seda:vendor1", "seda:vendor2",
"seda:vendor3");

from("seda:vendor1").to("bean:vendor1").to("seda:quoteAggregator");

```

```

from("seda:vendor2").to("bean:vendor2").to("seda:quoteAggregator");
from("seda:vendor3").to("bean:vendor3").to("seda:quoteAggregator");


from("seda:quoteAggregator")
    .aggregate(header("quoteRequestId"), new
LowestQuoteAggregationStrategy()).to("mock:result")

```

7.14. LOOP

Loop

The *loop* pattern enables you to process a message multiple times. It is used mainly for testing.



DEFAULT MODE

Notice by default the loop uses the same exchange throughout the looping. So the result from the previous iteration is used for the next (eg [Pipes and Filters](#)). From **Camel 2.8** onwards you can enable copy mode instead. See the options table for more details.

Exchange properties

On each loop iteration, two exchange properties are set, which can optionally be read by any processors included in the loop.

Property	Description
<code>CamelLoopSize</code>	Apache Camel 2.0: Total number of loops
<code>CamelLoopIndex</code>	Apache Camel 2.0: Index of the current iteration (0 based)

Java DSL examples

The following examples show how to take a request from the **direct:x** endpoint and then send the message repeatedly to **mock:result**. The number of loop iterations is specified either as an argument to **loop()** or by evaluating an expression at run time, where the expression *must* evaluate to an **int** (or else a **RuntimeCamelException** is thrown).

The following example passes the loop count as a constant:

```
from("direct:a").loop(8).to("mock:result");
```

The following example evaluates a simple expression to determine the loop count:

```
from("direct:b").loop(header("loop")).to("mock:result");
```

The following example evaluates an XPath expression to determine the loop count:

```
from("direct:c").loop().xpath("/hello/@times").to("mock:result");
```

XML configuration example

You can configure the same routes in Spring XML.

The following example passes the loop count as a constant:

```
<route>
  <from uri="direct:a"/>
  <loop>
    <constant>8</constant>
    <to uri="mock:result"/>
  </loop>
</route>
```

The following example evaluates a simple expression to determine the loop count:

```
<route>
  <from uri="direct:b"/>
  <loop>
    <header>loop</header>
    <to uri="mock:result"/>
  </loop>
</route>
```

Using copy mode

Now suppose we send a message to **direct:start** endpoint containing the letter A. The output of processing this route will be that, each **mock:loop** endpoint will receive AB as message.

```
from("direct:start")
  // instruct loop to use copy mode, which mean it will use a copy of
  // the input exchange
  // for each loop iteration, instead of keep using the same exchange
  all over
  .loop(3).copy()
    .transform(body().append("B"))
    .to("mock:loop")
  .end()
  .to("mock:result");
```

However if we do *not* enable copy mode then **mock:loop** will receive AB, ABB, AB BB messages.

```
from("direct:start")
  // by default loop will keep using the same exchange so on the 2nd
  // and 3rd iteration its
  // the same exchange that was previous used that are being looped all
  over
  .loop(3)
    .transform(body().append("B"))
    .to("mock:loop")
  .end()
```



```
.to("mock:result");
```

The equivalent example in XML DSL in copy mode is as follows:

```
<route>
  <from uri="direct:start"/>
  <!-- enable copy mode for loop eip -->
  <loop copy="true">
    <constant>3</constant>
    <transform>
      <simple>${body}B</simple>
    </transform>
    <to uri="mock:loop"/>
  </loop>
  <to uri="mock:result"/>
</route>
```

Options

The **loop** DSL command supports the following options:

Name	Default Value	Description
copy	false	Camel 2.8: Whether or not copy mode is used. If false then the same Exchange is being used throughout the looping. So the result from the previous iteration will be <i>visible</i> for the next iteration. Instead you can enable copy mode, and then each iteration is <i>restarting</i> with a fresh copy of the input Exchange .

7.15. SAMPLING

Sampling Throttler

A sampling throttler allows you to extract a sample of exchanges from the traffic through a route. It is configured with a sampling period during which only a *single* exchange is allowed to pass through. All other exchanges will be stopped.

By default, the sample period is 1 second.

Java DSL example

Use the **sample()** DSL command to invoke the sampler as follows:

```
// Sample with default sampling period (1 second)
from("direct:sample")
  .sample()
```

```

        .to("mock:result");

// Sample with explicitly specified sample period
from("direct:sample-configured")
    .sample(1, TimeUnit.SECONDS)
    .to("mock:result");

// Alternative syntax for specifying sampling period
from("direct:sample-configured-via-dsl")
    .sample().samplePeriod(1).timeUnits(TimeUnit.SECONDS)
    .to("mock:result");

from("direct:sample-messageFrequency")
    .sample(10)
    .to("mock:result");

from("direct:sample-messageFrequency-via-dsl")
    .sample().sampleMessageFrequency(5)
    .to("mock:result");

```

Spring XML example

In Spring XML, use the `sample` element to invoke the sampler, where you have the option of specifying the sampling period using the **samplePeriod** and **units** attributes:

```

<route>
  <from uri="direct:sample"/>
  <sample samplePeriod="1" units="seconds">
    <to uri="mock:result"/>
  </sample>
</route>
<route>
  <from uri="direct:sample-messageFrequency"/>
  <sample messageFrequency="10">
    <to uri="mock:result"/>
  </sample>
</route>
<route>
  <from uri="direct:sample-messageFrequency-via-dsl"/>
  <sample messageFrequency="5">
    <to uri="mock:result"/>
  </sample>
</route>

```

Options

The **sample** DSL command supports the following options:

Name	Default Value	Description
------	---------------	-------------

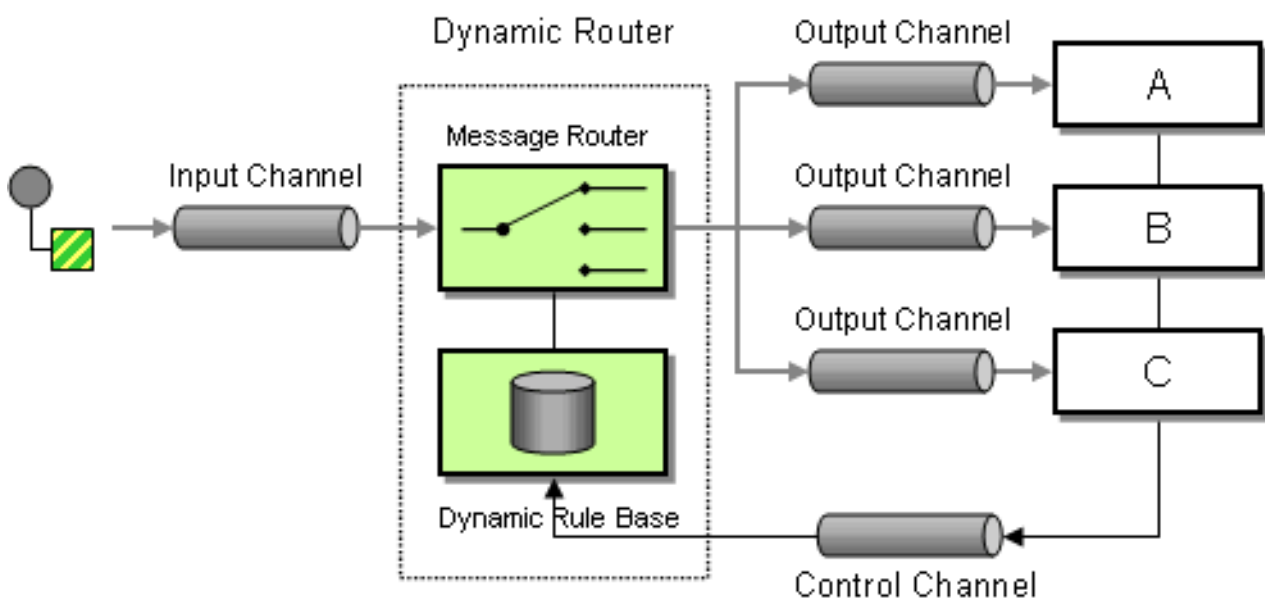
messageFrequency		Samples the message every N'th message. You can only use either frequency or period.
samplePeriod	1	Samples the message every N'th period. You can only use either frequency or period.
units	SECOND	Time unit as an enum of java.util.concurrent.TimeUnit from the JDK.

7.16. DYNAMIC ROUTER

Dynamic Router

The [Dynamic Router](#) pattern, as shown in [Figure 7.13, “Dynamic Router Pattern”](#), enables you to route a message consecutively through a series of processing steps, where the sequence of steps is not known at design time. The list of endpoints through which the message should pass is calculated dynamically *at run time*. Each time the message returns from an endpoint, the dynamic router calls back on a bean to discover the next endpoint in the route.

Figure 7.13. Dynamic Router Pattern



In **Camel 2.5** we introduced a **dynamicRouter** in the DSL, which is like a [dynamicRouting Slip](#) that evaluates the slip *on-the-fly*.

**BEWARE**

You must ensure the expression used for the **dynamicRouter** (such as a bean), returns **null** to indicate the end. Otherwise, the **dynamicRouter** will continue in an endless loop.

Dynamic Router in Camel 2.5 onwards

From Camel 2.5, the [Dynamic Router](#) updates the exchange property, **Exchange.SLIP_ENDPOINT**, with the current endpoint as it advances through the slip. This enables you to find out how far the exchange has progressed through the slip. (It's a slip because the [Dynamic Router](#) implementation is based on [Routing Slip](#)).

Java DSL

In Java DSL you can use the **dynamicRouter** as follows:

```
from("direct:start")
    // use a bean as the dynamic router
    .dynamicRouter(bean(DynamicRouterTest.class, "slip"));
```

Which will leverage a [Bean](#) to compute the slip *on-the-fly*, which could be implemented as follows:

```
// Java
/**
 * Use this method to compute dynamic where we should route next.
 *
 * @param body the message body
 * @return endpoints to go, or <tt>null</tt> to indicate the end
 */
public String slip(String body) {
    bodies.add(body);
    invoked++;

    if (invoked == 1) {
        return "mock:a";
    } else if (invoked == 2) {
        return "mock:b,mock:c";
    } else if (invoked == 3) {
        return "direct:foo";
    } else if (invoked == 4) {
        return "mock:result";
    }

    // no more so return null
    return null;
}
```

**NOTE**

The preceding example is *not* thread safe. You would have to store the state on the **Exchange** to ensure thread safety.

Spring XML

The same example in Spring XML would be:

```
<bean id="mySlip" class="org.apache.camel.processor.DynamicRouterTest"/>
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <dynamicRouter>
      <!-- use a method call on a bean as dynamic router -->
      <method ref="mySlip" method="slip"/>
    </dynamicRouter>
  </route>

  <route>
    <from uri="direct:foo"/>
    <transform><constant>Bye World</constant></transform>
    <to uri="mock:foo"/>
  </route>
</camelContext>
```

Options

The **dynamicRouter** DSL command supports the following options:

Name	Default Value	Description
uriDelimiter	,	Delimiter used if the Expression returned multiple endpoints.
ignoreInvalidEndpoints	false	If an endpoint uri could not be resolved, should it be ignored. Otherwise Camel will throw an exception stating the endpoint uri is not valid.

@DynamicRouter annotation

You can also use the **@DynamicRouter** annotation. For example:

```
// Java
public class MyDynamicRouter {

    @Consume(uri = "activemq:foo")
```

```
@DynamicRouter
public String route(@XPath("/customer/id") String customerId,
@Header("Location") String location, Document body) {
    // query a database to find the best match of the endpoint based
    on the input parameteres
    // return the next endpoint uri, where to go. Return null to
    indicate the end.
    }
}
```

The **route** method is invoked repeatedly as the message progresses through the slip. The idea is to return the endpoint URI of the next destination. Return **null** to indicate the end. You can return multiple endpoints if you like, just as the [Routing Slip](#), where each endpoint is separated by a delimiter.

CHAPTER 8. MESSAGE TRANSFORMATION

Abstract


The message transformation patterns describe how to modify the contents of messages for various purposes.

8.1. CONTENT ENRICHER

Overview

The *content enricher* pattern describes a scenario where the message destination requires more data than is present in the original message. In this case, you would use a content enricher to pull in the extra data from an external resource.

Figure 8.1. Content Enricher Pattern



Models of content enrichment

Apache Camel supports two kinds of content enricher, as follows:

- **enrich()**—obtains additional data from the resource by sending a copy of the current exchange to a *producer* endpoint and then using the data from the resulting reply (the exchange created by the enricher is always an *InOut* exchange).
- **pollEnrich()**—obtains the additional data by polling a *consumer* endpoint for data. Effectively, the consumer endpoint from the main route and the consumer endpoint in **pollEnrich()** are coupled, such that exchanges incoming on the main route trigger a poll of the **pollEnrich()** endpoint.

Content enrichment using `enrich()`

```

AggregationStrategy aggregationStrategy = ...

from("direct:start")
  .enrich("direct:resource", aggregationStrategy)
  .to("direct:result");
  
```

```
from("direct:resource")
...
```

The content enricher (**enrich**) retrieves additional data from a *resource endpoint* in order to enrich an incoming message (contained in the *original exchange*). An aggregation strategy combines the original exchange and the *resource exchange*. The first parameter of the **AggregationStrategy.aggregate(Exchange, Exchange)** method corresponds to the the original exchange, and the second parameter corresponds to the resource exchange. The results from the resource endpoint are stored in the resource exchange's *Out* message. Here is a sample template for implementing your own aggregation strategy class:

```
public class ExampleAggregationStrategy implements AggregationStrategy {
    public Exchange aggregate(Exchange original, Exchange resource) {
        Object originalBody = original.getIn().getBody();
        Object resourceResponse = resource.getOut().getBody();
        Object mergeResult = ... // combine original body and resource
response
        if (original.getPattern().isOutCapable()) {
            original.getOut().setBody(mergeResult);
        } else {
            original.getIn().setBody(mergeResult);
        }
        return original;
    }
}
```

Using this template, the original exchange can have any exchange pattern. The resource exchange created by the enricher is always an *InOut* exchange.

Spring XML enrich example

The preceding example can also be implemented in Spring XML:

```
<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <enrich uri="direct:resource" strategyRef="aggregationStrategy"/>
    <to uri="direct:result"/>
  </route>
  <route>
    <from uri="direct:resource"/>
    ...
  </route>
</camelContext>

<bean id="aggregationStrategy" class="..." />
```

Default aggregation strategy

The aggregation strategy is optional. If you do not provide it, Apache Camel will use the body obtained from the resource by default. For example:


```
from("direct:start")
  .enrich("direct:resource")
  .to("direct:result");
```

In the preceding route, the message sent to the **direct:result** endpoint contains the output from the **direct:resource**, because this example does not use any custom aggregation.

In XML DSL, just omit the **strategyRef** attribute, as follows:

```
<route>
  <from uri="direct:start"/>
  <enrich uri="direct:resource"/>
  <to uri="direct:result"/>
</route>
```

Enrich Options

The **enrich** DSL command supports the following options:

Name	Default Value	Description
uri		The endpoint uri for the external service to enrich from. You must use either uri or ref .
ref		Refers to the endpoint for the external service to enrich from. You must use either uri or ref .
strategyRef		Refers to an AggregationStrategy to be used to merge the reply from the external service, into a single outgoing message. By default Camel will use the reply from the external service as outgoing message.

Content enrich using pollEnrich

The **pollEnrich** command treats the resource endpoint as a *consumer*. Instead of sending an exchange to the resource endpoint, it *polls* the endpoint. By default, the poll returns immediately, if there is no exchange available from the resource endpoint. For example, the following route reads a file whose name is extracted from the header of an incoming JMS message:

```
from("activemq:queue:order")
  .pollEnrich("file://order/data/additional?fileName=orderId")
  .to("bean:processOrder");
```

And if you want to wait at most 20 seconds for the file to be ready, you can use a timeout as follows:

```
from("activemq:queue:order")
    .pollEnrich("file://order/data/additional?fileName=orderId", 20000) //
    timeout is in milliseconds
    .to("bean:processOrder");
```

You can also specify an aggregation strategy for **pollEnrich**, as follows:

```
.pollEnrich("file://order/data/additional?fileName=orderId", 20000,
aggregationStrategy)
```



NOTE

The resource exchange passed to the aggregation strategy's **aggregate()** method might be **null**, if the poll times out before an exchange is received.



DATA FROM CURRENT EXCHANGE NOT USED

pollEnrich does *not* access any data from the current Exchange, so that, when polling, it cannot use any of the existing headers you may have set on the Exchange. For example, you cannot set a filename in the **Exchange.FILE_NAME** header and use **pollEnrich** to consume only that file. For that, you must set the filename in the endpoint URI.

Polling methods used by pollEnrich()

In general, the **pollEnrich()** enricher polls the consumer endpoint using one of the following polling methods:

- **receiveNoWait()** (*used by default*)
- **receive()**
- **receive(long timeout)**

The **pollEnrich()** command's timeout argument (specified in milliseconds) determines which method gets called, as follows:

- Timeout is **0** or not specified, **receiveNoWait** is called.
- Timeout is negative, **receive** is called.
- Otherwise, **receive(timeout)** is called.

pollEnrich example

In this example we enrich the message by loading the content from the file named `inbox/data.txt`.

```
from("direct:start")
  .pollEnrich("file:inbox?fileName=data.txt")
  .to("direct:result");
```

And in XML DSL you do:

```
<route>
  <from uri="direct:start"/>
  <pollEnrich uri="file:inbox?fileName=data.txt"/>
  <to uri="direct:result"/>
</route>
```

If there is no file then the message is empty. We can use a timeout to either wait (potential forever) until a file exists, or use a timeout to wait a period. For example to wait up til 5 seconds you can do:

```
<route>
  <from uri="direct:start"/>
  <pollEnrich uri="file:inbox?fileName=data.txt" timeout="5000"/>
  <to uri="direct:result"/>
</route>
```

PollEnrich Options

The `pollEnrich` DSL command supports the following options:

Name	Default Value	Description
<code>uri</code>		The endpoint uri for the external service to enrich from. You must use either uri or ref .
<code>ref</code>		Refers to the endpoint for the external service to enrich from. You must use either uri or ref .
<code>strategyRef</code>		Refers to an AggregationStrategy to be used to merge the reply from the external service, into a single outgoing message. By default Camel will use the reply from the external service as outgoing message.


<code>timeout</code>	<code>0</code>	Timeout in millis to use when polling from the external service. See below for important details about the timeout.
----------------------	----------------	---

8.2. CONTENT FILTER

Overview

The *content filter* pattern describes a scenario where you need to filter out extraneous content from a message before delivering it to its intended recipient. For example, you might employ a content filter to strip out confidential information from a message.

Figure 8.2. Content Filter Pattern



A common way to filter messages is to use an expression in the DSL, written in one of the supported scripting languages (for example, XSLT, XQuery or JoSQL).

Implementing a content filter

A content filter is essentially an application of a message processing technique for a particular purpose. To implement a content filter, you can employ any of the following message processing techniques:

- *Message translator*—see [message translators](#).
- *Processors*—see [Chapter 41, Implementing a Processor](#).
- [Bean integration](#).

XML configuration example

The following example shows how to configure the same route in XML:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="activemq:My.Queue"/>
    <to uri="xslt:classpath:com/acme/content_filter.xsl"/>
    <to uri="activemq:Another.Queue"/>
  </route>
</camelContext>
```

Using an XPath filter

You can also use XPath to filter out part of the message you are interested in:

```
<route>
  <from uri="activemq:Input"/>
  <setBody><xpath resultType="org.w3c.dom.Document">//foo:bar</xpath>
</setBody>
  <to uri="activemq:Output"/>
</route>
```


8.3. NORMALIZER

Overview

The *normalizer* pattern is used to process messages that are semantically equivalent, but arrive in different formats. The normalizer transforms the incoming messages into a common format.

In Apache Camel, you can implement the normalizer pattern by combining a [content-based router](#), which detects the incoming message's format, with a collection of different [message translators](#), which transform the different incoming formats into a common format.

Figure 8.3. Normalizer Pattern



Java DSL example

This example shows a Message Normalizer that converts two types of XML messages into a common format. Messages in this common format are then filtered.

Using the [Fluent Builders](#)

```
// we need to normalize two types of incoming messages
from("direct:start")
  .choice()
    .when().xpath("/employee").to("bean:normalizer?
method=employeeToPerson")
    .when().xpath("/customer").to("bean:normalizer?
```

```
method=customerToPerson")
    .end()
    .to("mock:result");
```

In this case we're using a Java bean as the normalizer. The class looks like this

```
// Java
public class MyNormalizer {
    public void employeeToPerson(Exchange exchange,
    @XPath("/employee/name/text()") String name) {
        exchange.getOut().setBody(createPerson(name));
    }

    public void customerToPerson(Exchange exchange,
    @XPath("/customer/@name") String name) {
        exchange.getOut().setBody(createPerson(name));
    }

    private String createPerson(String name) {
        return "<person name=\"" + name + "\"/>";
    }
}
```

XML configuration example

The same example in the XML DSL

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <choice>
      <when>
        <xpath>/employee</xpath>
        <to uri="bean:normalizer?method=employeeToPerson"/>
      </when>
      <when>
        <xpath>/customer</xpath>
        <to uri="bean:normalizer?method=customerToPerson"/>
      </when>
    </choice>
    <to uri="mock:result"/>
  </route>
</camelContext>

<bean id="normalizer" class="org.apache.camel.processor.MyNormalizer"/>
```

8.4. CLAIM CHECK

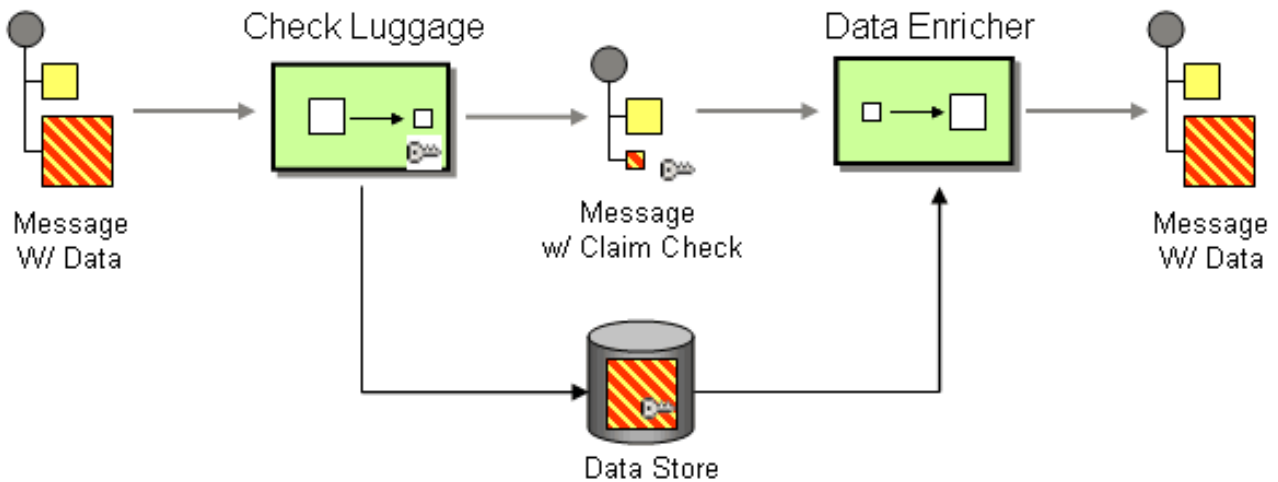
Claim Check

The *claim check* pattern, shown in [Figure 8.4, “Claim Check Pattern”](#), allows you to replace message content with a claim check (a unique key), which can be used to retrieve the message content at a later time. The message content is stored temporarily in a persistent

store like a database or file system. This pattern is very useful when message content is very large (thus it would be expensive to send around) and not all components require all information.

It can also be useful in situations where you cannot trust the information with an outside party; in this case, you can use the Claim Check to hide the sensitive portions of data.

Figure 8.4. Claim Check Pattern



Java DSL example

The following example shows how to replace a message body with a claim check and restore the body at a later step.

```
from("direct:start").to("bean:checkLuggage", "mock:testCheckpoint",
    "bean:dataEnricher", "mock:result");
```

The next step in the pipeline is the **mock:testCheckpoint** endpoint, which checks that the message body has been removed, the claim check added, and so on.

XML DSL example

The preceding example can also be written in XML, as follows:

```
<route>
  <from uri="direct:start"/>
  <pipeline>
    <to uri="bean:checkLuggage"/>
    <to uri="mock:testCheckpoint"/>
    <to uri="bean:dataEnricher"/>
    <to uri="mock:result"/>
  </pipeline>
</route>
```

checkLuggage bean

The message is first sent to the **checkLuggage** bean which is implemented as follows:

```
public static final class CheckLuggageBean {
```

```

    public void checkLuggage(Exchange exchange, @Body String body,
    @XPath("/order/@custId") String custId) {
        // store the message body into the data store, using the custId as
the claim check
        datastore.put(custId, body);
        // add the claim check as a header
        exchange.getIn().setHeader("claimCheck", custId);
        // remove the body from the message
        exchange.getIn().setBody(null);
    }
}

```

This bean stores the message body into the data store, using the **custId** as the claim check. In this example, we are using a **HashMap** to store the message body; in a real application you would use a database or the file system. The claim check is added as a message header for later use and, finally, we remove the body from the message and pass it down the pipeline.

testCheckpoint endpoint

The example route is just a [Pipeline](#). In a real application, you would substitute some other steps for the **mock:testCheckpoint** endpoint.

dataEnricher bean

To add the message body back into the message, we use the **dataEnricher** bean, which is implemented as follows:

```

public static final class DataEnricherBean {
    public void addDataBackIn(Exchange exchange, @Header("claimCheck")
String claimCheck) {
        // query the data store using the claim check as the key and add
the data
        // back into the message body
        exchange.getIn().setBody(dataStore.get(claimCheck));
        // remove the message data from the data store
        datastore.remove(claimCheck);
        // remove the claim check header
        exchange.getIn().removeHeader("claimCheck");
    }
}

```

This bean queries the data store, using the claim check as the key, and then adds the recovered data back into the message body. The bean then deletes the message data from the data store and removes the **claimCheck** header from the message.

8.5. SORT

Sort

The *sort* pattern is used to sort the contents of a message body, assuming that the message body contains a list of items that can be sorted.

By default, the contents of the message are sorted using a default comparator that handles

numeric values or strings. You can provide your own comparator and you can specify an expression that returns the list to be sorted (the expression must be convertible to `java.util.List`).

Java DSL example

The following example generates the list of items to sort by tokenizing on the line break character:

```
from("file://inbox").sort(body().tokenize("\n")).to("bean:MyServiceBean.processLine");
```

You can pass in your own comparator as the second argument to `sort()`:

```
from("file://inbox").sort(body().tokenize("\n"), new MyReverseComparator()).to("bean:MyServiceBean.processLine");
```

XML configuration example

You can configure the same routes in Spring XML.

The following example generates the list of items to sort by tokenizing on the line break character:

```
<route>
  <from uri="file://inbox"/>
  <sort>
    <simple>body</simple>
  </sort>
  <beanRef ref="myServiceBean" method="processLine"/>
</route>
```

And to use a custom comparator, you can reference it as a Spring bean:

```
<route>
  <from uri="file://inbox"/>
  <sort comparatorRef="myReverseComparator">
    <simple>body</simple>
  </sort>
  <beanRef ref="MyServiceBean" method="processLine"/>
</route>

<bean id="myReverseComparator" class="com.mycompany.MyReverseComparator"/>
```

Besides `<simple>`, you can supply an expression using any language you like, so long as it returns a list.

Options

The `sort` DSL command supports the following options:

Name	Default Value	Description
------	---------------	-------------

comparatorRef		Refers to a custom java.util.Comparator to use for sorting the message body. Camel will by default use a comparator which does a A..Z sorting.
----------------------	--	---

8.6. VALIDATE

Overview

The validate pattern provides a convenient syntax to check whether the content of a message is valid. The validate DSL command takes a predicate expression as its sole argument: if the predicate evaluates to **true**, the route continues processing normally; if the predicate evaluates to **false**, a **PredicateValidationException** is thrown.

Java DSL example

The following route validates the body of the current message using a regular expression:

```
from("jms:queue:incoming")
  .validate(body(String.class).regex("^\\w{10}\\,\\d{2}\\,\\w{24}$"))
  .to("bean:MyServiceBean.processLine");
```

You can also validate a message header—for example:

```
from("jms:queue:incoming")
  .validate(header("bar").isGreaterThan(100))
  .to("bean:MyServiceBean.processLine");
```

And you can use validate with the [simple](#) expression language:

```
from("jms:queue:incoming")
  .validate(simple("${in.header.bar} == 100"))
  .to("bean:MyServiceBean.processLine");
```

XML DSL example

To use validate in the XML DSL, the recommended approach is to use the [simple](#) expression language:

```
<route>
  <from uri="jms:queue:incoming"/>
  <validate>
    <simple>${body} regex ^\\w{10}\\,\\d{2}\\,\\w{24}$</simple>
  </validate>
  <beanRef ref="myServiceBean" method="processLine"/>
</route>

<bean id="myServiceBean" class="com.mycompany.MyServiceBean"/>
```

You can also validate a message header—for example:

```
<route>
  <from uri="jms:queue:incoming"/>
  <validate>
    <simple>${in.header.bar} == 100</simple>
  </validate>
  <beanRef ref="myServiceBean" method="processLine"/>
</route>

<bean id="myServiceBean" class="com.mycompany.MyServiceBean"/>
```

CHAPTER 9. MESSAGING ENDPOINTS

Abstract

The messaging endpoint patterns describe various features and qualities of service that can be configured on an endpoint.

9.1. MESSAGING MAPPER

Overview

The *messaging mapper* pattern describes how to map domain objects to and from a canonical message format, where the message format is chosen to be as platform neutral as possible. The chosen message format should be suitable for transmission through a [message bus](#), where the message bus is the backbone for integrating a variety of different systems, some of which might not be object-oriented.

Many different approaches are possible, but not all of them fulfill the requirements of a messaging mapper. For example, an obvious way to transmit an object is to use *object serialization*, which enables you to write an object to a data stream using an unambiguous encoding (supported natively in Java). However, this is *not* a suitable approach to use for the messaging mapper pattern, however, because the serialization format is understood only by Java applications. Java object serialization creates an impedance mismatch between the original application and the other applications in the messaging system.

The requirements for a messaging mapper can be summarized as follows:

- The canonical message format used to transmit domain objects should be suitable for consumption by non-object oriented applications.
- The mapper code should be implemented separately from both the domain object code and the messaging infrastructure. Apache Camel helps fulfill this requirement by providing hooks that can be used to insert mapper code into a route.
- The mapper might need to find an effective way of dealing with certain object-oriented concepts such as inheritance, object references, and object trees. The complexity of these issues varies from application to application, but the aim of the mapper implementation must always be to create messages that can be processed effectively by non-object-oriented applications.

Finding objects to map

You can use one of the following mechanisms to find the objects to map:

- *Find a registered bean.* — For singleton objects and small numbers of objects, you could use the **CamelContext** registry to store references to beans. For example, if a bean instance is instantiated using Spring XML, it is automatically entered into the registry, where the bean is identified by the value of its **id** attribute.
- *Select objects using the JoSQL language.* — If all of the objects you want to access are already instantiated at runtime, you could use the JoSQL language to locate a specific object (or objects). For example, if you have a class,

`org.apache.camel.builder.sql.Person`, with a `name` bean property and the incoming message has a `UserName` header, you could select the object whose `name` property equals the value of the `UserName` header using the following code:

```
import static org.apache.camel.builder.sql.SqlBuilder.sql;
import org.apache.camel.Expression;
...
Expression expression = sql("SELECT * FROM
org.apache.camel.builder.sql.Person where name = :UserName");
Object value = expression.evaluate(exchange);
```

Where the syntax, `:HeaderName`, is used to substitute the value of a header in a JoSQL expression.

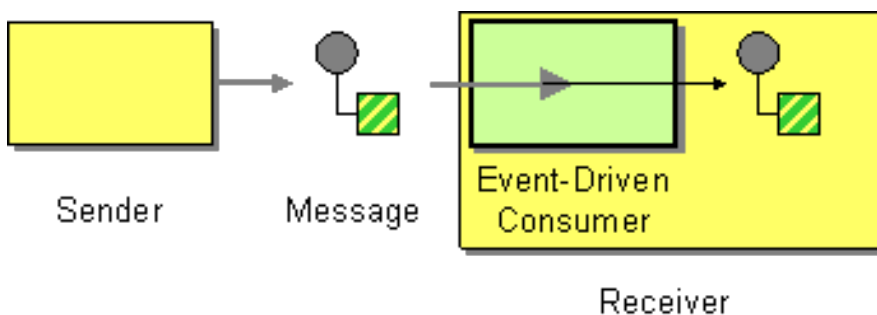
- *Dynamic* — For a more scalable solution, it might be necessary to read object data from a database. In some cases, the existing object-oriented application might already provide a finder object that can load objects from the database. In other cases, you might have to write some custom code to extract objects from a database, and in these cases the JDBC component and the SQL component might be useful.

9.2. EVENT DRIVEN CONSUMER

Overview

The *event-driven consumer* pattern, shown in [Figure 9.1, “Event Driven Consumer Pattern”](#), is a pattern for implementing the consumer endpoint in a Apache Camel component, and is only relevant to programmers who need to develop a custom component in Apache Camel. Existing components already have a consumer implementation pattern hard-wired into them.

Figure 9.1. Event Driven Consumer Pattern



Consumers that conform to this pattern provide an event method that is automatically called by the messaging channel or transport layer whenever an incoming message is received. One of the characteristics of the event-driven consumer pattern is that the consumer endpoint itself does not provide any threads to process the incoming messages. Instead, the underlying transport or messaging channel implicitly provides a processor thread when it invokes the exposed event method (which blocks for the duration of the message processing).

For more details about this implementation pattern, see [Section 44.1.3, “Consumer Patterns and Threading”](#) and [Chapter 47, *Consumer Interface*](#).

9.3. POLLING CONSUMER


Overview

The *polling consumer* pattern, shown in [Figure 9.2, “Polling Consumer Pattern”](#), is a pattern for implementing the consumer endpoint in a Apache Camel component, so it is only relevant to programmers who need to develop a custom component in Apache Camel. Existing components already have a consumer implementation pattern hard-wired into them.

Consumers that conform to this pattern expose polling methods, `receive()`, `receive(long timeout)`, and `receiveNoWait()` that return a new exchange object, if one is available from the monitored resource. A polling consumer implementation must provide its own thread pool to perform the polling.

For more details about this implementation pattern, see [Section 44.1.3, “Consumer Patterns and Threading”](#), [Chapter 47, *Consumer Interface*](#), and [Section 43.2, “Using the Consumer Template”](#).

Figure 9.2. Polling Consumer Pattern



Scheduled poll consumer

Many of the Apache Camel consumer endpoints employ a scheduled poll pattern to receive messages at the start of a route. That is, the endpoint appears to implement an event-driven consumer interface, but internally a scheduled poll is used to monitor a resource that provides the incoming messages for the endpoint.

See [Section 47.2, “Implementing the Consumer Interface”](#) for details of how to implement this pattern.

Quartz component

You can use the quartz component to provide scheduled delivery of messages using the *Quartz* enterprise scheduler. See [and Quartz Component](#) for details.

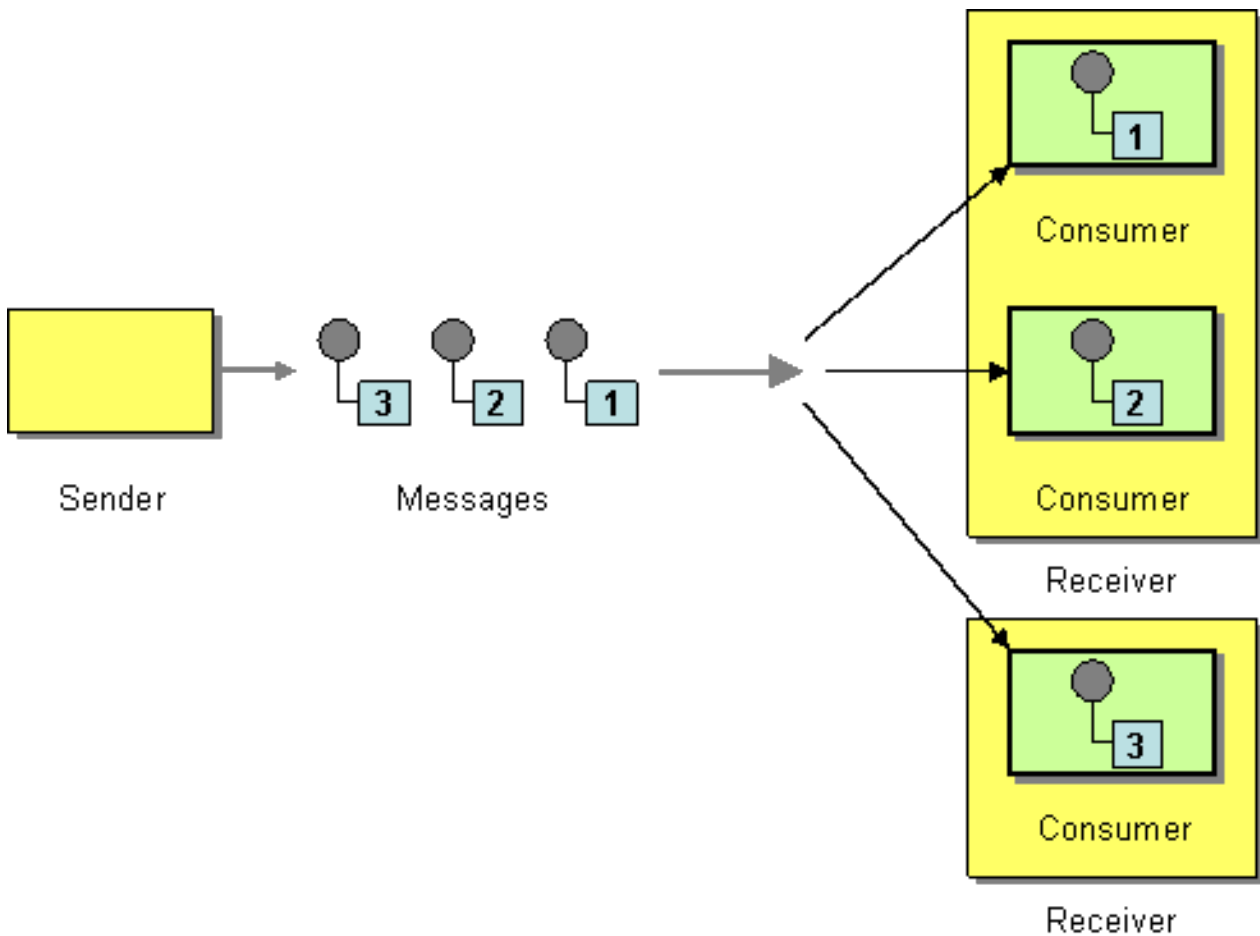
9.4. COMPETING CONSUMERS

Overview

The *competing consumers* pattern, shown in [Figure 9.3, “Competing Consumers Pattern”](#), enables multiple consumers to pull messages from the same queue, with the guarantee that *each message is consumed once only*. This pattern can be used to replace serial

message processing with concurrent message processing (bringing a corresponding reduction in response latency).

Figure 9.3. Competing Consumers Pattern



The following components demonstrate the competing consumers pattern:

- [the section called “JMS based competing consumers”](#)
- [the section called “SEDA based competing consumers”](#)

JMS based competing consumers

A regular JMS queue implicitly guarantees that each message can only be consumed at once. Hence, a JMS queue automatically supports the competing consumers pattern. For example, you could define three competing consumers that pull messages from the JMS queue, **HighVolumeQ**, as follows:

```
from("jms:HighVolumeQ").to("cxf:bean:replica01");
from("jms:HighVolumeQ").to("cxf:bean:replica02");
from("jms:HighVolumeQ").to("cxf:bean:replica03");
```

Where the CXF (Web services) endpoints, **replica01**, **replica02**, and **replica03**, process messages from the **HighVolumeQ** queue in parallel.

Alternatively, you can set the JMS query option, **concurrentConsumers**, to create a thread pool of competing consumers. For example, the following route creates a pool of three competing threads that pick messages from the specified queue:

-

```
from("jms:HighVolumeQ?concurrentConsumers=3").to("cxf:bean:replica01");
```

And the **concurrentConsumers** option can also be specified in XML DSL, as follows:

```
<route>
  <from uri="jms:HighVolumeQ?concurrentConsumers=3"/>
  <to uri="cxf:bean:replica01"/>
</route>
```



NOTE

JMS topics *cannot* support the competing consumers pattern. By definition, a JMS topic is intended to send multiple copies of the same message to different consumers. Therefore, it is not compatible with the competing consumers pattern.

SEDA based competing consumers

The purpose of the SEDA component is to simplify concurrent processing by breaking the computation into stages. A SEDA endpoint essentially encapsulates an in-memory blocking queue (implemented by **java.util.concurrent.BlockingQueue**). Therefore, you can use a SEDA endpoint to break a route into stages, where each stage might use multiple threads. For example, you can define a SEDA route consisting of two stages, as follows:

```
// Stage 1: Read messages from file system.
from("file://var/messages").to("seda:fanout");

// Stage 2: Perform concurrent processing (3 threads).
from("seda:fanout").to("cxf:bean:replica01");
from("seda:fanout").to("cxf:bean:replica02");
from("seda:fanout").to("cxf:bean:replica03");
```


Where the first stage contains a single thread that consumes message from a file endpoint, **file://var/messages**, and routes them to a SEDA endpoint, **seda:fanout**. The second stage contains three threads: a thread that routes exchanges to **cxf:bean:replica01**, a thread that routes exchanges to **cxf:bean:replica02**, and a thread that routes exchanges to **cxf:bean:replica03**. These three threads compete to take exchange instances from the SEDA endpoint, which is implemented using a blocking queue. Because the blocking queue uses locking to prevent more than one thread from accessing the queue at a time, you are guaranteed that each exchange instance can only be consumed once.

For a discussion of the differences between a SEDA endpoint and a thread pool created by **thread()**, see .

9.5. MESSAGE DISPATCHER

Overview

The *message dispatcher* pattern, shown in [Figure 9.4, “Message Dispatcher Pattern”](#), is used to consume messages from a channel and then distribute them locally to *performers*, which are responsible for processing the messages. In a Apache Camel application, performers are usually represented by in-process endpoints, which are used to transfer messages to another section of the route.

Figure 9.4. Message Dispatcher Pattern

You can implement the message dispatcher pattern in Apache Camel using one of the following approaches:

- [the section called “JMS selectors”](#)
- [the section called “JMS selectors in ActiveMQ”](#)
- [the section called “Content-based router”](#)

JMS selectors

If your application consumes messages from a JMS queue, you can implement the message dispatcher pattern using *JMS selectors*. A JMS selector is a predicate expression involving JMS headers and JMS properties. If the selector evaluates to **true**, the JMS message is allowed to reach the consumer, and if the selector evaluates to **false**, the JMS message is blocked. In many respects, a JMS selector is like a [filter processor](#), but it has the additional advantage that the filtering is implemented inside the JMS provider. This means that a JMS selector can block messages before they are transmitted to the Apache Camel application. This provides a significant efficiency advantage.

In Apache Camel, you can define a JMS selector on a consumer endpoint by setting the **selector** query option on a JMS endpoint URI. For example:

```
from("jms:dispatcher?selector=CountryCode='US'").to("cxf:bean:replica01");
from("jms:dispatcher?selector=CountryCode='IE'").to("cxf:bean:replica02");
from("jms:dispatcher?selector=CountryCode='DE'").to("cxf:bean:replica03");
```

Where the predicates that appear in a selector string are based on a subset of the SQL92 conditional expression syntax (for full details, see the [JMS specification](#)). The identifiers

appearing in a selector string can refer either to JMS headers or to JMS properties. For example, in the preceding routes, the sender sets a JMS property called **CountryCode**.

If you want to add a JMS property to a message from within your Apache Camel application, you can do so by setting a message header (either on *In* message or on *Out* messages). When reading or writing to JMS endpoints, Apache Camel maps JMS headers and JMS properties to, and from, its native message headers.

Technically, the selector strings must be URL encoded according to the **application/x-www-form-urlencoded** MIME format (see the [HTML specification](#)). In practice, the **&** (ampersand) character might cause difficulties because it is used to delimit each query option in the URI. For more complex selector strings that might need to embed the **&** character, you can encode the strings using the **java.net.URLEncoder** utility class. For example:

```
from("jms:dispatcher?selector=" +
    java.net.URLEncoder.encode("CountryCode='US' ", "UTF-8"))
    to("cxf:bean:replica01");
```

Where the UTF-8 encoding must be used.

JMS selectors in ActiveMQ

You can also define JMS selectors on ActiveMQ endpoints. For example:

```
from("activemq:dispatcher?
selector=CountryCode='US'")
    .to("cxf:bean:replica01");
from("activemq:dispatcher?
selector=CountryCode='IE'")
    .to("cxf:bean:replica02");
from("activemq:dispatcher?
selector=CountryCode='DE'")
    .to("cxf:bean:replica03");
```

For more details, see [ActiveMQ: JMS Selectors](#) and [ActiveMQ Message Properties](#).

Content-based router

The essential difference between the content-based router pattern and the message dispatcher pattern is that a content-based router dispatches messages to physically separate destinations (remote endpoints), and a message dispatcher dispatches messages locally, within the same process space. In Apache Camel, the distinction between these two patterns is determined by the target endpoint. The same router logic is used to implement both a content-based router and a message dispatcher. When the target endpoint is remote, the route defines a content-based router. When the target endpoint is in-process, the route defines a message dispatcher.

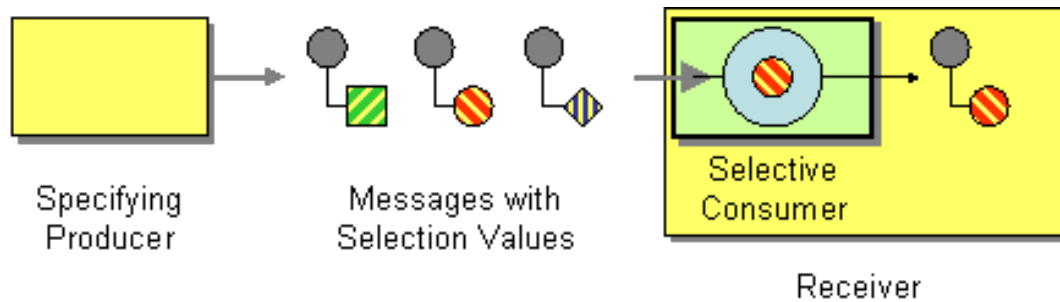
For details and examples of how to use the content-based router pattern see [Section 7.1, “Content-Based Router”](#).

9.6. SELECTIVE CONSUMER

Overview

The *selective consumer* pattern, shown in [Figure 9.5, “Selective Consumer Pattern”](#), describes a consumer that applies a filter to incoming messages, so that only messages meeting specific selection criteria are processed.

Figure 9.5. Selective Consumer Pattern



You can implement the selective consumer pattern in Apache Camel using one of the following approaches:


- [the section called “JMS selector”](#)
- [the section called “JMS selector in ActiveMQ”](#)
- [the section called “Message filter”](#)

JMS selector

A JMS selector is a predicate expression involving JMS headers and JMS properties. If the selector evaluates to **true**, the JMS message is allowed to reach the consumer, and if the selector evaluates to **false**, the JMS message is blocked. For example, to consume messages from the queue, **selective**, and select only those messages whose country code property is equal to **US**, you can use the following Java DSL route:

```
from("jms:selective?selector=" +
    java.net.URLEncoder.encode("CountryCode='US'", "UTF-8"))
    to("cxf:bean:replica01");
```

Where the selector string, **CountryCode='US'**, must be URL encoded (using UTF-8 characters) to avoid trouble with parsing the query options. This example presumes that the JMS property, **CountryCode**, is set by the sender. For more details about JMS selectors, see [the section called “JMS selectors”](#).



NOTE

If a selector is applied to a JMS queue, messages that are not selected remain on the queue and are potentially available to other consumers attached to the same queue.

JMS selector in ActiveMQ

You can also define JMS selectors on ActiveMQ endpoints. For example:

```
from("activemq:selective?selector=" +
    java.net.URLEncoder.encode("CountryCode='US'", "UTF-8"))
    to("cxf:bean:replica01");
```

For more details, see [ActiveMQ: JMS Selectors](#) and [ActiveMQ Message Properties](#).

Message filter

If it is not possible to set a selector on the consumer endpoint, you can insert a filter processor into your route instead. For example, you can define a selective consumer that processes only messages with a US country code using Java DSL, as follows:

```
from("seda:a").filter(header("CountryCode").isEqualTo("US")).process(myProcessor);
```

The same route can be defined using XML configuration, as follows:

```
<camelContext id="buildCustomProcessorWithFilter"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="seda:a"/>
    <filter>
      <xpath>$CountryCode = 'US'</xpath>
      <process ref="#myProcessor"/>
    </filter>
  </route>
</camelContext>
```

For more information about the Apache Camel filter processor, see [Message Filter](#).



WARNING

Be careful about using a message filter to select messages from a JMS *queue*. When using a filter processor, blocked messages are simply discarded. Hence, if the messages are consumed from a queue (which allows each message to be consumed only once—see [Section 9.4, “Competing Consumers”](#)), then blocked messages are not processed at all. This might not be the behavior you want.

9.7. DURABLE SUBSCRIBER

Overview

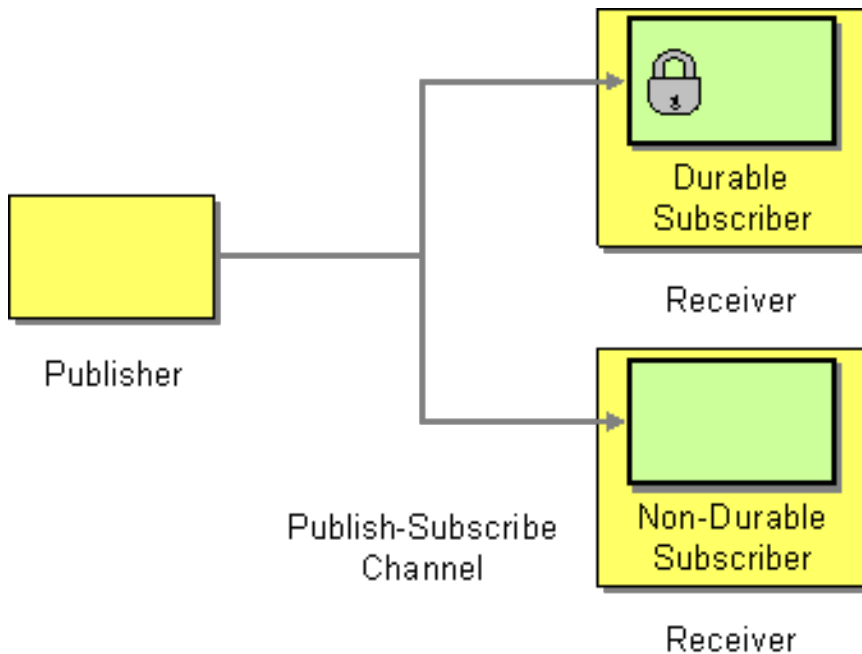
A *durable subscriber*, as shown in [Figure 9.6, “Durable Subscriber Pattern”](#), is a consumer that wants to receive all of the messages sent over a particular [publish-subscribe](#) channel, including messages sent while the consumer is disconnected from the messaging system. This requires the messaging system to store messages for later replay to the disconnected consumer. There also has to be a mechanism for a consumer to indicate that it wants to establish a durable subscription. Generally, a publish-subscribe channel (or topic) can have both durable and non-durable subscribers, which behave as follows:

- **non-durable subscriber**—Can have two states: *connected* and *disconnected*. While a non-durable subscriber is connected to a topic, it receives all of the topic's

messages in real time. However, a non-durable subscriber never receives messages sent to the topic while the subscriber is disconnected.

- **durable subscriber** —Can have two states: *connected* and *inactive*. The inactive state means that the durable subscriber is disconnected from the topic, but wants to receive the messages that arrive in the interim. When the durable subscriber reconnects to the topic, it receives a replay of all the messages sent while it was inactive.

Figure 9.6. Durable Subscriber Pattern



JMS durable subscriber

The JMS component implements the durable subscriber pattern. In order to set up a durable subscription on a JMS endpoint, you must specify a *client ID*, which identifies this particular connection, and a *durable subscription name*, which identifies the durable subscriber. For example, the following route sets up a durable subscription to the JMS topic, **news**, with a client ID of **conn01** and a durable subscription name of **John.Doe**:

```
from("jms:topic:news?clientId=conn01&durableSubscriptionName=John.Doe").
  to("cxf:bean:newsprocessor");
```

You can also set up a durable subscription using the ActiveMQ endpoint:

```
from("activemq:topic:news?
clientId=conn01&durableSubscriptionName=John.Doe").
  to("cxf:bean:newsprocessor");
```

If you want to process the incoming messages concurrently, you can use a SEDA endpoint to fan out the route into multiple, parallel segments, as follows:

```
from("jms:topic:news?clientId=conn01&durableSubscriptionName=John.Doe").
  to("seda:fanout");

from("seda:fanout").to("cxf:bean:newsproc01");
from("seda:fanout").to("cxf:bean:newsproc02");
```

```
from("seda:fanout").to("cxf:bean:newsproc03");
```

Where each message is processed only once, because the SEDA component supports the [competing consumers](#) pattern.

Alternative example

Another alternative is to combine the [Message Dispatcher](#) or [Content-Based Router](#) with [File component](#) or [JPA component](#) components for durable subscribers then something like [SEDA component](#) for non-durable.

Here is a simple example of creating durable subscribers to a [topic](#)

Using the Fluent Builders

```
from("direct:start").to("activemq:topic:foo");

from("activemq:topic:foo?
clientId=1&durableSubscriptionName=bar1").to("mock:result1");

from("activemq:topic:foo?
clientId=2&durableSubscriptionName=bar2").to("mock:result2");
```

Using the Spring XML Extensions

```
<route>
  <from uri="direct:start"/>
  <to uri="activemq:topic:foo"/>
</route>

<route>
  <from uri="activemq:topic:foo?
clientId=1&durableSubscriptionName=bar1"/>
  <to uri="mock:result1"/>
</route>

<route>
  <from uri="activemq:topic:foo?
clientId=2&durableSubscriptionName=bar2"/>
  <to uri="mock:result2"/>
</route>
```

Here is another example of [JMS](#) durable subscribers, but this time using [virtual topics](#) (recommended by AMQ over durable subscriptions)

Using the Fluent Builders

```
from("direct:start").to("activemq:topic:VirtualTopic.foo");

from("activemq:queue:Consumer.1.VirtualTopic.foo").to("mock:result1");

from("activemq:queue:Consumer.2.VirtualTopic.foo").to("mock:result2");
```

Using the Spring XML Extensions

-

```

<route>
  <from uri="direct:start"/>
  <to uri="activemq:topic:VirtualTopic.foo"/>
</route>

<route>
  <from uri="activemq:queue:Consumer.1.VirtualTopic.foo"/>
  <to uri="mock:result1"/>
</route>

<route>
  <from uri="activemq:queue:Consumer.2.VirtualTopic.foo"/>
  <to uri="mock:result2"/>
</route>

```

9.8. IDEMPOTENT CONSUMER

Overview

The *idempotent consumer* pattern is used to filter out duplicate messages. For example, consider a scenario where the connection between a messaging system and a consumer endpoint is abruptly lost due to some fault in the system. If the messaging system was in the middle of transmitting a message, it might be unclear whether or not the consumer received the last message. To improve delivery reliability, the messaging system might decide to redeliver such messages as soon as the connection is re-established. Unfortunately, this entails the risk that the consumer might receive duplicate messages and, in some cases, the effect of duplicating a message may have undesirable consequences (such as debiting a sum of money twice from your account). In this scenario, an idempotent consumer could be used to weed out undesired duplicates from the message stream.

Camel provides the following Idempotent Consumer implementations:

- [MemoryIdempotentRepository](#)
- [File](#)
- [HazelcastIdempotentRepository](#)
- [JdbcMessageIdRepository](#)
- [JpaMessageIdRepository](#)

Idempotent consumer with in-memory cache

In Apache Camel, the idempotent consumer pattern is implemented by the `idempotentConsumer()` processor, which takes two arguments:

- **messageIdExpression** — An expression that returns a message ID string for the current message.
- **messageIdRepository** — A reference to a message ID repository, which stores the IDs of all the messages received.

As each message comes in, the idempotent consumer processor looks up the current

message ID in the repository to see if this message has been seen before. If yes, the message is discarded; if no, the message is allowed to pass and its ID is added to the repository.

The code shown in [Example 9.1, “Filtering Duplicate Messages with an In-memory Cache”](#) uses the **TransactionID** header to filter out duplicates.

Example 9.1. Filtering Duplicate Messages with an In-memory Cache

```
import static
org.apache.camel.processor.idempotent.MemoryMessageIdRepository.memoryMe
ssageIdRepository;
...
RouteBuilder builder = new RouteBuilder() {
    public void configure() {
        from("seda:a")
            .idempotentConsumer(
                header("TransactionID"),
                memoryMessageIdRepository(200)
            ).to("seda:b");
    }
};
```

Where the call to **memoryMessageIdRepository(200)** creates an in-memory cache that can hold up to 200 message IDs.

You can also define an idempotent consumer using XML configuration. For example, you can define the preceding route in XML, as follows:

```
<camelContext id="buildIdempotentConsumer"
xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="seda:a"/>
        <idempotentConsumer messageIdRepositoryRef="MsgIDRepos">
            <simple>header.TransactionID</simple>
            <to uri="seda:b"/>
        </idempotentConsumer>
    </route>
</camelContext>

<bean id="MsgIDRepos"
class="org.apache.camel.processor.idempotent.MemoryMessageIdRepository">
    <!-- Specify the in-memory cache size. -->
    <constructor-arg type="int" value="200"/>
</bean>
```

Idempotent consumer with JPA repository

The in-memory cache suffers from the disadvantages of easily running out of memory and not working in a clustered environment. To overcome these disadvantages, you can use a Java Persistent API (JPA) based repository instead. The JPA message ID repository uses an object-oriented database to store the message IDs. For example, you can define a route that uses a JPA repository for the idempotent consumer, as follows:


```

import org.springframework.orm.jpa.JpaTemplate;

import org.apache.camel.spring.SpringRouteBuilder;
import static
org.apache.camel.processor.idempotent.jpa.JpaMessageIdRepository.jpaMessage
eIdRepository;
...
RouteBuilder builder = new SpringRouteBuilder() {
    public void configure() {
        from("seda:a").idempotentConsumer(
            header("TransactionID"),
            jpaMessageIdRepository(bean(JpaTemplate.class),
"myProcessorName")
        ).to("seda:b");
    }
};

```

The JPA message ID repository is initialized with two arguments:

- **JpaTemplate** instance—Provides the handle for the JPA database.
- processor name—Identifies the current idempotent consumer processor.

The **SpringRouteBuilder.bean()** method is a shortcut that references a bean defined in the Spring XML file. The **JpaTemplate** bean provides a handle to the underlying JPA database. See the JPA documentation for details of how to configure this bean.

For more details about setting up a JPA repository, see [JPA Component](#) documentation, the [Spring JPA](#) documentation, and the sample code in the [Camel JPA unit test](#)

Spring XML example

The following example uses the **myMessageId** header to filter out duplicates:

```

<!-- repository for the idempotent consumer -->
<bean id="myRepo"
class="org.apache.camel.processor.idempotent.MemoryIdempotentRepository"/>

<camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="direct:start"/>
        <idempotentConsumer messageIdRepositoryRef="myRepo">
            <!-- use the messageId header as key for identifying duplicate
messages -->
            <header>messageId</header>
            <!-- if not a duplicate send it to this mock endpoint -->
            <to uri="mock:result"/>
        </idempotentConsumer>
    </route>
</camelContext>

```

Idempotent consumer with JDBC repository

A JDBC repository is also supported for storing message IDs in the idempotent consumer pattern. The implementation of the JDBC repository is provided by the SQL component, so if

you are using the Maven build system, add a dependency on the **camel-sql** artifact.

You can use the **SingleConnectionDataSource** JDBC wrapper class from the Spring persistence API in order to instantiate the connection to a SQL database. For example, to instantiate a JDBC connection to a [HyperSQL](#) database instance, you could define the following JDBC data source:

```
<bean id="dataSource"
class="org.springframework.jdbc.datasource.SingleConnectionDataSource">
  <property name="driverClassName" value="org.hsqldb.jdbcDriver"/>
  <property name="url" value="jdbc:hsqldb:mem:camel_jdbc"/>
  <property name="username" value="sa"/>
  <property name="password" value=""/>
</bean>
```



NOTE

The preceding JDBC data source uses the HyperSQL **mem** protocol, which creates a memory-only database instance. This is a toy implementation of the HyperSQL database which is *not* actually persistent.

Using the preceding data source, you can define an idempotent consumer pattern that uses the JDBC message ID repository, as follows:

```
<bean id="messageIdRepository"
class="org.apache.camel.processor.idempotent.jdbc.JdbcMessageIdRepository"
>
  <constructor-arg ref="dataSource" />
  <constructor-arg value="myProcessorName" />
</bean>

<camel:camelContext>
  <camel:errorHandler id="deadLetterChannel" type="DeadLetterChannel"
deadLetterUri="mock:error">
    <camel:redeliveryPolicy maximumRedeliveries="0"
maximumRedeliveryDelay="0" logStackTrace="false" />
  </camel:errorHandler>

  <camel:route id="JdbcMessageIdRepositoryTest"
errorHandlerRef="deadLetterChannel">
    <camel:from uri="direct:start" />
    <camel:idempotentConsumer messageIdRepositoryRef="messageIdRepository">
      <camel:header>messageId</camel:header>
      <camel:to uri="mock:result" />
    </camel:idempotentConsumer>
  </camel:route>
</camel:camelContext>
```

How to handle duplicate messages in the route

Available as of Camel 2.8

You can now set the **skipDuplicate** option to **false** which instructs the idempotent consumer to route duplicate messages as well. However the duplicate message has been

marked as duplicate by having a property on the `Exchange` set to true. We can leverage this fact by using a `Content-Based Router` or `Message Filter` to detect this and handle duplicate messages.

For example in the following example we use the `Message Filter` to send the message to a duplicate endpoint, and then stop continue routing that message.

```
from("direct:start")
    // instruct idempotent consumer to not skip duplicates as we will
    filter then our self

.idempotentConsumer(header("messageId")).messageIdRepository(repo).skipDuplicate(false)
    .filter(property(Exchange.DUPLICATE_MESSAGE).isEqualTo(true))
    // filter out duplicate messages by sending them to someplace
else and then stop
    .to("mock:duplicate")
    .stop()
.end()
// and here we process only new messages (no duplicates)
.to("mock:result");
```

The sample example in XML DSL would be:

```
<!-- idempotent repository, just use a memory based for testing -->
<bean id="myRepo"
class="org.apache.camel.processor.idempotent.MemoryIdempotentRepository"/>

<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <!-- we do not want to skip any duplicate messages -->
    <idempotentConsumer messageIdRepositoryRef="myRepo"
skipDuplicate="false">
      <!-- use the messageId header as key for identifying
duplicate messages -->
      <header>messageId</header>
      <!-- we will to handle duplicate messages using a filter -->
      <filter>
        <!-- the filter will only react on duplicate messages,
if this property is set on the Exchange -->
        <property>CamelDuplicateMessage</property>
        <!-- and send the message to this mock, due its part of
an unit test -->
        <!-- but you can of course do anything as its part of the
route -->
        <to uri="mock:duplicate"/>
        <!-- and then stop -->
        <stop/>
      </filter>
      <!-- here we route only new messages -->
      <to uri="mock:result"/>
    </idempotentConsumer>
  </route>
</camelContext>
```

How to handle duplicate message in a clustered environment with a data grid

If you have running Camel in a clustered environment, a in memory idempotent repository doesn't work (see above). You can setup either a central database or use the idempotent consumer implementation based on the [Hazelcast](#) data grid. Hazelcast finds the nodes over multicast (which is default - configure Hazelcast for tcp-ip) and creates automatically a map based repository:

```
HazelcastIdempotentRepository idempotentRepo = new
HazelcastIdempotentRepository("myrepo");

from("direct:in").idempotentConsumer(header("messageId"),
idempotentRepo).to("mock:out");
```

You have to define how long the repository should hold each message id (default is to delete it never). To avoid that you run out of memory you should create an eviction strategy based on the [Hazelcast configuration](#). For additional information see [camel-hazelcast](#).

See this [little tutorial](#), how setup such an idempotent repository on two cluster nodes using Apache Karaf.

Options

The Idempotent Consumer has the following options:

Option	Default	Description
eager	true	Camel 2.0: Eager controls whether Camel adds the message to the repository before or after the exchange has been processed. If enabled before then Camel will be able to detect duplicate messages even when messages are currently in progress. By disabling Camel will only detect duplicates when a message has successfully been processed.
messageIdRepositoryRef	null	A reference to a IdempotentRepository to lookup in the registry. This option is mandatory when using XML DSL.


skipDuplicate	true	Camel 2.8: Sets whether to skip duplicate messages. If set to false then the message will be continued. However the Exchange has been marked as a duplicate by having the Exchange.DUPLICATE_MESSAGE exchange property set to a Boolean.TRUE value.
----------------------	-------------	---

9.9. TRANSACTIONAL CLIENT

Overview

The *transactional client* pattern, shown in [Figure 9.7, “Transactional Client Pattern”](#), refers to messaging endpoints that can participate in a transaction. Apache Camel supports transactions using [Spring transaction management](#).

Figure 9.7. Transactional Client Pattern



Transaction oriented endpoints

Not all Apache Camel endpoints support transactions. Those that do are called *transaction oriented endpoints* (or TOEs). For example, both the JMS component and the ActiveMQ component support transactions.

To enable transactions on a component, you must perform the appropriate initialization before adding the component to the **CamelContext**. This entails writing code to initialize your transactional components explicitly.

References

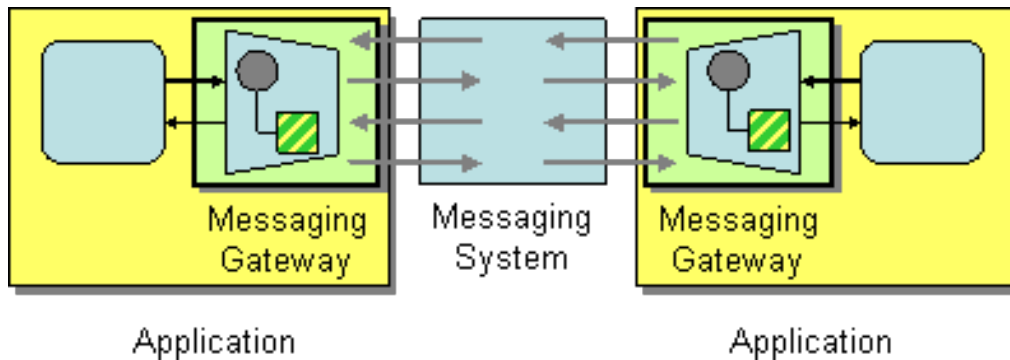
The details of configuring transactions in Apache Camel are beyond the scope of this guide. For full details of how to use transactions, see the *Apache Camel Transaction Guide*.

9.10. MESSAGING GATEWAY

Overview

The *messaging gateway* pattern, shown in [Figure 9.8, “Messaging Gateway Pattern”](#), describes an approach to integrating with a messaging system, where the messaging system's API remains hidden from the programmer at the application level. One of the more common example is when you want to translate synchronous method calls into request/reply message exchanges, without the programmer being aware of this.

Figure 9.8. Messaging Gateway Pattern



The following Apache Camel components provide this kind of integration with the messaging system:


-
-

9.11. SERVICE ACTIVATOR

Overview

The *service activator* pattern, shown in [Figure 9.9, “Service Activator Pattern”](#), describes the scenario where a service's operations are invoked in response to an incoming request message. The service activator identifies which operation to call and extracts the data to use as the operation's parameters. Finally, the service activator invokes an operation using the data extracted from the message. The operation invocation can be either oneway (request only) or two-way (request/reply).

Figure 9.9. Service Activator Pattern



In many respects, a service activator resembles a conventional remote procedure call (RPC), where operation invocations are encoded as messages. The main difference is that a service activator needs to be more flexible. An RPC framework standardizes the request and reply message encodings (for example, Web service operations are encoded as SOAP messages), whereas a service activator typically needs to improvise the mapping between the messaging system and the service's operations.

Bean integration

The main mechanism that Apache Camel provides to support the service activator pattern is *bean integration*. [Bean integration](#) provides a general framework for mapping incoming messages to method invocations on Java objects. For example, the Java fluent DSL provides the processors `bean()` and `beanRef()` that you can insert into a route to invoke methods on a registered Java bean. The detailed mapping of message data to Java method parameters is determined by the *bean binding*, which can be implemented by adding annotations to the bean class.

For example, consider the following route which calls the Java method, `BankBean.getUserAccBalance()`, to service requests incoming on a JMS/ActiveMQ queue:

```
from("activemq:BalanceQueries")
    .setProperty("userid",
xpath("/Account/BalanceQuery/UserID").stringResult())
    .beanRef("bankBean", "getUserAccBalance")
    .to("velocity:file:src/scripts/acc_balance.vm")
    .to("activemq:BalanceResults");
```

The messages pulled from the ActiveMQ endpoint, `activemq:BalanceQueries`, have a simple XML format that provides the user ID of a bank account. For example:

```
<?xml version='1.0' encoding='UTF-8'?>
<Account>
  <BalanceQuery>
    <UserID>James.Strachan</UserID>
  </BalanceQuery>
</Account>
```

The first processor in the route, `setProperty()`, extracts the user ID from the *in* message and stores it in the `userid` exchange property. This is preferable to storing it in a header, because the *in* headers are not available after invoking the bean.

The service activation step is performed by the `beanRef()` processor, which binds the incoming message to the `getUserAccBalance()` method on the Java object identified by the `bankBean` bean ID. The following code shows a sample implementation of the `BankBean` class:

```
package tutorial;

import org.apache.camel.language.XPath;

public class BankBean {
    public int getUserAccBalance(@XPath("/Account/BalanceQuery/UserID")
String user) {
        if (user.equals("James.Strachan")) {
            return 1200;
        }
    }
}
```

```

        }
        else {
            return 0;
        }
    }
}

```

Where the binding of message data to method parameter is enabled by the `@XPath` annotation, which injects the content of the `UserID` XML element into the `user` method parameter. On completion of the call, the return value is inserted into the body of the `Out` message which is then copied into the `In` message for the next step in the route. In order for the bean to be accessible to the `beanRef()` processor, you must instantiate an instance in Spring XML. For example, you can add the following lines to the `META-INF/spring/camel-context.xml` configuration file to instantiate the bean:

```

<?xml version="1.0" encoding="UTF-8"?>
<beans ... >
    ...
    <bean id="bankBean" class="tutorial.BankBean"/>
</beans>

```

Where the bean ID, `bankBean`, identifies this bean instance in the registry.

The output of the bean invocation is injected into a Velocity template, to produce a properly formatted result message. The Velocity endpoint, `velocity:file:src/scripts/acc_balance.vm`, specifies the location of a velocity script with the following contents:

```

<?xml version='1.0' encoding='UTF-8'?>
<Account>
    <BalanceResult>
        <UserID>${exchange.getProperty("userid")}</UserID>
        <Balance>${body}</Balance>
    </BalanceResult>
</Account>

```

The exchange instance is available as the Velocity variable, `exchange`, which enables you to retrieve the `userid` exchange property, using `${exchange.getProperty("userid")}`. The body of the current `In` message, `${body}`, contains the result of the `getUserAccBalance()` method invocation.

CHAPTER 10. SYSTEM MANAGEMENT

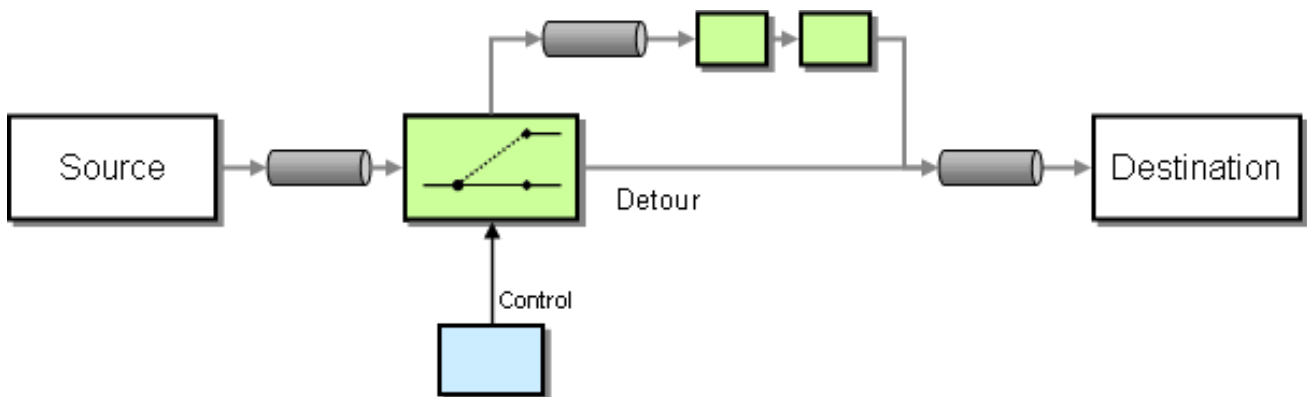
Abstract

The system management patterns describe how to monitor, test, and administer a messaging system.

10.1. DETOUR

Detour

The [Detour](#) from the [Introducing Enterprise Integration Patterns](#) allows you to send messages through additional steps if a control condition is met. It can be useful for turning on extra validation, testing, debugging code when needed.



Example

In this example we essentially have a route like `from("direct:start").to("mock:result")` with a conditional detour to the `mock:detour` endpoint in the middle of the route..

```
from("direct:start").choice()
    .when().method("controlBean", "isDetour").to("mock:detour").end()
    .to("mock:result");
```

Using the [Spring XML Extensions](#)

```
<route>
  <from uri="direct:start"/>
  <choice>
    <when>
      <method bean="controlBean" method="isDetour"/>
    </when>
    <to uri="mock:detour"/>
  </choice>
  <to uri="mock:result"/>
</split>
</route>
```

whether the detour is turned on or off is decided by the **ControlBean**. So, when the detour is on the message is routed to **mock:detour** and then **mock:result**. When the detour is off, the message is routed to **mock:result**.

For full details, check the example source here:

[camel-core/src/test/java/org/apache/camel/processor/DetourTest.java](https://github.com/apache/camel-core/blob/master/src/test/java/org/apache/camel/processor/DetourTest.java)

10.2. LOGEIP

Overview

Apache Camel provides several ways to perform logging in a route:

- Using the **log** DSL command.
- Using the **Log** component, which can log the message content.
- Using the Tracer, which traces message flow.
- Using a **Processor** or a **Bean** endpoint to perform logging in Java.



DIFFERENCE BETWEEN THE LOG DSL COMMAND AND THE LOG COMPONENT

The **log** DSL is much lighter and meant for logging human logs such as **Starting to do . . .**. It can only log a message based on the **Simple** language. In contrast, the **Log** component is a fully featured logging component. The **Log** component is capable of logging the message itself and you have many URI options to control the logging.

Java DSL example

Since **Apache Camel 2.2**, you can use the **log** DSL command to construct a log message at run time using the Simple expression language. For example, you can create a log message within a route, as follows:

```
from("direct:start").log("Processing ${id}").to("bean:foo");
```

This route constructs a **String** format message at run time. The log message will be logged at **INFO** level, using the route ID as the log name. By default, routes are named consecutively, **route-1**, **route-2** and so on. But you can use the DSL command, **routeId("myCoolRoute")**, to specify a custom route ID.

The log DSL also provides variants that enable you to set the logging level and the log name explicitly. For example, to set the logging level explicitly to **LogLevel.DEBUG**, you can invoke the log DSL as follows:

has overloaded methods to set the logging level and/or name as well.

```
from("direct:start").log(LogLevel.DEBUG, "Processing  
${id}").to("bean:foo");
```

To set the log name to **fileRoute**, you can invoke the log DSL as follows:

```
from("file://target/files").log(LoggingLevel.DEBUG, "fileRoute",
"Processing file ${file:name}").to("bean:foo");
```

XML DSL example

In XML DSL, the log DSL is represented by the **log** element and the log message is specified by setting the **message** attribute to a Simple expression, as follows:

```
<route id="foo">
  <from uri="direct:foo"/>
  <log message="Got ${body}"/>
  <to uri="mock:foo"/>
</route>
```

The **log** element supports the **message**, **loggingLevel** and **logName** attributes. For example:



```
<route id="baz">
  <from uri="direct:baz"/>
  <log message="Me Got ${body}" loggingLevel="FATAL" logName="cool"/>
  <to uri="mock:baz"/>
</route>
```

10.3. WIRE TAP

Wire Tap

The *wire tap* pattern, as shown in [Figure 10.1, “Wire Tap Pattern”](#), enables you to route a copy of the message to a separate tap location, while the original message is forwarded to the ultimate destination.

Figure 10.1. Wire Tap Pattern



STREAMS

If you [Wire Tap](#) a stream message body, you should consider enabling [Stream Caching](#) to ensure the message body can be re-read. See more details at [Stream Caching](#)

WireTap node

Apache Camel 2.0 introduces the **wireTap** node for doing wire taps. The **wireTap** node copies the original exchange to a tapped exchange, whose exchange pattern is set to *InOnly*, because the tapped exchange should be propagated in *oneway* style. The tapped exchange is processed in a separate thread, so that it can run concurrently with the main route.

The **wireTap** supports two different approaches to tapping an exchange:

- Tap a copy of the original exchange.
- Tap a new exchange instance, enabling you to customize the tapped exchange.

Tap a copy of the original exchange

Using the Java DSL:

```
from("direct:start")
    .to("log:foo")
    .wireTap("direct:tap")
    .to("mock:result");
```

Using Spring XML extensions:

```
<route>
  <from uri="direct:start"/>
  <to uri="log:foo"/>
  <wireTap uri="direct:tap"/>
  <to uri="mock:result"/>
</route>
```

Tap and modify a copy of the original exchange

Using the Java DSL, Apache Camel supports using either a processor or an expression to modify a copy of the original exchange. Using a processor gives you full power over how the exchange is populated, because you can set properties, headers and so on. The expression approach can only be used to modify the *In* message body.

For example, to modify a copy of the original exchange using the *processor* approach:

```
from("direct:start")
    .wireTap("direct:foo", new Processor() {
        public void process(Exchange exchange) throws Exception {
            exchange.getIn().setHeader("foo", "bar");
        }
    }).to("mock:result");

from("direct:foo").to("mock:foo");
```

And to modify a copy of the original exchange using the *expression* approach:

```
from("direct:start")
    .wireTap("direct:foo", constant("Bye World"))
```

```

        .to("mock:result");
    from("direct:foo").to("mock:foo");

```

Using the Spring XML extensions, you can modify a copy of the original exchange using the *processor* approach, where the **processorRef** attribute references a spring bean with the **myProcessor** ID:

```

<route>
  <from uri="direct:start2"/>
  <wireTap uri="direct:foo" processorRef="myProcessor"/>
  <to uri="mock:result"/>
</route>

```

And to modify a copy of the original exchange using the *expression* approach:

```

<route>
  <from uri="direct:start"/>
  <wireTap uri="direct:foo">
    <body><constant>Bye World</constant></body>
  </wireTap>
  <to uri="mock:result"/>
</route>

```

Tap a new exchange instance

You can define a wiretap with a new exchange instance by setting the copy flag to **false** (the default is **true**). In this case, an initially empty exchange is created for the wiretap.

For example, to create a new exchange instance using the *processor* approach:

```

from("direct:start")
    .wireTap("direct:foo", false, new Processor() {
        public void process(Exchange exchange) throws Exception {
            exchange.getIn().setBody("Bye World");
            exchange.getIn().setHeader("foo", "bar");
        }
    }).to("mock:result");

from("direct:foo").to("mock:foo");

```

Where the second **wireTap** argument sets the copy flag to **false**, indicating that the original exchange is *not* copied and an empty exchange is created instead.

To create a new exchange instance using the *expression* approach:

```

from("direct:start")
    .wireTap("direct:foo", false, constant("Bye World"))
    .to("mock:result");

from("direct:foo").to("mock:foo");

```

Using the Spring XML extensions, you can indicate that a new exchange is to be created by setting the `wireTap` element's `copy` attribute to `false`.

To create a new exchange instance using the *processor* approach, where the `processorRef` attribute references a spring bean with the `myProcessor` ID, as follows:

```
<route>
  <from uri="direct:start2"/>
  <wireTap uri="direct:foo" processorRef="myProcessor" copy="false"/>
  <to uri="mock:result"/>
</route>
```

And to create a new exchange instance using the *expression* approach:

```
<route>
  <from uri="direct:start"/>
  <wireTap uri="direct:foo" copy="false">
    <body><constant>Bye World</constant></body>
  </wireTap>
  <to uri="mock:result"/>
</route>
```

Sending a new **Exchange** and set headers in DSL

Available as of Camel 2.8

If you send a new messages using the **Wire Tap** then you could only set the message body using an **Expression** from the DSL. If you also need to set new headers you would have to use a **Processor** for that. So in Camel 2.8 onwards we have improved this situation so you can now set headers as well in the DSL.

The following example sends a new message which has

- "Bye World" as message body
- a header with key "id" with the value 123
- a header with key "date" which has current date as value

Java DSL

```
from("direct:start")
  // tap a new message and send it to direct:tap
  // the new message should be Bye World with 2 headers
  .wireTap("direct:tap")
    // create the new tap message body and headers
    .newExchangeBody(constant("Bye World"))
    .newExchangeHeader("id", constant(123))
    .newExchangeHeader("date", simple("${date:now:yyyyMMdd}"))
  .end()
  // here we continue routing the original messages
  .to("mock:result");
```

```
// this is the tapped route
from("direct:tap")
  .to("mock:tap");
```

XML DSL

The XML DSL is slightly different than Java DSL as how you configure the message body and headers. In XML you use `<body>` and `<setHeader>` as shown:

```
<route>
  <from uri="direct:start"/>
  <!-- tap a new message and send it to direct:tap -->
  <!-- the new message should be Bye World with 2 headers -->
  <wireTap uri="direct:tap">
    <!-- create the new tap message body and headers -->
    <body><constant>Bye World</constant></body>
    <setHeader headerName="id"><constant>123</constant></setHeader>
    <setHeader headerName="date"><simple>${date:now:yyyyMMdd}
  </simple></setHeader>
  </wireTap>
  <!-- here we continue routing the original message -->
  <to uri="mock:result"/>
</route>
```

Using `onPrepare` to execute custom logic when preparing messages

Available as of Camel 2.8

For details, see [Multicast](#).

Options

The `wireTap` DSL command supports the following options:

Name	Default Value	Description
<code>uri</code>		The endpoint uri where to send the wire tapped message. You should use either <code>uri</code> or <code>ref</code> .
<code>ref</code>		Refers to the endpoint where to send the wire tapped message. You should use either <code>uri</code> or <code>ref</code> .
<code>executorServiceRef</code>		Refers to a custom Thread Pool to be used when processing the wire tapped messages. If not set then Camel uses a default thread pool.

processorRef		Refers to a custom Processor to be used for creating a new message (eg the send a new message mode). See below.
copy	true	Camel 2.3: Should a copy of the Exchange to used when wire tapping the message.
onPrepareRef		Camel 2.8: Refers to a custom Processor to prepare the copy of the Exchange to be wire tapped. This allows you to do any custom logic, such as deep-cloning the message payload if that's needed etc.

APPENDIX A. MIGRATING FROM SERVICEMIX EIP

Abstract

If you are currently an Apache ServiceMix 3.x user, you might already have implemented some Enterprise Integration Patterns using the *ServiceMix EIP module*. It is recommended that you migrate these legacy patterns to Apache Camel, which has more extensive support for Enterprise Integration Patterns. After migrating, you can deploy your patterns into a Red Hat JBoss Fuse container.

A.1. MIGRATING ENDPPOINTS

Overview

A typical ServiceMix EIP route exposes a service that consumes exchanges from the NMR. The route also defines one or more target destinations, to which exchanges are sent. In the Apache Camel environment, the exposed ServiceMix service maps to a *consumer endpoint* and the ServiceMix target destinations map to *producer endpoints*. The Apache Camel consumer endpoints and producer endpoints are both defined using *endpoint URIs*.

When migrating endpoints from ServiceMix EIP to Apache Camel, you must express the ServiceMix services/endpoints as Apache Camel endpoint URIs. You can adopt one of the following approaches:

- Connect to an existing ServiceMix service/endpoint through the ServiceMix Camel module (which integrates Apache Camel with the NMR).
- If the existing ServiceMix service/endpoint represents a ServiceMix binding component, you can replace the ServiceMix binding component with an equivalent Apache Camel component (thus bypassing the NMR).

The ServiceMix Camel module

The integration between Apache Camel and ServiceMix is provided by the **servicemix-camel** module. This module is provided with ServiceMix, but actually implements a plug-in for the Apache Camel product: the *JBI component* (see [and JBI Component](#)).

To access the JBI component from Apache Camel, make sure that the **servicemix-camel** JAR file is included on your Classpath or, if you are using Maven, include a dependency on the **servicemix-camel** artifact in your project POM. You can then access the JBI component by defining Apache Camel endpoint URIs with the **jbi:** component prefix.

Translating ServiceMix URIs into Apache Camel endpoint URIs

ServiceMix defines a flexible format for defining URIs, which is described in detail in [ServiceMix URIs](#). To translate a ServiceMix URI into a Apache Camel endpoint URI, perform the following steps:

1. If the ServiceMix URI contains a namespace prefix, replace the prefix by its corresponding namespace.

For example, after modifying the ServiceMix URI, **service:test:messageFilter**, where **test** corresponds to the namespace, **http://progress.com/demos/test**, you get **service:http://progress.com/demos/test:messageFilter**.

2. Modify the separator character, depending on what kind of namespace appears in the URI:

- If the namespace starts with **http://**, use the **/** character as the separator between namespace, service name, and endpoint name (if present).

For example, the URI,

service:http://progress.com/demos/test:messageFilter, would be modified to **service:http://progress.com/demos/test/messageFilter**.

- If the namespace starts with **urn:**, use the **:** character as the separator between namespace, service name, and endpoint name (if present).

For example, **service:urn:progress:com:demos:test:messageFilter**.

3. Create a JBI endpoint URI by adding the **jbi:** prefix.

For example, **jbi:service:http://progress.com/demos/test/messageFilter**.

Example of mapping ServiceMix URIs

For example, consider the following configuration of the [static recipient list](#) pattern in ServiceMix EIP. The **eip:exchange-target** elements define some targets using the ServiceMix URI format.

```
<beans xmlns:sm="http://servicemix.apache.org/config/1.0"
      xmlns:eip="http://servicemix.apache.org/eip/1.0"
      xmlns:test="http://progress.com/demos/test" >
  ...
  <eip:static-recipient-list service="test:recipients"
endpoint="endpoint">
    <eip:recipients>
      <eip:exchange-target uri="service:test:messageFilter" />
      <eip:exchange-target uri="service:test:trace4" />
    </eip:recipients>
  </eip:static-recipient-list>
  ...
</beans>
```

When the preceding ServiceMix configuration is mapped to an equivalent Apache Camel configuration, you get the following route:

```
<route>
  <from
uri="jbi:endpoint:http://progress.com/demos/test/recipients/endpoint"/>
    <to uri="jbi:service:http://progress.com/demos/test/messageFilter"/>
    <to uri="jbi:service:http://progress.com/demos/test/trace4"/>
  </route>
```

Replacing ServiceMix bindings with Apache Camel components

Instead of using the Apache Camel JBI component to route all your messages through the ServiceMix NMR, you can use one of the many supported Apache Camel components to connect directly to a consumer or a producer endpoint. In particular, when sending messages to an external endpoint, it is more efficient to send the messages directly through a Apache Camel component than sending them through the NMR and a ServiceMix binding.

For details of all the Apache Camel components that are available, see [and Apache Camel Components](#).

A.2. COMMON ELEMENTS

Overview

When configuring ServiceMix EIP patterns in a ServiceMix configuration file, there are some common elements that occur in many of the pattern schemas. This section provides a brief overview of these common elements and explains how they can be mapped to equivalent constructs in Apache Camel.

Exchange target

All of the patterns supported by ServiceMix EIP use the **eip:exchange-target** element to specify JBI target endpoints. [Table A.1, “Mapping the Exchange Target Element”](#) shows examples of how to map sample **eip:exchange-target** elements to Apache Camel endpoint URIs, where it is assumed that the **test** prefix maps to the **http://progress.com/demos/test** namespace.

Table A.1. Mapping the Exchange Target Element

ServiceMix EIP Target	Apache Camel Endpoint URI
<code><eip:exchange-target interface="HelloWorld" /></code>	<code>jbi:interface:HelloWorld</code>
<code><eip:exchange-target service="test:HelloWorldService" /></code>	<code>jbi:service:http://progress.com/demos/test/HelloWorldService</code>
<code><eip:exchange-target service="test:HelloWorldService" endpoint="secure" /></code>	<code>jbi:service:http://progress.com/demos/test/HelloWorldService/secure</code>
<code><eip:exchange-target uri="service:test:HelloWorldService" /></code>	<code>jbi:service:http://progress.com/demos/test/HelloWorldService</code>

Predicates

The ServiceMix EIP component allows you to define predicate expressions in the XPath language. For example, XPath predicates can appear in **eip:xpath-predicate** elements or in **eip:xpath-splitter** elements, where the XPath predicate is specified using an **xpath** attribute.

ServiceMix XPath predicates can easily be migrated to equivalent constructs in Apache Camel: that is, either the **xpath** element (in XML configuration) or the **xpath()** command (in Java DSL). For example, the message filter pattern in Apache Camel can incorporate an XPath predicate as follows:

```
<route>
  <from
uri="jbi:endpoint:http://progress.com/demos/test/messageFilter/endpoint">
  <filter>
    <xpath>count(/test:world) = 1</xpath>
    <to uri="jbi:service:http://progress.com/demos/test/trace3"/>
  </filter>
</route>
```

Where the **xpath** element specifies that only messages containing the **test:world** element will pass through the filter.



NOTE

Apache Camel also supports a wide range of other scripting languages including XQuery, PHP, Python, and Ruby, which can be used to define predicates. For details of all the supported predicate languages, see [Expression and Predicate Languages](#).

Namespace contexts

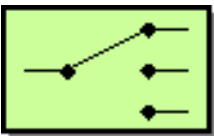
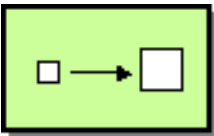
When using XPath predicates in the ServiceMix EIP configuration, it is necessary to define a namespace context using the **eip:namespace-context** element. The namespace is then referenced using a **namespaceContext** attribute.

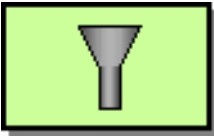

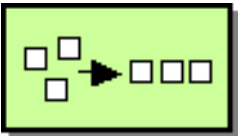
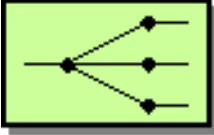

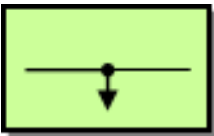
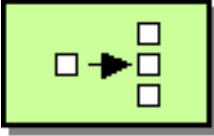
When ServiceMix EIP configuration is migrated to Apache Camel, there is no need to define namespace contexts, because Apache Camel allows you to define XPath predicates without referencing a namespace context. You can simply drop the **eip:namespace-context** elements when you migrate to Apache Camel.

A.3. SERVICEMIX EIP PATTERNS

The patterns supported by ServiceMix EIP are shown in [Table A.2, “ServiceMix EIP Patterns”](#).

Table A.2. ServiceMix EIP Patterns


	<p>Content-Based Router</p>	<p>How we handle a situation where the implementation of a single logical function (e.g., inventory check) is spread across multiple physical systems.</p>
	<p>Content Enricher</p>	<p>How we communicate with another system if the message originator does not have all the required data items available.</p>

	Message Filter	How a component avoids receiving uninteresting messages.
	Pipeline	How we perform complex processing on a message while maintaining independence and flexibility.
	Resequencer	How we get a stream of related but out-of-sequence messages back into the correct order.
	Static Recipient List	How we route a message to a list of specified recipients.
	Static Routing Slip	How we route a message consecutively through a series of processing steps.
	Wire Tap	How you inspect messages that travel on a point-to-point channel.
	XPath Splitter	How we process a message if it contains multiple elements, each of which may have to be processed in a different way.

A.4. CONTENT-BASED ROUTER

Overview

A *content-based router* enables you to route messages to the appropriate destination, where the routing decision is based on the message contents. This pattern maps to the corresponding [content-based router](#) pattern in Apache Camel.

Figure A.1. Content-based Router Pattern**Example ServiceMix EIP route**

Example A.1, “[ServiceMix EIP Content-based Route](#)” shows how to define a content-based router using the ServiceMix EIP component. If a **test:echo** element is present in the message body, the message is routed to the **http://test/pipeline/endpoint** endpoint. Otherwise, the message is routed to the **test:recipients** endpoint.

Example A.1. ServiceMix EIP Content-based Route

```
<eip:content-based-router service="test:router" endpoint="endpoint">
  <eip:rules>
    <eip:routing-rule>
      <eip:predicate>
        <eip:xpath-predicate xpath="count(/test:echo) = 1"
namespaceContext="#nsContext" />
      </eip:predicate>
      <eip:target>
        <eip:exchange-target uri="endpoint:test:pipeline:endpoint" />
      </eip:target>
    </eip:routing-rule>
    <eip:routing-rule>
      <!-- There is no predicate, so this is the default destination --
    >
      <eip:target>
        <eip:exchange-target service="test:recipients" />
      </eip:target>
    </eip:routing-rule>
  </eip:rules>
</eip:content-based-router>
```

Equivalent Apache Camel XML route

Example A.2, “[Apache Camel Content-based Router Using XML Configuration](#)” shows how to define an equivalent route using Apache Camel XML configuration.

Example A.2. Apache Camel Content-based Router Using XML Configuration

```
<route>
  <from
uri="jbi:endpoint:http://progress.com/demos/test/router/endpoint"/>
  <choice>
    <when>
```

```

    <xpath>count(/test:echo) = 1</xpath>
    <to
uri="jbi:endpoint:http://progress.com/demos/test/pipeline/endpoint"/>
    </when>
    <otherwise>
    <!-- This is the default destination -->
    <to uri="jbi:service:http://progress.com/demos/test/recipients"/>
    </otherwise>
    </choice>
</route>

```

Equivalent Apache Camel Java DSL route

Example A.3, “[Apache Camel Content-based Router Using Java DSL](#)” shows how to define an equivalent route using the Apache Camel Java DSL.

Example A.3. Apache Camel Content-based Router Using Java DSL

```

from("jbi:endpoint:http://progress.com/demos/test/router/endpoint").
    choice().when(xpath("count(/test:echo) =
1")).to("jbi:endpoint:http://progress.com/demos/test/pipeline/endpoint")
.
otherwise().to("jbi:service:http://progress.com/demos/test/recipients");


```

A.5. CONTENT ENRICHER

Overview

A *content enricher*, shown in [Figure A.2, “Content Enricher Pattern”](#), is a pattern for augmenting a message with missing information. The ServiceMix EIP content enricher is roughly equivalent to a pipeline that adds missing data as the message passes through an enricher target. Consequently, when migrating to Apache Camel, you can re-implement the ServiceMix content enricher as a Apache Camel pipeline.

Figure A.2. Content Enricher Pattern



Example ServiceMix EIP route

[Example A.4, “ServiceMix EIP Content Enricher”](#) shows how to define a content enricher using the ServiceMix EIP component. Incoming messages pass through the enricher target, **test:additionalInformationExtractor**, which adds missing data to the message. The message is then sent on to its ultimate destination, **test:myTarget**.

Example A.4. ServiceMix EIP Content Enricher

```
<eip:content-enricher service="test:contentEnricher"
endpoint="endpoint">
  <eip:enricherTarget>
    <eip:exchange-target service="test:additionalInformationExtractor"
/>
  </eip:enricherTarget>
  <eip:target>
    <eip:exchange-target service="test:myTarget" />
  </eip:target>
</eip:content-enricher>
```

Equivalent Apache Camel XML route

[Example A.5, “Apache Camel Content Enricher using XML Configuration”](#) shows how to define an equivalent route using Apache Camel XML configuration.

Example A.5. Apache Camel Content Enricher using XML Configuration

```
<route>
  <from>
    uri="jbi:endpoint:http://progress.com/demos/test/contentEnricher/endpoi
nt"/>
  <to>
    uri="jbi:service:http://progress.com/demos/test/additionalInformationExt
racter"/>
    <to uri="jbi:service:http://progress.com/demos/test/myTarget"/>
  </route>
```

Equivalent Apache Camel Java DSL route

[Example A.6, “Apache Camel Content Enricher using Java DSL”](#) shows how to define an equivalent route using the Apache Camel Java DSL:

Example A.6. Apache Camel Content Enricher using Java DSL


```
from("jbi:endpoint:http://progress.com/demos/test/contentEnricher/endpoi
nt").
to("jbi:service:http://progress.com/demos/test/additionalInformationExt
racter").
  to("jbi:service:http://progress.com/demos/test/myTarget");
```


A.6. MESSAGE FILTER

Overview

A *message filter*, shown in [Figure A.3, “Message Filter Pattern”](#), is a processor that eliminates undesired messages based on specific criteria. Filtering is controlled by specifying a predicate in the filter: when the predicate is **true**, the incoming message is allowed to pass; otherwise, it is blocked. This pattern maps to the corresponding [message filter](#) pattern in Apache Camel.

Figure A.3. Message Filter Pattern



Example ServiceMix EIP route

[Example A.7, “ServiceMix EIP Message Filter”](#) shows how to define a message filter using the ServiceMix EIP component. Incoming messages are passed through a filter mechanism that blocks messages that lack a **test:world** element.

Example A.7. ServiceMix EIP Message Filter

```
<eip:message-filter service="test:messageFilter" endpoint="endpoint">
  <eip:target>
    <eip:exchange-target service="test:trace3" />
  </eip:target>
  <eip:filter>
    <eip:xpath-predicate xpath="count(/test:world) = 1"
      namespaceContext="#nsContext"/>
  </eip:filter>
</eip:message-filter>
```

Equivalent Apache Camel XML route

[Example A.8, “Apache Camel Message Filter Using XML”](#) shows how to define an equivalent route using Apache Camel XML configuration.

Example A.8. Apache Camel Message Filter Using XML

```
<route>
  <from
    uri="jbi:endpoint:http://progress.com/demos/test/messageFilter/endpoint"
  >
    <filter>
      <xpath>count(/test:world) = 1</xpath>
```

```

<to uri="jbi:service:http://progress.com/demos/test/trace3"/>
</filter>
</route>

```

Equivalent Apache Camel Java DSL route

Example A.9, “[Apache Camel Message Filter Using Java DSL](#)” shows how to define an equivalent route using the Apache Camel Java DSL.

Example A.9. Apache Camel Message Filter Using Java DSL

```

from("jbi:endpoint:http://progress.com/demos/test/messageFilter/endpoint")
    .
    filter(xpath("count(/test:world) = 1")).
    to("jbi:service:http://progress.com/demos/test/trace3");


```

A.7. PIPELINE

Overview

The ServiceMix EIP *pipeline* pattern, shown in [Figure A.4, “Pipes and Filters Pattern”](#), is used to pass messages through a single transformer endpoint, where the transformer's input is taken from the source endpoint and the transformer's output is routed to the target endpoint. This pattern is thus a special case of the more general Apache Camel [pipes and filters](#) pattern, which enables you to pass an *n* message through *multiple* transformer endpoints.

Figure A.4. Pipes and Filters Pattern



Example ServiceMix EIP route

Example A.10, “[ServiceMix EIP Pipeline](#)” shows how to define a pipeline using the ServiceMix EIP component. Incoming messages are passed into the transformer endpoint, **test:decrypt**, and the output from the transformer endpoint is then passed into the target endpoint, **test:plaintextOrder**.

Example A.10. ServiceMix EIP Pipeline

```

<eip:pipeline service="test:pipeline" endpoint="endpoint">
  <eip:transformer>
    <eip:exchange-target service="test:decrypt" />
  </eip:transformer>
  <eip:target>

```

```

    <eip:exchange-target service="test:plaintextOrder" />
  </eip:target>
</eip:pipeline>

```

Equivalent Apache Camel XML route

Example A.11, “[Apache Camel Pipeline Using XML](#)” shows how to define an equivalent route using Apache Camel XML configuration.

Example A.11. Apache Camel Pipeline Using XML

```

<route>
  <from
uri="jbi:endpoint:http://progress.com/demos/test/pipeline/endpoint"/>
  <to uri="jbi:service:http://progress.com/demos/test/decrypt"/>
  <to uri="jbi:service:http://progress.com/demos/test/plaintextOrder"/>
</route>

```

Equivalent Apache Camel Java DSL route

Example A.12, “[Apache Camel Pipeline Using Java DSL](#)” shows how to define an equivalent route using the Apache Camel Java DSL.

Example A.12. Apache Camel Pipeline Using Java DSL

```

from("jbi:endpoint:http://progress.com/demos/test/pipeline/endpoint").
  pipeline("jbi:service:http://progress.com/demos/test/decrypt",
"jbi:service:http://progress.com/demos/test/plaintextOrder");


```

A.8. RESEQUENCER

Overview

The *resequencer* pattern, shown in [Figure A.5, “Resequencer Pattern”](#), enables you to resequence messages according to the sequence number stored in an NMR property. The ServiceMix EIP resequencer pattern maps to the Apache Camel [resequencer](#) configured with the *stream resequencing* algorithm.

Figure A.5. Resequencer Pattern



Sequence number property

The sequence of messages emitted from the resequencer is determined by the value of the sequence number property: messages with a low sequence number are emitted first and messages with a higher number are emitted later. By default, the sequence number is read from the `org.apache.servicemix.eip.sequence.number` property in ServiceMix, but you can customize the name of this property using the `eip:default-comparator` element in ServiceMix.

The equivalent concept in Apache Camel is a *sequencing expression*, which can be any message-dependent expression. When migrating from ServiceMix EIP, you normally define an expression that extracts the sequence number from a header (a Apache Camel header is equivalent to an NMR message property). For example, to extract a sequence number from a `seqnum` header, you can use the simple expression, `header.seqnum`.

Example ServiceMix EIP route

[Example A.13, “ServiceMix EIP Resequencer”](#) shows how to define a resequencer using the ServiceMix EIP component.

Example A.13. ServiceMix EIP Resequencer

```
<eip:resequencer
  service="sample:Resequencer"
  endpoint="ResequencerEndpoint"
  comparator="#comparator"
  capacity="100"
  timeout="2000">
  <eip:target>
    <eip:exchange-target service="sample:SampleTarget" />
  </eip:target>
</eip:resequencer>

<!-- Configure default comparator with custom sequence number property -
->
<eip:default-comparator xml:id="comparator"
  sequenceNumberKey="seqnum"/>
```

Equivalent Apache Camel XML route

[Example A.14, “Apache Camel Resequencer Using XML”](#) shows how to define an equivalent route using Apache Camel XML configuration.

Example A.14. Apache Camel Resequencer Using XML

```
<route>
  <from uri="jbi:endpoint:sample:Resequencer:ResequencerEndpoint"/>
  <resequencer>
    <simple>header.seqnum</simple>
    <to uri="jbi:service:sample:SampleTarget" />
    <stream-config capacity="100" timeout="2000"/>
  </resequencer>
</route>
```

Equivalent Apache Camel Java DSL route

Example A.15, “Apache Camel Resequencer Using Java DSL” shows how to define an equivalent route using the Apache Camel Java DSL.

Example A.15. Apache Camel Resequencer Using Java DSL

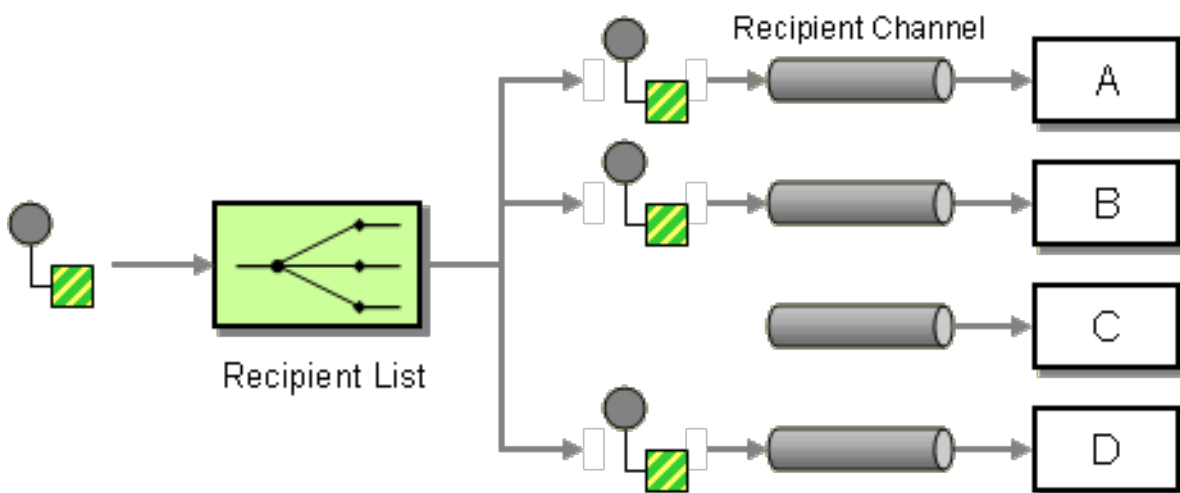
```
from("jbi:endpoint:sample:Resequencer:ResequencerEndpoint").
    resequencer(header("seqnum")).
    stream(new StreamResequencerConfig(100, 2000L)).
    to("jbi:service:sample:SampleTarget");
```

A.9. STATIC RECIPIENT LIST

Overview

A *recipient list*, shown in Figure A.6, “Static Recipient List Pattern”, is a type of router that sends each incoming message to multiple different destinations. The ServiceMix EIP recipient list is restricted to processing *InOnly* and *RobustInOnly* exchange patterns. Moreover, the list of recipients must be static. This pattern maps to the *recipient list* with fixed destination pattern in Apache Camel.

Figure A.6. Static Recipient List Pattern



Example ServiceMix EIP route

Example A.16, “ServiceMix EIP Static Recipient List” shows how to define a static recipient list using the ServiceMix EIP component. Incoming messages are copied to the `test:messageFilter` endpoint and to the `test:trace4` endpoint.

Example A.16. ServiceMix EIP Static Recipient List

```
<eip:static-recipient-list service="test:recipients"
    endpoint="endpoint">
    <eip:recipients>
```

```

    <eip:exchange-target service="test:messageFilter" />
    <eip:exchange-target service="test:trace4" />
  </eip:recipients>
</eip:static-recipient-list>

```

Equivalent Apache Camel XML route

Example A.17, “Apache Camel Static Recipient List Using XML” shows how to define an equivalent route using Apache Camel XML configuration.

Example A.17. Apache Camel Static Recipient List Using XML

```

<route>
  <from
uri="jbi:endpoint:http://progress.com/demos/test/recipients/endpoint"/>
  <to uri="jbi:service:http://progress.com/demos/test/messageFilter"/>
  <to uri="jbi:service:http://progress.com/demos/test/trace4"/>
</route>

```



NOTE

The preceding route configuration appears to have the same syntax as a Apache Camel pipeline pattern. The key difference is that the preceding route is intended for processing *InOnly* message exchanges, which are processed in a different way. See [Section 4.4, “Pipes and Filters”](#) for more details.

Equivalent Apache Camel Java DSL route

Example A.18, “Apache Camel Static Recipient List Using Java DSL” shows how to define an equivalent route using the Apache Camel Java DSL.

Example A.18. Apache Camel Static Recipient List Using Java DSL

```

from("jbi:endpoint:http://progress.com/demos/test/recipients/endpoint").
  to("jbi:service:http://progress.com/demos/test/messageFilter",
    "jbi:service:http://progress.com/demos/test/trace4");

```

A.10. STATIC ROUTING SLIP

Overview

The *static routing slip* pattern in the ServiceMix EIP component is used to route an *InOut* message exchange through a series of endpoints. Semantically, it is equivalent to the [pipeline](#) pattern in Apache Camel.

Example ServiceMix EIP route

Example A.19, “ServiceMix EIP Static Routing Slip” shows how to define a static routing slip

using the ServiceMix EIP component. Incoming messages pass through each of the endpoints, **test:procA**, **test:procB**, and **test:procC**, where the output of each endpoint is connected to the input of the next endpoint in the chain. The final endpoint, **test:procC**, sends its output (*Out* message) back to the caller.

Example A.19. ServiceMix EIP Static Routing Slip

```
<eip:static-routing-slip service="test:routingSlip"
                        endpoint="endpoint">
  <eip:targets>
    <eip:exchange-target service="test:procA" />
    <eip:exchange-target service="test:procB" />
    <eip:exchange-target service="test:procC" />
  </eip:targets>
</eip:static-routing-slip>
```

Equivalent Apache Camel XML route

[Example A.20, “Apache Camel Static Routing Slip Using XML”](#) shows how to define an equivalent route using Apache Camel XML configuration.

Example A.20. Apache Camel Static Routing Slip Using XML

```
<route>
  <from
uri="jbi:endpoint:http://progress.com/demos/test/routingSlip/endpoint"/>
  <to uri="jbi:service:http://progress.com/demos/test/procA"/>
  <to uri="jbi:service:http://progress.com/demos/test/procB"/>
  <to uri="jbi:service:http://progress.com/demos/test/procC"/>
</route>
```

Equivalent Apache Camel Java DSL route

[Example A.21, “Apache Camel Static Routing Slip Using Java DSL”](#) shows how to define an equivalent route using the Apache Camel Java DSL.

Example A.21. Apache Camel Static Routing Slip Using Java DSL


```
from("jbi:endpoint:http://progress.com/demos/test/routingSlip/endpoint")
    .
    pipeline("jbi:service:http://progress.com/demos/test/procA",
            "jbi:service:http://progress.com/demos/test/procB",
            "jbi:service:http://progress.com/demos/test/procC");
```

A.11. WIRE TAP

Overview

The *wire tap* pattern, shown in [Figure A.7, “Wire Tap Pattern”](#), allows you to route messages to a separate tap location before it is forwarded to the ultimate destination. The ServiceMix EIP wire tap pattern maps to the [wire tap](#) pattern in Apache Camel.

Figure A.7. Wire Tap Pattern



Example ServiceMix EIP route

[Example A.22, “ServiceMix EIP Wire Tap”](#) shows how to define a wire tap using the ServiceMix EIP component. The *In* message from the source endpoint is copied to the *In*-listener endpoint, before being forwarded on to the target endpoint. If you want to monitor any returned *Out* messages or *Fault* messages from the target endpoint, you also must define an *Out* listener (using the `eip:outListener` element) and a *Fault* listener (using the `eip:faultListener` element).

Example A.22. ServiceMix EIP Wire Tap

```

<eip:wire-tap service="test:wireTap" endpoint="endpoint">
  <eip:target>
    <eip:exchange-target service="test:target" />
  </eip:target>
  <eip:inListener>
    <eip:exchange-target service="test:trace1" />
  </eip:inListener>
</eip:wire-tap>

```

Equivalent Apache Camel XML route

[Example A.23, “Apache Camel Wire Tap Using XML”](#) shows how to define an equivalent route using Apache Camel XML configuration.

Example A.23. Apache Camel Wire Tap Using XML

```

<route>
  <from
    uri="jbi:endpoint:http://progress.com/demos/test/wireTap/endpoint"/>
    <to uri="jbi:service:http://progress.com/demos/test/trace1"/>
    <to uri="jbi:service:http://progress.com/demos/test/target"/>
  </route>

```


Equivalent Apache Camel Java DSL route

Example A.24, “Apache Camel Wire Tap Using Java DSL” shows how to define an equivalent route using the Apache Camel Java DSL.

Example A.24. Apache Camel Wire Tap Using Java DSL


```
from("jbi:endpoint:http://progress.com/demos/test/wireTap/endpoint")
    .to("jbi:service:http://progress.com/demos/test/trace1",
        "jbi:service:http://progress.com/demos/test/target");
```

A.12. XPATH SPLITTER

Overview

A *splitter*, shown in Figure A.8, “XPath Splitter Pattern”, is a type of router that splits an incoming message into a series of outgoing messages, where each of the messages contains a piece of the original message. The ServiceMix EIP XPath splitter pattern is restricted to using the *InOnly* and *RobustInOnly* exchange patterns. The expression that defines how to split up the original message is defined in the XPath language. The XPath splitter pattern maps to the [splitter](#) pattern in Apache Camel.

Figure A.8. XPath Splitter Pattern



Forwarding NMR attachments and properties

The `eip:xpath-splitter` element supports a `forwardAttachments` attribute and a `forwardProperties` attribute, either of which can be set to `true`, if you want the splitter to copy the incoming message's attachments or properties to the outgoing messages. The corresponding splitter pattern in Apache Camel does not support any such attributes. By default, the incoming message's headers are copied to each of the outgoing messages by the Apache Camel splitter.

Example ServiceMix EIP route

Example A.25, “ServiceMix EIP XPath Splitter” shows how to define a splitter using the ServiceMix EIP component. The specified XPath expression, `/*/*`, causes an incoming message to split at every occurrence of a nested XML element (for example, the `/foo/bar` and `/foo/car` elements are split into distinct messages).

Example A.25. ServiceMix EIP XPath Splitter

```

<eip:xpath-splitter service="test:xpathSplitter"
                  endpoint="endpoint"
                  xpath="/*/*"
                  namespaceContext="#nsContext">
  <eip:target>
    <eip:exchange-target uri="service:http://test/router" />
  </eip:target>
</eip:xpath-splitter>

```

Equivalent Apache Camel XML route

[Example A.26, “Apache Camel XPath Splitter Using XML”](#) shows how to define an equivalent route using Apache Camel XML configuration.

Example A.26. Apache Camel XPath Splitter Using XML

```

<route>
  <from
    uri="jbi:endpoint:http://progress.com/demos/test/xpathSplitter/endpoint"
  />
  <splitter>
    <xpath>/*/*</xpath>
    <to uri="jbi:service:http://test/router"/>
  </splitter>
</route>

```

Equivalent Apache Camel Java DSL route

[Example A.27, “Apache Camel XPath Splitter Using Java DSL”](#) shows how to define an equivalent route using the Apache Camel Java DSL.

Example A.27. Apache Camel XPath Splitter Using Java DSL

```

from("jbi:endpoint:http://progress.com/demos/test/xpathSplitter/endpoint")
    .splitter(xpath("/*/*")).to("jbi:service:http://test/router");

```

PART II. ROUTING EXPRESSION AND PREDICATE LANGUAGES

Abstract

This guide describes the basic syntax used by the evaluative languages supported by Apache Camel.

CHAPTER 11. INTRODUCTION

Abstract

This chapter provides an overview of all the expression languages supported by Apache Camel.

11.1. OVERVIEW OF THE LANGUAGES

Table of expression and predicate languages

Table 11.1, “Expression and Predicate Languages” gives an overview of the different syntaxes for invoking expression and predicate languages.

Table 11.1. Expression and Predicate Languages

Language	Static Method	Fluent DSL Method	XML Element	Annotation	Artifact
Section 2.4, “Bean Integration”	<code>bean()</code>	<code>EIP().method()</code>	<code>method</code>	<code>@Bean</code>	<i>Camel core</i>
Constant	<code>constant()</code>	<code>EIP().constant()</code>	<code>constant</code>	<code>@Constant</code>	<i>Camel core</i>
EL	<code>el()</code>	<code>EIP().el()</code>	<code>el</code>	<code>@EL</code>	<code>camel-juel</code>
Groovy	<code>groovy()</code>	<code>EIP().groovy()</code>	<code>groovy</code>	<code>@Groovy</code>	<code>camel-groovy</code>
Header	<code>header()</code>	<code>EIP().header()</code>	<code>header</code>	<code>@Header</code>	<i>Camel core</i>
JavaScript	<code>javascript()</code>	<code>EIP().javascript()</code>	<code>javascript</code>	<code>@JavaScript</code>	<code>camel-script</code>
JoSQL	<code>sql()</code>	<code>EIP().sql()</code>	<code>sql</code>	<code>@SQL</code>	<code>camel-josql</code>
XPath	<i>None</i>	<code>EIP().xpath()</code>	<code>xpath</code>	<code>@XPath</code>	<code>camel-jxpath</code>
MVEL	<code>mvel()</code>	<code>EIP().mvel()</code>	<code>mvel</code>	<code>@MVEL</code>	<code>camel-mvel</code>
OGNL	<code>ognl()</code>	<code>EIP().ognl()</code>	<code>ognl</code>	<code>@OGNL</code>	<code>camel-ognl</code>

Language	Static Method	Fluent DSL Method	XML Element	Annotation	Artifact
PHP	<code>php()</code>	<code>EIP().php()</code>	<code>php</code>	<code>@PHP</code>	<code>camel-script</code>
Property	<code>property()</code>	<code>EIP().property()</code>	<code>property</code>	<code>@Property</code>	<i>Camel core</i>
Python	<code>python()</code>	<code>EIP().python()</code>	<code>python</code>	<code>@Python</code>	<code>camel-script</code>
Ref	<code>ref()</code>	<code>EIP().ref()</code>	<code>ref</code>	<i>N/A</i>	<i>Camel core</i>
Ruby	<code>ruby()</code>	<code>EIP().ruby()</code>	<code>ruby</code>	<code>@Ruby</code>	<code>camel-script</code>
Simple/File	<code>simple()</code>	<code>EIP().simple()</code>	<code>simple</code>	<code>@Simple</code>	<i>Camel core</i>
SpEL	<code>spel()</code>	<code>EIP().spel()</code>	<code>spel</code>	<code>@SpEL</code>	<code>camel-spring</code>
XPath	<code>xpath()</code>	<code>EIP().xpath()</code>	<code>xpath</code>	<code>@XPath</code>	<i>Camel core</i>
XQuery	<code>xquery()</code>	<code>EIP().xquery()</code>	<code>xquery</code>	<code>@XQuery</code>	<code>camel-saxon</code>

11.2. HOW TO INVOKE AN EXPRESSION LANGUAGE

Prerequisites

Before you can use a particular expression language, you must ensure that the required JAR files are available on the classpath. If the language you want to use is not included in the Apache Camel core, you must add the relevant JARs to your classpath.

If you are using the Maven build system, you can modify the build-time classpath simply by adding the relevant dependency to your POM file. For example, if you want to use the Ruby language, add the following dependency to your POM file:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-script</artifactId>
  <!-- Use the same version as your Camel core version -->
  <version>${camel.version}</version>
</dependency>
```

If you are going to deploy your application in a Red Hat JBoss Fuse OSGi container, you also

need to ensure that the relevant language features are installed (features are named after the corresponding Maven artifact). For example, to use the Groovy language in the OSGi container, you must first install the **camel-groovy** feature by entering the following OSGi console command:

```
karaf@root> features:install camel-groovy
```

Approaches to invoking

As shown in [Table 11.1, “Expression and Predicate Languages”](#), there are several different syntaxes for invoking an expression language, depending on the context in which it is used. You can invoke an expression language:

- the section called “As a static method”.
- the section called “As a fluent DSL method”.
- the section called “As an XML element”.
- the section called “As an annotation”.

As a static method

Most of the languages define a static method that can be used in *any* context where an **org.apache.camel.Expression** type or an **org.apache.camel.Predicate** type is expected. The static method takes a string expression (or predicate) as its argument and returns an **Expression** object (which is usually also a **Predicate** object).

For example, to implement a content-based router that processes messages in XML format, you could route messages based on the value of the **/order/address/countryCode** element, as follows:

```
from("SourceURL")
  .choice
    .when(xpath("/order/address/countryCode = 'us'"))
      .to("file://countries/us/")
    .when(xpath("/order/address/countryCode = 'uk'"))
      .to("file://countries/uk/")
    .otherwise()
      .to("file://countries/other/")
  .to("TargetURL");
```

As a fluent DSL method

The Java fluent DSL supports another style of invoking expression languages. Instead of providing the expression as an argument to an Enterprise Integration Pattern (EIP), you can provide the expression as a sub-clause of the DSL command. For example, instead of invoking an XPath expression as **filter(xpath("Expression"))**, you can invoke the expression as, **filter().xpath("Expression")**.

For example, the preceding content-based router can be re-implemented in this style of invocation, as follows:

```
from("SourceURL")
```

```

.choice
  .when().xpath("/order/address/countryCode = 'us'")
    .to("file://countries/us/")
  .when().xpath("/order/address/countryCode = 'uk'")
    .to("file://countries/uk/")
  .otherwise()
    .to("file://countries/other/")
.to("TargetURL");

```

As an XML element

You can also invoke an expression language in XML, by putting the expression string inside the relevant XML element.

For example, the XML element for invoking XPath in XML is **xpath** (which belongs to the standard Apache Camel namespace). You can use XPath expressions in a XML DSL content-based router, as follows:

```

<from uri="file://input/orders"/>
<choice>
  <when>
    <xpath>/order/address/countryCode = 'us'</xpath>
    <to uri="file://countries/us/">
  </when>
  <when>
    <xpath>/order/address/countryCode = 'uk'</xpath>
    <to uri="file://countries/uk/">
  </when>
  <otherwise>
    <to uri="file://countries/other/">
  </otherwise>
</choice>

```

Alternatively, you can specify a language expression using the **language** element, where you specify the name of the language in the **language** attribute. For example, you can define an XPath expression using the **language** element as follows:

```

<language language="xpath">/order/address/countryCode = 'us'</language>

```

As an annotation

Language annotations are used in the context of bean integration (see [Section 2.4, “Bean Integration”](#)). The annotations provide a convenient way of extracting information from a message or header and then injecting the extracted data into a bean's method parameters.

For example, consider the bean, **myBeanProc**, which is invoked as a predicate of the **filter()** EIP. If the bean's **checkCredentials** method returns **true**, the message is allowed to proceed; but if the method returns **false**, the message is blocked by the filter. The filter pattern is implemented as follows:

```

// Java
MyBeanProcessor myBeanProc = new MyBeanProcessor();

```

```
from("SourceURL")
  .filter().method(myBeanProc, "checkCredentials")
  .to("TargetURL");
```

The implementation of the **MyBeanProcessor** class exploits the **@XPath** annotation to extract the **username** and **password** from the underlying XML message, as follows:

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

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

The **@XPath** annotation is placed just before the parameter into which it gets injected. Notice how the XPath expression *explicitly* selects the text node, by appending **text()** to the path, which ensures that just the content of the element is selected, not the enclosing tags.

As a Camel endpoint URI

Using the Camel Language component, you can invoke a supported language in an endpoint URI. There are two alternative syntaxes.

To invoke a language script stored in a file (or other resource type defined by **Scheme**), use the following URI syntax:

```
language://LanguageName:resource:Scheme:Location[?Options]
```

Where the scheme can be **file:**, **classpath:**, or **http:**.

For example, the following route executes the **mysimplescript.txt** from the classpath:

```
from("direct:start")
  .to("language:simple:classpath:org/apache/camel/component/language/mysimplescript.txt")
  .to("mock:result");
```

To invoke an embedded language script, use the following URI syntax:

```
language://LanguageName[:Script][?Options]
```

For example, to run the Simple language script stored in the **script** string:

```
String script = URLEncoder.encode("Hello ${body}", "UTF-8");
from("direct:start")
  .to("language:simple:" + script)
```



```
    .to("mock:result");
```

CHAPTER 12. CONSTANT

OVERVIEW

The constant language is a trivial built-in language that is used to specify a plain text string. This makes it possible to provide a plain text string in any context where an expression type is expected.

XML EXAMPLE

In XML, you can set the **username** header to the value, **Jane Doe** as follows:

```
<camelContext>
  <route>
    <from uri="SourceURL" />
    <setHeader headerName="username">
      <constant>Jane Doe</constant>
    </setHeader>
    <to uri="TargetURL" />
  </route>
</camelContext>
```

JAVA EXAMPLE

In Java, you can set the **username** header to the value, **Jane Doe** as follows:

```
from("SourceURL")
  .setHeader("username", constant("Jane Doe"))
  .to("TargetURL");
```

CHAPTER 13. EL

OVERVIEW

The Unified Expression Language (EL) was originally specified as part of the JSP 2.1 standard (JSR-245), but it is now available as a standalone language. Apache Camel integrates with JUEL (<http://juel.sourceforge.net/>), which is an open source implementation of the EL language.

ADDING JUEL PACKAGE

To use EL in your routes you need to add a dependency on **camel-juel** to your project as shown in [Example 13.1, “Adding the camel-juel dependency”](#).

Example 13.1. Adding the camel-juel dependency

```
<!-- Maven POM File -->
<properties>
  <camel-version>2.12.0.redhat-610379</camel-version>
  ...
</properties>

<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-juel</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

STATIC IMPORT

To use the `el()` static method in your application code, include the following import statement in your Java source files:

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

VARIABLES

[Table 13.1, “EL variables”](#) lists the variables that are accessible when using EL.

Table 13.1. EL variables

Variable	Type	Value
<code>exchange</code>	<code>org.apache.camel.Exchange</code>	The current Exchange

Variable	Type	Value
<code>in</code>	<code>org.apache.camel.Message</code>	The IN message
<code>out</code>	<code>org.apache.camel.Message</code>	The OUT message

EXAMPLE

Example 13.2, “Routes using EL” shows two routes that use EL.

Example 13.2. Routes using EL

```
<camelContext>
  <route>
    <from uri="seda:foo"/>
    <filter>
      <language language="el">${in.headers.foo == 'bar'}</language>
      <to uri="seda:bar"/>
    </filter>
  </route>
  <route>
    <from uri="seda:foo2"/>
    <filter>
      <language language="el">${in.headers['My Header'] ==
'bar'}</language>
      <to uri="seda:bar"/>
    </filter>
  </route>
</camelContext>
```

CHAPTER 14. THE FILE LANGUAGE

Abstract

The file language is an extension to the simple language, not an independent language in its own right. But the file language extension can only be used in conjunction with File or FTP endpoints.

14.1. WHEN TO USE THE FILE LANGUAGE

Overview

The file language is an extension to the simple language which is not always available. You can use it under the following circumstances:

- [the section called “In a File or FTP consumer endpoint”](#)
- [the section called “On exchanges created by a File or FTP consumer”](#)



NOTE

The escape character, `\`, is not available in the file language.

In a File or FTP consumer endpoint

There are several URI options that you can set on a File or FTP consumer endpoint, which take a file language expression as their value. For example, in a File consumer endpoint URI you can set the **fileName**, **move**, **preMove**, **moveFailed**, and **sortBy** options using a file expression.

In a File consumer endpoint, the **fileName** option acts as a filter, determining which file will actually be read from the starting directory. If a plain text string is specified (for example, **fileName=report.txt**), the File consumer reads the same file each time it is updated. You can make this option more dynamic, however, by specifying a simple expression. For example, you could use a counter bean to select a different file each time the File consumer polls the starting directory, as follows:

```
file://target/filelanguage/bean/?
fileName=${bean:counter.next}.txt&delete=true
```

Where the **\${bean:counter.next}** expression invokes the **next()** method on the bean registered under the ID, **counter**.

The **move** option is used to move files to a backup location after they have been read by a File consumer endpoint. For example, the following endpoint moves files to a backup directory, after they have been processed:

```
file://target/filelanguage/?
move=backup/${date:now:yyyyMMdd}/${file:name.noext}.bak&recursive=false
```

Where the **\${file:name.noext}.bak** expression modifies the original file name, replacing the file extension with **.bak**.

You can use the **sortBy** option to specify the order in which file should be processed. For example, to process files according to the alphabetical order of their file name, you could use the following File consumer endpoint:

```
file://target/filelanguage/?sortBy=file:name
```

To process file according to the order in which they were last modified, you could use the following File consumer endpoint:

```
file://target/filelanguage/?sortBy=file:modified
```

You can reverse the order by adding the **reverse:** prefix—for example:

```
file://target/filelanguage/?sortBy=reverse:file:modified
```

On exchanges created by a File or FTP consumer

When an exchange originates from a File or FTP consumer endpoint, it is possible to apply file language expressions to the exchange throughout the route (as long as the original message headers are not erased). For example, you could define a content-based router, which routes messages according to their file extension, as follows:

```
<from uri="file://input/orders"/>
<choice>
  <when>
    <simple>${file:ext} == 'txt'</simple>
    <to uri="bean:orderService?method=handleTextFiles"/>
  </when>
  <when>
    <simple>${file:ext} == 'xml'</simple>
    <to uri="bean:orderService?method=handleXmlFiles"/>
  </when>
  <otherwise>
    <to uri="bean:orderService?method=handleOtherFiles"/>
  </otherwise>
</choice>
```

14.2. FILE VARIABLES

Overview

File variables can be used whenever a route starts with a File or FTP consumer endpoint, which implies that the underlying message body is of **java.io.File** type. The file variables enable you to access various parts of the file pathname, almost as if you were invoking the methods of the **java.io.File** class (in fact, the file language extracts the information it needs from message headers that have been set by the File or FTP endpoint).

Starting directory

Some of file variables return paths that are defined relative to a *starting directory*, which is just the directory that is specified in the File or FTP endpoint. For example, the following File consumer endpoint has the starting directory, **./filetransfer** (a relative path):

```
file:filetransfer
```

The following FTP consumer endpoint has the starting directory, `./ftptransfer` (a relative path):

```
ftp://myhost:2100/ftptransfer
```

Naming convention of file variables

In general, the file variables are named after corresponding methods on the `java.io.File` class. For example, the `file:absolute` variable gives the value that would be returned by the `java.io.File.getAbsolute()` method.



NOTE

This naming convention is not strictly followed, however. For example, there is *no* such method as `java.io.File.getSize()`.

Table of variables

Table 14.1, “Variables for the File Language” shows all of the variable supported by the file language.

Table 14.1. Variables for the File Language

Variable	Type	Description
<code>file:name</code>	<code>String</code>	The pathname relative to the starting directory.
<code>file:name.ext</code>	<code>String</code>	The file extension (characters following the last <code>.</code> character in the pathname).
<code>file:name.noext</code>	<code>String</code>	The pathname relative to the starting directory, omitting the file extension.
<code>file:onlyname</code>	<code>String</code>	The final segment of the pathname. That is, the file name without the parent directory path.
<code>file:onlyname.noext</code>	<code>String</code>	The final segment of the pathname, omitting the file extension.
<code>file:ext</code>	<code>String</code>	The file extension (same as <code>file:name.ext</code>).

Variable	Type	Description
file:parent	String	The pathname of the parent directory, including the starting directory in the path.
file:path	String	The file pathname, including the starting directory in the path.
file:absolute	Boolean	true , if the starting directory was specified as an absolute path; false , otherwise.
file:absolute.path	String	The absolute pathname of the file.
file:length	Long	The size of the referenced file.
file:size	Long	Same as file:length .
file:modified	java.util.Date	Date last modified.

14.3. EXAMPLES

Relative pathname

Consider a File consumer endpoint, where the starting directory is specified as a *relative pathname*. For example, the following File endpoint has the starting directory, **./filelanguage**:

```
file://filelanguage
```

Now, while scanning the **filelanguage** directory, suppose that the endpoint has just consumed the following file:

```
./filelanguage/test/hello.txt
```

And, finally, assume that the **filelanguage** directory itself has the following absolute location:

```
/workspace/camel/camel-core/target/filelanguage
```

Given the preceding scenario, the file language variables return the following values, when applied to the current exchange:

Expression	Result
<code>file:name</code>	<code>test/hello.txt</code>
<code>file:name.ext</code>	<code>txt</code>
<code>file:name.noext</code>	<code>test/hello</code>
<code>file:onlyname</code>	<code>hello.txt</code>
<code>file:onlyname.noext</code>	<code>hello</code>
<code>file:ext</code>	<code>txt</code>
<code>file:parent</code>	<code>filelanguage/test</code>
<code>file:path</code>	<code>filelanguage/test/hello.txt</code>
<code>file:absolute</code>	<code>false</code>
<code>file:absolute.path</code>	<code>/workspace/camel/camel-core/target/filelanguage/test/hello.txt</code>

Absolute pathname

Consider a File consumer endpoint, where the starting directory is specified as an *absolute pathname*. For example, the following File endpoint has the starting directory, `/workspace/camel/camel-core/target/filelanguage`:

```
file:///workspace/camel/camel-core/target/filelanguage
```

Now, while scanning the `filelanguage` directory, suppose that the endpoint has just consumed the following file:

```
./filelanguage/test/hello.txt
```

Given the preceding scenario, the file language variables return the following values, when applied to the current exchange:

Expression	Result
<code>file:name</code>	<code>test/hello.txt</code>
<code>file:name.ext</code>	<code>txt</code>
<code>file:name.noext</code>	<code>test/hello</code>

Expression	Result
<code>file:onlyname</code>	<code>hello.txt</code>
<code>file:onlyname.noext</code>	<code>hello</code>
<code>file:ext</code>	<code>txt</code>
<code>file:parent</code>	<code>/workspace/camel/camel-core/target/filelanguage/test</code>
<code>file:path</code>	<code>/workspace/camel/camel-core/target/filelanguage/test/hello.txt</code>
<code>file:absolute</code>	<code>true</code>
<code>file:absolute.path</code>	<code>/workspace/camel/camel-core/target/filelanguage/test/hello.txt</code>

CHAPTER 15. GROOVY

OVERVIEW

Groovy is a Java-based scripting language that allows quick parsing of object. The Groovy support is part of the `camel-script` module.

ADDING THE SCRIPT MODULE

To use Groovy in your routes you need to add a dependency on `camel-script` to your project as shown in [Example 15.1, “Adding the camel-script dependency”](#).

Example 15.1. Adding the camel-script dependency

```
<!-- Maven POM File -->
<properties>
  <camel-version>2.12.0.redhat-610379</camel-version>
  ...
</properties>

<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-script</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

STATIC IMPORT

To use the `groovy()` static method in your application code, include the following import statement in your Java source files:

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

BUILT-IN ATTRIBUTES

[Table 15.1, “Groovy attributes”](#) lists the built-in attributes that are accessible when using Groovy.

Table 15.1. Groovy attributes

Attribute	Type	Value
<code>context</code>	<code>org.apache.camel.CamelContext</code>	The Camel Context

Attribute	Type	Value
exchange	org.apache.camel.Exchange	The current Exchange
request	org.apache.camel.Message	The IN message
response	org.apache.camel.Message	The OUT message
properties	org.apache.camel.builder.script.PropertiesFunction	Function with a resolve method to make it easier to use the properties component inside scripts.

The attributes all set at **ENGINE_SCOPE**.

EXAMPLE

Example 15.2, “Routes using Groovy” shows two routes that use Groovy scripts.

Example 15.2. Routes using Groovy

```
<camelContext>
  <route>
    <from uri=""mock:mock0" />
    <filter>
      <language language="groovy">request.lineItems.any { i -> i.value
> 100 }</language>
      <to uri=""mock:mock1" />
    </filter>
  </route>
  <route>
    <from uri=""direct:in"/>
    <setHeader headerName=""firstName">
      <language language="groovy">$user.firstName
$user.lastName</language>
    </setHeader>
    <to uri=""seda:users"/>
  </route>
</camelContext>
```

USING THE PROPERTIES COMPONENT

To access a property value from the properties component, invoke the **resolve** method on the built-in **properties** attribute, as follows:

```
.setHeader("myHeader").groovy("properties.resolve(PropKey)")
```

Where *PropKey* is the key of the property you want to resolve, where the key value is of **String** type.

For more details about the properties component, see .

CHAPTER 16. HEADER

OVERVIEW

The header language provides a convenient way of accessing header values in the current message. When you supply a header name, the header language performs a case-insensitive lookup and returns the corresponding header value.

The header language is part of **camel-core**.

XML EXAMPLE

For example, to resequence incoming exchanges according to the value of a **SequenceNumber** header (where the sequence number must be a positive integer), you can define a route as follows:

```
<camelContext xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="SourceURL"/>
    <resequence>
      <language language="header">SequenceNumber</language>
    </resequence>
    <to uri="TargetURL"/>
  </route>
</camelContext>
```

JAVA EXAMPLE

The same route can be defined in Java, as follows:

```
from("SourceURL")
  .resequence(header("SequenceNumber"))
  .to("TargetURL");
```

CHAPTER 17. JAVASCRIPT

OVERVIEW

JavaScript, also known as ECMAScript is a Java-based scripting language that allows quick parsing of object. The JavaScript support is part of the **camel-script** module.

ADDING THE SCRIPT MODULE

To use JavaScript in your routes you need to add a dependency on **camel-script** to your project as shown in [Example 17.1, “Adding the camel-script dependency”](#).

Example 17.1. Adding the camel-script dependency

```
<!-- Maven POM File -->
<properties>
  <camel-version>2.12.0.redhat-610379</camel-version>
  ...
</properties>

<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-script</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

STATIC IMPORT

To use the `javaScript()` static method in your application code, include the following import statement in your Java source files:

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

BUILT-IN ATTRIBUTES

[Table 17.1, “JavaScript attributes”](#) lists the built-in attributes that are accessible when using JavaScript.

Table 17.1. JavaScript attributes

Attribute	Type	Value
<code>context</code>	<code>org.apache.camel.CamelContext</code>	The Camel Context

Attribute	Type	Value
exchange	<code>org.apache.camel.Exchange</code>	The current Exchange
request	<code>org.apache.camel.Message</code>	The IN message
response	<code>org.apache.camel.Message</code>	The OUT message
properties	<code>org.apache.camel.builder.script.PropertiesFunction</code>	Function with a resolve method to make it easier to use the properties component inside scripts.

The attributes all set at **ENGINE_SCOPE**.

EXAMPLE

Example 17.2, “Route using JavaScript” shows a route that uses JavaScript.

Example 17.2. Route using JavaScript

```
<camelContext>
  <route>
    <from uri="direct:start" />
    <choice>
      <when>
        <language language="javascript">request.headers.get('user') ==
'admin'</language>
        <to uri="seda:adminQueue" />
      </when>
      <otherwise>
        <to uri="seda:regularQueue" />
      </otherwise>
    </choice>
  </route>
</camelContext>
```

USING THE PROPERTIES COMPONENT

To access a property value from the properties component, invoke the **resolve** method on the built-in **properties** attribute, as follows:

```
.setHeader("myHeader").javascript("properties.resolve(PropKey)")
```

Where *PropKey* is the key of the property you want to resolve, where the key value is of **String** type.

For more details about the properties component, see .

CHAPTER 18. JOSQL

OVERVIEW

The JoSQL (SQL for Java objects) language enables you to evaluate predicates and expressions in Apache Camel. JoSQL employs a SQL-like query syntax to perform selection and ordering operations on data from in-memory Java objects—however, JoSQL is *not* a database. 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/>.

ADDING THE JOSQL MODULE

To use JoSQL in your routes you need to add a dependency on `camel-josql` to your project as shown in [Example 18.1, “Adding the camel-josql dependency”](#).

Example 18.1. Adding the camel-josql dependency

```
<!-- Maven POM File -->
...
<dependencies>
...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-josql</artifactId>
    <version>${camel-version}</version>
  </dependency>
...
</dependencies>
```

STATIC IMPORT

To use the `sql()` static method in your application code, include the following import statement in your Java source files:

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

VARIABLES

[Table 18.1, “SQL variables”](#) lists the variables that are accessible when using JoSQL.

Table 18.1. SQL variables

Name	Type	Description
<code>exchange</code>	<code>org.apache.camel.Exchange</code>	The current Exchange

Name	Type	Description
in	org.apache.camel.Message	The IN message
out	org.apache.camel.Message	The OUT message
<i>property</i>	Object	the Exchange property whose key is <i>property</i>
<i>header</i>	Object	the IN message header whose key is <i>header</i>
<i>variable</i>	Object	the variable whose key is <i>variable</i>

EXAMPLE

Example 18.2, “Route using JoSQL” shows a route that uses JoSQL.

Example 18.2. Route using JoSQL

```
<camelContext>
  <route>
    <from uri="direct:start"/>
    <setBody>
      <language language="sql">select * from MyType</language>
    </setBody>
    <to uri="seda:regularQueue"/>
  </route>
</camelContext>
```

CHAPTER 19. JXPath

OVERVIEW

The JXPath language enables you to invoke Java beans using the [Apache Commons JXPath](#) language. The JXPath language has a similar syntax to XPath, but instead of selecting element or attribute nodes from an XML document, it invokes methods on an object graph of Java beans. If one of the bean attributes returns an XML document (a DOM/JDOM instance), however, the remaining portion of the path is interpreted as an XPath expression and is used to extract an XML node from the document. In other words, the JXPath language provides a hybrid of object graph navigation and XML node selection.

ADDING JXPath PACKAGE

To use JXPath in your routes you need to add a dependency on **camel-jxpath** to your project as shown in [Example 19.1, “Adding the camel-jxpath dependency”](#).

Example 19.1. Adding the camel-jxpath dependency

```
<!-- Maven POM File -->
<properties>
  <camel-version>2.12.0.redhat-610379</camel-version>
  ...
</properties>

<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-jxpath</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

VARIABLES

[Table 19.1, “JXPath variables”](#) lists the variables that are accessible when using JXPath.

Table 19.1. JXPath variables

Variable	Type	Value
this	org.apache.camel.Exchange	The current Exchange
in	org.apache.camel.Message	The IN message

Variable	Type	Value
out	org.apache.camel.Message	The OUT message

EXAMPLE

Example 19.2, “Routes using JXPath” shows a route that use JXPath.

Example 19.2. Routes using JXPath

```
<camelContext>
  <route>
    <from uri="activemq:MyQueue"/>
    <filter>
      <jxpath>in/body/name = 'James'</xpath>
      <to uri="mqseries:SomeOtherQueue"/>
    </filter>
  </route>
</camelContext>
```

CHAPTER 20. MVEL

OVERVIEW

MVEL (<http://mvel.codehaus.org/>) is a Java-based dynamic language that is similar to OGNL, but is reported to be much faster. The MVEL support is in the `camel-mvel` module.

SYNTAX

You use the MVEL dot syntax to invoke Java methods, for example:

```
getRequest().getBody().getFamilyName()
```

Because MVEL is dynamically typed, it is unnecessary to cast the message body instance (of `Object` type) before invoking the `getFamilyName()` method. You can also use an abbreviated syntax for invoking bean attributes, for example:

```
request.body.familyName
```

ADDING THE MVEL MODULE

To use MVEL in your routes you need to add a dependency on `camel-mvel` to your project as shown in [Example 20.1, “Adding the camel-mvel dependency”](#).

Example 20.1. Adding the camel-mvel dependency

```
<!-- Maven POM File -->
<properties>
  <camel-version>2.12.0.redhat-610379</camel-version>
  ...
</properties>

<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-mvel</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

BUILT-IN VARIABLES

[Table 20.1, “MVEL variables”](#) lists the built-in variables that are accessible when using MVEL.

Table 20.1. MVEL variables

Name	Type	Description
this	org.apache.camel.Exchange	The current Exchange
exchange	org.apache.camel.Exchange	The current Exchange
exception	Throwable	the Exchange exception (if any)
exchangeID	String	the Exchange ID
fault	org.apache.camel.Message	The Fault message(if any)
request	org.apache.camel.Message	The IN message
response	org.apache.camel.Message	The OUT message
properties	Map	The Exchange properties
property(<i>name</i>)	Object	The value of the named Exchange property
property(<i>name</i>, <i>type</i>)	Type	The typed value of the named Exchange property

EXAMPLE

Example 20.2, “Route using MVEL” shows a route that uses MVEL.

Example 20.2. Route using MVEL

```
<camelContext>
  <route>
    <from uri="seda:foo"/>
    <filter>
      <language language="mvel">request.headers.foo == 'bar'</language>
      <to uri="seda:bar"/>
    </filter>
  </route>
</camelContext>
```

CHAPTER 21. THE OBJECT-GRAPH NAVIGATION LANGUAGE(OGNL)

OVERVIEW

OGNL (<http://www.opensymphony.com/ognl/>) is an expression language for getting and setting properties of Java objects. You use the same expression for both getting and setting the value of a property. The OGNL support is in the **camel-ognl** module.

ADDING THE OGNL MODULE

To use OGNL in your routes you need to add a dependency on **camel-ognl** to your project as shown in [Example 21.1, “Adding the camel-ognl dependency”](#).

Example 21.1. Adding the camel-ognl dependency

```
<!-- Maven POM File -->
...
<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-ognl</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

STATIC IMPORT

To use the **ognl()** static method in your application code, include the following import statement in your Java source files:

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

BUILT-IN VARIABLES

[Table 21.1, “OGNL variables”](#) lists the built-in variables that are accessible when using OGNL.

Table 21.1. OGNL variables

Name	Type	Description
this	org.apache.camel.Exchange	The current Exchange
exchange	org.apache.camel.Exchange	The current Exchange

Name	Type	Description
exception	Throwable	the Exchange exception (if any)
exchangeID	String	the Exchange ID
fault	org.apache.camel.Message	The Fault message(if any)
request	org.apache.camel.Message	The IN message
response	org.apache.camel.Message	The OUT message
properties	Map	The Exchange properties
property(<i>name</i>)	Object	The value of the named Exchange property
property(<i>name</i>, <i>type</i>)	Type	The typed value of the named Exchange property

EXAMPLE

Example 21.2, “Route using OGNL” shows a route that uses OGNL.

Example 21.2. Route using OGNL

```
<camelContext>
  <route>
    <from uri="seda:foo"/>
    <filter>
      <language language="ognl">request.headers.foo == 'bar'</language>
      <to uri="seda:bar"/>
    </filter>
  </route>
</camelContext>
```

CHAPTER 22. PHP

OVERVIEW

PHP is a widely-used general-purpose scripting language that is especially suited for Web development. The PHP support is part of the `camel-script` module.

ADDING THE SCRIPT MODULE

To use PHP in your routes you need to add a dependency on `camel-script` to your project as shown in [Example 22.1, “Adding the camel-script dependency”](#).

Example 22.1. Adding the camel-script dependency

```
<!-- Maven POM File -->
...
<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-script</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

STATIC IMPORT

To use the `php()` static method in your application code, include the following import statement in your Java source files:

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

BUILT-IN ATTRIBUTES

[Table 22.1, “PHP attributes”](#) lists the built-in attributes that are accessible when using PHP.

Table 22.1. PHP attributes

Attribute	Type	Value
<code>context</code>	<code>org.apache.camel.CamelContext</code>	The Camel Context
<code>exchange</code>	<code>org.apache.camel.Exchange</code>	The current Exchange
<code>request</code>	<code>org.apache.camel.Message</code>	The IN message

Attribute	Type	Value
response	org.apache.camel.Message	The OUT message
properties	org.apache.camel.builder.script.PropertiesFunction	Function with a resolve method to make it easier to use the properties component inside scripts.

The attributes all set at **ENGINE_SCOPE**.

EXAMPLE

Example 22.2, “Route using PHP” shows a route that uses PHP.

Example 22.2. Route using PHP

```
<camelContext>
  <route>
    <from uri="direct:start"/>
    <choice>
      <when>
        <language language="php">strpos(request.headers.get('user'),
'admin')!= FALSE</language>
        <to uri="seda:adminQueue"/>
      </when>
      <otherwise>
        <to uri="seda:regularQueue"/>
      </otherwise>
    </choice>
  </route>
</camelContext>
```

USING THE PROPERTIES COMPONENT

To access a property value from the properties component, invoke the **resolve** method on the built-in **properties** attribute, as follows:

```
.setHeader("myHeader").php("properties.resolve(PropKey)")
```

Where *PropKey* is the key of the property you want to resolve, where the key value is of **String** type.

For more details about the properties component, see .

CHAPTER 23. PROPERTY

OVERVIEW

The property language provides a convenient way of accessing *exchange properties*. When you supply a key that matches one of the exchange property names, the property language returns the corresponding value.

The property language is part of **camel-core**.

XML EXAMPLE

For example, to implement the recipient list pattern when the **listOfEndpoints** exchange property contains the recipient list, you could define a route as follows:

```
<camelContext>
  <route>
    <from uri="direct:a"/>
    <recipientList>
      <property>listOfEndpoints</property>
    </recipientList>
  </route>
</camelContext>
```

JAVA EXAMPLE

The same recipient list example can be implemented in Java as follows:

```
from("direct:a").recipientList(property("listOfEndpoints"));
```

CHAPTER 24. PYTHON

OVERVIEW

Python is a remarkably powerful dynamic programming language that is used in a wide variety of application domains. Python is often compared to Tcl, Perl, Ruby, Scheme or Java. The Python support is part of the **camel-script** module.

ADDING THE SCRIPT MODULE

To use Python in your routes you need to add a dependency on **camel-script** to your project as shown in [Example 24.1, “Adding the camel-script dependency”](#).

Example 24.1. Adding the camel-script dependency

```
<!-- Maven POM File -->
...
<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-script</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

STATIC IMPORT

To use the **python()** static method in your application code, include the following import statement in your Java source files:

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

BUILT-IN ATTRIBUTES

[Table 24.1, “Python attributes”](#) lists the built-in attributes that are accessible when using Python.

Table 24.1. Python attributes

Attribute	Type	Value
context	org.apache.camel.CamelContext	The Camel Context
exchange	org.apache.camel.Exchange	The current Exchange

Attribute	Type	Value
request	org.apache.camel.Message	The IN message
response	org.apache.camel.Message	The OUT message
properties	org.apache.camel.builder.script.PropertiesFunction	Function with a resolve method to make it easier to use the properties component inside scripts.

The attributes all set at **ENGINE_SCOPE**.

EXAMPLE

[Example 24.2, “Route using Python”](#) shows a route that uses Python.

Example 24.2. Route using Python

```
<camelContext>
  <route>
    <from uri="direct:start"/>
    <choice>
      <when>
        <language language="python">if request.headers.get('user') =
'admin'</language>
        <to uri="seda:adminQueue"/>
      </when>
      <otherwise>
        <to uri="seda:regularQueue"/>
      </otherwise>
    </choice>
  </route>
</camelContext>
```

USING THE PROPERTIES COMPONENT

To access a property value from the properties component, invoke the **resolve** method on the built-in **properties** attribute, as follows:

```
.setHeader("myHeader").python("properties.resolve(PropKey)")
```

Where *PropKey* is the key of the property you want to resolve, where the key value is of **String** type.

For more details about the properties component, see .

CHAPTER 25. REF

OVERVIEW

The Ref expression language is really just a way to look up a custom [Expression](#) from the [Registry](#). This is particular convenient to use in the XML DSL.

The Ref language is part of **camel-core**.

STATIC IMPORT

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

```
import static org.apache.camel.language.simple.RefLanguage.ref;
```

XML EXAMPLE

For example, the splitter pattern can reference a custom expression using the Ref language, as follows:

```
<beans ...>
  <bean id="myExpression" class="com.mycompany.MyCustomExpression"/>
  ...
  <camelContext>
    <route>
      <from uri="seda:a"/>
      <split>
        <ref>myExpression</ref>
        <to uri="mock:b"/>
      </split>
    </route>
  </camelContext>
</beans>
```

JAVA EXAMPLE

The preceding route can also be implemented in the Java DSL, as follows:

```
from("seda:a")
  .split().ref("myExpression")
  .to("seda:b");
```

CHAPTER 26. RUBY

OVERVIEW

Ruby is a dynamic, open source programming language with a focus on simplicity and productivity. It has an elegant syntax that is natural to read and easy to write. The Ruby support is part of the `camel-script` module.

ADDING THE SCRIPT MODULE

To use Ruby in your routes you need to add a dependency on `camel-script` to your project as shown in [Example 26.1, “Adding the camel-script dependency”](#).

Example 26.1. Adding the camel-script dependency

```
<!-- Maven POM File -->
...
<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-script</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

STATIC IMPORT

To use the `ruby()` static method in your application code, include the following import statement in your Java source files:

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

BUILT-IN ATTRIBUTES

[Table 26.1, “Ruby attributes”](#) lists the built-in attributes that are accessible when using Ruby.

Table 26.1. Ruby attributes

Attribute	Type	Value
<code>context</code>	<code>org.apache.camel.CamelContext</code>	The Camel Context
<code>exchange</code>	<code>org.apache.camel.Exchange</code>	The current Exchange

Attribute	Type	Value
request	org.apache.camel.Message	The IN message
response	org.apache.camel.Message	The OUT message
properties	org.apache.camel.builder.script.PropertiesFunction	Function with a resolve method to make it easier to use the properties component inside scripts.

The attributes all set at **ENGINE_SCOPE**.

EXAMPLE

Example 26.2, “Route using Ruby” shows a route that uses Ruby.

Example 26.2. Route using Ruby

```
<camelContext>
  <route>
    <from uri="direct:start"/>
    <choice>
      <when>
        <language language="ruby">$request.headers['user'] ==
'admin'</language>
        <to uri="seda:adminQueue"/>
      </when>
      <otherwise>
        <to uri="seda:regularQueue"/>
      </otherwise>
    </choice>
  </route>
</camelContext>
```

USING THE PROPERTIES COMPONENT

To access a property value from the properties component, invoke the **resolve** method on the built-in **properties** attribute, as follows:

```
.setHeader("myHeader").ruby("properties.resolve(PropKey)")
```

Where *PropKey* is the key of the property you want to resolve, where the key value is of **String** type.

For more details about the properties component, see .

CHAPTER 27. THE SIMPLE LANGUAGE

Abstract

The simple language is a language that was developed in Apache Camel specifically for the purpose of accessing and manipulating the various parts of an exchange object. The language is not quite as simple as when it was originally created and it now features a comprehensive set of logical operators and conjunctions.

27.1. JAVA DSL

Simple expressions in Java DSL

In the Java DSL, there are two styles for using the `simple()` command in a route. You can either pass the `simple()` command as an argument to a processor, as follows:

```
from("seda:order")
  .filter(simple("${in.header.foo}"))
  .to("mock:fooOrders");
```

Or you can call the `simple()` command as a sub-clause on the processor, for example:

```
from("seda:order")
  .filter()
  .simple("${in.header.foo}")
  .to("mock:fooOrders");
```

Embedding in a string

If you are embedding a simple expression inside a plain text string, you must use the placeholder syntax, `${Expression}`. For example, to embed the `in.header.name` expression in a string:

```
simple("Hello ${in.header.name}, how are you?")
```

Customizing the start and end tokens

From Java, you can customize the start and end tokens (`{` and `}`, by default) by calling the `changeFunctionStartToken` static method and the `changeFunctionEndToken` static method on the `SimpleLanguage` object.

For example, you can change the start and end tokens to `[` and `]` in Java, as follows:

```
// Java
import org.apache.camel.language.simple.SimpleLanguage;
...
SimpleLanguage.changeFunctionStartToken("[");
SimpleLanguage.changeFunctionEndToken("]");
```

**NOTE**

Customizing the start and end tokens affects all Apache Camel applications that share the same **camel-core** library on their classpath. For example, in an OSGi server this might affect many applications; whereas in a Web application (WAR file) it would affect only the Web application itself.

27.2. XML DSL

Simple expressions in XML DSL

In the XML DSL, you can use a simple expression by putting the expression inside a **simple** element. For example, to define a route that performs filtering based on the contents of the **foo** header:

```
<route id="simpleExample">
  <from uri="seda:orders"/>
  <filter>
    <simple>${in.header.foo}</simple>
    <to uri="mock:fooOrders"/>
  </filter>
</route>
```

Alternative placeholder syntax

Sometimes—for example, if you have enabled Spring property placeholders or OSGi blueprint property placeholders—you might find that the **\${Expression}** syntax clashes with another property placeholder syntax. In this case, you can disambiguate the placeholder using the alternative syntax, **\$simple{Expression}**, for the simple expression. For example:

```
<simple>Hello $simple{in.header.name}, how are you?</simple>
```

Customizing the start and end tokens

From XML configuration, you can customize the start and end tokens (**{** and **}**, by default) by overriding the **SimpleLanguage** instance. For example, to change the start and end tokens to **[** and **]**, define a new **SimpleLanguage** bean in your XML configuration file, as follows:

```
<bean id="simple" class="org.apache.camel.language.simple.SimpleLanguage">
  <constructor-arg name="functionStartToken" value="["/>
  <constructor-arg name="functionEndToken" value="]"/>
</bean>
```

**NOTE**

Customizing the start and end tokens affects all Apache Camel applications that share the same **camel-core** library on their classpath. For example, in an OSGi server this might affect many applications; whereas in a Web application (WAR file) it would affect only the Web application itself.

Whitespace and auto-trim in XML DSL

By default, whitespace preceding and following a simple expression is automatically trimmed in XML DSL. So this expression with surrounding whitespace:

```
<transform>
  <simple>
    data=${body}
  </simple>
</transform>
```

is automatically trimmed, so that it is equivalent to this expression (no surrounding whitespace):

```
<transform>
  <simple>data=${body}</simple>
</transform>
```

If you want to include newlines before or after the expression, you can either explicitly add a newline character, as follows:

```
<transform>
  <simple>data=${body}\n</simple>
</transform>
```

or you can switch off auto-trimming, by setting the **trim** attribute to **false**, as follows:

```
<transform trim="false">
  <simple>data=${body}
</simple>
</transform>
```

27.3. INVOKING AN EXTERNAL SCRIPT

Overview

It is possible to execute Simple scripts that are stored in an external resource, as described here.

Syntax for script resource

Use the following syntax to access a Simple script that is stored as an external resource:

```
resource:Scheme:Location
```

Where the **Scheme:** can be either **classpath:**, **file:**, or **http:**.


For example, to read the **mysimple.txt** script from the classpath,

```
simple("resource:classpath:mysimple.txt")
```

27.4. EXPRESSIONS

Overview

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.



NOTE

Since Apache Camel 2.9, you must *always* use the placeholder syntax, `${Expression}`, to return a variable value. It is never permissible to omit the enclosing tokens (`{` and `}`).

Contents of a single variable

You can use the simple language to define string expressions, based on the variables provided. For example, you can use a variable of the form, `in.header.HeaderName`, to obtain the value of the `HeaderName` header, as follows:

```
simple("${in.header.foo}")
```

Variables embedded in a string

You can embed simple variables in a string expression—for example:

```
simple("Received a message from ${in.header.user} on  
$(date:in.header.date:yyyyMMdd).")
```

date and bean variables

As well as providing variables that access all of the different parts of an exchange (see [Table 27.1, “Variables for the Simple Language”](#)), the simple language also provides special variables for formatting dates, `date:command:pattern`, and for calling bean methods, `bean:beanRef`. For example, you can use the date and the bean variables as follows:

```
simple("Todays date is ${date:now:yyyyMMdd}")  
simple("The order type is ${bean:orderService?method=getOrderType}")
```

Specifying the result type

You can specify the result type of an expression explicitly. This is mainly useful for converting the result type to a boolean or numerical type.

In the Java DSL, specify the result type as an extra argument to `simple()`. For example, to return a boolean result, you could evaluate a simple expression as follows:

```
...  
.setHeader("cool", simple("true", Boolean.class))
```

In the XML DSL, specify the result type using the `resultType` attribute. For example:

```
<setHeader headerName="cool">
  <!-- use resultType to indicate that the type should be a
  java.lang.Boolean -->
  <simple resultType="java.lang.Boolean">true</simple>
</setHeader>
```

Nested expressions

Simple expressions can be nested—for example:

```
simple("${header.${bean:headerChooser?method=whichHeader}}")
```

Accessing constants or enums

You can access a bean's constant or enum fields using the following syntax:

```
type:ClassName.Field
```

For example, consider the following Java **enum** type:

```
package org.apache.camel.processor;
...
public enum Customer {
    GOLD, SILVER, BRONZE
}
```

You can access the **Customer** enum fields, as follows:

```
from("direct:start")
  .choice()
    .when().simple("${header.customer} ==
    ${type:org.apache.camel.processor.Customer.GOLD}")
    .to("mock:gold")
    .when().simple("${header.customer} ==
    ${type:org.apache.camel.processor.Customer.SILVER}")
    .to("mock:silver")
    .otherwise()
    .to("mock:other");
```

OGNL expressions

The [Object Graph Navigation Language](#) (OGNL) is a notation for invoking bean methods in a chain-like fashion. If a message body contains a Java bean, you can easily access its bean properties using OGNL notation. For example, if the message body is a Java object with a **getAddress()** accessor, you can access the **Address** object and the **Address** object's properties as follows:

```
simple("${body.address}")
simple("${body.address.street}")
simple("${body.address.zip}")
simple("${body.address.city}")
```

Where the notation, `${body.address.street}`, is shorthand for `${body.getAddress.getStreet}`.

OGNL null-safe operator

You can use the null-safe operator, `?.`, to avoid encountering null-pointer exceptions, in case the body does *not* have an address. For example:

```
simple("${body?.address?.street}")
```

If the body is a `java.util.Map` type, you can look up a value in the map with the key `foo`, using the following notation:

```
simple("${body[foo]?.name}")
```

OGNL list element access

You can also use square brackets notation, `[k]`, to access the elements of a list. For example:

```
simple("${body.address.lines[0]}")
simple("${body.address.lines[1]}")
simple("${body.address.lines[2]}")
```

The `last` keyword returns the index of the last element of a list. For example, you can access the *second last* element of a list, as follows:

```
simple("${body.address.lines[last-1]}")
```

You can use the `size` method to query the size of a list, as follows:

```
simple("${body.address.lines.size}")
```

OGNL array length access

You can access the length of a Java array through the `length` method, as follows:

```
String[] lines = new String[]{"foo", "bar", "cat"};
exchange.getIn().setBody(lines);

simple("There are ${body.length} lines")
```

27.5. PREDICATES

Overview

You can 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`.

Syntax


You can also test various parts of an exchange (headers, message body, and so on) using simple predicates. Simple predicates have the following general syntax:

```
{LHSVariable} Op RHSValue
```

Where the variable on the left hand side, *LHSVariable*, is one of the variables shown in [Table 27.1, “Variables for the Simple Language”](#) and the value on the right hand side, *RHSValue*, is one of the following:

- Another variable, `{RHSVariable}`.
- A string literal, enclosed in single quotes, `' '`.
- A numeric constant, enclosed in single quotes, `' '`.
- The null object, `null`.

The simple language always attempts to convert the RHS value to the type of the LHS value.



NOTE

While the simple language will attempt to convert the RHS, depending on the operator the LHS may need to be cast into the appropriate Type before the comparison is made.

Examples

For example, you can perform simple string comparisons and numerical comparisons as follows:

```
simple("${in.header.user} == 'john'")
simple("${in.header.number} > '100'") // String literal can be converted
to integer
```

You can test whether the left hand side is a member of a comma-separated list, as follows:

```
simple("${in.header.type} in 'gold,silver'")
```

You can test whether the left hand side matches a regular expression, as follows:

```
simple("${in.header.number} regex '\d{4}')
```

You can test the type of the left hand side using the **is** operator, as follows:

```
simple("${in.header.type} is 'java.lang.String'")
simple("${in.header.type} is 'String'") // You can abbreviate java.lang.
types
```

You can test whether the left hand side lies in a specified numerical range (where the range is inclusive), as follows:


```
simple("${in.header.number} range '100..199'")
```

Conjunctions

You can also combine predicates using the logical conjunctions, **&&** and **||**.

For example, here is an expression using the **&&** conjunction (logical and):

```
simple("${in.header.title} contains 'Camel' && ${in.header.type} == 'gold'")
```

And here is an expression using the **||** conjunction (logical inclusive or):

```
simple("${in.header.title} contains 'Camel' || ${in.header.type} == 'gold'")
```

27.6. VARIABLE REFERENCE

Table of variables

Table 27.1, “Variables for the Simple Language” shows all of the variables supported by the simple language.

Table 27.1. Variables for the Simple Language

Variable	Type	Description
camelContext	Object	The Camel context. <i>Supports OGNL expressions.</i>
camelId	String	The Camel context's ID value.
exchangeId	String	The exchange's ID value.
id	String	The <i>In</i> message ID value.
body	Object	The <i>In</i> message body. <i>Supports OGNL expressions.</i>
in.body	Object	The <i>In</i> message body. <i>Supports OGNL expressions.</i>
out.body	Object	The <i>Out</i> message body.

Variable	Type	Description
bodyAs (Type)	<i>Type</i>	The <i>In</i> message body, converted to the specified type. All types, <i>Type</i> , must be specified using their fully-qualified Java name, except for the types: byte[] , String , Integer , and Long . The converted body can be null.
mandatoryBodyAs (Type)	<i>Type</i>	The <i>In</i> message body, converted to the specified type. All types, <i>Type</i> , must be specified using their fully-qualified Java name, except for the types: byte[] , String , Integer , and Long . The converted body is expected to be non-null.
header .HeaderName	Object	The <i>In</i> message's <i>HeaderName</i> header. Supports OGNL expressions.
header [HeaderName]	Object	The <i>In</i> message's <i>HeaderName</i> header (alternative syntax).
headers .HeaderName	Object	The <i>In</i> message's <i>HeaderName</i> header.
headers [HeaderName]	Object	The <i>In</i> message's <i>HeaderName</i> header (alternative syntax).
in .header .HeaderName	Object	The <i>In</i> message's <i>HeaderName</i> header. Supports OGNL expressions.
in .header [HeaderName]	Object	The <i>In</i> message's <i>HeaderName</i> header (alternative syntax).
in .headers .HeaderName	Object	The <i>In</i> message's <i>HeaderName</i> header. Supports OGNL expressions.
in .headers [HeaderName]	Object	The <i>In</i> message's <i>HeaderName</i> header (alternative syntax).

Variable	Type	Description
<code>out.header.HeaderName</code>	Object	The <i>Out</i> message's <i>HeaderName</i> header.
<code>out.header[HeaderName]</code>	Object	The <i>Out</i> message's <i>HeaderName</i> header (alternative syntax).
<code>out.headers.HeaderName</code>	Object	The <i>Out</i> message's <i>HeaderName</i> header.
<code>out.headers[HeaderName]</code>	Object	The <i>Out</i> message's <i>HeaderName</i> header (alternative syntax).
<code>headerAs(Key, Type)</code>	<i>Type</i>	The <i>Key</i> header, converted to the specified type. All types, <i>Type</i> , must be specified using their fully-qualified Java name, except for the types: byte[] , String , Integer , and Long . The converted value can be null.
<code>headers</code>	Map	All of the <i>In</i> headers (as a java.util.Map type).
<code>in.headers</code>	Map	All of the <i>In</i> headers (as a java.util.Map type).
<code>property.PropertyName</code>	Object	The <i>PropertyName</i> property on the exchange.
<code>property[PropertyName]</code>	Object	The <i>PropertyName</i> property on the exchange (alternative syntax).
<code>sys.SysPropertyName</code>	String	The <i>SysPropertyName</i> Java system property.
<code>sysenv.SysEnvVar</code>	String	The <i>SysEnvVar</i> system environment variable.
<code>exception</code>	String	Either the exception object from Exchange.getException() or, if this value is null, the caught exception from the Exchange.EXCEPTION_CAUGHT property; otherwise null. <i>Supports OGNL expressions.</i>

Variable	Type	Description
exception.message	String	If an exception is set on the exchange, returns the value of Exception.getMessage() ; otherwise, returns null .
exception.stacktrace	String	If an exception is set on the exchange, returns the value of Exception.getStackTrace() ; otherwise, returns null . Note: The simple language first tries to retrieve an exception from Exchange.getException() . If that property is not set, it checks for a caught exception, by calling Exchange.getProperty(Exchange.CAUGHT_EXCEPTION) .
date:command:pattern	String	A date formatted using a java.text.SimpleDateFormat pattern. The following commands are supported: now , for the current date and time; header.HeaderName , or in.header.HeaderName to use a java.util.Date object in the <i>HeaderName</i> header from the <i>In</i> message; out.header.HeaderName to use a java.util.Date object in the <i>HeaderName</i> header from the <i>Out</i> message;
bean:beanID.Method	Object	Invokes a method on the referenced bean and <i>returns the result of the method invocation</i> . To specify a method name, you can either use the beanID.Method syntax; or you can use the beanID?method=methodName syntax.

Variable	Type	Description
<code>ref:beanID</code>	Object	Looks up the bean with the ID, <i>beanID</i> , in the registry and returns a reference to the bean itself. For example, if you are using the splitter EIP, you could use this variable to reference the bean that implements the splitting algorithm.
<code>properties:Key</code>	String	The value of the <i>Key</i> property placeholder (see Section 2.7, “Property Placeholders”).
<code>properties:Location:Key</code>	String	The value of the <i>Key</i> property placeholder, where the location of the properties file is given by <i>Location</i> (see Section 2.7, “Property Placeholders”).
<code>threadName</code>	String	The name of the current thread.
<code>routeId</code>	String	Returns the ID of the current route through which the Exchange is being routed.
<code>type:Name[.Field]</code>	Object	References a type or field by its Fully-Qualified-Name (FQN). To refer to a field, append .Field . For example, you can refer to the FILE_NAME constant field from the Exchange class as type:org.apache.camel.Exchange.FILE_NAME

27.7. OPERATOR REFERENCE

Binary operators

The binary operators for simple language predicates are shown in [Table 27.2, “Binary Operators for the Simple Language”](#).

Table 27.2. Binary Operators for the Simple Language

Operator	Description
<code>==</code>	Equals.
<code>></code>	Greater than.
<code>>=</code>	Greater than or equals.
<code><</code>	Less than.
<code><=</code>	Less than or equals.
<code>!=</code>	Not equal to.
<code>contains</code>	Test if LHS string contains RHS string.
<code>not contains</code>	Test if LHS string does <i>not</i> contain RHS string.
<code>regex</code>	Test if LHS string matches RHS regular expression.
<code>not regex</code>	Test if LHS string does <i>not</i> match RHS regular expression.
<code>in</code>	Test if LHS string appears in the RHS comma-separated list.
<code>not in</code>	Test if LHS string does <i>not</i> appear in the RHS comma-separated list.
<code>is</code>	Test if LHS is an instance of RHS Java type (using Java <code>instanceof</code> operator).
<code>not is</code>	Test if LHS is <i>not</i> an instance of RHS Java type (using Java <code>instanceof</code> operator).
<code>range</code>	Test if LHS number lies in the RHS range (where range has the format, ' <code>min...max</code> ').
<code>not range</code>	Test if LHS number does <i>not</i> lie in the RHS range (where range has the format, ' <code>min...max</code> ').

Unary operators and character escapes

The binary operators for simple language predicates are shown in [Table 27.3, “Unary Operators for the Simple Language”](#).

Table 27.3. Unary Operators for the Simple Language

Operator	Description
<code>++</code>	Increment a number by 1.
<code>--</code>	Decrement a number by 1.
<code>\n</code>	The newline character.
<code>\r</code>	The carriage return character.
<code>\t</code>	The tab character.
<code>\</code>	<i>(Obsolete)</i> Since Camel version 2.11, the backslash escape character is not supported.

Combining predicates

The conjunctions shown in [Table 27.4, “Conjunctions for Simple Language Predicates”](#) can be used to combine two or more simple language predicates.

Table 27.4. Conjunctions for Simple Language Predicates

Operator	Description
<code>&&</code>	Combine two predicates with logical <i>and</i> .
<code> </code>	Combine two predicates with logical <i>inclusive or</i> .
<code>and</code>	<i>Deprecated.</i> Use <code>&&</code> instead.
<code>or</code>	<i>Deprecated.</i> Use <code> </code> instead.

CHAPTER 28. SPEL

OVERVIEW

The [Spring Expression Language \(SpEL\)](#) is an object graph navigation language provided with Spring 3, which can be used to construct predicates and expressions in a route. A notable feature of SpEL is the ease with which you can access beans from the registry.

SYNTAX

The SpEL expressions must use the placeholder syntax, `#{SpELExpression}`, so that they can be embedded in a plain text string (in other words, SpEL has expression templating enabled).

SpEL can also look up beans in the registry (typically, the Spring registry), using the `@BeanID` syntax. For example, given a bean with the ID `headerUtils`, and the method, `count()` (which counts the number of headers on the current message), you could use the `headerUtils` bean in an SpEL predicate, as follows:

```
#{@headerUtils.count > 4}
```

ADDING SPEL PACKAGE

To use SpEL in your routes you need to add a dependency on `camel-spring` to your project as shown in [Example 28.1, “Adding the camel-spring dependency”](#).

Example 28.1. Adding the camel-spring dependency

```
<!-- Maven POM File -->
<properties>
  <camel-version>2.12.0.redhat-610379</camel-version>
  ...
</properties>

<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-spring</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

VARIABLES

[Table 28.1, “SpEL variables”](#) lists the variables that are accessible when using SpEL.

Table 28.1. SpEL variables

Variable	Type	Description
this	Exchange	The current exchange is the root object.
exchange	Exchange	The current exchange.
exchangeId	String	The current exchange's ID.
exception	Throwable	The exchange exception (if any).
fault	Message	The fault message (if any).
request	Message	The exchange's <i>In</i> message.
response	Message	The exchange's <i>Out</i> message (if any).
properties	Map	The exchange properties.
property(<i>Name</i>)	Object	The exchange property keyed by <i>Name</i> .
property(<i>Name</i>, <i>Type</i>)	Type	The exchange property keyed by <i>Name</i> , converted to the type, <i>Type</i> .

XML EXAMPLE

For example, to select only those messages whose **Country** header has the value **USA**, you can use the following SpEL expression:

```
<route>
  <from uri="SourceURL"/>
  <filter>
    <spel>#{request.headers['Country'] == 'USA'}}</spel>
    <to uri="TargetURL"/>
  </filter>
</route>
```

JAVA EXAMPLE

You can define the same route in the Java DSL, as follows:

```
from("SourceURL")
  .filter().spel("#{request.headers['Country'] == 'USA'}")
  .to("TargetURL");
```

The following example shows how to embed SpEL expressions within a plain text string:

```
from("SourceURL")
  .setBody(spel("Hello #{request.body}! What a beautiful #
{request.headers['dayOrNight']}"))
  .to("TargetURL");
```

CHAPTER 29. THE XPATH LANGUAGE

Abstract

When processing XML messages, the XPath language enables you to select part of a message, by specifying an XPath expression that acts on the message's Document Object Model (DOM). You can also define XPath predicates to test the contents of an element or an attribute.

29.1. JAVA DSL

Basic expressions

You can use `xpath("Expression")` to evaluate an XPath expression on the current exchange (where the XPath expression is applied to the body of the current *In* message). The result of the `xpath()` expression is an XML node (or node set, if more than one node matches).

For example, to extract the contents of the `/person/name` element from the current *In* message body and use it to set a header named `user`, you could define a route like the following:

```
from("queue:foo")
  .setHeader("user", xpath("/person/name/text()"))
  .to("direct:tie");
```

Instead of specifying `xpath()` as an argument to `setHeader()`, you can use the fluent builder `xpath()` command—for example:

```
from("queue:foo")
  .setHeader("user").xpath("/person/name/text()")
  .to("direct:tie");
```

If you want to convert the result to a specific type, specify the result type as the second argument of `xpath()`. For example, to specify explicitly that the result type is `String`:

```
xpath("/person/name/text()", String.class)
```

Namespaces

Typically, XML elements belong to a schema, which is identified by a namespace URI. When processing documents like this, it is necessary to associate namespace URIs with prefixes, so that you can identify element names unambiguously in your XPath expressions. Apache Camel provides the helper class, `org.apache.camel.builder.xml.Namespaces`, which enables you to define associations between namespaces and prefixes.

For example, to associate the prefix, `cust`, with the namespace, `http://acme.com/customer/record`, and then extract the contents of the element, `/cust:person/cust:name`, you could define a route like the following:

```
import org.apache.camel.builder.xml.Namespaces;
```

```

...
Namespaces ns = new Namespaces("cust", "http://acme.com/customer/record");

from("queue:foo")
    .setHeader("user", xpath("/cust:person/cust:name/text()", ns))
    .to("direct:tie");

```

Where you make the namespace definitions available to the `xpath()` expression builder by passing the `Namespaces` object, `ns`, as an additional argument. If you need to define multiple namespaces, use the `Namespace.add()` method, as follows:

```

import org.apache.camel.builder.xml.Namespaces;
...
Namespaces ns = new Namespaces("cust", "http://acme.com/customer/record");
ns.add("inv", "http://acme.com/invoice");
ns.add("xsi", "http://www.w3.org/2001/XMLSchema-instance");

```

If you need to specify the result type *and* define namespaces, you can use the three-argument form of `xpath()`, as follows:

```

xpath("/person/name/text()", String.class, ns)

```

Auditing namespaces

One of the most frequent problems that can occur when using XPath expressions is that there is a mismatch between the namespaces appearing in the incoming messages and the namespaces used in the XPath expression. To help you troubleshoot this kind of problem, the XPath language supports an option to dump all of the namespaces from all of the incoming messages into the system log.

To enable namespace logging at the **INFO** log level, enable the `LogNamespaces` option in the Java DSL, as follows:

```

xpath("/foo:person/@id", String.class).logNamespaces()

```

Alternatively, you could configure your logging system to enable **TRACE** level logging on the `org.apache.camel.builder.xml.XPathBuilder` logger.

When namespace logging is enabled, you will see log messages like the following for each processed message:

```

2012-01-16 13:23:45,878 [stSaxonWithFlag] INFO XPathBuilder -
Namespaces discovered in message: {xmlns:a=[http://apache.org/camel],
DEFAULT=[http://apache.org/default],
xmlns:b=[http://apache.org/camelA, http://apache.org/camelB]}

```

29.2. XML DSL

Basic expressions

To evaluate an XPath expression in the XML DSL, put the XPath expression inside an `xpath` element. The XPath expression is applied to the body of the current *In* message and returns an XML node (or node set). Typically, the returned XML node is automatically converted to

a string.

For example, to extract the contents of the `/person/name` element from the current/*n* message body and use it to set a header named `user`, you could define a route like the following:

```
<beans ...>

  <camelContext xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="queue:foo"/>
      <setHeader headerName="user">
        <xpath>/person/name/text()</xpath>
      </setHeader>
      <to uri="direct:tie"/>
    </route>
  </camelContext>

</beans>
```

If you want to convert the result to a specific type, specify the result type by setting the `resultType` attribute to a Java type name (where you must specify the fully-qualified type name). For example, to specify explicitly that the result type is `java.lang.String` (you can omit the `java.lang.` prefix here):

```
<xpath resultType="String">/person/name/text()</xpath>
```

Namespaces

When processing documents whose elements belong to one or more XML schemas, it is typically necessary to associate namespace URIs with prefixes, so that you can identify element names unambiguously in your XPath expressions. It is possible to use the standard XML mechanism for associating prefixes with namespace URIs. That is, you can set an attribute like this: `xmlns:Prefix="NamespaceURI"`.

For example, to associate the prefix, `cust`, with the namespace, `http://acme.com/customer/record`, and then extract the contents of the element, `/cust:person/cust:name`, you could define a route like the following:

```
<beans ...>

  <camelContext xmlns="http://camel.apache.org/schema/spring"
                xmlns:cust="http://acme.com/customer/record" >
    <route>
      <from uri="queue:foo"/>
      <setHeader headerName="user">
        <xpath>/cust:person/cust:name/text()</xpath>
      </setHeader>
      <to uri="direct:tie"/>
    </route>
  </camelContext>

</beans>
```

Auditing namespaces

One of the most frequent problems that can occur when using XPath expressions is that there is a mismatch between the namespaces appearing in the incoming messages and the namespaces used in the XPath expression. To help you troubleshoot this kind of problem, the XPath language supports an option to dump all of the namespaces from all of the incoming messages into the system log.

To enable namespace logging at the **INFO** log level, enable the **logNamespaces** option in the XML DSL, as follows:

```
<xpath logNamespaces="true" resultType="String">/foo:person/@id</xpath>
```

Alternatively, you could configure your logging system to enable **TRACE** level logging on the **org.apache.camel.builder.xml.XPathBuilder** logger.

When namespace logging is enabled, you will see log messages like the following for each processed message:

```
2012-01-16 13:23:45,878 [stSaxonWithFlag] INFO XPathBuilder -
Namespaces discovered in message: {xmlns:a=[http://apache.org/camel],
DEFAULT=[http://apache.org/default],
xmlns:b=[http://apache.org/camelA, http://apache.org/camelB]}
```

29.3. XPATH INJECTION

Parameter binding annotation

When using Apache Camel bean integration to invoke a method on a Java bean, you can use the **@XPath** annotation to extract a value from the exchange and bind it to a method parameter.

For example, consider the following route fragment, which invokes the **credit** method on an **AccountService** object:

```
from("queue:payments")
    .beanRef("accountService","credit")
    ...
```

The **credit** method uses parameter binding annotations to extract relevant data from the message body and inject it into its parameters, as follows:

```
public class AccountService {
    ...
    public void credit(
        @XPath("/transaction/transfer/receiver/text()") String name,
        @XPath("/transaction/transfer/amount/text()") String amount
    )
    {
        ...
    }
    ...
}
```

For more information about bean integration, see [Section 2.4, “Bean Integration”](#).

Namespaces

[Table 29.1, “Predefined Namespaces for @XPath”](#) shows the namespaces that are predefined for XPath. You can use these namespace prefixes in the **XPath** expression that appears in the **@XPath** annotation.

Table 29.1. Predefined Namespaces for @XPath

Namespace URI	Prefix
http://www.w3.org/2001/XMLSchema	xsd
http://www.w3.org/2003/05/soap-envelope	soap

Custom namespaces

You can use the **@NamespacePrefix** annotation to define custom XML namespaces. Invoke the **@NamespacePrefix** annotation to initialize the **namespaces** argument of the **@XPath** annotation. The namespaces defined by **@NamespacePrefix** can then be used in the **@XPath** annotation's expression value.

For example, to associate the prefix, **ex**, with the custom namespace, <http://fusesource.com/examples>, invoke the **@XPath** annotation as follows:

```
public class AccountService {
    ...
    public void credit(
        @XPath(
            value = "/ex:transaction/ex:transfer/ex:receiver/text()",
            namespaces = @NamespacePrefix(
                prefix = "ex",
                uri = "http://fusesource.com/examples"
            )
        ) String name,
        @XPath(
            value = "/ex:transaction/ex:transfer/ex:amount/text()",
            namespaces = @NamespacePrefix(
                prefix = "ex",
                uri = "http://fusesource.com/examples"
            )
        ) String amount,
    )
    {
        ...
    }
    ...
}
```

29.4. XPATH BUILDER

Overview

The `org.apache.camel.builder.xml.XPathBuilder` class enables you to evaluate XPath expressions independently of an exchange. That is, if you have an XML fragment from any source, you can use `XPathBuilder` to evaluate an XPath expression on the XML fragment.

Matching expressions

Use the `matches()` method to check whether one or more XML nodes can be found that match the given XPath expression. The basic syntax for matching an XPath expression using `XPathBuilder` is as follows:

```
boolean matches = XPathBuilder
    .xpath("Expression")
    .matches(CamelContext, "XMLString");
```

Where the given expression, *Expression*, is evaluated against the XML fragment, *XMLString*, and the result is true, if at least one node is found that matches the expression. For example, the following example returns `true`, because the XPath expression finds a match in the `xyz` attribute.

```
boolean matches = XPathBuilder
    .xpath("/foo/bar/@xyz")
    .matches(getContext(), "<foo><bar xyz='cheese' />
</foo>");
```

Evaluating expressions

Use the `evaluate()` method to return the contents of the first node that matches the given XPath expression. The basic syntax for evaluating an XPath expression using `XPathBuilder` is as follows:

```
String nodeValue = XPathBuilder
    .xpath("Expression")
    .evaluate(CamelContext, "XMLString");
```

You can also specify the result type by passing the required type as the second argument to `evaluate()`—for example:

```
String name = XPathBuilder
    .xpath("foo/bar")
    .evaluate(context, "<foo><bar>cheese</bar></foo>",
String.class);
Integer number = XPathBuilder
    .xpath("foo/bar")
    .evaluate(context, "<foo><bar>123</bar></foo>",
Integer.class);
Boolean bool = XPathBuilder
    .xpath("foo/bar")
    .evaluate(context, "<foo><bar>>true</bar></foo>",
Boolean.class);
```

29.5. ENABLING SAXON

Prerequisites

A prerequisite for using the Saxon parser is that you add a dependency on the **camel-saxon** artifact (either adding this dependency to your Maven POM, if you use Maven, or adding the **camel-saxon-6.1.0.redhat-379.jar** file to your classpath, otherwise).

Using the Saxon parser in Java DSL

In Java DSL, the simplest way to enable the Saxon parser is to call the **saxon()** fluent builder method. For example, you could invoke the Saxon parser as shown in the following example:

```
// Java
// create a builder to evaluate the xpath using saxon
XPathBuilder builder = XPathBuilder.xpath("tokenize(/foo/bar, '_' )
[2]").saxon();

// evaluate as a String result
String result = builder.evaluate(context, "<foo><bar>abc_def_ghi</bar>
</foo>");
```

Using the Saxon parser in XML DSL

In XML DSL, the simplest way to enable the Saxon parser is to set the **saxon** attribute to true in the **xpath** element. For example, you could invoke the Saxon parser as shown in the following example:

```
<xpath saxon="true" resultType="java.lang.String">current-dateTime()
</xpath>
```

Programming with Saxon

If you want to use the Saxon XML parser in your application code, you can create an instance of the Saxon transformer factory explicitly using the following code:

```
// Java
import javax.xml.transform.TransformerFactory;
import net.sf.saxon.TransformerFactoryImpl;
...
TransformerFactory saxonFactory = new
net.sf.saxon.TransformerFactoryImpl();
```

On the other hand, if you prefer to use the generic JAXP API to create a transformer factory instance, you *must* first set the **javax.xml.transform.TransformerFactory** property in the **ESBInstall/etc/system.properties** file, as follows:

```
javax.xml.transform.TransformerFactory=net.sf.saxon.TransformerFactoryImpl
```

You can then instantiate the Saxon factory using the generic JAXP API, as follows:

```
// Java
import javax.xml.transform.TransformerFactory;
...

```

```
TransformerFactory factory = TransformerFactory.newInstance();
```

If your application depends on any third-party libraries that use Saxon, it might be necessary to use the second, generic approach.



NOTE

The Saxon library must be installed in the container as the OSGi bundle, **net.sf.saxon/saxon9he** (normally installed by default). In versions of Fuse ESB prior to 7.1, it is not possible to load Saxon using the generic JAXP API.

29.6. EXPRESSIONS

Result type

By default, an XPath expression returns a list of one or more XML nodes, of **org.w3c.dom.NodeList** type. You can use the type converter mechanism to convert the result to a different type, however. In the Java DSL, you can specify the result type in the second argument of the **xpath()** command. For example, to return the result of an XPath expression as a **String**:

```
xpath("/person/name/text()", String.class)
```

In the XML DSL, you can specify the result type in the **resultType** attribute, as follows:

```
<xpath resultType="java.lang.String">/person/name/text()</xpath>
```

Patterns in location paths

You can use the following patterns in XPath location paths:

/people/person

The basic location path specifies the nested location of a particular element. That is, the preceding location path would match the person element in the following XML fragment:

```
<people>
  <person>...</person>
</people>
```

Note that this basic pattern can match *multiple* nodes—for example, if there is more than one **person** element inside the **people** element.

/name/text()

If you just want to access the *text* inside by the element, append **text()** to the location path, otherwise the node includes the element's start and end tags (and these tags would be included when you convert the node to a string).

/person/telephone/@isDayTime

To select the value of an attribute, *AttributeName*, use the syntax **@AttributeName**. For example, the preceding location path returns **true** when applied to the following XML fragment:

```
<person>
  <telephone isDayTime="true">1234567890</telephone>
</person>
```

*

A wildcard that matches all elements in the specified scope. For example, `/people/person/*` matches all the child elements of `person`.

@*

A wildcard that matches all attributes of the matched elements. For example, `/person/name/@*` matches all attributes of every matched `name` element.

//

Match the location path at every nesting level. For example, the `//name` pattern matches every `name` element highlighted in the following XML fragment:

```
<invoice>
  <person>
    <name .../>
  </person>
</invoice>
<person>
  <name .../>
</person>
<name .../>
```

..

Selects the parent of the current context node. Not normally useful in the Apache Camel XPath language, because the current context node is the document root, which has no parent.

node()

Match any kind of node.

text()

Match a text node.

comment()

Match a comment node.

processing-instruction()

Match a processing-instruction node.

Predicate filters

You can filter the set of nodes matching a location path by appending a predicate in square brackets, `[Predicate]`. For example, you can select the N^{th} node from the list of matches by appending `[N]` to a location path. The following expression selects the first matching `person` element:

```
/people/person[1]
```

The following expression selects the second-last **person** element:

```
/people/person[last()-1]
```

You can test the value of attributes in order to select elements with particular attribute values. The following expression selects the **name** elements, whose **surname** attribute is either Strachan or Davies:

```
/person/name[@surname="Strachan" or @surname="Davies"]
```

You can combine predicate expressions using any of the conjunctions **and**, **or**, **not()**, and you can compare expressions using the comparators, **=**, **!=**, **>**, **>=**, **<**, **<=** (in practice, the less-than symbol must be replaced by the **<** entity). You can also use XPath functions in the predicate filter.

Axes

When you consider the structure of an XML document, the root element contains a sequence of children, and some of those child elements contain further children, and so on. Looked at in this way, where nested elements are linked together by the *child-of* relationship, the whole XML document has the structure of a *tree*. Now, if you choose a particular node in this element tree (call it the *context node*), you might want to refer to different parts of the tree relative to the chosen node. For example, you might want to refer to the children of the context node, to the parent of the context node, or to all of the nodes that share the same parent as the context node (*sibling nodes*).

An *XPath axis* is used to specify the scope of a node match, restricting the search to a particular part of the node tree, relative to the current context node. The axis is attached as a prefix to the node name that you want to match, using the syntax, **AxisType::MatchingNode**. For example, you can use the **child::** axis to search the children of the current context node, as follows:

```
/invoice/items/child::item
```

The context node of **child::item** is the **items** element that is selected by the path, **/invoice/items**. The **child::** axis restricts the search to the children of the context node, **items**, so that **child::item** matches the children of **items** that are named **item**. As a matter of fact, the **child::** axis is the default axis, so the preceding example can be written equivalently as:

```
/invoice/items/item
```

But there several other axes (13 in all), some of which you have already seen in abbreviated form: **@** is an abbreviation of **attribute::**, and **//** is an abbreviation of **descendant-or-self::**. The full list of axes is as follows (for details consult the reference below):

- **ancestor**
- **ancestor-or-self**

- **attribute**
- **child**
- **descendant**
- **descendant-or-self**
- **following**
- **following-sibling**
- **namespace**
- **parent**
- **preceding**
- **preceding-sibling**
- **self**

Functions

XPath provides a small set of standard functions, which can be useful when evaluating predicates. For example, to select the last matching node from a node set, you can use the `last()` function, which returns the index of the last node in a node set, as follows:

```
/people/person[last()]
```

Where the preceding example selects the last **person** element in a sequence (in document order).

For full details of all the functions that XPath provides, consult the reference below.

Reference

For full details of the XPath grammar, see the [XML Path Language, Version 1.0](#) specification.

29.7. PREDICATES

Basic predicates

You can use **xpath** in the Java DSL or the XML DSL in a context where a predicate is expected—for example, as the argument to a **filter()** processor or as the argument to a **when()** clause.

For example, the following route filters incoming messages, allowing a message to pass, only if the **/person/city** element contains the value, **London**:

```
from("direct:tie")
    .filter().xpath("/person/city =
'London').to("file:target/messages/uk");
```

The following route evaluates the XPath predicate in a **when()** clause:

```
from("direct:tie")
  .choice()
    .when(xpath("/person/city =
'London')).to("file:target/messages/uk")
    .otherwise().to("file:target/messages/others");
```

XPath predicate operators

The XPath language supports the standard XPath predicate operators, as shown in [Table 29.2, “Operators for the XPath Language”](#).

Table 29.2. Operators for the XPath Language

Operator	Description
=	Equals.
!=	Not equal to.
>	Greater than.
>=	Greater than or equals.
<	Less than.
<=	Less than or equals.
or	Combine two predicates with logical <i>and</i> .
and	Combine two predicates with logical <i>inclusive or</i> .
not()	Negate predicate argument.

29.8. USING VARIABLES AND FUNCTIONS

Evaluating variables in a route

When evaluating XPath expressions inside a route, you can use XPath variables to access the contents of the current exchange, as well as O/S environment variables and Java system properties. The syntax to access a variable value is **\$VarName** or **\$Prefix:VarName**, if the variable is accessed through an XML namespace.

For example, you can access the *In* message's body as **\$in:body** and the *In* message's header value as **\$in:HeaderName**. O/S environment variables can be accessed as **\$env:EnvVar** and Java system properties can be accessed as **\$system:SysVar**.

In the following example, the first route extracts the value of the **/person/city** element

and inserts it into the **city** header. The second route filters exchanges using the XPath expression, `$in:city = 'London'`, where the `$in:city` variable is replaced by the value of the **city** header.

```
from("file:src/data?noop=true")
    .setHeader("city").xpath("/person/city/text()")
    .to("direct:tie");

from("direct:tie")
    .filter().xpath("$in:city = 'London'").to("file:target/messages/uk");
```

Evaluating functions in a route

In addition to the standard XPath functions, the XPath language defines additional functions. These additional functions (which are listed in [Table 29.4, “XPath Custom Functions”](#)) can be used to access the underlying exchange, to evaluate a simple expression or to look up a property in the Apache Camel property placeholder component.

For example, the following example uses the `in:header()` function and the `in:body()` function to access a head and the body from the underlying exchange:

```
from("direct:start").choice()
    .when().xpath("in:header('foo') = 'bar'").to("mock:x")
    .when().xpath("in:body() = '<two/>'").to("mock:y")
    .otherwise().to("mock:z");
```

Notice the similarity between these functions and the corresponding `in:HeaderName` or `in:body` variables. The functions have a slightly different syntax however: `in:header('HeaderName')` instead of `in:HeaderName`; and `in:body()` instead of `in:body`.

Evaluating variables in XPathBuilder

You can also use variables in expressions that are evaluated using the `XPathBuilder` class. In this case, you cannot use variables such as `$in:body` or `$in:HeaderName`, because there is no exchange object to evaluate against. But you *can* use variables that are defined inline using the `variable(Name, Value)` fluent builder method.

For example, the following `XPathBuilder` construction evaluates the `$test` variable, which is defined to have the value, **London**:

```
String var = XPathBuilder.xpath("$test")
    .variable("test", "London")
    .evaluate(getContext(), "<name>foo</name>");
```

Note that variables defined in this way are automatically entered into the global namespace (for example, the variable, `$test`, uses no prefix).

29.9. VARIABLE NAMESPACES

Table of namespaces

[Table 29.3, “XPath Variable Namespaces”](#) shows the namespace URIs that are associated with the various namespace prefixes.

Table 29.3. XPath Variable Namespaces

Namespace URI	Prefix	Description
http://camel.apache.org/schema/spring	<i>None</i>	Default namespace (associated with variables that have no namespace prefix).
http://camel.apache.org/xml/in/	in	Used to reference header or body of the current exchange's <i>In</i> message.
http://camel.apache.org/xml/out/	out	Used to reference header or body of the current exchange's <i>Out</i> message.
http://camel.apache.org/xml/functions/	functions	Used to reference some custom functions.
http://camel.apache.org/xml/variables/environment-variables	env	Used to reference O/S environment variables.
http://camel.apache.org/xml/variables/system-properties	system	Used to reference Java system properties.
http://camel.apache.org/xml/variables/exchange-property	<i>Undefined</i>	Used to reference exchange properties. You must define your own prefix for this namespace.

29.10. FUNCTION REFERENCE

Table of custom functions

[Table 29.4, “XPath Custom Functions”](#) shows the custom functions that you can use in Apache Camel XPath expressions. These functions can be used in addition to the standard XPath functions.

Table 29.4. XPath Custom Functions

Function	Description
in:body()	Returns the <i>In</i> message body.
in:header(<i>HeaderName</i>)	Returns the <i>In</i> message header with name, <i>HeaderName</i> .
out:body()	Returns the <i>Out</i> message body.

Function	Description
out:header (<i>HeaderName</i>)	Returns the <i>Out</i> message header with name, <i>HeaderName</i> .
function:properties (<i>PropKey</i>)	Looks up a property with the key, <i>PropKey</i> (see Section 2.7, "Property Placeholders").
function:simple (<i>SimpleExp</i>)	Evaluates the specified simple expression, <i>SimpleExp</i> .

CHAPTER 30. XQUERY

OVERVIEW

XQuery was originally devised as a query language for data stored in XML form in a database. The XQuery language enables you to select parts of the current message, when the message is in XML format. XQuery is a superset of the XPath language; hence, any valid XPath expression is also a valid XQuery expression.

JAVA SYNTAX

You can pass an XQuery expression to `xquery()` in several ways. For simple expressions, you can pass the XQuery expressions as a string (`java.lang.String`). For longer XQuery expressions, 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 the result. Depending on the context, the return value is interpreted either as a predicate (where an empty node set is interpreted as false) or as an expression.

ADDING THE SAXON MODULE

To use XQuery in your routes you need to add a dependency on `camel-saxon` to your project as shown in [Example 30.1, “Adding the camel-saxon dependency”](#).

Example 30.1. Adding the camel-saxon dependency

```
<!-- Maven POM File -->
...
<dependencies>
  ...
  <dependency>
    <groupId>org.apache.camel</groupId>
    <artifactId>camel-saxon</artifactId>
    <version>${camel-version}</version>
  </dependency>
  ...
</dependencies>
```

STATIC IMPORT

To use the `xquery()` static method in your application code, include the following import statement in your Java source files:

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

VARIABLES

[Table 30.1, “XQuery variables”](#) lists the variables that are accessible when using XQuery.

Table 30.1. XQuery variables

Variable	Type	Description
<code>exchange</code>	Exchange	The current Exchange
<code>in.body</code>	Object	The body of the IN message
<code>out.body</code>	Object	The body of the OUT message
<code>in.headers.key</code>	Object	The IN message header whose key is <code>key</code>
<code>out.headers.key</code>	Object	The OUT message header whose key is <code>key</code>
<code>key</code>	Object	The Exchange property whose key is <code>key</code>

EXAMPLE

[Example 30.2, “Route using XQuery”](#) shows a route that uses XQuery.

Example 30.2. Route using XQuery

```
<camelContext>
  <route>
    <from uri="activemq:MyQueue"/>
    <filter>
      <language language="xquery">/foo:person[@name='James ' ]</language>
      <to uri="mqseries:SomeOtherQueue"/>
    </filter>
  </route>
</camelContext>
```

PART III. WEB SERVICES AND ROUTING WITH CAMEL CXF

Abstract

This guide describes how to use Apache Camel's CXF component to create Web services or wrap existing functionality in Web service facades.

CHAPTER 31. DEMONSTRATION CODE FOR CAMEL/CXF

Abstract

This chapter explains how to install, build, and run the demonstrations that accompany this guide.

31.1. DOWNLOADING AND INSTALLING THE DEMONSTRATIONS

Overview

Most of the examples discussed in this guide are based on working demonstrations, which you can download and try out for yourself. The examples can easily be run by deploying them into a Red Hat JBoss Fuse container, as described here.

Prerequisites

For building and running the demonstration code, you must have the following prerequisites installed:

- *Java platform*—the demonstrations can run on Java 6 or Java 7.
- *Apache Maven build tool*—to build the demonstration, you require Apache [Maven 3.0.4](#).
- *Internet connection*—Maven requires an Internet connection in order to download required dependencies from remote repositories while performing a build.
- *Red Hat JBoss Fuse*—the demonstrations are deployed into the JBoss Fuse container.



NOTE

For more details of the requirements for installing and working with JBoss Fuse, see "[Installation Guide](#)".

Downloading the demonstration package

The source code for the demonstrations is packaged as a Zip file, **cx-f-webinars-jboss-fuse-6.1.zip**, and is available from the following location:

- <https://github.com/FuseByExample/cxf-webinars/archive/jboss-fuse-6.1.zip>

31.2. RUNNING THE DEMONSTRATIONS

Building the demonstrations

Use Apache Maven to build the demonstrations. Open a new command prompt, change directory to **cx-f-webinars-jboss-fuse-6.1**, and enter the following commands:

-

```
cd parent
mvn install
cd ..
mvn install
```

The first invocation of `mvn install` is to install the parent POM in your local Maven repository (so that it is available for the main build step).

The second invocation of `mvn install` is the main build step. This command builds all of the demonstrations under the `cxp-webinars-jboss-fuse-6.1` directory (where the demonstrations are defined to be submodules of the `cxp-webinars-jboss-fuse-6.1/pom.xml` project). While Maven is building the demonstration code, it downloads whatever dependencies it needs from the Internet and installs them in the local Maven repository.

Starting and configuring the Red Hat JBoss Fuse container

Start and configure the Red Hat JBoss Fuse container as follows:

1. (Optional) If your local Maven repository is in a non-standard location, you might need to edit the JBoss Fuse configuration to specify your custom location. Edit the `InstallDir/etc/org.ops4j.pax.url.mvn.cfg` file and set the `org.ops4j.pax.url.mvn.localRepository` property to the location of your local Maven repository:

```
#
# Path to the local maven repository which is used to avoid
# downloading
# artifacts when they already exist locally.
# The value of this property will be extracted from the settings.xml
# file
# above, or defaulted to:
#   System.getProperty( "user.home" ) + "/.m2/repository"
#
#org.ops4j.pax.url.mvn.localRepository=
org.ops4j.pax.url.mvn.localRepository=file:E:/Data/.m2/repository
```

2. Launch the JBoss Fuse container. Open a new command prompt, change directory to `InstallDir/bin`, and enter the following command:

```
./fuse
```

Running the customer-ws-osgi-bundle demonstration

It is now a relatively straightforward task to run each of the demonstrations by installing the relevant OSGi bundles.

For example, to start up the WSDL-first Web service (discussed in [Chapter 33, WSDL-First Service Implementation](#)), enter the following console commands:

```
JBossFuse:karaf@root> install -s mvn:com.fusesource.byexample.cxf-
webinars/customer-ws-osgi-bundle
```

To see the Web service in action, start up the sample Web service client (discussed in [Chapter 34, Implementing a WS Client](#)), by entering the following console command:

```
JBossFuse:karaf@root> install -s mvn:com.fusesource.byexample.cxf-
webinars/customer-ws-client
```

The bundle creates a thread that invokes the Web service once a second and logs the response. View the log by entering the following console command:

```
JBossFuse:karaf@root> log:tail -n 4
```

You should see log output like the following:

```
18:03:58,609 | INFO | qtp5581640-231 | CustomerServiceImpl
| ? ? |
218 - org.fusesource.sparks.fuse-webinars.cxf-webinars.customer-ws-osgi-
bundle - 1.1.4 | Getting status for custome
r 1234
18:03:58,687 | INFO | invoker thread. | ClientInvoker
| ? ? |
219 - org.fusesource.sparks.fuse-webinars.cxf-webinars.customer-ws-client
- 1.1.4 | Got back: status = Active, stat
usMessage = In the park, playing with my frisbee.
18:04:00,687 | INFO | qtp5581640-232 | CustomerServiceImpl
| ? ? |
218 - org.fusesource.sparks.fuse-webinars.cxf-webinars.customer-ws-osgi-
bundle - 1.1.4 | Getting status for custome
r 1234
18:04:00,703 | INFO | invoker thread. | ClientInvoker
| ? ? |
219 - org.fusesource.sparks.fuse-webinars.cxf-webinars.customer-ws-client
- 1.1.4 | Got back: status = Active, stat
usMessage = In the park, playing with my frisbee.
```

To stop viewing the log, type the interrupt character (usually Ctrl-C).

To stop the client, first discover the client's bundle ID using the **osgi:list** console command. For example:

```
JBossFuse:karaf@root> list | grep customer-ws-client
[ 219] [Active ] [ ] [Started] [ 60] customer-ws-client
(1.1.4)
```

You can then stop the client using the **osgi:stop** console command. For example:

```
JBossFuse:karaf@root> stop 219
```

To shut down the container completely, enter the following console command:

```
JBossFuse:karaf@root> shutdown -f
```

Running the other demonstrations

The remaining demonstrations are all based on the Camel CXF component. You can only run *one* of these demonstrations at a time, because they all use the same Web service port and would clash, if started at the same time:

- **customer-ws-camel-cxf-pojo**
- **customer-ws-camel-cxf-payload**
- **customer-ws-camel-cxf-provider**

The preceding demonstrations all require the Camel CXF component and some of them require the Camel Velocity component as well. Before you run the demonstrations, you must install the requisite features for these Camel components, as follows:

```
JBossFuse:karaf@root> features:install camel-cxf  
JBossFuse:karaf@root> features:install camel-velocity
```

You can test these demonstrations using the provided **customer-ws-client** client or using the third-party [SoapUI](#) utility.

CHAPTER 32. JAVA-FIRST SERVICE IMPLEMENTATION

32.1. JAVA-FIRST OVERVIEW

Overview

The Java-first approach is a convenient way to get started with Web services, if you are unfamiliar with WSDL syntax. Using this approach, you can define the Web service interface using an ordinary Java interface and then use the provided Apache CXF utilities to generate the corresponding WSDL contract from the Java interface.



NOTE

There is no demonstration code to accompany this example.

Service Endpoint Interface (SEI)

An SEI is an ordinary Java interface. In order to use the standard JAX-WS frontend, the SEI must be annotated with the `@WebService` annotation.^[1]

In the Java-first approach, the SEI is the starting point for implementing the Web service and it plays a central role in the development of the Web service implementation. The SEI is used in the following ways:

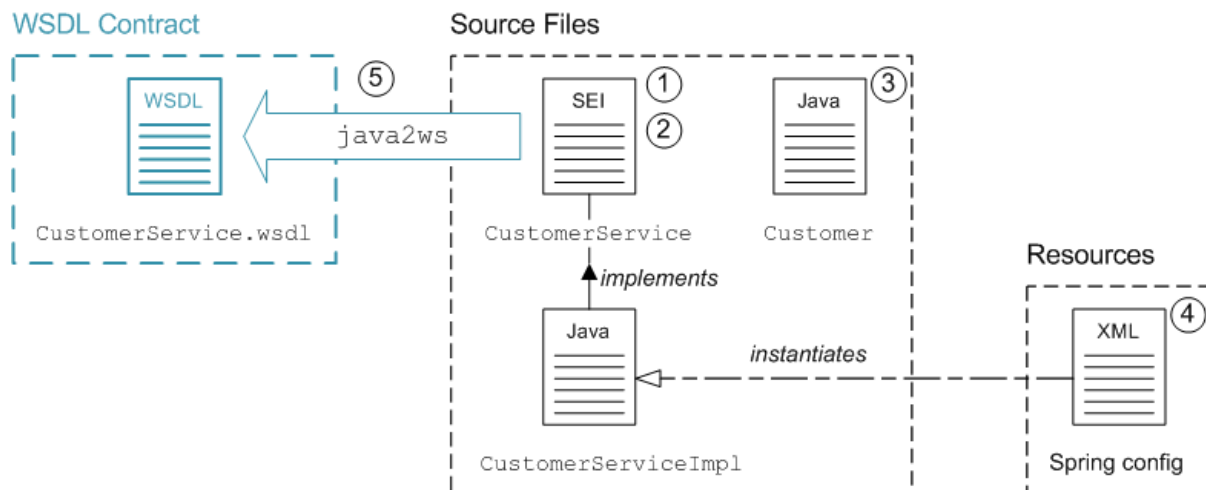
- *Base type of the Web service implementation (server side)*—you define the Web service by implementing the SEI.
- *Proxy type (client side)*—on the client side, you use the SEI to invoke operations on the client proxy object.
- *Basis for generating the WSDL contract*—in the Java-first approach, you generate the WSDL contract by converting the SEI to WSDL.

WSDL contract

The WSDL contract is a platform-neutral and language-neutral description of the Web service interface. When you want to make the Web service available to third-party clients, you should publish the WSDL contract to some well-known location. The WSDL contract contains all of the metadata required by WS clients.

The CustomerService demonstration

Figure 32.1, “Building a Java-First Web Service” shows an overview of the files required to implement and build the `CustomerService` Web service using the Java-first approach.

Figure 32.1. Building a Java-First Web Service

Implementing and building the service

To implement and build the Java-first example shown in [Figure 32.1, “Building a Java-First Web Service”](#), you would perform the following steps:

1. Implement the SEI, which constitutes the basic definition of the Web service's interface.
2. Annotate the SEI (you can use the annotations to influence the ultimate form of the generated WSDL contract).
3. Implement any other requisite Java classes. In particular, implement the following:
 - Any data types referenced by the SEI—for example, the **Customer** class.
 - The implementation of the SEI, **CustomerServiceImpl**.
4. Instantiate the Web service endpoint, by adding the appropriate code to a Spring XML file.
5. Generate the WSDL contract using a Java-to-WSDL converter.

32.2. DEFINE SEI AND RELATED CLASSES

Overview

The Service Endpoint Interface (SEI) is the starting point for implementing a Web service in the Java-first approach. The SEI represents the Web service in Java and it is ultimately used as the basis for generating the WSDL contract. This section describes how to create a sample SEI, the **CustomerService** interface, which enables you to access the details of a customer's account.

The CustomerService SEI

A JAX-WS service endpoint interface (SEI) is essentially an ordinary Java interface, augmented by certain annotations (which are discussed in the next section). For example, consider the following **CustomerService** interface, which defines methods for accessing the **Customer** data type:

```
// Java
package com.fusesource.demo.wsdl.customerservice;

// NOT YET ANNOTATED!
public interface CustomerService {

    public com.fusesource.demo.customer.Customer lookupCustomer(
        java.lang.String customerId
    );

    public void updateCustomer(
        com.fusesource.demo.customer.Customer cust
    );

    public void getCustomerStatus(
        java.lang.String customerId,
        javax.xml.ws.Holder<java.lang.String> status,
        javax.xml.ws.Holder<java.lang.String> statusMessage
    );
}
```

After adding the requisite annotations to the **CustomerService** interface, this interface provides the basis for defining the **CustomerService** Web service.

javax.xml.ws.Holder<?> types

The **getCustomerStatus** method from the **CustomerService** interface has parameters declared to be of **javax.xml.ws.Holder<String>** type. These so-called *holder types* are needed in order to declare the **OUT** or **INOUT** parameters of a WSDL operation.

The syntax of WSDL operations allows you to define any number of OUT or INOUT parameters, which means that the parameters are used to *return* a value to the caller. This kind of parameter passing is *not* natively supported by the Java language. Normally, the only way that Java allows you to return a value is by declaring it as the return value of a method. You can work around this language limitation, however, by declaring parameters to be *holder* types.

For example, consider the definition of the following method, **getStringValues()**, which takes a holder type as its second parameter:

```
// Java
public void getStringValues(
    String wrongWay,
    javax.xml.ws.Holder<String> rightWay
) {
    wrongWay = "Caller will never see this string!";
    rightWay.value = "But the caller *can* see this string.";
}
```

The caller can access the value of the returned **rightWay** string as **rightWay.value**. For example:

```
// Java
String wrongWay = "This string never changes";
javax.xml.ws.Holder<String> rightWay.value = "This value *can* change.";
```

```
sampleObject.getStringValues(wrongWay, rightWay);

System.out.println("Unchanged string: " + wrongWay);
System.out.println("Changed string: " + rightWay.value);
```

It is, perhaps, slightly unnatural to use **Holder**<> types in a Java-first example, because this is not a normal Java idiom. But it is interesting to include OUT parameters in the example, so that you can see how a Web service processes this kind of parameter.

Related classes

When you run the Java-to-WSDL compiler on the SEI, it converts not only the SEI, but also the classes referenced as parameters or return values. The parameter types must be convertible to XML, otherwise it would not be possible for WSDL operations to send or to receive those data types. In fact, when you run the Java-to-WSDL compiler, it is typically necessary to convert an entire tree of related classes to XML using the standard JAX-B encoding.

Normally, as long as the related classes do not require any exotic language features, the JAX-B encoding should be quite straightforward.

Default constructor for related classes

There is one simple rule, however, that you need to keep in mind when implementing related classes: each related class *must* have a default constructor (that is, a constructor without arguments). If you do not define any constructor for a class, the Java language automatically adds a default constructor. But if you define a class's constructors explicitly, you must ensure that one of them is a default constructor.

The Customer class

For example, the **Customer** class appears as a related class in the definition of the **CustomerService** SEI ([the section called "The CustomerService SEI"](#)). The **Customer** class consists of a collection of **String** fields and the only special condition it needs to satisfy is that it includes a default constructor:

```
// Java
package com.fusesource.demo.customer;

public class Customer {
    protected String firstName;
    protected String lastName;
    protected String phoneNumber;
    protected String id;

    // Default constructor, required by JAX-WS
    public Customer() { }

    public Customer(String firstName, String lastName, String phoneNumber,
        String id) {
        super();
        this.firstName = firstName;
        this.lastName = lastName;
        this.phoneNumber = phoneNumber;
    }
}
```

```

        this.id = id;
    }

    public String getFirstName() {
        return firstName;
    }

    public void setFirstName(String value) {
        this.firstName = value;
    }

    public String getLastName() {
        return lastName;
    }

    public void setLastName(String value) {
        this.lastName = value;
    }

    public String getPhoneNumber() {
        return phoneNumber;
    }

    public void setPhoneNumber(String value) {
        this.phoneNumber = value;
    }

    public String getId() {
        return id;
    }

    public void setId(String value) {
        this.id = value;
    }
}

```

32.3. ANNOTATE SEI FOR JAX-WS

Overview

To use the JAX-WS frontend, an SEI must be annotated using standardised JAX-WS annotations. The annotations signal to the Web services tooling that the SEI uses JAX-WS and the annotations are also used to customize the mapping from Java to WSDL. Here we only cover the most basic annotations—for complete details of JAX-WS annotations, see *Developing Applications Using JAX-WS* from the Apache CXF library.



NOTE

It sometimes makes sense also to annotate the service implementation class (the class that implements the SEI)—for example, if you want to associate an implementation class with a specific WSDL **serviceName** and **portName** (there can be more than one implementation of a given SEI).

Minimal annotation

The minimal annotation required for an SEI using the JAX-WS frontend is to prefix the interface declaration with `@WebService`. For example, the `CustomerService` SEI could be minimally annotated as follows:

```
// Java
package com.fusesource.demo.wsdl.customerservice;

import javax.jws.WebService;

@WebService
public interface CustomerService {
    ...
}
```

If you run the Java-to-WSDL utility on this interface, it will generate a complete WSDL contract using the standard default style of code conversion.

@WebService annotation

Although it is sufficient to specify the `@WebService` annotation without any attributes, it is usually better to specify some attributes to provide a more descriptive WSDL service name and WSDL port name. You will also usually want to specify the XML target namespace. For this, you can specify the following optional attributes of the `@WebService` annotation:

name

Specifies the name of the WSDL contract (appearing in the `wsdl:definitions` element).

serviceName

Specifies the name of the WSDL service (a SOAP service is defined by default in the generated contract).

portName

Specifies the name of the WSDL port (a SOAP/HTTP port is defined by default in the generated contract).

targetNamespace

The XML schema namespace that is used, by default, to qualify the elements and types defined in the contract.

@WebParam annotation

You can add the `@WebParam` annotation to method arguments in the SEI. The `@WebParam` annotation is optional, but there are a couple of good reasons for adding it:

- By default, JAX-WS maps Java arguments to parameters with names like `arg0`, ..., `argN`. Messages are much easier to read, however, when the parameters have meaningful names.
- It is a good idea to define parameter elements without a namespace. This makes the XML encoding of requests and responses more compact.
- To enable support for WSDL OUT and INOUT parameters.

You can add `@WebParam` annotations with the following attributes:

name

Specifies the mapped name of the parameter.

targetNamespace

Specifies the namespace of the mapped parameter. Set this to a blank string for a more compact XML encoding.

mode

Can have one of the following values:

- **WebParam.Mode.IN**—(*default*) parameter is passed from client to service (in request).
- **WebParam.Mode.INOUT**—parameter is passed from client to service (request) and from the service back to the client (in reply).
- **WebParam.Mode.OUT**—parameter is passed from service back to the client (in reply).

OUT and INOUT parameters

In WSDL, OUT and INOUT parameters represent values that can be sent from the service back to the client (where the INOUT parameter is sent in both directions).

In Java syntax, the only value that can ordinarily be returned from a method is the method's return value. In order to support OUT or INOUT parameters in Java (which are effectively like additional return values), you must:

- Declare the corresponding Java argument using a `javax.xml.ws.Holder<ParamType>` type, where *ParamType* is the type of the parameter you want to send.
- Annotate the Java argument with `@WebParam`, setting either `mode = WebParam.Mode.OUT` or `mode = WebParam.Mode.INOUT`.

Annotated CustomerService SEI

The following example shows the `CustomerService` SEI after it has been annotated. Many other annotations are possible, but this level of annotation is usually adequate for a WSDL-first project.

```
// Java
package com.fusesource.demo.wsdl.customerservice;

import javax.jws.WebParam;
import javax.jws.WebService;

@WebService(
    targetNamespace = "http://demo.fusesource.com/wsdl/CustomerService/",
    name = "CustomerService",
    serviceName = "CustomerService",
```

```

    portName = "SOAPOverHTTP"
)
public interface CustomerService {

    public com.fusesource.demo.customer.Customer lookupCustomer(
        @WebParam(name = "customerId", targetNamespace = "")
        java.lang.String customerId
    );

    public void updateCustomer(
        @WebParam(name = "cust", targetNamespace = "")
        com.fusesource.demo.customer.Customer cust
    );

    public void getCustomerStatus(
        @WebParam(name = "customerId", targetNamespace = "")
        java.lang.String customerId,
        @WebParam(mode = WebParam.Mode.OUT, name = "status", targetNamespace =
        "")
        javax.xml.ws.Holder<java.lang.String> status,
        @WebParam(mode = WebParam.Mode.OUT, name = "statusMessage",
        targetNamespace = "")
        javax.xml.ws.Holder<java.lang.String> statusMessage
    );
}

```

32.4. INSTANTIATE THE WS ENDPOINT

Overview

In Apache CXF, you create a WS endpoint by defining a **jaxws:endpoint** element in XML. The WS endpoint is effectively the runtime representation of the Web service: it opens an IP port to listen for SOAP/HTTP requests, is responsible for marshalling and unmarshalling messages (making use of the generated Java stub code), and routes incoming requests to the relevant methods on the implementor class.

In other words, creating a Web service in Spring XML consists essentially of the following two steps:

1. Create an instance of the implementor class, using the Spring **bean** element.
2. Create a WS endpoint, using the **jaxws:endpoint** element.

The **jaxws:endpoint** element

You can instantiate a WS endpoint using the **jaxws:endpoint** element in a Spring file, where the **jaxws:** prefix is associated with the **http://cxf.apache.org/jaxws** namespace.



NOTE

Take care not to confuse the **jaxws:endpoint** element with the **cxf:cxfEndpoint** element, which you meet later in this guide: the **jaxws:endpoint** element is used to integrate a WS endpoint with a Java implementation class; whereas the **cxf:cxfEndpoint** is used to integrate a WS endpoint with a Camel route.

Define JAX-WS endpoint in XML

The following sample Spring file shows how to define a JAX-WS endpoint in XML, using the **jaxws:endpoint** element.

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xmlns:jaxws="http://cxf.apache.org/jaxws"
  xmlns:soap="http://cxf.apache.org/bindings/soap"
  xsi:schemaLocation="
http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
http://cxf.apache.org/bindings/soap
http://cxf.apache.org/schemas/configuration/soap.xsd
http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd">

  <jaxws:endpoint
    xmlns:customer="http://demo.fusesource.com/wsdl/CustomerService/"
    id="customerService"
    address="/Customer"
    serviceName="customer:CustomerService"
    endpointName="customer:SOAPOverHTTP"
    implementor="#customerServiceImpl">
  </jaxws:endpoint>

  <bean id="customerServiceImpl"
    class="com.fusesource.customer.ws.CustomerServiceImpl"/>

</beans>
```

Address for the Jetty container

Apache CXF deploys the WS endpoint into a [Jetty](#) servlet container instance and the **address** attribute of **jaxws:endpoint** is therefore used to configure the addressing information for the endpoint in the Jetty container.

Apache CXF supports the notion of a *default servlet container* instance. The way the default servlet container is initialized and configured depends on the particular mode of deployment that you choose. For example the Red Hat JBoss Fuse container and Web containers (such as Tomcat) provide a default servlet container.

There are two different syntaxes you can use for the endpoint address, where the syntax that you use effectively determines whether or not the endpoint is deployed into the default servlet container, as follows:

- *Address syntax for default servlet container*—to use the default servlet container,

specify only the servlet context for this endpoint. Do *not* specify the protocol, host, and IP port in the address. For example, to deploy the endpoint to the `/Customers` servlet context in the default servlet container:

```
address="/Customers"
```

- *Address syntax for custom servlet container*—to instantiate a custom Jetty container for the endpoint, specify a complete HTTP URL, including the host and IP port (the value of the IP port effectively identifies the target Jetty container). Typically, for a Jetty container, you specify the host as `0.0.0.0`, which is interpreted as a wildcard that matches every IP network interface on the local machine (that is, if deployed on a multi-homed host, Jetty opens a listening port on every network card). For example, to deploy the endpoint to the custom Jetty container listening on IP port, **8083**:

```
address="http://0.0.0.0:8083/Customers"
```



NOTE

If you want to configure a secure endpoint (secured by SSL), you would specify the **https:** scheme in the address.

Referencing the service implementation

The `implementor` attribute of the `jaxws:endpoint` element references the implementation of the WS service. The value of this attribute can either be the name of the implementation class or (as in this example) a bean reference in the format, `#BeanID`, where the `#` character indicates that the following identifier is the name of a bean in the bean registry.

32.5. JAVA-TO-WSDL MAVEN PLUG-IN

Overview

To generate a WSDL contract from your SEI, you can use either the `java2ws` command-line utility or the `cxf-java2ws-plugin` Maven plug-in. The plug-in approach is ideal for Maven-based projects: after you paste the requisite plug-in configuration into your POM file, the WSDL code generation step is integrated into your build.

Configure the Java-to-WSDL Maven plug-in

Configuring the Java-to-WSDL Maven plug-in is relatively easy, because most of the default configuration settings can be left as they are. After copying and pasting the sample `plugin` element into your project's POM file, there are just a few basic settings that need to be customized, as follows:

- *CXF version*—make sure that the plug-in's dependencies are using the latest version of Apache CXF.
- *SEI class name*—specify the fully-qualified class name of the SEI in the `configuration/className` element.
- *Location of output*—specify the location of the generated WSDL file in the `configuration/outputFile` element.

For example, the following POM fragment shows how to configure the `cxf-java2ws-plugin` plug-in to generate WSDL from the `CustomerService` SEI:

```

<project ...>
  ...
  <properties>
    <cxf.version>2.7.0.redhat-610379</cxf.version>
  </properties>

  <build>
    <defaultGoal>install</defaultGoal>
    <plugins>
      ...
      <plugin>
        <groupId>org.apache.cxf</groupId>
        <artifactId>cxf-java2ws-plugin</artifactId>
        <version>${cxf.version}</version>
        <dependencies>
          <dependency>
            <groupId>org.apache.cxf</groupId>
            <artifactId>cxf-rt-frontend-jaxws</artifactId>
            <version>${cxf.version}</version>
          </dependency>
          <dependency>
            <groupId>org.apache.cxf</groupId>
            <artifactId>cxf-rt-frontend-simple</artifactId>
            <version>${cxf.version}</version>
          </dependency>
        </dependencies>
        <executions>
          <execution>
            <id>process-classes</id>
            <phase>process-classes</phase>
            <configuration>

<className>org.fusesource.demo.camelcxf.ws.server.CustomerService</classNa
me>

<outputFile>${basedir}/../src/main/resources/wsdL/CustomerService.wsdL</ou
tputFile>
          <genWsdL>true</genWsdL>
          <verbose>true</verbose>
        </configuration>
          <goals>
            <goal>java2ws</goal>
          </goals>
        </execution>
      </executions>
    </plugin>
  </plugins>
</build>

</project>

```

Generated WSDL

When using the Java-first approach to defining a Web service, there are typically other parts of your application (for example, WS clients) that depend on the generated WSDL file. For this reason, it is generally a good idea to output the generated WSDL file to a common location, which is accessible to other projects in your application, using the **outputFile** configuration element.

If you do not specify the **outputFile** configuration element, the generated WSDL is sent to the following location, by default:

```
BaseDir/target/generated/wsd\SEIClassName.wsdl
```

Reference

For full details of how to configure the Java-to-WSDL plug-in, see the Maven [Java2WS plug-in reference page](#).

[1] If the SEI is left without annotations, Apache CXF defaults to using the [simple frontend](#). This is a non-standard frontend, which is *not* recommended for most applications.

CHAPTER 33. WSDL-FIRST SERVICE IMPLEMENTATION

33.1. WSDL-FIRST OVERVIEW

Overview

If you are familiar with the syntax of WSDL and you want to have ultimate control over the layout and conventions applied to the WSDL contract, you will probably prefer to develop your Web service using the WSDL-first approach. In this approach, you start with the WSDL contract and then use the provided Apache CXF utilities to generate the requisite Java stub files from the WSDL contract.

Demonstration location

The code presented in this chapter is taken from the following demonstration:

```
cxf-webinars-jboss-fuse-6.1/customer-ws-osgi-bundle
```

For details of how to download and install the demonstration code, see [Chapter 31, *Demonstration Code for Camel/CXF*](#)

WSDL contract

The WSDL contract is a platform-neutral and language-neutral description of the Web service interface. In the WSDL-first approach, the WSDL contract is the starting point for implementing the Web service. You can use it to generate Java stub code, which provides the basis for implementing the Web service on the server side.

Service Endpoint Interface (SEI)


The most important piece of the generated stub code is the SEI, which is an ordinary Java interface that represents the Web service interface in the Java language.

The SEI is used in the following ways:

- *Base type of the Web service implementation (server side)*—you define the Web service by implementing the SEI.
- *Proxy type (client side)*—on the client side, you use the SEI to invoke operations on the client proxy object.

The CustomerService demonstration

[Figure 33.1, “Building a WSDL-First Web Service”](#) shows an overview of the files required to implement and build the **CustomerService** Web service using the WSDL-first approach.

Figure 33.1. Building a WSDL-First Web Service

Implementing and building the service

To implement and build the WSDL-first example shown in [Figure 33.1, “Building a WSDL-First Web Service”](#), starting from scratch, you would perform the following steps:

1. Create the WSDL contract.
2. Generate the Java stub code from the WSDL contract using a WSDL-to-Java converter, `ws2java`. This gives you the SEI, `CustomerService`, and its related classes, such as `Customer`.
3. Write the implementation of the SEI, `CustomerServiceImpl`.
4. Instantiate the Web service endpoint, by adding the appropriate code to a Spring XML file.

33.2. CUSTOMERSERVICE WSDL CONTRACT

Sample WSDL contract

The WSDL contract used in this demonstration is the `CustomerService` WSDL contract, which is available in the following location:

```
cxfr-wbinars-jboss-fuse-6.1/src/main/resources
```

Because the WSDL contract is a fairly verbose format, it is not shown in here in full. The main point you need to be aware of is that the `CustomerService` WSDL contract exposes the following operations:

lookupCustomer

Given a customer ID, the operation returns the corresponding `Customer` data object.

updateCustomer

Stores the given `Customer` data object against the given customer ID.

getCustomerStatus

Returns the status of the customer with the given customer ID.

Parts of the WSDL contract

A WSDL contract has the following main parts:

- the section called “Port type”.
- the section called “WSDL binding”.
- the section called “WSDL port”.

Port type

The port type is defined in the WSDL contract by the `wsdl:portType` element. It is analogous to an interface and it defines the operations that can be invoked on the Web service.

For example, the following WSDL fragment shows the `wsdl:portType` definition from the `CustomerService` WSDL contract:

```

<wsdl:definitions name="CustomerService"
  targetNamespace="http://demo.fusesource.com/wsdl/CustomerService/"
  ...>
  ...
  <wsdl:portType name="CustomerService">
    <wsdl:operation name="lookupCustomer">
      <wsdl:input message="tns:lookupCustomer"></wsdl:input>
      <wsdl:output message="tns:lookupCustomerResponse">
</wsdl:output>
    </wsdl:operation>
    <wsdl:operation name="updateCustomer">
      <wsdl:input message="tns:updateCustomer"></wsdl:input>
      <wsdl:output message="tns:updateCustomerResponse">
</wsdl:output>
    </wsdl:operation>
    <wsdl:operation name="getCustomerStatus">
      <wsdl:input message="tns:getCustomerStatus"></wsdl:input>
      <wsdl:output message="tns:getCustomerStatusResponse">
</wsdl:output>
    </wsdl:operation>
  </wsdl:portType>
  ...
</wsdl:definitions>

```

WSDL binding

A WSDL binding describes how to *encode* all of the operations and data types associated with a particular port type. A binding is specific to a particular protocol—for example, SOAP or JMS.

WSDL port

A WSDL port *specifies the transport protocol* and contains addressing data that enables clients to locate and connect to a remote server endpoint.

For example, the **CustomerService** WSDL contract defines the following WSDL port:

```
<wsdl:definitions ...>
  ...
  <wsdl:service name="CustomerService">
    <wsdl:port name="SOAPoverHTTP" binding="tns:CustomerServiceSOAP">
      <soap:address location="http://0.0.0.0:8183/CustomerService"
    />
    </wsdl:port>
  </wsdl:service>
</wsdl:definitions>
```

The address specified by the **soap:address** element's **location** attribute in the original WSDL contract is typically overridden at run time, however.

The **getCustomerStatus** operation

Because a WSDL contract is fairly verbose, it can be a bit difficult to see what the parameters of an operation are. Typically, for each operation, you can find data types in the XML schema section that represent the operation request and the operation response. For example, the **getCustomerStatus** operation has its request parameters (IN parameters) encoded by the **getCustomerStatus** element and its response parameters (OUT parameters) encoded by the **getCustomerStatusResponse** element, as follows:

```
<wsdl:definitions name="CustomerService"
  targetNamespace="http://demo.fusesource.com/wsdl/CustomerService/"
  ...>
  <wsdl:types>
    <xsd:schema ...>
      ...
      <xsd:element name="getCustomerStatus">
        <xsd:complexType>
          <xsd:sequence>
            <xsd:element name="customerId"
type="xsd:string"/>
          </xsd:sequence>
        </xsd:complexType>
      </xsd:element>
      <xsd:element name="getCustomerStatusResponse">
        <xsd:complexType>
          <xsd:sequence>
            <xsd:element name="status" type="xsd:string"/>
            <xsd:element name="statusMessage"
type="xsd:string"/>
          </xsd:sequence>
        </xsd:complexType>
      </xsd:element>
    </xsd:schema>
```



```

    </wsdl:types>
    ...
</wsdl:definitions>

```

References

For more details about the format of WSDL contracts and how to create your own WSDL contracts, see *Writing WSDL Contracts* and the [Eclipse JAX-WS Tools Component](#).

33.3. WSDL-TO-JAVA MAVEN PLUG-IN

Overview

In contrast to the Java-first approach, which starts with a Java interface and then generates the WSDL contract, the WSDL-first approach needs to generate Java stub code from the WSDL contract.

To generate Java stub code from the WSDL contract, you can use either the **ws2java** command-line utility or the **cxf-codegen-plugin** Maven plug-in. The plug-in approach is ideal for Maven-based projects: after you paste the requisite plug-in configuration into your POM file, the WSDL-to-Java code generation step is integrated into your build.

Configure the WSDL-to-Java Maven plug-in

Configuring the WSDL-to-Java Maven plug-in is relatively easy, because most of the default configuration settings can be left as they are. After copying and pasting the sample **plugin** element into your project's POM file, there are just a few basic settings that need to be customized, as follows:

- *CXF version*—make sure that the plug-in's dependencies are using the latest version of Apache CXF.
- *WSDL file location*—specify the WSDL file location in the **configuration/wsdlOptions/wsdlOption/wsdl** element.
- *Location of output*—specify the root directory of the generated Java source files in the **configuration/sourceRoot** element.

For example, the following POM fragment shows how to configure the **cxf-codegen-plugin** plug-in to generate Java stub code from the **CustomerService.wsdl** WSDL file:

```

<project ...>
  ...
  <parent>
    <groupId>com.fusesource.byexample.cxf-webinars</groupId>
    <artifactId>cxf-webinars</artifactId>
    <version>1.0-SNAPSHOT</version>
  </parent>

  <build>
    <defaultGoal>install</defaultGoal>
    <plugins>
      ...
      <plugin>

```

```

<groupId>org.apache.cxf</groupId>
<artifactId>cxf-codegen-plugin</artifactId>
<version>${cxf-version}</version>
<executions>
  <execution>
    <id>generate-sources</id>
    <phase>generate-sources</phase>
    <configuration>
      <!-- Maven auto-compiles any source files under
target/generated-sources/ -->
      <sourceRoot>${basedir}/target/generated-
sources/jaxws</sourceRoot>
      <wsdlOptions>
        <wsdlOption>

<wsdl>${basedir}/../src/main/resources/wsdl/CustomerService.wsdl</wsdl>
        </wsdlOption>
      </wsdlOptions>
    </configuration>
    <goals>
      <goal>wsdl2java</goal>
    </goals>
  </execution>
</executions>
</plugin>

</plugins>
</build>

</project>

```

Generated Java source code

With the sample configuration shown here, the generated Java source code is written under the **target/generated-sources/jaxws** directory. Note that the Web service implementation is dependent on this generated stub code—for example, the service implementation class must implement the generated **CustomerService** SEI.

Adding the generated source to an IDE

If you are using an IDE such as Eclipse or IntelliJ's IDEA, you need to make sure that the IDE is aware of the generated Java code. For example, in Eclipse it is necessary to add the **target/generated-sources/jaxws** directory to the project as a source code directory.

Compiling the generated code

You must ensure that the generated Java code is compiled and added to the deployment package. By convention, Maven automatically compiles any source files that it finds under the following directory:

```
BaseDir/target/generated-sources/
```

Hence, if you configure the output directory as shown in the preceding POM fragment, the generated code is automatically compiled by Maven.

Reference

For full details of how to configure the Java-to-WSDL plug-in, see the [Maven cxf-codegen-plugin](#) reference page.

33.4. INSTANTIATE THE WS ENDPOINT

Overview

In Apache CXF, you create a WS endpoint by defining a `jaxws:endpoint` element in XML. The WS endpoint is effectively the runtime representation of the Web service: it opens an IP port to listen for SOAP/HTTP requests, is responsible for marshalling and unmarshalling messages (making use of the generated Java stub code), and routes incoming requests to the relevant methods on the implementor class.

In other words, creating a Web service in Spring XML consists essentially of the following two steps:

1. Create an instance of the implementor class, using the Spring `bean` element.
2. Create a WS endpoint, using the `jaxws:endpoint` element.

Define JAX-WS endpoint in XML

The following sample Spring file shows how to define a JAX-WS endpoint in XML, using the `jaxws:endpoint` element.

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

  <import resource="classpath:META-INF/cxf/cxf.xml" />

  <jaxws:endpoint
    xmlns:customer="http://demo.fusesource.com/wsdl/CustomerService/"
    id="customerService"
    address="/Customer"
    serviceName="customer:CustomerService"
    endpointName="customer:SOAPOverHTTP"
    implementor="#customerServiceImpl">
  </jaxws:endpoint>

  <bean id="customerServiceImpl"
    class="com.fusesource.customer.ws.CustomerServiceImpl"/>

</beans>
```

Address for the Jetty container

In the preceding example, the **address** attribute of the **jaxws:endpoint** element specifies the servlet context for this endpoint, relative to the Jetty container in which it is deployed.

For more details about the options for specifying the endpoint address, see [the section called "Address for the Jetty container"](#).

Referencing the service implementation

The **implementor** attribute of the **jaxws:endpoint** element is used to reference the implementation of the WS service. The value of this attribute can either be the name of the implementation class or (as in this example) a bean reference in the format, **#BeanID**, where the **#** character indicates that the following identifier is the name of a bean in the bean registry.

33.5. DEPLOY TO AN OSGI CONTAINER

Overview

One of the options for deploying the Web service is to package it as an OSGi bundle and deploy it into an OSGi container such as Red Hat JBoss Fuse. Some of the advantages of an OSGi deployment include:

- Bundles are a relatively lightweight deployment option (because dependencies can be shared between deployed bundles).
- OSGi provides sophisticated dependency management, ensuring that only version-consistent dependencies are added to the bundle's classpath.

Using the Maven bundle plug-in

The Maven bundle plug-in is used to package your project as an OSGi bundle, in preparation for deployment into the OSGi container. There are two essential modifications to make to your project's **pom.xml** file:

1. Change the packaging type to **bundle** (by editing the value of the **project/packaging** element in the POM).
2. Add the Maven bundle plug-in to your POM file and configure it as appropriate.

Configuring the Maven bundle plug-in is quite a technical task (although the default settings are often adequate). For full details of how to customize the plug-in configuration, consult *Deploying into the OSGi Container* and *Managing OSGi Dependencies*.

Sample bundle plug-in configuration

The following POM fragment shows a sample configuration of the Maven bundle plug-in, which is appropriate for the current example.

```
<?xml version="1.0"?>
<project ...>
  ...
  <groupId>com.fusesource.byexample.cxf-webinars</groupId>
```

```

<artifactId>customer-ws-osgi-bundle</artifactId>
<name>customer-ws-osgi-bundle</name>
<url>http://www.fusesource.com</url>
<packaging>bundle</packaging>
...
<build>
  <plugins>
    ...
    <plugin>
      <groupId>org.apache.felix</groupId>
      <artifactId>maven-bundle-plugin</artifactId>
      <version>${version.maven-bundle-plugin}</version>
      <extensions>>true</extensions>
      <configuration>
        <instructions>
          <Export-Package>
            !com.fusesource.customer.ws,
            !com.fusesource.demo.customer,
            !com.fusesource.demo.wsdل.customerservice
          </Export-Package>
          <Import-Package>
            META-INF.cxf,
            META-INF.cxf.osgi,
            *
          </Import-Package>
          <DynamicImport-Package>
            org.apache.cxf.*,
            org.springframework.beans.*
          </DynamicImport-Package>
        </instructions>
      </configuration>
    </plugin>
    ...
  </plugins>
</build>
</project>

```

Dynamic imports

The Java packages from Apache CXF and the Spring API are imported using dynamic imports (specified using the **DynamicImport-Package** element). This is a pragmatic way of dealing with the fact that Spring XML files are not terribly well integrated with the Maven bundle plug-in. At build time, the Maven bundle plug-in is *not* able to figure out which Java classes are required by the Spring XML code. By listing wildcarded package names in the **DynamicImport-Package** element, however, you allow the OSGi container to figure out which Java classes are needed by the Spring XML code at *run* time.

**NOTE**

In general, using **DynamicImport-Package** headers is not recommended in OSGi, because it short-circuits OSGi version checking. Normally, what should happen is that the Maven bundle plug-in lists the Java packages used at *build* time, along with their versions, in the **Import-Package** header. At *deploy* time, the OSGi container then checks that the available Java packages are compatible with the build-time versions listed in the **Import-Package** header. With dynamic imports, this version checking cannot be performed.

Build and deploy the service bundle

After you have configured the POM file, you can build the Maven project and install it in your local repository by entering the following command:

```
mvn install
```

To deploy the service bundle, enter the following command at the command console:

```
karaf@root> install -s mvn:com.fusesource.byexample.cxf-webinars/customer-  
ws-osgi-bundle
```

**NOTE**

If your local Maven repository is stored in a non-standard location, you might need to customize the value of the **org.ops4j.pax.url.mvn.localRepository** property in the **EsbInstallDir/etc/org.ops4j.pax.url.mvn.cfg** file, before you can use the **mvn:** scheme to access Maven artifacts.

Red Hat JBoss Fuse default servlet container

Red Hat JBoss Fuse has a default Jetty container which, by default, listens for HTTP requests on port 8181. Moreover, WS endpoints in this container are implicitly deployed under the servlet context **cx/**. Hence, any WS endpoint whose **address** attribute is configured in the **jaxws:endpoint** element as **/EndpointContext** will have the following effective address:

```
http://Hostname:8181/cx/EndpointContext
```

You can optionally customize the default servlet container by editing settings in the following file:

```
InstallDir/etc/org.ops4j.pax.web.cfg
```

Full details of the properties you can set in this file are given in the [Ops4j Pax Web configuration reference](#).

Check that the service is running

A simple way of checking that the service is running is to point your browser at the following URL:

```
http://localhost:8181/cx/Customer?wsdl
```

- This query should return a copy of the WS endpoint's WSDL contract.

CHAPTER 34. IMPLEMENTING A WS CLIENT

34.1. WS CLIENT OVERVIEW

Overview

The key object in a WS client is the WS client proxy object, which enables you to access the remote Web service by invoking methods on the SEI. The proxy object itself can easily be instantiated using the `jaxws:client` element in Spring XML.

Demonstration location

The code presented in this chapter is taken from the following demonstration:

```
cxf-webinars-jboss-fuse-6.1/customer-ws-client
```

For details of how to download and install the demonstration code, see [Chapter 31, Demonstration Code for Camel/CXF](#)

WSDL contract

The WSDL contract is a platform-neutral and language-neutral description of the Web service interface. It contains all of the metadata that a client needs to find a Web service and invoke its operations. You can generate Java stub code from the WSDL contract, which provides an API that makes it easy to invoke the remote WSDL operations.

Service Endpoint Interface (SEI)

The most important piece of the generated stub code is the SEI, which is an ordinary Java interface that represents the Web service interface in the Java language.

WS client proxy

The WS client proxy is an object that converts Java method invocations to remote procedure calls, sending and receiving messages to a remote instance of the Web service across the network. The methods of the proxy are exposed through the SEI.



NOTE

The proxy type is generated *dynamically* by Apache CXF at run time. That is, there is no class in the stub code that corresponds to the implementation of the proxy (the only relevant entity is the SEI, which defines the proxy's interface).

The CustomerService client

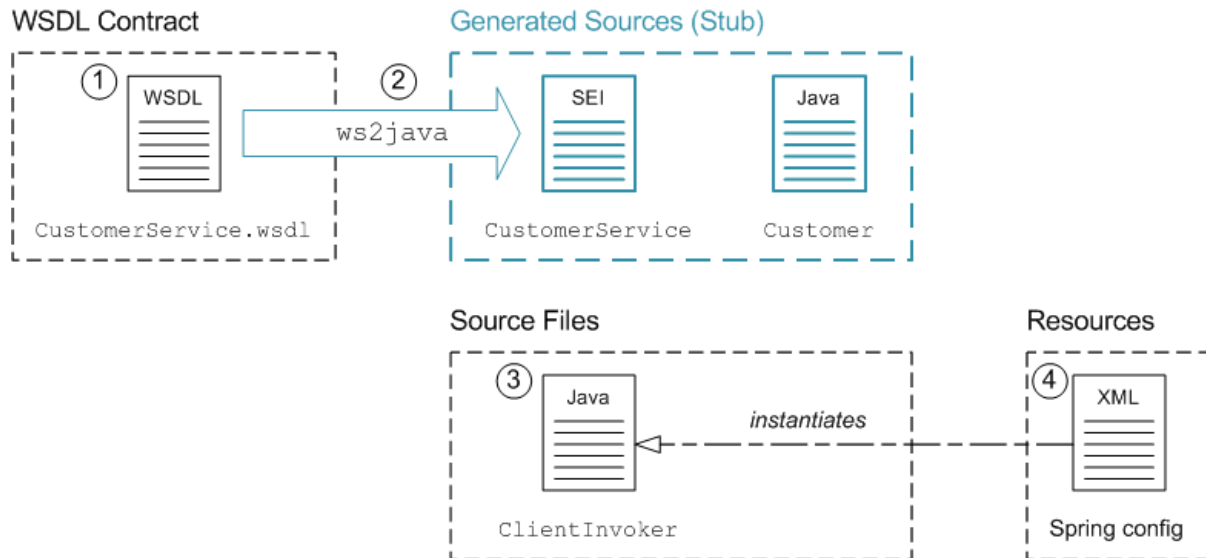
To take a specific example, consider the `customer-ws-client` demonstration, which is available from the following location:

```
cxf-webinars-jboss-fuse-6.1/customer-ws-client
```

[Figure 34.1, “Building a WS Client”](#) shows an overview of the files required to implement

and build the WS client.

Figure 34.1. Building a WS Client



Implementing and building the WS client

To implement and build the sample WS client shown in [Figure 34.1, “Building a WS Client”](#), starting from scratch, you would perform the following steps:

1. Obtain a copy of the WSDL contract.
2. Generate the Java stub code from the WSDL contract using a WSDL-to-Java converter, **ws2java**. This gives you the SEI, **CustomerService**, and its related classes, such as **Customer**.
3. Implement the main client class, **ClientInvoker**, which invokes the Web service operations. In this class define a bean property of type, **CustomerService**, so that the client class can receive a reference to the WS client proxy by property injection.
4. In a Spring XML file, instantiate the WS client proxy and inject it into the main client class, **ClientInvoker**.

34.2. WSDL-TO-JAVA MAVEN PLUG-IN

Overview

To generate Java stub code from the WSDL contract, you can use either the **ws2java** command-line utility or the **cxfr-codegen-plugin** Maven plug-in. When using Maven, the plug-in approach is ideal: after you paste the requisite plug-in configuration into your POM file, the WSDL-to-Java code generation step is integrated into your build.

Configure the WSDL-to-Java Maven plug-in

Configuring the WSDL-to-Java Maven plug-in is relatively easy, because most of the default configuration settings can be left as they are. After copying and pasting the sample **plugin** element into your project's POM file, there are just a few basic settings that need to be customized, as follows:

- *CXF version*—make sure that the plug-in's dependencies are using the latest version

of Apache CXF.

- *WSDL file location*—specify the WSDL file location in the **configuration/wsdloptions/wsdloption/wsdL** element.
- *Location of output*—specify the root directory of the generated Java source files in the **configuration/sourceRoot** element.

For example, the following POM fragment shows how to configure the **cxf-codegen-plugin** plug-in to generate Java stub code from the **CustomerService.wsdl** WSDL file:

```
<project ...>
  ...
  <parent>
    <groupId>com.fusesource.byexample.cxf-webinars</groupId>
    <artifactId>cxf-webinars</artifactId>
    <version>1.0-SNAPSHOT</version>
  </parent>

  <build>
    <defaultGoal>install</defaultGoal>
    <plugins>
      ...
      <plugin>
        <groupId>org.apache.cxf</groupId>
        <artifactId>cxf-codegen-plugin</artifactId>
        <version>${cxf-version}</version>
        <executions>
          <execution>
            <id>generate-sources</id>
            <phase>generate-sources</phase>
            <configuration>
              <sourceRoot>${basedir}/target/generated-
sources/jaxws</sourceRoot>
              <wsdlOptions>
                <wsdlOption>
<wsdl>${basedir}/../src/main/resources/wsdl/CustomerService.wsdl</wsdl>
                </wsdlOption>
              </wsdlOptions>
            </configuration>
            <goals>
              <goal>wsdl2java</goal>
            </goals>
          </execution>
        </executions>
      </plugin>

    </plugins>
  </build>
</project>
```

Generated Java source code

With the sample configuration shown here, the generated Java source code is written under the **target/generated-sources/jaxws** directory. Note that the client implementation is dependent on this generated stub code—for example, the client invokes the proxy using the generated **CustomerService** SEI.

Add generated source to IDE

If you are using an IDE such as Eclipse or IntelliJ's IDEA, you need to make sure that the IDE is aware of the generated Java code. For example, in Eclipse it is necessary to add the **target/generated-sources/jaxws** directory to the project as a source code directory.

Compiling the generated code

You must ensure that the generated Java code is compiled and added to the deployment package. By convention, Maven automatically compiles any source files that it finds under the following directory:

```
BaseDir/target/generated-sources/
```

Hence, if you configure the output directory as shown in the preceding POM fragment, the generated code is automatically compiled by Maven.

Reference

For full details of how to configure the Java-to-WSDL plug-in, see the [Maven cxf-codegen-plugin](#) reference page.

34.3. INSTANTIATE THE WS CLIENT PROXY

Overview

The WS client proxy is the most important kind of object in a WS client, because it provides a simple way of invoking operations on a remote Web service. The proxy enables you to access a Web service by invoking methods locally on a Java interface. The methods invoked on the proxy object are then translated into remote procedure calls on the Web service.

You can instantiate a WS client proxy straightforwardly using the **jaxws:client** element.

Define the WS client in XML

The following Spring XML fragment shows how to instantiate a client proxy bean using the **jaxws:client** element.

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xmlns:jaxws="http://cxf.apache.org/jaxws"
       xmlns:soap="http://cxf.apache.org/bindings/soap"
       xsi:schemaLocation="
http://www.springframework.org/schema/beans
http://www.springframework.org/schema/beans/spring-beans.xsd
http://cxf.apache.org/bindings/soap
http://cxf.apache.org/schemas/configuration/soap.xsd
```

```

http://cxf.apache.org/jaxws http://cxf.apache.org/schemas/jaxws.xsd">

  <import resource="classpath:META-INF/cxf/cxf.xml" />

  <jaxws:client
    id="customerServiceProxy"
    address="http://localhost:8181/cxf/Customer"

serviceClass="com.fusesource.demo.wsdl.customerservice.CustomerService"
  />

  <bean id="customerServiceClient"
    class="com.fusesource.customer.client.ClientInvoker"
    init-method="init" destroy-method="destroy">
    <property name="customerService" ref="customerServiceProxy"/>
  </bean>

</beans>

```

The `jaxws:client` element

The `jaxws:client` element creates a client proxy *dynamically* (that is, there is no dedicated class that represents a proxy implementation in the Java stub code). The following attributes are used to define the proxy:

`id`

The ID that you specify here is entered in the bean registry and can be used to reference the proxy instance from other beans.

`address`

The full address of the remote Web service that this proxy connects to.

`serviceClass`

The fully-qualified class name of the Web service's SEI (you invoke methods on the proxy through the SEI).

Injecting with the proxy reference

To access the proxy instance, simply inject the proxy into one or more other beans defined in XML. Given that the proxy ID has the value, `customerServiceProxy`, you can inject it into a bean property using the Spring `property` element, as follows:

```

<bean ...>
  <property name="customerService" ref="customerServiceProxy"/>
</bean>

```

The bean class that is being injected must have a corresponding `setCustomerService` setter method—for example:

```

// Java
...
public class ClientInvoker implements Runnable {

```

```

...
public void setCustomerService(CustomerService customerService) {
    this.customerService = customerService;
}
}

```

34.4. INVOKE WS OPERATIONS

Proxy interface is SEI interface

The proxy implements the SEI. Hence, to make remote procedure calls on the Web service, simply invoke the SEI methods on the proxy instance.

Invoking the `lookupCustomer` operation

For example, the `CustomerService` SEI exposes the `lookupCustomer` method, which takes a customer ID as its argument and returns a `Customer` data object. Using the proxy instance, `customerService`, you can invoke the `lookupCustomer` operation as follows:

```

// Java
com.fusesource.demo.customer.Customer response
= customerService.lookupCustomer("1234");

log.info("Got back " + response.getFirstName() + " "
+ response.getLastName()
+ ", ph:" + response.getPhoneNumber() );

```

The `ClientInvoker` class

In the `cxfr-webinars-jboss-fuse-6.1/customer-ws-client` project, there is a `ClientInvoker` class (located in `src/main/java/com/fusesource/customer/client`), which defines a continuous loop that invokes the `lookupCustomer` operation.

When you are experimenting with the demonstration code in the latter chapters of this guide, you might need to modify the `ClientInvoker` class, possibly adding operation invocations.

34.5. DEPLOY TO AN OSGI CONTAINER

Overview

One of the options for deploying the WS client is to package it as an OSGi bundle and deploy it into an OSGi container such as Red Hat JBoss Fuse. Some of the advantages of an OSGi deployment include:

- Bundles are a relatively lightweight deployment option (because dependencies can be shared between deployed bundles).
- OSGi provides sophisticated dependency management, ensuring that only version-consistent dependencies are added to the bundle's classpath.

Using the Maven bundle plug-in

The Maven bundle plug-in is used to package your project as an OSGi bundle, in preparation for deployment into the OSGi container. There are two essential modifications to make to your project's `pom.xml` file:

1. Change the packaging type to **bundle** (by editing the value of the `project/packaging` element in the POM).
2. Add the Maven bundle plug-in to your POM file and configure it as appropriate.

Configuring the Maven bundle plug-in is quite a technical task (although the default settings are often adequate). For full details of how to customize the plug-in configuration, consult *Deploying into the OSGi Container* and *Managing OSGi Dependencies*.

Sample bundle plug-in configuration

The following POM fragment shows a sample configuration of the Maven bundle plug-in, which is appropriate for the current example.

```
<?xml version="1.0"?>
<project ...>
  ...
  <groupId>com.fusesource.byexample.cxf-webinars</groupId>
  <artifactId>customer-ws-client</artifactId>
  <name>customer-ws-client</name>
  <packaging>bundle</packaging>
  ...
  <build>
    <plugins>
      ...
      <plugin>
        <groupId>org.apache.felix</groupId>
        <artifactId>maven-bundle-plugin</artifactId>
        <extensions>>true</extensions>
        <configuration>
          <instructions>
            <Export-Package>
              !com.fusesource.customer.client,
              !com.fusesource.demo.customer,
              !com.fusesource.demo.wsdل.customerservice
            </Export-Package>
            <Import-Package>
              META-INF.cxf,
              *
            </Import-Package>
            <DynamicImport-Package>
              org.apache.cxf.*,
              org.springframework.beans.*
            </DynamicImport-Package>
          </instructions>
        </configuration>
      </plugin>
      ...
    </plugins>
  </build>
</project>
```

```

    </plugins>
  </build>
</project>

```

Dynamic imports

The Java packages from Apache CXF and the Spring API are imported using dynamic imports (specified using the **DynamicImport-Package** element). This is a pragmatic way of dealing with the fact that Spring XML files are not terribly well integrated with the Maven bundle plug-in. At build time, the Maven bundle plug-in is *not* able to figure out which Java classes are required by the Spring XML code. By listing wildcarded package names in the **DynamicImport-Package** element, however, you allow the OSGi container to figure out which Java classes are needed by the Spring XML code at *run* time.



NOTE

In general, using **DynamicImport-Package** headers is not recommended in OSGi, because it short-circuits OSGi version checking. Normally, what should happen is that the Maven bundle plug-in lists the Java packages used at *build* time, along with their versions, in the **Import-Package** header. At *deploy* time, the OSGi container then checks that the available Java packages are compatible with the build time versions listed in the **Import-Package** header. With dynamic imports, this version checking cannot be performed.

Build and deploy the client bundle

After you have configured the POM file, you can build the Maven project and install it in your local repository by entering the following command:

```
mvn install
```

To deploy the client bundle, enter the following command at the containers command console:

```
karaf@root> install -s mvn:com.fusesource.byexample.cxf-webinars/customer-
ws-client
```



NOTE

If your local Maven repository is stored in a non-standard location, you might need to customize the value of the **org.ops4j.pax.url.mvn.localRepository** property in the **EsbInstallDir/etc/org.ops4j.pax.url.mvn.cfg** file, before you can use the **mvn:** scheme to access Maven artifacts.

Check that the client is running

Assuming that you have already deployed the corresponding Web service into the OSGi container, you can verify that the client is successfully invoking WSDL operations by checking the log, as follows:

```
karaf@root> log:display -n 10
```

The client invokes an operation on the Web service once every second.

CHAPTER 35. POJO-BASED ROUTE

35.1. PROCESSING MESSAGES IN POJO FORMAT

Overview

By default, the Camel CXF component marshals incoming Web service requests into the POJO data form, where the *In* message body is encoded as a list of Java objects (one for each operation parameter). The POJO data format has advantages and disadvantages, as follows:

- The big advantage of the POJO data format is that the operation parameters are encoded using the JAX-B standard, which makes them easy to manipulate in Java.
- The downside of the POJO data format, on the other hand, is that it requires that the WSDL metadata is converted to Java in advance (as defined by the JAX-WS and JAX-B mappings) and compiled into your application. This means that a POJO-based route is not very dynamic.

Demonstration location

The code presented in this chapter is taken from the following demonstration:

```
cxf-webinars-jboss-fuse-6.1/customer-ws-camel-cxf-pojo
```

For details of how to download and install the demonstration code, see [Chapter 31, Demonstration Code for Camel/CXF](#)

Camel CXF component

The Camel CXF component is an Apache CXF component that integrates Web services with routes. You can use it either to instantiate *consumer endpoints* (at the start of a route), which behave like Web service instances, or to instantiate *producer endpoints* (at any other points in the route), which behave like WS clients.



NOTE

Camel CXF endpoints—which are instantiated using the `cxf:cxfEndpoint` XML element and are implemented by the Apache Camel project—are not to be confused with the Apache CXF JAX-WS endpoints—which are instantiated using the `jaxws:endpoint` XML element and are implemented by the Apache CXF project.

POJO data format

POJO data format is the *default* data format used by the Camel CXF component and it has the following characteristics:

- JAX-WS and JAX-B stub code (as generated from the WSDL contract) *must* be provided.
- The SOAP body is marshalled into a list of Java objects.

- One Java object for each part or parameter of the corresponding WSDL operation.
- The type of the message body is `org.apache.cxf.message.MessageContentsList`.
- The SOAP headers are converted into headers in the exchange's *In* message.

Implementing and building a POJO route

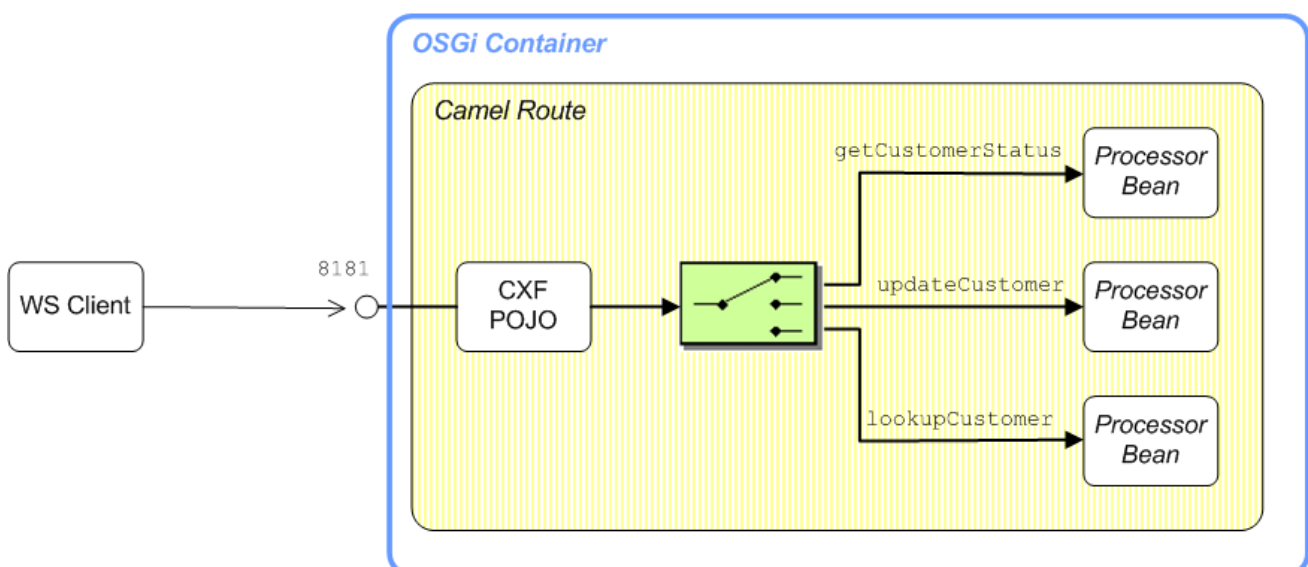
To implement and build the demonstration POJO-based route, starting from scratch, you would perform the following steps:

1. Obtain a copy of the WSDL contract that is to be integrated into the route.
2. Generate the Java stub code from the WSDL contract using a WSDL-to-Java converter. This gives you the SEI, `CustomerService`, and its related classes, such as `Customer`.
3. Instantiate the Camel CXF endpoint in Spring, using the `cxf:cxfEndpoint` element.
4. Implement the route in XML, where you can use the content-based router to sort requests by operation name.
5. Implement the operation processor beans, which are responsible for processing each operation. When implementing these beans, the message contents must be accessed in POJO data format.

Sample POJO route

Figure 35.1, “Sample POJO Route” shows an outline of the route that is used to process the operations of the `CustomerService` Web service using the POJO data format. After sorting the request messages by operation name, an operation-specific processor bean reads the incoming request parameters and then generates a response in the POJO data format.

Figure 35.1. Sample POJO Route



35.2. WSDL-TO-JAVA MAVEN PLUG-IN

Overview

To generate Java stub code from the WSDL contract, you can use either the **ws2java** command-line utility or the **cxf-codegen-plugin** Maven plug-in. When using Maven, the plug-in approach is ideal: after you paste the requisite plug-in configuration into your POM file, the WSDL-to-Java code generation step is integrated into your build.

Configure the WSDL-to-Java Maven plug-in

Configuring the WSDL-to-Java Maven plug-in is relatively easy, because most of the default configuration settings can be left as they are. After copying and pasting the sample **plugin** element into your project's POM file, there are just a few basic settings that need to be customized, as follows:

- *CXF version*—make sure that the plug-in's dependencies are using the latest version of Apache CXF.
- *WSDL file location*—specify the WSDL file location in the **configuration/wsdloptions/wsdloption/wsd** element.
- *Location of output*—specify the root directory of the generated Java source files in the **configuration/sourceRoot** element.

For example, the following POM fragment shows how to configure the **cxf-codegen-plugin** plug-in to generate Java stub code from the **CustomerService.wsdl** WSDL file:

```
<project ...>
  ...
  <parent>
    <groupId>com.fusesource.byexample.cxf-webinars</groupId>
    <artifactId>cxf-webinars</artifactId>
    <version>1.0-SNAPSHOT</version>
  </parent>

  <build>
    <defaultGoal>install</defaultGoal>
    <plugins>
      ...
      <plugin>
        <groupId>org.apache.cxf</groupId>
        <artifactId>cxf-codegen-plugin</artifactId>
        <version>${cxf-version}</version>
        <executions>
          <execution>
            <id>generate-sources</id>
            <phase>generate-sources</phase>
            <configuration>
              <sourceRoot>${basedir}/target/generated-
sources/jaxws</sourceRoot>
              <wsdlOptions>
                <wsdlOption>
<wsdl>${basedir}/../src/main/resources/wsdl/CustomerService.wsdl</wsdl>
                </wsdlOption>
              </wsdlOptions>
            </configuration>
          </execution>
        </executions>
      </plugin>
    </plugins>
  </build>
</project>
```

```
        <goals>
          <goal>wsdl2java</goal>
        </goals>
      </execution>
    </executions>
  </plugin>

</plugins>
</build>

</project>
```

Generated Java source code

With the sample configuration shown here, the generated Java source code is written under the **target/generated-sources/jaxws** directory. Note that the route is dependent on this generated stub code—for example, when processing the POJO parameters, the parameter processor uses the **Customer** data type from the stub code.

Add generated code to IDE

If you are using an IDE such as Eclipse or IntelliJ's IDEA, you need to make sure that the IDE is aware of the generated Java code. For example, in Eclipse it is necessary to add the **target/generated-sources/jaxws** directory to the project as a source code directory.

Compiling the generated code

You must ensure that the generated Java code is compiled and added to the deployment package. By convention, Maven automatically compiles any source files that it finds under the following directory:

```
BaseDir/target/generated-sources/
```

Hence, if you configure the output directory as shown in the preceding POM fragment, the generated code is automatically compiled by Maven.

Reference

For full details of how to configure the Java-to-WSDL plug-in, see the [Maven cxf-codegen-plugin](#) reference page.

35.3. INSTANTIATE THE WS ENDPOINT

Overview

In Apache Camel, the Camel CXF component is the key to integrating routes with Web services. You can use the Camel CXF component to create a CXF endpoint, which can be used in either of the following ways:

- *Consumer*—(at the start of a route) represents a Web service instance, which integrates with the route. The type of payload injected into the route depends on the value of the endpoint's **dataFormat** option.

- *Producer*—(at other points in the route) represents a WS client proxy, which converts the current exchange object into an operation invocation on a remote Web service. The format of the current exchange must match the endpoint's **dataFormat** setting.

In the current demonstration, we are interested in creating a Camel CXF consumer endpoint, with the **dataFormat** option set to POJO.

Maven dependency

The Camel CXF component requires you to add a dependency on the **camel-cxf** component in your Maven POM. For example, the **pom.xml** file from the **customer-ws-camel-cxf-pojo** demonstration project includes the following dependency:

```
<dependency>
  <groupId>org.apache.camel</groupId>
  <artifactId>camel-cxf</artifactId>
  <version>${camel-version}</version>
</dependency>
```

The cxf:bean: URI syntax

The **cxf:bean:** URI is used to bind an Apache CXF endpoint to a route and has the following general syntax:

```
cxf:bean:CxfEndpointID[?Options]
```

Where **CxfEndpointID** is the ID of a bean created using the **cxf:cxfEndpoint** element, which configures the details of the WS endpoint. You can append options to this URI (where the options are described in detail in). If you do not specify any additional options, *the endpoint uses the POJO data format by default.*

For example, to start a route with a Apache CXF endpoint that is configured by the bean with ID, **customer-ws**, define the route as follows:

```
<route>
  <from uri="cxf:bean:customer-ws"/>
  ...
</route>
```



NOTE

There is an alternative URI syntax, **cxf://WsAddress[?Options]**, which enables you to specify *all* of the WS endpoint details in the URI (so there is no need to reference a bean instance). This typically results in a long and cumbersome URI, but is useful in some cases.

The cxf:cxfEndpoint element

The **cxf:cxfEndpoint** element is used to define a WS endpoint that binds either to the start (consumer endpoint) or the end (producer endpoint) of a route. For example, to define the **customer-ws** WS endpoint referenced in the preceding route, you would define a **cxf:cxfEndpoint** element as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ...
  xmlns:cxf="http://camel.apache.org/schema/cxf" ...>
  ...
  <cxf:cxfEndpoint id="customer-ws"
    address="/Customer"
    endpointName="c:SOAPOverHTTP"
    serviceName="c:CustomerService"

serviceClass="com.fusesource.demo.wsdl.customerservice.CustomerService"
  xmlns:c="http://demo.fusesource.com/wsdl/CustomerService/">
  ...
</beans>
```



IMPORTANT

Remember that the **cxf:cxfEndpoint** element and the **jaxws:endpoint** element use *different* XML schemas (although the syntax looks superficially similar). These elements bind a WS endpoint in different ways: the **cxf:cxfEndpoint** element instantiates and binds a WS endpoint to an Apache Camel route, whereas the **jaxws:endpoint** element instantiates and binds a WS endpoint to a Java class using the JAX-WS mapping.

Address for the Jetty container

Apache CXF deploys the WS endpoint into a [Jetty](#) servlet container instance and the **address** attribute of **cxf:cxfEndpoint** is therefore used to configure the addressing information for the endpoint in the Jetty container.

Apache CXF supports the notion of a *default servlet container* instance. The way the default servlet container is initialized and configured depends on the particular mode of deployment that you choose. For example the Red Hat JBoss Fuse container and Web containers (such as Tomcat) provide a default servlet container.

There are two different syntaxes you can use for the endpoint address, where the syntax that you use effectively determines whether or not the endpoint is deployed into the default servlet container, as follows:

- *Address syntax for default servlet container*—to use the default servlet container, specify only the servlet context for this endpoint. Do *not* specify the protocol, host, and IP port in the address. For example, to deploy the endpoint to the **/Customer** servlet context in the default servlet container:

```
address="/Customer"
```

- *Address syntax for custom servlet container*—to instantiate a custom Jetty container for this endpoint, specify a complete HTTP URL, including the host and IP port (the value of the IP port effectively identifies the target Jetty container). Typically, for a Jetty container, you specify the host as **0.0.0.0**, which is interpreted as a wildcard that matches every IP network interface on the local machine (that is, if deployed on a multi-homed host, Jetty opens a listening port on every network card). For example, to deploy the endpoint to the custom Jetty container listening on IP port, **8083**:

```
address="http://0.0.0.0:8083/Customer"
```



NOTE

If you want to configure a secure endpoint (secured by SSL), you would specify the **https:** scheme in the address.

Referencing the SEI

The **serviceClass** attribute of the **cxf:cxfEndpoint** element references the SEI of the Web service, which in this case is the **CustomerService** interface.

35.4. SORT MESSAGES BY OPERATION NAME

The operationName header

When the WS endpoint parses an incoming operation invocation in POJO mode, it automatically sets the **operationName** header to the name of the invoked operation. You can then use this header to sort messages by operation name.

Sorting by operation name

For example, the **customer-ws-camel-cxf-pojo** demonstration defines the following route, which uses the content-based router pattern to sort incoming messages, based on the operation name. The **when** predicates check the value of the **operationName** header using simple language expressions, sorting messages into invocations on the **updateCustomer** operation, the **lookupCustomer** operation, or the **getCustomerStatus** operation.

```
<beans ...>
  ...
  <camelContext id="camel"
xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="cxf:bean:customer-ws"/>
    <choice>
      <when>
        <simple>${in.header.operationName} ==
'updateCustomer'</simple>
        <to uri="updateCustomer"/>
      </when>
      <when>
        <simple>${in.header.operationName} ==
'lookupCustomer'</simple>
        <to uri="lookupCustomer"/>
      </when>
      <when>
        <simple>${in.header.operationName} ==
'getCustomerStatus'</simple>
        <to uri="getCustomerStatus"/>
      </when>
    </choice>
  </route>
</camelContext>
```

```

    <bean id="updateCustomer"
class="com.fusesource.customerwscamelcxfpojo.UpdateCustomerProcessor"/>
    <bean id="getCustomerStatus"
class="com.fusesource.customerwscamelcxfpojo.GetCustomerStatusProcessor"/>
    <bean id="lookupCustomer"
class="com.fusesource.customerwscamelcxfpojo.LookupCustomerProcessor"/>
</beans>

```

Beans as endpoints

Note how the preceding route uses a convenient shortcut to divert each branch of the **choice** DSL to a different processor bean. The DSL for sending exchanges to producer endpoints (for example, `<to uri="Destination"/>`) is integrated with the bean registry: if the *Destination* does not resolve to an endpoint or a component, the *Destination* is used as a bean ID to look up the bean registry. In this example, the exchange is routed to processor beans (which implement the `org.apache.camel.Processor` interface).

35.5. PROCESS OPERATION PARAMETERS

Overview

The most important characteristic of using Camel CXF in POJO mode is that the exchange's message body contains a list of Java objects, representing the parameters of the WSDL operation. The types of the Java objects are defined by the standard JAX-B mapping and the implementations of these parameter types are provided by the Java stub code.

Contents of request message body

In POJO mode, the body of the request message is an `org.apache.cxf.message.MessageContentsList` object. You can also obtain the message body as an `Object[]` array (where type conversion is automatic).

When the body is obtained as an `Object[]` array, the array contains the list of all the operation's IN, INOUT, and OUT parameters *in exactly the same order as defined in the WSDL contract* (and in the same order as the corresponding operation signature of the SEI). The parameter mode affects the content as follows:

IN


Contains a parameter value from the client.

INOUT

Contains a **Holder** object containing a parameter value from the client.

OUT

Contains an *empty* **Holder** object, which is a placeholder for the response.

**NOTE**

Unlike OUT parameters, there is no placeholder in the request's **Object[]** array to represent a return value.

Contents of response message body

In POJO mode, the body of the response message can be either an **org.apache.cxf.message.MessageContentsList** object or an **Object[]** array.


When setting the response body as an **Object[]** array, the array should contain *only* the operation's INOUT and OUT parameters in the same order as defined in the WSDL contract, omitting the IN parameters. The parameter mode affects the content as follows:

INOUT

Contains a **Holder** object, which you must set to a response value. The **Holder** object used here *must be exactly the Holder object for the corresponding parameter that was extracted from the request Object[] array*. Creating and inserting a new **Holder** object into the **Object[]** array does *not* work.

OUT

Contains a **Holder** object, which you must initialize with a response value. The **Holder** object used here *must be exactly the Holder object for the corresponding parameter that was extracted from the request Object[] array*. Creating and inserting a new **Holder** object into the **Object[]** array does *not* work.

**NOTE**

If you defined the Web service interface using the Java-first approach, note that the return value (if any) must be set as the *first* element in the response's **Object[]** array. The return type is set as a plain object: it does *not* use a **Holder** object.

Example: getCustomerStatus operation

For example, the **getCustomerStatus** operation takes three parameters: IN, OUT, and OUT, respectively. The corresponding method signature in the SEI is, as follows:

```
// Java
public void getCustomerStatus(
    @WebParam(name = "customerId", targetNamespace = "")
    java.lang.String customerId,

    @WebParam(mode = WebParam.Mode.OUT, name = "status", targetNamespace =
    "")
    javax.xml.ws.Holder<java.lang.String> status,

    @WebParam(mode = WebParam.Mode.OUT, name = "statusMessage",
    targetNamespace = "")
    javax.xml.ws.Holder<java.lang.String> statusMessage
);
```

Example: request and response bodies

For the `getCustomerStatus` operation, the bodies of the request message and the response message have the following contents:

- *Request message*—as an `Object[]` array type, the contents are: { `String customerId`, `Holder<String> status`, `Holder<String> statusMessage` }.
- *Response message*—as an `Object[]` array type, the contents are: {`Holder<String> status`, `Holder<String> statusMessage` }

Example: processing `getCustomerStatus`

The `GetCustomerStatusProcessor` class is responsible for processing incoming `getCustomerStatus` invocations. The following sample implementation for POJO mode shows how to read the request parameters from the *In* message body and then set the response parameters in the *Out* message body.

```
// Java
package com.fusesource.customerwscamelcxfpojo;

import javax.xml.ws.Holder;
import org.apache.camel.Exchange;
import org.apache.camel.Processor;
import org.slf4j.Logger;
import org.slf4j.LoggerFactory;

public class GetCustomerStatusProcessor implements Processor {
    public static final Logger log =
        LoggerFactory.getLogger(GetCustomerStatusProcessor.class);

    public void process(Exchange exchng) throws Exception {
        Object[] args = exchng.getIn().getBody(Object[].class);

        String id = (String) args[0];
        Holder<String> status = (Holder<String>) args[1];
        Holder<String> statusMsg = (Holder<String>) args[2];

        log.debug("Getting status for customer '" + id + "'");

        // This is where you'd actually do the work! Setting
        // the holder values to constants for the sake of brevity.
        //
        status.value = "Offline";
        statusMsg.value = "Going to sleep now!";

        exchng.getOut().setBody(new Object[] {status , statusMsg});
    }
}
```

35.6. DEPLOY TO OSGI

Overview

One of the options for deploying the POJO-based route is to package it as an OSGi bundle and deploy it into an OSGi container such as Red Hat JBoss Fuse. Some of the advantages of an OSGi deployment include:

- Bundles are a relatively lightweight deployment option (because dependencies can be shared between deployed bundles).
- OSGi provides sophisticated dependency management, ensuring that only version-consistent dependencies are added to the bundle's classpath.

Using the Maven bundle plug-in

The Maven bundle plug-in is used to package your project as an OSGi bundle, in preparation for deployment into the OSGi container. There are two essential modifications to make to your project's `pom.xml` file:

1. Change the packaging type to **bundle** (by editing the value of the `project/packaging` element in the POM).
2. Add the Maven bundle plug-in to your POM file and configure it as appropriate.

Configuring the Maven bundle plug-in is quite a technical task (although the default settings are often adequate). For full details of how to customize the plug-in configuration, consult *Deploying into the OSGi Container* and *Managing OSGi Dependencies*.

Sample bundle plug-in configuration

The following POM fragment shows a sample configuration of the Maven bundle plug-in, which is appropriate for the current example.

```
<?xml version="1.0"?>
<project ...>
  ...
  <groupId>com.fusesource.byexample.cxf-webinars</groupId>
  <artifactId>customer-ws-camel-cxf-pojo</artifactId>

  <name>customer-ws-camel-cxf-pojo</name>
  <packaging>bundle</packaging>
  ...
  <build>
    <plugins>
      ...
      <plugin>
        <groupId>org.apache.felix</groupId>
        <artifactId>maven-bundle-plugin</artifactId>
        <extensions>true</extensions>
        <configuration>
          <instructions>
            <Import-Package>
              META-INF.cxf,
              META-INF.cxf.osgi,
              *
            </Import-Package>
            <DynamicImport-Package>
```

```

        org.apache.cxf.*,
        org.springframework.beans.*
    </DynamicImport-Package>
</instructions>
</configuration>
</plugin>
...
</plugins>
</build>
</project>

```

Dynamic imports

The Java packages from Apache CXF and the Spring API are imported using dynamic imports (specified using the **DynamicImport-Package** element). This is a pragmatic way of dealing with the fact that Spring XML files are not terribly well integrated with the Maven bundle plug-in. At build time, the Maven bundle plug-in is *not* able to figure out which Java classes are required by the Spring XML code. By listing wildcarded package names in the **DynamicImport-Package** element, however, you allow the OSGi container to figure out which Java classes are needed by the Spring XML code at *run* time.



NOTE

In general, using **DynamicImport-Package** headers is not recommended in OSGi, because it short-circuits OSGi version checking. Normally, what should happen is that the Maven bundle plug-in lists the Java packages used at *build* time, along with their versions, in the **Import-Package** header. At *deploy* time, the OSGi container then checks that the available Java packages are compatible with the build time versions listed in the **Import-Package** header. With dynamic imports, this version checking cannot be performed.

Build and deploy the POJO route bundle

After you have configured the POM file, you can build the Maven project and install it in your local repository by entering the following command:

```
mvn install
```

To deploy the route bundle, enter the following command at the JBoss Fuse console:

```
karaf@root> install -s mvn:com.fusesource.byexample.cxf-webinars/customer-
ws-camel-cxf-pojo
```



NOTE

If your local Maven repository is stored in a non-standard location, you might need to customize the value of the **org.ops4j.pax.url.mvn.localRepository** property in the **EsbInstallDir/etc/org.ops4j.pax.url.mvn.cfg** file, before you can use the **mvn:** scheme to access Maven artifacts.

CHAPTER 36. PAYLOAD-BASED ROUTE

36.1. PROCESSING MESSAGES IN PAYLOAD FORMAT

Overview

Select the PAYLOAD format, if you want to access the SOAP message body in XML format, encoded as a DOM object (that is, of `org.w3c.dom.Node` type). One of the advantages of the PAYLOAD format is that no JAX-WS and JAX-B stub code is required, which allows your application to be dynamic, potentially handling many different WSDL interfaces.

Having a message body in XML format enables you to parse the request using XML languages such as XPath and to generate responses using templating languages, such as Velocity.



NOTE

The DOM format is not the optimal type to use for *large* XML message bodies. For large messages, consider using the techniques described in [Chapter 37, Provider-Based Route](#).

Demonstration location


The code presented in this chapter is taken from the following demonstration:

```
cxf-webinars-jboss-fuse-6.1/customer-ws-camel-cxf-payload
```

For details of how to download and install the demonstration code, see [Chapter 31, Demonstration Code for Camel/CXF](#)

Camel CXF component

The Camel CXF component is an Apache CXF component that integrates Web services with routes. You can use it either to instantiate *consumer endpoints* (at the start of a route), which behave like Web service instances, or to instantiate *producer endpoints* (at any other points in the route), which behave like WS clients.



NOTE

Camel CXF endpoints—which are instantiated using the `cxf:cxfEndpoint` XML element and are implemented by the Apache Camel project—are not to be confused with the Apache CXF JAX-WS endpoints—which are instantiated using the `jaxws:endpoint` XML element and are implemented by the Apache CXF project.

PAYLOAD data format

The PAYLOAD data format is selected by setting the `dataFormat=PAYLOAD` option on a Camel CXF endpoint URI and it has the following characteristics:

- Enables you to access the message body as a DOM object (XML payload).
- No JAX-WS or JAX-B stub code required.

- The SOAP body is marshalled as follows:
 - The message body is effectively an XML payload of `org.w3c.dom.Node` type (wrapped in a `CxfPayload` object).
 - The type of the message body is `org.apache.camel.component.cxf.CxfPayload`.
- The SOAP headers are converted into headers in the exchange's *In* message, of `org.apache.cxf.binding.soap.SoapHeader` type.

Implementing and building a PAYLOAD route


To implement and build the demonstration PAYLOAD-based route, starting from scratch, you would perform the following steps:

1. Instantiate the Camel CXF endpoint in Spring, using the `cxf:cxfEndpoint` element.
2. Implement the route in XML, where you can use the content-based router to sort requests by operation name.
3. For each operation, define a processor bean to process the request.
4. Define velocity templates for generating the response messages.

Sample PAYLOAD route

Figure 36.1, “Sample PAYLOAD Route” shows an outline of the route that is used to process the operations of the `CustomerService` Web service using the PAYLOAD data format. After sorting the request messages by operation name, an operation-specific processor bean reads the incoming request parameters. Finally, the response messages are generated using Velocity templates.

Figure 36.1. Sample PAYLOAD Route



36.2. INSTANTIATE THE WS ENDPOINT

Overview

In Apache Camel, the CXF component is the key to integrating routes with Web services. You can use the CXF component to create two different kinds of endpoint:

- *Consumer endpoint*—(at the start of a route) represents a Web service instance, which integrates with the route. The type of payload injected into the route depends on the value of the endpoint's **dataFormat** option.
- *Producer endpoint*—represents a special kind of WS client proxy, which converts the current exchange object into an operation invocation on a remote Web service. The format of the current exchange must match the endpoint's **dataFormat** setting.

The `cxf:bean`: URI syntax

The `cxf:bean`: URI is used to bind an Apache CXF endpoint to a route and has the following general syntax:

```
cxf:bean:CxfEndpointID[?Options]
```

Where **CxfEndpointID** is the ID of a bean created using the `cxf:cxfEndpoint` element, which configures the details of the WS endpoint. You can append options to this URI (where the options are described in detail in). To enable payload mode, you must set the URI option, **dataFormat=PAYLOAD**.

For example, to start a route with an endpoint in PAYLOAD mode, where the endpoint is configured by the **customer-ws** bean, define the route as follows:

```
<route>
  <from uri="cxf:bean:customer-ws?dataFormat=PAYLOAD"/>
  ...
</route>
```

The `cxf:cxfEndpoint` element

The `cxf:cxfEndpoint` element is used to define a WS endpoint that binds either to the start (consumer endpoint) or the end (producer endpoint) of a route. For example, to define the **customer-ws** WS endpoint in PAYLOAD mode, you define a `cxf:cxfEndpoint` element as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ...>
  ...
  <cxf:cxfEndpoint id="customer-ws"
    address="/Customer"
    endpointName="c:SOAPoverHTTP"
    serviceName="c:CustomerService"
    wsdlURL="wsdl/CustomerService.wsdl"
    xmlns:c="http://demo.fusesource.com/wsdl/CustomerService/">
  ...
</beans>
```



NOTE

In the case of PAYLOAD mode, you do *not* need to reference the SEI and you *must* specify the WSDL location instead. In fact, in PAYLOAD mode, you do not require any Java stub code at all.

Address for the Jetty container

Apache CXF deploys the WS endpoint into a [Jetty](#) servlet container instance and the **address** attribute of **cxf:cxfEndpoint** is therefore used to configure the addressing information for the endpoint in the Jetty container.

Apache CXF supports the notion of a *default servlet container* instance. The way the default servlet container is initialized and configured depends on the particular mode of deployment that you choose. For example the OSGi container and Web containers (such as Tomcat) provide a default servlet container.

There are two different syntaxes you can use for the endpoint address, where the syntax that you use effectively determines whether or not the endpoint is deployed into the default servlet container, as follows:

- *Address syntax for default servlet container*—to use the default servlet container, specify only the servlet context for this endpoint. Do *not* specify the protocol, host, and IP port in the address. For example, to deploy the endpoint to the **/Customer** servlet context in the default servlet container:

```
address="/Customer"
```

- *Address syntax for custom servlet container*—to instantiate a custom Jetty container for this endpoint, specify a complete HTTP URL, including the host and IP port (the value of the IP port effectively identifies the target Jetty container). Typically, for a Jetty container, you specify the host as **0.0.0.0**, which is interpreted as a wildcard that matches every IP network interface on the local machine (that is, if deployed on a multi-homed host, Jetty opens a listening port on every network card). For example, to deploy the endpoint to the custom Jetty container listening on IP port, **8083**:

```
address="http://0.0.0.0:8083/Customer"
```



NOTE

If you want to configure a secure endpoint (secured by SSL), you would specify the **https:** scheme in the address.

Specifying the WSDL location

The **wsdlURL** attribute of the **cxf:cxfEndpoint** element is used to specify the location of the WSDL contract for this endpoint. The WSDL contract is used exclusively as the source of metadata for this endpoint: there is need to specify an SEI in PAYLOAD mode.

36.3. SORT MESSAGES BY OPERATION NAME

The operationName header

When the WS endpoint parses an incoming operation invocation in PAYLOAD mode, it automatically sets the **operationName** header to the name of the invoked operation. You can then use this header to sort messages by operation name.

Sorting by operation name

For example, the **customer-ws-camel-cxf-payload** demonstration defines the following route, which uses the content-based router pattern to sort incoming messages, based on the operation name. The **when** predicates check the value of the **operationName** header using simple language expressions, sorting messages into invocations on the **updateCustomer** operation, the **lookupCustomer** operation, or the **getCustomerStatus** operation.

```
<beans ...>
  ...
  <camelContext id="camel"
xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="cxf:bean:customer-ws?dataFormat=PAYLOAD"/>
      <choice>
        <when>
          <simple>${in.header.operationName} ==
'updateCustomer'</simple>
          ...
        </when>
        <when>
          <simple>${in.header.operationName} ==
'lookupCustomer'</simple>
          ...
        </when>
        <when>
          <simple>${in.header.operationName} ==
'getCustomerStatus'</simple>
          ...
        </when>
      </choice>
    </route>
  </camelContext>
</beans>
```

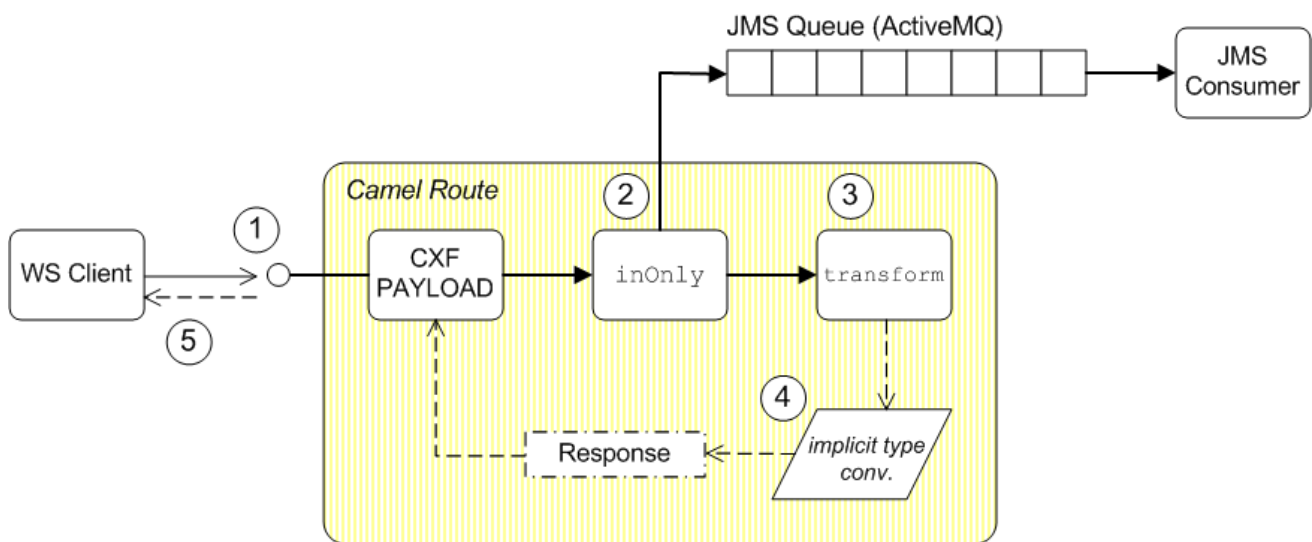
36.4. SOAP/HTTP-TO-JMS BRIDGE USE CASE

Overview

In this section, we consider a SOAP/HTTP-to-JMS bridge use case: that is, you want to create a route that transforms a synchronous operation invocation (over SOAP/HTTP) into an asynchronous message delivery (by pushing the message onto a JMS queue). In this way, it becomes possible to process the incoming operation invocations at a later time, by pulling messages off the JMS queue.

Of course, an alternative solution would be to modify the WSDL contract directly to declare the operation as *OneWay*, thus making the operation asynchronous. Unfortunately, it is often impractical to modify existing WSDL contracts in the real world, because this can have an impact on third-party applications.

Figure 36.2, “SOAP/HTTP-to-JMS Bridge” shows the general outline of a bridge that can transform synchronous SOAP/HTTP invocations into asynchronous JMS message deliveries.

Figure 36.2. SOAP/HTTP-to-JMS Bridge

Transforming RPC operations to One Way

As shown in [Figure 36.2, “SOAP/HTTP-to-JMS Bridge”](#), the route for transforming synchronous SOAP/HTTP to asynchronous JMS works as follows:

1. The WS client invokes a synchronous operation on the Camel CXF endpoint at the start of the route. The Camel CXF endpoint then creates an initial *InOut* exchange at the start of the route, where the body of the exchange message contains a payload in XML format.
2. The **inOnly** DSL command pushes a copy of the XML payload onto a JMS queue, so that it can be processed offline at some later time.
3. The **transform** DSL command constructs an immediate response to send back to the client, where the response has the form of an XML string.
4. The Camel CXF component supports implicit type conversion of the XML string to payload format.
5. The response is sent back to the WS client, thus completing the synchronous operation invocation.

Evidently, this transformation can only work, if the original operation invocation has no return value. Otherwise, it would be impossible to generate a response message before the request has been processed.

Creating a broker instance

You can use Apache ActiveMQ as the JMS implementation. A convenient approach to use in this demonstration is to embed the Apache ActiveMQ broker in the bridge bundle. Simply define an **amq:broker** element in the Spring XML file, as follows:

```
<beans xmlns="http://www.springframework.org/schema/beans"
  ...
  xmlns:amq="http://activemq.apache.org/schema/core"
  ...>

  <amq:broker brokerName="CxfPayloadDemo" persistent="false">
```

```

    <amq:transportConnectors>
      <amq:transportConnector name="openwire"
uri="tcp://localhost:51616"/>
      <amq:transportConnector name="vm" uri="vm:local"/>
    </amq:transportConnectors>
  </amq:broker>
  ...
</beans>

```

**NOTE**

This broker instance is created with the **persistent** attribute set to **false**, so that the messages are stored *only* in memory.

Configuring the JMS component

Because the broker is co-located with the bridge route (in the same JVM), the most efficient way to connect to the broker is to use the VM (Virtual Machine) transport. Configure the Apache ActiveMQ component as follows, to connect to the co-located broker using the VM protocol:

```

<beans ...>
  ...
  <bean id="activemq"
class="org.apache.activemq.camel.component.ActiveMQComponent">
    <property name="brokerURL" value="vm:local"/>
  </bean>
  ...
</beans>

```

**NOTE**

By defining the bean with an **id** value of **activemq**, you are implicitly overriding the component associated with the endpoint URI prefix, **activemq:**. In other words, your custom **ActiveMQComponent** instance is used instead of the default **ActiveMQComponent** instance from the **camel-activemq** JAR file.

Sample SOAP/HTTP-to-JMS route

For example, you could define a route that implements the SOAP/HTTP-to-JMS bridge specifically for the **updateCustomer** operation from the **CustomerService** SEI, as follows:

```

<when>
  <simple>${in.header.operationName} == 'updateCustomer'</simple>
  <log message="Placing update customer message onto queue."/>
  <inOnly uri="activemq:queue:CustomerUpdates?jmsMessageType=Text"/>
  <transform>
    <constant>
      <![CDATA[
<ns2:updateCustomerResponse
xmlns:ns2="http://demo.fusesource.com/wsdl/CustomerService/">
      ]]>

```

```

        </constant>
    </transform>
</when>

```

Sending to the JMS endpoint in `inOnly` mode

Note how the message payload is sent to the JMS queue using the `inOnly` DSL command instead of the `to` DSL command. When you send a message using the `to` DSL command, the default behavior is to use the same invocation mode as the current exchange. But the current exchange has an *InOut* MEP, which means that the `to` DSL command would wait forever for a response message from JMS.

The invocation mode we want to use when sending the payload to the JMS queue is *InOnly* (asynchronous), and we can force this mode by inserting the `inOnly` DSL command into the route.



NOTE

By specifying the option, `jmsMessageType=Text`, Camel CXF implicitly converts the message payload to an XML string before pushing it onto the JMS queue.

Returning a literal response value

The `transform` DSL command uses an expression to set the body of the exchange's *Out* message and this message is then used as the response to the client. Your first impulse when defining a response in XML format might be to use a DOM API, but in this example, the response is specified as a string literal. This approach has the advantage of being both efficient and very easy to program.

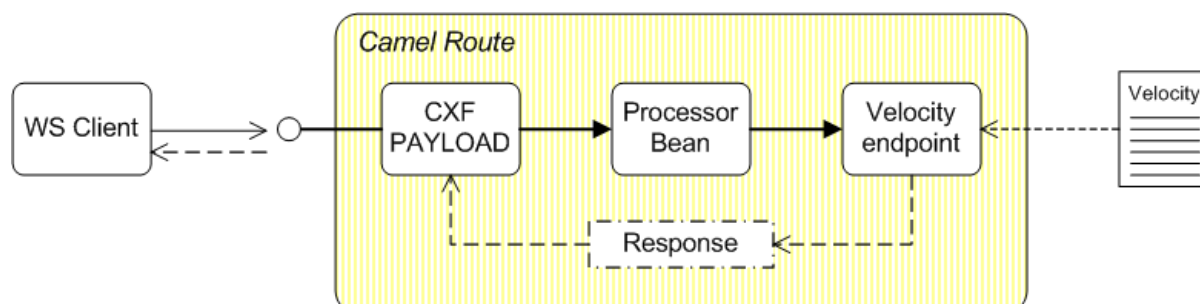
The final step of processing, which consists of converting the XML string to a DOM object, is performed by Apache Camel's implicit type conversion mechanism.

36.5. GENERATING RESPONSES USING TEMPLATES

Overview

One of the simplest and quickest approaches to generating a response message is to use a velocity template. [Figure 36.3, "Response Generated by Velocity"](#) shows the outline of a general template-based route. At the start of the route is a Camel CXF endpoint in `PAYLOAD` mode, which is the appropriate mode to use for processing the message as an XML document. After doing the work required to process the message and stashing some intermediate results in message headers, the route generates the response message using a Velocity template.

Figure 36.3. Response Generated by Velocity



Sample template-based route

For example, you could define a template-based route specifically for the `getCustomerStatus` operation, as follows:

```

...
<when>
  <simple>${in.header.operationName} ==
'getCustomerStatus'</simple>
  <convertBodyTo type="org.w3c.dom.Node"/>
  <setHeader headerName="customerId">
    <xpath>/cus:getCustomerStatus/customerId/text()</xpath>
  </setHeader>
  <to uri="getCustomerStatus"/>
  <to uri="velocity:getCustomerStatusResponse.vm"/>
</when>
</choice>
</route>
</camelContext>
...
<bean id="getCustomerStatus"
  class="com.fusesource.customerwscamelcxfpayload.GetCustomerStatus"/>

```

Route processing steps

Given the preceding route definition, any message whose operation name matches `getCustomerStatus` would be processed as follows:

1. To facilitate processing the payload body, the first step uses `convertBodyTo` to convert the body type from `org.apache.camel.component.cxf.CxfPayload` (the default payload type) to `org.w3c.dom.Node`.
2. The route then applies an XPath expression to the message in order to extract the customer ID value and then stashes it in the `customerId` header.
3. The next step sends the message to the `getCustomerStatus` bean, which does whatever processing is required to get the customer status for the specified customer ID. The results from this step are stashed in message headers.
4. Finally, a response is generated using a velocity template.



NOTE

A common pattern when implementing Apache Camel routes is to use message headers as a temporary stash to hold intermediate results (you could also use exchange properties in the same way).

Converting XPath result to a string

Because the default return type of XPath is a node list, you must explicitly convert the result to a string, if you want to obtain the string contents of an element. There are two alternative ways of obtaining the string value of an element:

- Specify the result type explicitly using the `resultType` attribute, as follows:



```
<xpath
  resultType="java.lang.String">/cus:getCustomerStatus/customerId</xpath>
```

- Modify the expression so that it returns a **text()** node, which automatically converts to string:

```
<xpath>/cus:getCustomerStatus/customerId/text()</xpath>
```

getCustomerStatus processor bean

The **getCustomerStatus** processor bean is an instance of the **GetCustomerStatus** processor class, which is defined as follows:

```
// Java
package com.fusesource.customerwscamelcxfpayload;

import org.apache.camel.Exchange;
import org.apache.camel.Processor;

public class GetCustomerStatus implements Processor
{
    public void process(Exchange exchng) throws Exception {
        String id = exchng.getIn().getHeader("customerId", String.class);

        // Maybe do some kind of lookup here!
        //

        exchng.getIn().setHeader("status", "Away");
        exchng.getIn().setHeader("statusMessage", "Going to sleep.");
    }
}
```

The implementation shown here is just a placeholder. In a realistic application you would perform some sort of checks or database lookup to obtain the customer status. In the demonstration code, however, the **status** and **statusMessage** are simply set to constant values and stashed in message headers.

In the preceding code, we make the modifications directly to the *In* message. When the exchange's *Out* message is **null**, the next processor in the route gets a copy of the current *In* message instead



NOTE

An exceptional case occurs when the message exchange pattern is *inOnly*, in which case the *Out* message value is *always* copied into the *In* message, even if it is **null**.

getCustomerStatusResponse.vm Velocity template

You can generate a response message very simply using a Velocity template. The Velocity template consists of a message in plain text, where specific pieces of data can be inserted using expressions—for example, the expression **`\${header.HeaderName}** substitutes the value of a named header.

The Velocity template for generating the `getCustomerStatus` response is located in the `customer-ws-camel-cxf-payload/src/main/resources` directory and it contains the following template script:

```
<ns2:getCustomerStatusResponse
xmlns:ns2="http://demo.fusesource.com/wsd/CustomerService/">
  <status>${headers.status}</status>
  <statusMessage>${headers.statusMessage}</statusMessage>
</ns2:getCustomerStatusResponse>
```

36.6. TYPECONVERTER FOR CXFPAYLOAD

Overview

Apache Camel supports a type converter mechanism, which is used to perform implicit and explicit type conversions of message bodies and message headers. The payload demonstration requires a customer type converter that can convert **String** objects to **CXFPayload** objects. This type converter automatically gets invoked at the end of the Camel route, when the generated response message (which is a **String** type) gets converted into a **CXFPayload** object.

String to CXFPayload type converter

The **String** to **CXFPayload** type converter is implemented in the `AdditionalCxfPayloadConverters` class, as follows:

```
// Java
package com.fusesource.customerwscamelcxfpayload;

import java.io.ByteArrayInputStream;
import java.io.StringWriter;
import java.util.ArrayList;
import java.util.List;
import javax.xml.transform.OutputKeys;
import javax.xml.transform.Transformer;
import javax.xml.transform.TransformerFactory;
import javax.xml.transform.dom.DOMSource;
import javax.xml.transform.stream.StreamResult;
import org.apache.camel.Converter;
import org.apache.camel.component.cxf.CxfPayload;
import org.apache.cxf.binding.soap.SoapHeader;
import org.slf4j.Logger;
import org.slf4j.LoggerFactory;
import org.w3c.dom.Document;
import org.w3c.dom.Element;
import org.w3c.dom.Node;

@Converter
public class AdditionalCxfPayloadConverters {
    ...
    @Converter
    public static CxfPayload<SoapHeader> toCxfPayload(String xml) {
        // System.out.println("To CxfPayload " + xml);
        List<Element> elements = new ArrayList<Element>();
```

```

        try {
            Document doc = b.newDocumentBuilder().parse(new
ByteArrayInputStream(xml.getBytes()));
            elements.add(doc.getDocumentElement());
        } catch (Exception ex) {
            log.warn("Exception while converting String payload to
CxfPayload; resulting payload will be empty.");
        }
        // The CxfPayload is changed to use Source object under layer, the
elements API only work if we already setup the list before creating the
CxfPayload
        CxfPayload<SoapHeader> ret = new CxfPayload<SoapHeader>(null,
elements);
        return ret;
    }
    ...
}

```

Reference

For full details of the type converter mechanism in Apache Camel, see [Section 40.3, “Built-In Type Converters”](#) and [Chapter 42, *Type Converters*](#).

36.7. DEPLOY TO OSGI

Overview

One of the options for deploying the payload-based route is to package it as an OSGi bundle and deploy it into an OSGi container such as Red Hat JBoss Fuse. Some of the advantages of an OSGi deployment include:

- Bundles are a relatively lightweight deployment option (because dependencies can be shared between deployed bundles).
- OSGi provides sophisticated dependency management, ensuring that only version-consistent dependencies are added to the bundle's classpath.

Using the Maven bundle plug-in

The Maven bundle plug-in is used to package your project as an OSGi bundle, in preparation for deployment into the OSGi container. There are two essential modifications to make to your project's `pom.xml` file:

1. Change the packaging type to **bundle** (by editing the value of the `project/packaging` element in the POM).
2. Add the Maven bundle plug-in to your POM file and configure it as appropriate.

Configuring the Maven bundle plug-in is quite a technical task (although the default settings are often adequate). For full details of how to customize the plug-in configuration, consult *Deploying into the OSGi Container* and *Managing OSGi Dependencies*.

Sample bundle plug-in configuration

The following POM fragment shows a sample configuration of the Maven bundle plug-in, which is appropriate for the current example.

```
<?xml version="1.0"?>
<project ...>
  ...
  <groupId>com.fusesource.byexample.cxf-webinars</groupId>
  <artifactId>customer-ws-camel-cxf-payload</artifactId>
  <name>customer-ws-camel-cxf-payload</name>
  <packaging>bundle</packaging>
  ...
  <build>
    <plugins>
      ...
      <plugin>
        <groupId>org.apache.felix</groupId>
        <artifactId>maven-bundle-plugin</artifactId>
        <extensions>>true</extensions>
        <configuration>
          <instructions>
            <Import-Package>
              org.apache.camel.component.velocity,
              META-INF.cxf,
              META-INF.cxf.osgi,
              javax.jws,
              javax.wsdl,
              javax.xml.bind,
              javax.xml.bind.annotation,
              javax.xml.namespace,
              javax.xml.ws,
              org.w3c.dom,
              <!-- Workaround to access DOM XPathFactory -->
              org.apache.xpath.jaxp,
              *
            </Import-Package>
            <DynamicImport-Package>
              org.apache.cxf.*,
              org.springframework.beans.*
            </DynamicImport-Package>
          </instructions>
        </configuration>
      </plugin>
      ...
    </plugins>
  </build>
</project>
```

Dynamic imports

The Java packages from Apache CXF and the Spring API are imported using dynamic imports (specified using the **DynamicImport-Package** element). This is a pragmatic way of dealing with the fact that Spring XML files are not terribly well integrated with the Maven bundle plug-in. At build time, the Maven bundle plug-in is *not* able to figure out which Java

classes are required by the Spring XML code. By listing wildcarded package names in the **DynamicImport-Package** element, however, you allow the OSGi container to figure out which Java classes are needed by the Spring XML code at *run* time.



NOTE

In general, using **DynamicImport-Package** headers is not recommended in OSGi, because it short-circuits OSGi version checking. Normally, what should happen is that the Maven bundle plug-in lists the Java packages used at *build* time, along with their versions, in the **Import-Package** header. At *deploy* time, the OSGi container then checks that the available Java packages are compatible with the build time versions listed in the **Import-Package** header. With dynamic imports, this version checking cannot be performed.

Build and deploy the client bundle

After you have configured the POM file, you can build the Maven project and install it in your local repository by entering the following command:

```
mvn install
```

Install the **camel-velocity** feature, which is needed for this example:

```
karaf@root> features:install camel-velocity
```

To deploy the route bundle, enter the following command at the console:

```
karaf@root> install -s mvn:com.fusesource.byexample.cxf-webinars/customer-  
ws-camel-cxf-payload
```



NOTE

If your local Maven repository is stored in a non-standard location, you might need to customize the value of the **org.ops4j.pax.url.mvn.localRepository** property in the **InstallDir/etc/org.ops4j.pax.url.mvn.cfg** file, before you can use the **mvn:** scheme to access Maven artifacts.

CHAPTER 37. PROVIDER-BASED ROUTE

37.1. PROVIDER-BASED JAX-WS ENDPOINT

Overview

Use the provider-based approach, if you need to process *very large* Web services messages. The provider-based approach is a variant of the PAYLOAD data format that enables you to encode the message body as an XML streaming type, such as **SAXSource**. Since the XMLstreaming types are more efficient than DOM objects, the provider-based approach is ideal for large XML messages.

Demonstration location

The code presented in this chapter is taken from the following demonstration:

```
cxf-webinars-jboss-fuse-6.1/customer-ws-camel-cxf-provider
```

For details of how to download and install the demonstration code, see [Chapter 31, Demonstration Code for Camel/CXF](#)

Camel CXF component

The Camel CXF component is an Apache CXF component that integrates Web services with routes. You can use it either to instantiate *consumer endpoints* (at the start of a route), which behave like Web service instances, or to instantiate *producer endpoints* (at any other points in the route), which behave like WS clients.



NOTE

Camel CXF endpoints—which are instantiated using the **cxf:cxfEndpoint** XML element and are implemented by the Apache Camel project—are not to be confused with the Apache CXF JAX-WS endpoints—which are instantiated using the **jaxws:endpoint** XML element and are implemented by the Apache CXF project.

Provider-based approach and the PAYLOAD data format

The provider-based approach is a variant of the PAYLOAD data format, which is enabled as follows:

- Define a custom **javax.xml.ws.Provider<StreamType>** class, where the *StreamType* type is an XML streaming type, such as **SAXSource**.
- The PAYLOAD data format is selected by an annotation on the custom **Provider<?>** class (see [the section called “The SAXSourceService provider class”](#)).
- The custom **Provider<?>** class is referenced by setting the **serviceClass** attribute of the **cxf:cxfEndpoint** element in XML configuration.

The provider-based approach has the following characteristics:

- Enables you to access the message body as a streamed XML type—for example, `javax.xml.transform.sax.SAXSource`.
- No JAX-WS or JAX-B stub code required.
- The SOAP body is marshalled into a stream-based **SAXSource** type.
- The SOAP headers are converted into headers in the exchange's *In* message, of `org.apache.cxf.binding.soap.SoapHeader` type.

Implementing and building a provider-based route

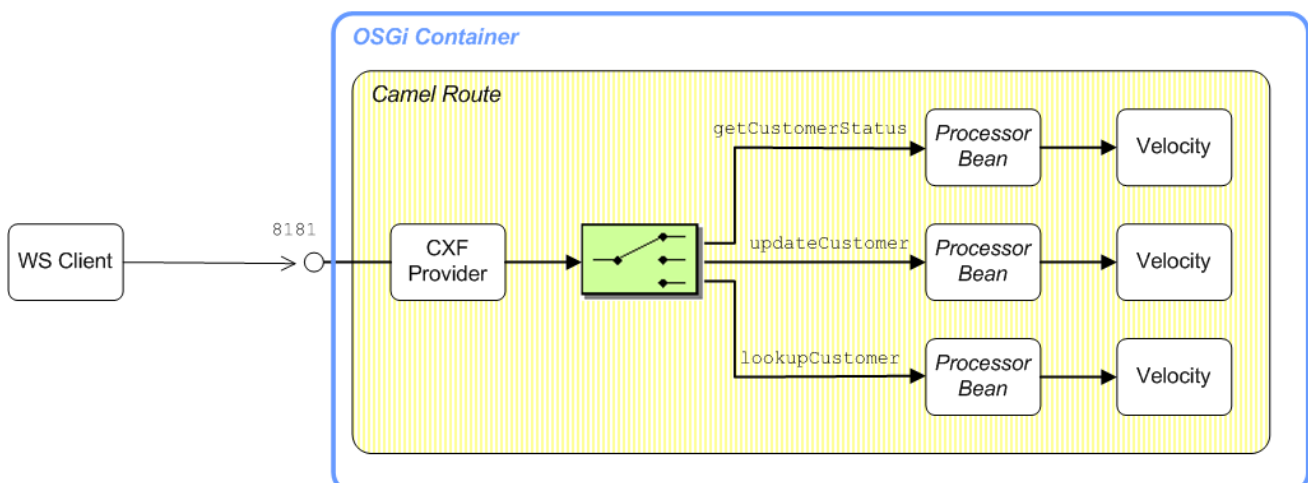
To implement and build the demonstration provider-based route, starting from scratch, you would perform the following steps:

1. Define a custom `javax.xml.ws.Provider<StreamType>` class (the current demonstration uses **SAXSource** as the *StreamType* type).
2. Instantiate the Camel CXF endpoint in Spring, using the `cxf:cxfEndpoint` element and reference the custom provider class (using the `serviceClass` attribute).
3. Implement the route in XML, where you can use the content-based router to sort requests by operation name.
4. For each operation, define a processor bean to process the request.
5. Define velocity templates for generating the response messages.
6. Define a custom type converter, to support converting a **String** message body to a **SAXSource** message body.

Sample provider-based route

Figure 37.1, “Sample Provider-Based Route” shows an outline of the route that is used to process the operations of the **CustomerService** Web service using the provider-based approach. After sorting the request messages by operation name, an operation-specific processor bean reads the incoming request parameters. Finally, the response messages are generated using Velocity templates.

Figure 37.1. Sample Provider-Based Route



37.2. CREATE A PROVIDER<?> IMPLEMENTATION CLASS

Overview

The fundamental prerequisite for using provider mode is to define a custom **Provider<>** class that implements the **invoke()** method. In fact, the sole purpose of this class is to provide runtime type information for Apache CXF: *the **invoke()** method never gets called!*

By implementing the provider class in the way shown here, you are merely indicating to the Apache CXF runtime that the WS endpoint should operate in in PAYLOAD mode and the type of the message PAYLOAD should be **SAXSource**.

The **SAXSourceService** provider class

The definition of the provider class is relatively short and the complete definition of the customer provider class, **SAXSourceService**, is as follows:

```
// Java
package com.fusesource.customerwscamelcxfprovider;

import javax.xml.transform.sax.SAXSource;
import javax.xml.ws.Provider;
import javax.xml.ws.Service.Mode;
import javax.xml.ws.ServiceMode;
import javax.xml.ws.WebServiceProvider;

@WebServiceProvider()
@ServiceMode(Mode.PAYLOAD)
public class SAXSourceService implements Provider<SAXSource>
{
    public SAXSource invoke(SAXSource t) {
        throw new UnsupportedOperationException("Not supported yet.");
    }
}
```

The customer provider class, **SAXSourceService**, must be annotated by the **@WebServiceProvider** annotation to mark it as a provider class and can be optionally annotated by the **@ServiceMode** annotation to select PAYLOAD mode.

37.3. INSTANTIATE THE WS ENDPOINT

Overview

In Apache Camel, the CXF component is the key to integrating routes with Web services. You can use the CXF component to create two different kinds of endpoint:

- *Consumer endpoint*—(at the start of a route) represents a Web service instance, which integrates with the route. The type of payload injected into the route depends on the value of the endpoint's **dataFormat** option.
- *Producer endpoint*—represents a special kind of WS client proxy, which converts the current exchange object into an operation invocation on a remote Web service. The format of the current exchange must match the endpoint's **dataFormat** setting.

The `cxf:bean:` URI syntax

The `cxf:bean:` URI is used to bind an Apache CXF endpoint to a route and has the following general syntax:

```
cxf:bean:CxfEndpointID[?Options]
```

Where **`CxfEndpointID`** is the ID of a bean created using the `cxf:cxfEndpoint` element, which configures the details of the WS endpoint. You can append options to this URI (where the options are described in detail in). Provider mode is essentially a variant of PAYLOAD mode: you could specify this mode on the URI (by setting `dataFormat=PAYLOAD`), but this is not necessary, because PAYLOAD mode is already selected by the `@ServiceMode` annotation on the custom **Provider** class.

For example, to start a route with an endpoint in provider mode, where the endpoint is configured by the `customer-ws` bean, define the route as follows:

```
<route>
  <from uri="cxf:bean:customer-ws"/>
  ...
</route>
```

The `cxf:cxfEndpoint` element

The `cxf:cxfEndpoint` element is used to define a WS endpoint that binds either to the start (consumer endpoint) or the end (producer endpoint) of a route. For example, to define the `customer-ws` WS endpoint in provider mode, you define `cxf:cxfEndpoint` element as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ...>
  ...
  <cxf:cxfEndpoint id="customer-ws"
    address="/Customer"
    endpointName="c:SOAPoverHTTP"
    serviceName="c:CustomerService"
    wsdlURL="wsdl/CustomerService.wsdl"

    serviceClass="com.fusesource.customerwscamelcxfprovider.SAXSourceService"
    xmlns:c="http://demo.fusesource.com/wsdl/CustomerService/" />
  ...
</beans>
```

Specifying the WSDL location

The `wsdlURL` attribute of the `cxf:cxfEndpoint` element is used to specify the location of the WSDL contract for this endpoint. The WSDL contract is used as the source of metadata for this endpoint.

Specifying the service class

A key difference between provider mode and ordinary PAYLOAD mode is that the `serviceClass` attribute must be set to the provider class, **`SAXSourceService`**.

37.4. SORT MESSAGES BY OPERATION NAME

The operationName header

When the WS endpoint parses an incoming operation invocation in PROVIDER mode, it automatically sets the **operationName** header to the name of the invoked operation. You can then use this header to sort messages by operation name.

Sorting by operation name

For example, the **customer-ws-camel-cxf-provider** demonstration defines the following route, which uses the content-based router pattern to sort incoming messages, based on the operation name. The **when** predicates check the value of the **operationName** header using simple language expressions, sorting messages into invocations on the **updateCustomer** operation, the **lookupCustomer** operation, or the **getCustomerStatus** operation.

```
<beans ...>
  ...
  <camelContext id="camel"
xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="cxf:bean:customer-ws"/>
      <choice>
        <when>
          <simple>${in.header.operationName} ==
'updateCustomer'</simple>
          ...
        </when>
        <when>
          <simple>${in.header.operationName} ==
'lookupCustomer'</simple>
          ...
        </when>
        <when>
          <simple>${in.header.operationName} ==
'getCustomerStatus'</simple>
          ...
        </when>
      </choice>
    </route>
  </camelContext>
  ...
</beans>
```

37.5. SOAP/HTTP-TO-JMS BRIDGE USE CASE

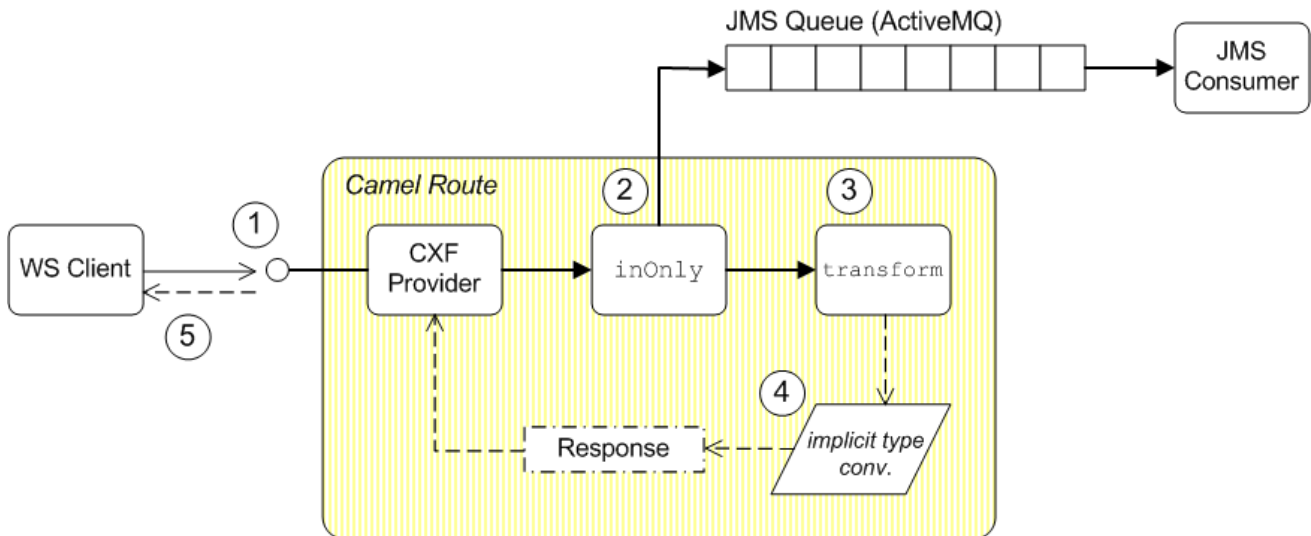
Overview

In this section, we consider a SOAP/HTTP-to-JMS bridge use case: that is, you want to create a route that transforms a synchronous operation invocation (over SOAP/HTTP) into an asynchronous message delivery (by pushing the message onto a JMS queue). In this way, it

becomes possible to process the incoming operation invocations at a later time, by pulling messages off the JMS queue.

Figure 37.2, “SOAP/HTTP-to-JMS Bridge” shows the general outline of a bridge that can transform synchronous SOAP/HTTP invocations into asynchronous JMS message deliveries.

Figure 37.2. SOAP/HTTP-to-JMS Bridge



Transforming RPC operations to One Way

As shown in Figure 37.2, “SOAP/HTTP-to-JMS Bridge”, the route for transforming synchronous SOAP/HTTP to asynchronous JMS works as follows:

1. The WS client invokes a synchronous operation on the Camel CXF endpoint at the start of the route. The Camel CXF endpoint then creates an initial *InOut* exchange at the start of the route, where the body of the exchange message contains a payload in XML format.
2. The **inOnly** DSL command pushes a copy of the XML payload onto a JMS queue, so that it can be processed offline at some later time.
3. The **transform** DSL command constructs an immediate response to send back to the client, where the response has the form of an XML string.
4. The route explicitly converts the XML string to the `javax.xml.transform.sax.SAXSource` type.
5. The response is sent back to the WS client, thus completing the synchronous operation invocation.

Evidently, this transformation can only work, if the original operation invocation has no return value. Otherwise, it would be impossible to generate a response message before the request has been processed.

Creating a broker instance

You can use Apache ActiveMQ as the JMS implementation. A convenient approach to use in this demonstration is to embed the Apache ActiveMQ broker in the bridge bundle. Simply define an `amq:broker` element in the Spring XML file, as follows:


```

<beans xmlns="http://www.springframework.org/schema/beans"
  ...
  xmlns:amq="http://activemq.apache.org/schema/core"
  ...>

  <amq:broker brokerName="CxfPayloadDemo" persistent="false">
    <amq:transportConnectors>
      <amq:transportConnector name="openwire"
uri="tcp://localhost:51616"/>
      <amq:transportConnector name="vm" uri="vm:local"/>
    </amq:transportConnectors>
  </amq:broker>
  ...
</beans>

```

**NOTE**

This broker instance is created with the **persistent** attribute set to **false**, so that the messages are stored *only* in memory.

Configuring the JMS component

Because the broker is co-located with the bridge route (in the same JVM), the most efficient way to connect to the broker is to use the VM (Virtual Machine) transport. Configure the Apache ActiveMQ component as follows, to connect to the co-located broker using the VM protocol:

```

<beans ...>
  ...
  <bean id="activemq"
class="org.apache.activemq.camel.component.ActiveMQComponent">
    <property name="brokerURL" value="vm:local"/>
  </bean>
  ...
</beans>

```

**NOTE**

By defining the bean with an **id** value of **activemq**, you are implicitly overriding the component associated with the endpoint URI prefix, **activemq:**. In other words, your custom **ActiveMQComponent** instance is used instead of the default **ActiveMQComponent** instance from the **camel-activemq** JAR file.

Sample SOAP/HTTP-to-JMS route

For example, you could define a route that implements the SOAP/HTTP-to-JMS bridge specifically for the **updateCustomer** operation from the **CustomerService** SEI, as follows:

```

<when>
  <simple>${in.header.operationName} == 'updateCustomer'</simple>
  <log message="Placing update customer message onto queue."/>
  <inOnly uri="activemq:queue:CustomerUpdates?jmsMessageType=Text"/>
  <transform>

```

```

        <constant>
        <![CDATA[
<ns2:updateCustomerResponse
xmlns:ns2="http://demo.fusesource.com/wsdl/CustomerService/" />
        ]]>
        </constant>
    </transform>
    <convertBodyTo type="javax.xml.transform.sax.SAXSource" />
</when>

```

Sending to the JMS endpoint in `inOnly` mode

Note how the message payload is sent to the JMS queue using the `inOnly` DSL command instead of the `to` DSL command. When you send a message using the `to` DSL command, the default behavior is to use the same invocation mode as the current exchange. But the current exchange has an *InOut* MEP, which means that the `to` DSL command would wait forever for a response message from JMS.

The invocation mode we want to use when sending the payload to the JMS queue is *InOnly* (asynchronous), and we can force this mode by inserting the `inOnly` DSL command into the route.



NOTE

By specifying the option, `jmsMessageType=Text`, Camel CXF implicitly converts the message payload to an XML string before pushing it onto the JMS queue.

Returning a literal response value

The `transform` DSL command uses an expression to set the body of the exchange's *Out* message and this message is then used as the response to the client. Your first impulse when defining a response in XML format might be to use a DOM API, but in this example, the response is specified as a string literal. This approach has the advantage of being both efficient and very easy to program.

Type conversion of the response message

In this example, the reply message (like the request message) is required to be of type, `javax.xml.transform.sax.SAXSource`. In the last step of the route, therefore, you must convert the message body from `String` type to `javax.xml.transform.sax.SAXSource` type, by invoking the `convertBodyTo` DSL command.

The implementation of the `String` to `SAXSource` conversion is provided by a custom type converter, as described in [Section 37.7, "TypeConverter for SAXSource"](#).

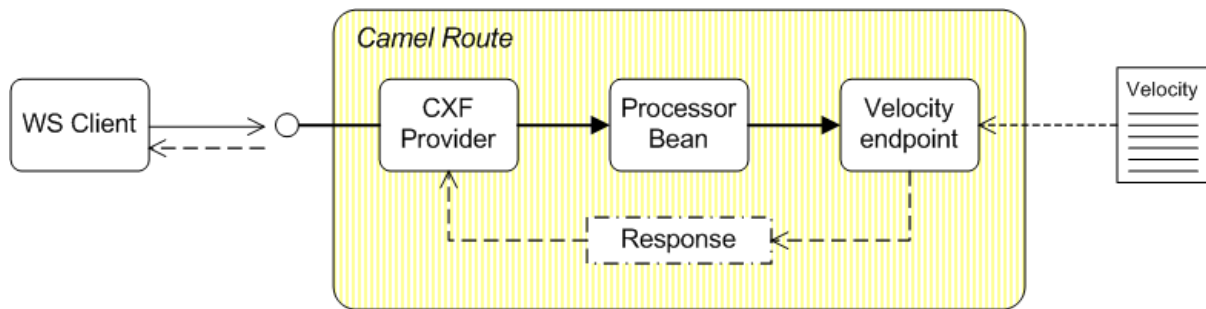
37.6. GENERATING RESPONSES USING TEMPLATES

Overview

One of the simplest and quickest approaches to generating a response message is to use a velocity template. [Figure 37.3, "Response Generated by Velocity"](#) shows the outline of a general template-based route. At the start of the route is a Camel CXF endpoint in provider mode, which is the appropriate mode to use for processing the message as an XML document. After doing the work required to process the message and stashing some

intermediate results in message headers, the route generates the response message using a Velocity template.

Figure 37.3. Response Generated by Velocity



Sample template-based route

For example, you could define a template-based route specifically for the `getCustomerStatus` operation, as follows:

```

...
<when>
  <simple>${in.header.operationName} ==
'getCustomerStatus'</simple>
  <setHeader headerName="customerId">
    <xpath
resultType="java.lang.String">/cus:getCustomerStatus/customerId</xpath>
  </setHeader>
  <to uri="getCustomerStatus"/>
  <to uri="velocity:getCustomerStatusResponse.vm"/>
  <convertBodyTo type="javax.xml.transform.sax.SAXSource"/>
</when>
</choice>
</route>
</camelContext
...
<bean id="getCustomerStatus"
  class="com.fusesource.customerwscamelcxfpayload.GetCustomerStatus"/>

```

Route processing steps

Given the preceding route definition, any message whose operation name matches `getCustomerStatus` would be processed as follows:

1. The route applies an XPath expression to the message in order to extract the customer ID value and then stashes it in the `customerId` header.
2. The next step sends the message to the `getCustomerStatus` bean, which does whatever processing is required to get the customer status for the specified customer ID. The results from this step are stashed in message headers.
3. A response is generated using a Velocity template.
4. Finally, the XML string generated by the Velocity template must be explicitly converted to the `javax.xml.transform.sax.SAXSource` type using `convertBodyTo` (which implicitly relies on a type converter).

**NOTE**

A common pattern when implementing Apache Camel routes is to use message headers as a temporary stash to hold intermediate results (you could also use exchange properties in the same way).

XPath expressions and SAXSource

XPath expressions can be applied directly to SAXSource objects. The XPath implementation has a pluggable architecture that supports a variety of XML parsers and when XPath encounters a SAXSource object, it automatically loads the plug-in required to support SAXSource parsing.

getCustomerStatus processor bean

The `getCustomerStatus` processor bean is an instance of the `GetCustomerStatus` processor class, which is defined as follows:

```
// Java
package com.fusesource.customerwscamelcxfpayload;

import org.apache.camel.Exchange;
import org.apache.camel.Processor;

public class GetCustomerStatus implements Processor
{
    public void process(Exchange exchng) throws Exception {
        String id = exchng.getIn().getHeader("customerId", String.class);

        // Maybe do some kind of lookup here!
        //

        exchng.getIn().setHeader("status", "Away");
        exchng.getIn().setHeader("statusMessage", "Going to sleep.");
    }
}
```

The implementation shown here is just a placeholder. In a realistic application you would perform some sort of checks or database lookup to obtain the customer status. In the demonstration code, however, the `status` and `statusMessage` are simply set to constant values and stashed in message headers.

getCustomerStatusResponse.vm Velocity template

You can generate a response message very simply using a Velocity template. The Velocity template consists of a message in plain text, where specific pieces of data can be inserted using expressions—for example, the expression `${header.HeaderName}` substitutes the value of a named header.

The Velocity template for generating the `getCustomerStatus` response is located in the `customer-ws-camel-cxf-provider/src/main/resources` directory and it contains the following template script:

```
<ns2:getCustomerStatusResponse
xmlns:ns2="http://demo.fusesource.com/wsd/CustomerService/">
```

```

<status>${headers.status}</status>
<statusMessage>${headers.statusMessage}</statusMessage>
</ns2:getCustomerStatusResponse>

```

37.7. TYPECONVERTER FOR SAXSOURCE

Overview

Apache Camel supports a type converter mechanism, which is used to perform implicit and explicit type conversions of message bodies and message headers. The type converter mechanism is extensible and it so happens that the provider demonstration requires a custom type converter that can convert **String** objects to **SAXSource** objects.

String to SAXSource type converter

The **String** to **SAXSource** type converter is implemented in the **AdditionalConverters** class, as follows:

```

// Java
package com.fusesource.customerwscamelcxprovider;

import java.io.ByteArrayInputStream;
import javax.xml.transform.sax.SAXSource;
import org.apache.camel.Converter;
import org.xml.sax.InputSource;

@Converter
public class AdditionalConverters {

    @Converter
    public static SAXSource toSAXSource(String xml) {
        return new SAXSource(new InputSource(new
ByteArrayInputStream(xml.getBytes())));
    }
}

```

Reference

For full details of the type converter mechanism in Apache Camel, see [Section 40.3, “Built-In Type Converters”](#) and [Chapter 42, *Type Converters*](#).

37.8. DEPLOY TO OSGI

Overview

One of the options for deploying the provider-based route is to package it as an OSGi bundle and deploy it into an OSGi container such as Red Hat JBoss Fuse. Some of the advantages of an OSGi deployment include:

- Bundles are a relatively lightweight deployment option (because dependencies can be shared between deployed bundles).

- OSGi provides sophisticated dependency management, ensuring that only version-consistent dependencies are added to the bundle's classpath.

Using the Maven bundle plug-in

The Maven bundle plug-in is used to package your project as an OSGi bundle, in preparation for deployment into the OSGi container. There are two essential modifications to make to your project's `pom.xml` file:

1. Change the packaging type to **bundle** (by editing the value of the `project/packaging` element in the POM).
2. Add the Maven bundle plug-in to your POM file and configure it as appropriate.

Configuring the Maven bundle plug-in is quite a technical task (although the default settings are often adequate). For full details of how to customize the plug-in configuration, consult *Deploying into the OSGi Container* and *Managing OSGi Dependencies*.

Sample bundle plug-in configuration

The following POM fragment shows a sample configuration of the Maven bundle plug-in, which is appropriate for the current example.

```
<?xml version="1.0"?>
<project ...>
  ...
  <groupId>com.fusesource.byexample.cxf-webinars</groupId>
  <artifactId>customer-ws-camel-cxf-provider</artifactId>
  <name>customer-ws-camel-cxf-provider</name>
  <packaging>bundle</packaging>
  ...
  <build>
    <plugins>
      ...
      <plugin>
        <groupId>org.apache.felix</groupId>
        <artifactId>maven-bundle-plugin</artifactId>
        <extensions>>true</extensions>
        <configuration>
          <instructions>
            <Import-Package>
              org.apache.camel.component.velocity,
              META-INF.cxf,
              META-INF.cxf.osgi,
              javax.jws,
              javax.wsdl,
              javax.xml.bind,
              javax.xml.bind.annotation,
              javax.xml.namespace,
              javax.xml.ws,
              org.w3c.dom,
              <!-- Workaround to access DOM XPathFactory -->
              org.apache.xpath.jaxp,
              *
            </Import-Package>
            <DynamicImport-Package>
```

```

        org.apache.cxf.*,
        org.springframework.beans.*
    </DynamicImport-Package>
</instructions>
</configuration>
</plugin>
...
</plugins>
</build>
</project>

```

Dynamic imports

The Java packages from Apache CXF and the Spring API are imported using dynamic imports (specified using the **DynamicImport-Package** element). This is a pragmatic way of dealing with the fact that Spring XML files are not terribly well integrated with the Maven bundle plug-in. At build time, the Maven bundle plug-in is *not* able to figure out which Java classes are required by the Spring XML code. By listing wildcarded package names in the **DynamicImport-Package** element, however, you allow the OSGi container to figure out which Java classes are needed by the Spring XML code at *run* time.



NOTE

In general, using **DynamicImport-Package** headers is not recommended in OSGi, because it short-circuits OSGi version checking. Normally, what should happen is that the Maven bundle plug-in lists the Java packages used at *build* time, along with their versions, in the **Import-Package** header. At *deploy* time, the OSGi container then checks that the available Java packages are compatible with the build time versions listed in the **Import-Package** header. With dynamic imports, this version checking cannot be performed.


Build and deploy the client bundle

After you have configured the POM file, you can build the Maven project and install it in your local repository by entering the following command:

```
mvn install
```

To deploy the route bundle, enter the following command at the container console:

```
karaf@root> install -s mvn:com.fusesource.byexample.cxf-webinars/customer-
ws-camel-cxf-provider
```



NOTE

If your local Maven repository is stored in a non-standard location, you might need to customize the value of the **org.ops4j.pax.url.mvn.localRepository** property in the **EsbInstallDir/etc/org.ops4j.pax.url.mvn.cfg** file, before you can use the **mvn:** scheme to access Maven artifacts.

CHAPTER 38. PROXYING A WEB SERVICE

Abstract

A common use case for the Camel CXF component is to use a route as a *proxy* for a Web service. That is, in order to perform additional processing of WS request and response messages, you interpose a route between the WS client and the original Web service.


38.1. PROXYING WITH HTTP

Overview

The simplest way to proxy a SOAP/HTTP Web service is to treat the request and reply messages as HTTP packets. This type of proxying can be used where there is no requirement to read or modify the messages passing through the route. For example, you could use this kind of proxying to apply various patterns of flow control on the WS messages.

Figure 38.1, “Proxy Route with Message in HTTP Format” shows an overview of how to proxy a Web service using an Apache Camel route, where the route treats the messages as HTTP packets. The key feature of this route is that both the consumer endpoint (at the start of the route) and the producer endpoint (at the end of the route) must be compatible with the HTTP packet format.

Figure 38.1. Proxy Route with Message in HTTP Format



Alternatives for the consumer endpoint

The following Apache Camel endpoints can be used as consumer endpoints for HTTP format messages:

- *Jetty endpoint*—is a lightweight Web server. You can use Jetty to handle messages for *any* HTTP-based protocol, including the commonly-used Web service SOAP/HTTP protocol.
- *Camel CXF endpoint in MESSAGE mode*—when a Camel CXF endpoint is used in MESSAGE mode, the body of the exchange message is the raw message received from the transport layer (which is HTTP). In other words, the Camel CXF endpoint in MESSAGE mode is equivalent to a Jetty endpoint in the case of HTTP-based protocols.

Consumer endpoint for HTTP

A Jetty endpoint has the general form, **jetty:HttpAddress**. To configure the Jetty endpoint to be a proxy for a Web service, use a **HttpAddress** value that is almost identical to the HTTP address the client connects to, except that Jetty's version of **HttpAddress** uses the

special hostname, `0.0.0.0` (which matches *all* of the network interfaces on the current machine).

```
<route>
  <from uri="jetty:http://0.0.0.0:9093/Customers?
matchOnUriPrefix=true"/>
  ...
</route>
```

matchOnUriPrefix option

Normally, a Jetty consumer endpoint accepts only an *exact* match on the context path. For example, a request that is sent to the address `http://localhost:9093/Customers` would be accepted, but a request sent to `http://localhost:9093/Customers/Foo` would be rejected. By setting `matchOnUriPrefix` to `true`, however, you enable a kind of wildcarding on the context path, so that any context path prefixed by `/Customers` is accepted.

Alternatives for the producer endpoint

The following Apache Camel endpoints can be used as producer endpoints for HTTP format messages:

- *Jetty HTTP client endpoint*—(recommended) the Jetty library implements a HTTP client. In particular, the Jetty HTTP client features support for `HttpClient` thread pools, which means that the Jetty implementation scales particularly well.
- *HTTP endpoint*—the HTTP endpoint implements a HTTP client based on the `HttpClient` 3.x API.
- *HTTP4 endpoint*—the HTTP endpoint implements a HTTP client based on the `HttpClient` 4.x API.

Producer endpoint for HTTP

To configure a Jetty HTTP endpoint to send HTTP requests to a remote SOAP/HTTP Web service, set the `uri` attribute of the `to` element at the end of the route to be the address of the remote Web service, as follows:

```
<route>
  ...
  <to uri="jetty:http://localhost:8083/Customers?
bridgeEndpoint=true&throwExceptionOnFailure=false"/>
</route>
```

bridgeEndpoint option

The HTTP component supports a `bridgeEndpoint` option, which you can enable on a HTTP producer endpoint to configure the endpoint appropriately for operating in a HTTP-to-HTTP bridge (as is the case in this demonstration). In particular, when `bridgeEndpoint=true`, the HTTP endpoint ignores the value of the `Exchange.HTTP_URI` header, using the HTTP address from the endpoint URI instead.

throwExceptionOnFailure option

Setting `throwExceptionOnFailure` to `false` ensures that any HTTP exceptions are relayed back to the original WS client, instead of being thrown within the route.

Handling message headers

When defining a HTTP bridge application, the `CamelHttp*` headers set by the consumer endpoint at the start of the route can affect the behavior of the producer endpoint. For this reason, in a bridge application it is advisable to remove the `CamelHttp*` headers before the message reaches the producer endpoint, as follows:

```
<route>
  <from uri="jetty:http:..."/>
  ...
  <removeHeaders pattern="CamelHttp*" />
  <to uri="jetty:http:..."/>
</route>
```

Outgoing HTTP headers

By default, any headers in the exchange that are *not* prefixed by `Camel` will be converted into HTTP headers and sent out over the wire by the HTTP producer endpoint. This could have adverse consequences on the behavior of your application, *so it is important to be aware of any headers that are set in the exchange object and to remove them, if necessary.*

For more details about dealing with headers, see [Section 38.4, “Handling HTTP Headers”](#).


38.2. PROXYING WITH POJO FORMAT

Overview

If you want to access the content of the Web services messages that pass through the route, you might prefer to process the messages in POJO format: that is, where the body of the exchange consists of a list of Java objects representing the WS operation parameters. The key advantage of using POJO format is that you can easily *process the contents of a message*, by accessing the operation parameters as Java objects.

[Figure 38.2, “Proxy Route with Message in POJO Format”](#) shows an overview of how to proxy a Web service using an Apache Camel route, where the route processes the messages in POJO format. The key feature of this route is that both the consumer endpoint (at the start of the route) and the producer endpoint (at the end of the route) must be compatible with the POJO data format.

Figure 38.2. Proxy Route with Message in POJO Format



Consumer endpoint for CXF/POJO

To parse incoming messages into POJO data format, the consumer endpoint at the start of

the route must be a Camel CXF endpoint that is configured to use POJO mode. Use the **cxf:bean:BeanID** URI format to reference the Camel CXF endpoint as follows (where the **dataFormat** option defaults to POJO):

```
<route>
  <from uri="cxf:bean:customerServiceProxy"/>
  ...
</route>
```

The bean with the ID, **customerServiceProxy**, is a Camel CXF/POJO endpoint, which is defined as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ...>
  ...
  <cxf:cxfEndpoint
    id="customerServiceProxy"
    xmlns:c="http://demo.fusesource.org/wsdl/camelcxf"
    address="/Customers"
    endpointName="c:SOAPoverHTTP"
    serviceName="c:CustomerService"
    wsdlURL="wsdl/CustomerService.wsdl"
    serviceClass="org.fusesource.demo.wsdl.camelcxf.CustomerService"
  />
  ...
</beans>
```

Producer endpoint for CXF/POJO

To convert the exchange body from POJO data format to a SOAP/HTTP message, the producer endpoint at the end of the route must be a Camel CXF endpoint configured to use POJO mode. Use the **cxf:bean:BeanID** URI format to reference the Camel CXF endpoint as follows (where the **dataFormat** option defaults to POJO):

```
<route>
  ...
  <to uri="cxf:bean:customerServiceReal"/>
</route>
```

The bean with the ID, **customerServiceReal**, is a Camel CXF/POJO endpoint, which is defined as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<beans ...>
  ...
  <cxf:cxfEndpoint
    id="customerServiceReal"
    xmlns:c="http://demo.fusesource.org/wsdl/camelcxf"
    address="http://localhost:8083/Customers"
    endpointName="c:SOAPoverHTTP"
    serviceName="c:CustomerService"
    wsdlURL="wsdl/CustomerService.wsdl"
    serviceClass="org.fusesource.demo.wsdl.camelcxf.CustomerService"
  />
  ...
</beans>
```

```

/>
...
</beans>

```


38.3. PROXYING WITH PAYLOAD FORMAT

Overview

If you want to access the content of the Web services messages that pass through the route, you might prefer to process the messages in the normal PAYLOAD format: that is, where the body of the exchange is accessible as an XML document (essentially, an `org.w3c.dom.Node` object). The key advantage of using PAYLOAD format is that you can easily *process the contents of a message*, by accessing the message body as an XML document.

Figure 38.3, “Proxy Route with Message in PAYLOAD Format” shows an overview of how to proxy a Web service using an Apache Camel route, where the route processes the messages in PAYLOAD format. The key feature of this route is that both the consumer endpoint (at the start of the route) and the producer endpoint (at the end of the route) must be compatible with the PAYLOAD data format.

Figure 38.3. Proxy Route with Message in PAYLOAD Format



Consumer endpoint for CXF/PAYLOAD

To parse incoming messages into PAYLOAD data format, the consumer endpoint at the start of the route must be a Camel CXF endpoint that is configured to use PAYLOAD mode. Use the `cxf:bean:BeanID` URI format to reference the Camel CXF endpoint as follows, where you *must* set the `dataFormat` option to PAYLOAD:

```

<route>
  <from uri="cxf:bean:customerServiceProxy?dataFormat=PAYLOAD" />
  ...
</route>

```

The bean with the ID, `customerServiceProxy`, is a Camel CXF/PAYLOAD endpoint, which is defined as follows:

```

<?xml version="1.0" encoding="UTF-8"?>
<beans ...>
  ...
  <cxf:cxfEndpoint
    id="customerServiceProxy"
    xmlns:c="http://demo.fusesource.org/wsdl/camelcxf"
    address="/Customers"
    endpointName="c:SOAPoverHTTP"
    serviceName="c:CustomerService"
    wsdlURL="wsdl/CustomerService.wsdl"
  >

```

```

    />
    ...
</beans>

```

Producer endpoint for CXF/PAYLOAD

To convert the exchange body from PAYLOAD data format to a SOAP/HTTP message, the producer endpoint at the end of the route must be a Camel CXF endpoint configured to use PAYLOAD mode. Use the `cxf:bean:BeanID` URI format to reference the Camel CXF endpoint as follows, where you *must* set the `dataFormat` option to PAYLOAD:

```

<route>
    ...
    <to uri="cxf:bean:customerServiceReal?dataFormat=PAYLOAD" />
</route>

```

The bean with the ID, `customerServiceReal`, is a Camel CXF/PAYLOAD endpoint, which is defined as follows:

```

<?xml version="1.0" encoding="UTF-8"?>
<beans ...>
    ...
    <cxf:cxfEndpoint
        id="customerServiceReal"
        xmlns:c="http://demo.fusesource.org/wsd/camelcxf"
        address="http://localhost:8083/Customers"
        endpointName="c:SOAPoverHTTP"
        serviceName="c:CustomerService"
        wsdlURL="wsdl/CustomerService.wsdl"
    />
    ...
</beans>

```

Outgoing HTTP headers

By default, any headers in the exchange that are *not* prefixed by `Camel` will be converted into HTTP headers and sent out over the wire by the Camel CXF producer endpoint. This could have adverse consequences on the behavior of your application, *so it is important to be aware of any headers that are set in the exchange object and to remove them, if necessary.*

For more details about dealing with headers, see [Section 38.4, “Handling HTTP Headers”](#).

38.4. HANDLING HTTP HEADERS

Overview

When building bridge applications using HTTP or HTTP-based components, it is important to be aware of how the HTTP-based endpoints process headers. In many cases, internal headers (prefixed by Camel) or other headers can cause unwanted side-effects on your application. It is often necessary to remove or filter out certain headings or classes of headings in your route, in order to ensure that your application behaves as expected.

HTTP-based components

The behavior described in this section affects not just the Camel HTTP component (`camel-http`), but also a number of other HTTP-based components, including:

```
camel-http
camel-http4
camel-jetty
camel-restlet
camel-cxf
```

HTTP consumer endpoint

When a HTTP consumer endpoint receives an incoming message, it creates an *In* message with the following headers:

CamelHttp* headers

Several headers with the `CamelHttp` prefix are created, which record the status of the incoming message. For details of these internal headers, see [HTTP](#).

HTTP headers

All of the HTTP headers from the original incoming message are mapped to headers on the exchange's *In* message.

URL options (*Jetty only*)

The URL options from the original HTTP request URL are mapped to headers on the exchange's *In* message. For example, given the client request with the URL, `http://myserver/myserver?orderid=123`, a Jetty consumer endpoint creates the `orderid` header with value `123`.

HTTP producer endpoint

When a HTTP producer endpoint receives an exchange and converts it to the target message format, it handles the *In* message headers as follows:

CamelHttp*

Headers prefixed by `CamelHttp` are used to control the behavior of the HTTP producer endpoint. Any headers of this kind are consumed by the HTTP producer endpoint and the endpoint behaves as directed.

Camel*

All other headers prefixed by `Camel` are presumed to be meant for internal use and are *not* mapped to HTTP headers in the target message (in other words, these headers are ignored).

*

All other headers are converted to HTTP headers in the target message, with the exception of the following headers, which are blocked (based on a case-insensitive match):

```
content-length
```

```

content-type
cache-control
connection
date
pragma
trailer
transfer-encoding
upgrade
via
warning

```

Implications for HTTP bridge applications

When defining a HTTP bridge application (that is, a route starting with a HTTP consumer endpoint and ending with a HTTP producer endpoint), the **CamelHttp*** headers set by the consumer endpoint at the start of the route can affect the behavior of the producer endpoint. For this reason, in a bridge application it is advisable to remove the **CamelHttp*** headers, as follows:

```

from("http://0.0.0.0/context/path")
  .removeHeaders("CamelHttp*")
  ...
  .to("http://remoteHost/context/path");

```

Setting a custom header filter

If you want to customize the way that a HTTP producer endpoint processes headers, you can define your own customer header filter by defining the **headerFilterStrategy** option on the endpoint URI. For example, to configure a producer endpoint with the **myHeaderFilterStrategy** filter, you could use a URI like the following:

```

http://remoteHost/context/path?
headerFilterStrategy=#myHeaderFilterStrategy

```

Where **myHeaderFilterStrategy** is the bean ID of your custom filter instance.

CHAPTER 39. FILTERING SOAP MESSAGE HEADERS

Abstract

The Camel CXF component supports a flexible header filtering mechanism, which enables you to process SOAP headers, applying different filters according to the header's XML namespace.

39.1. BASIC CONFIGURATION

Overview

When more than one CXF endpoint appears in a route, you need to decide whether or not to allow headers to propagate between the endpoints. By default, the headers are relayed back and forth between the endpoints, but in many cases it might be necessary to filter the headers or to block them altogether. You can control header propagation by applying filters to producer endpoints.

CxfHeaderFilterStrategy

Header filtering is controlled by the **CxfHeaderFilterStrategy** class. Basic configuration of the **CxfHeaderFilterStrategy** class involves setting one or more of the following options:

- [the section called “relayHeaders option”](#).
- [the section called “relayAllMessageHeaders option”](#).

relayHeaders option

The semantics of the **relayHeaders** option can be summarized as follows:

	In-band headers	Out-of-band headers
relayHeaders=true, dataFormat=PAYLOAD	<i>Filter</i>	<i>Filter</i>
relayHeaders=true, dataFormat=POJO	<i>Relay all</i>	<i>Filter</i>
relayHeaders=false	<i>Block</i>	<i>Block</i>

In-band headers

An *in-band header* is a header that is explicitly defined as part of the WSDL binding contract for an endpoint.

Out-of-band headers

An *out-of-band header* is a header that is serialized over the wire, but is not explicitly part of the WSDL binding contract. In particular, the SOAP binding permits out-of-band headers,

because the SOAP specification does *not* require headers to be defined in the WSDL contract.

Payload format

The CXF endpoint's payload format affects the filter behavior as follows:

POJO

(Default) Only out-of-band headers are available for filtering, because the in-band headers have already been processed and removed from the list by CXF. The in-band headers are incorporated into the **MessageContentList** in POJO mode. If you require access to headers in POJO mode, you have the option of implementing a custom CXF interceptor or JAX-WS handler.

PAYLOAD

In this mode, both in-band and out-of-band headers are available for filtering.

MESSAGE

Not applicable. (In this mode, the message remains in a raw format and the headers are not processed at all.)

Default filter

The default filter is of type, **SoapMessageHeaderFilter**, which removes only the SOAP headers that the SOAP specification expects an intermediate Web service to consume. For more details, see [the section called "SoapMessageHeaderFilter"](#).

Overriding the default filter

You can override the default **CxfHeaderFilterStrategy** instance by defining a new **CxfHeaderFilterStrategy** bean and associating it with a CXF endpoint.

Sample relayHeaders configuration

The following example shows how you can use the **relayHeaders** option to create a **CxfHeaderFilterStrategy** bean that blocks all message headers. The CXF endpoints in the route use the **headerFilterStrategy** option to install the filter strategy in the endpoint, where the **headerFilterStrategy** setting has the syntax, **headerFilterStrategy=#BeanID**.

```
<beans ...>
  ...
  <bean id="dropAllMessageHeadersStrategy"
class="org.apache.camel.component.cxf.common.header.CxfHeaderFilterStrateg
y">
    <!-- Set relayHeaders to false to drop all SOAP headers -->
    <property name="relayHeaders" value="false"/>
  </bean>

  <camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="cxf:bean:routerNoRelayEndpoint?"
```

```

headerFilterStrategy=#dropAllMessageHeadersStrategy"/>
    <to uri="cxf:bean:serviceNoRelayEndpoint?
headerFilterStrategy=#dropAllMessageHeadersStrategy"/>
    </route>
</camelContext>
...
</beans>

```

relayAllMessageHeaders option

The **relayAllMessageHeaders** option is used to propagate *all* SOAP headers, without applying any filtering (any installed filters would be bypassed). In order to enable this feature, you must set *both* **relayHeaders** and **relayAllMessageHeaders** to **true**.

Sample relayAllMessageHeaders configuration

The following example shows how to configure CXF endpoints to propagate *all* SOAP message headers. The **propagateAllMessages** filter strategy sets both **relayHeaders** and **relayAllMessageHeaders** to **true**.

```

<beans ...>
  ...
  <bean id="propagateAllMessages"
class="org.apache.camel.component.cxf.common.header.CxfHeaderFilterStrateg
y">
    <!-- Set both properties to true to propagate *all* SOAP headers --
  >
    <property name="relayHeaders" value="true"/>
    <property name="relayAllMessageHeaders" value="true"/>
  </bean>

  <camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
    <route>
      <from uri="cxf:bean:routerNoRelayEndpoint?
headerFilterStrategy=#propagateAllMessages"/>
      <to uri="cxf:bean:serviceNoRelayEndpoint?
headerFilterStrategy=#propagateAllMessages"/>
    </route>
  </camelContext>
  ...
</beans>

```

39.2. HEADER FILTERING

Overview

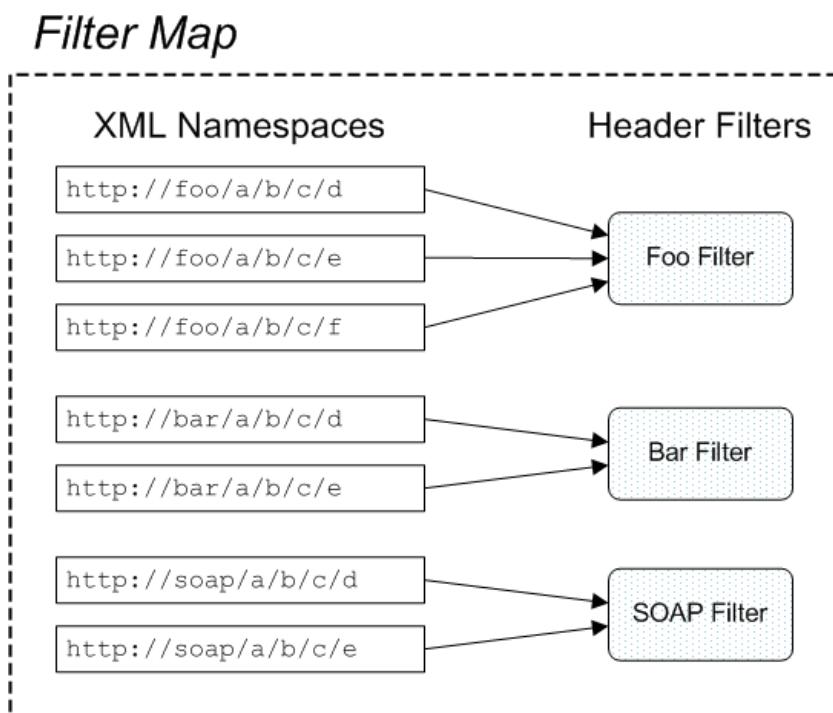
You can optionally install multiple headers in a **CxfHeaderFilterStrategy** instance. The filtering mechanism then uses the header's XML namespace to lookup a particular filter, which it then applies to the header.

Filter map

Figure 39.1, “Filter Map” shows an overview of the filter map that is contained within a

CxfHeaderFilterStrategy instance. For each filter that you install in **CxfHeaderFilterStrategy**, corresponding entries are made in the filter map, where one or more XML schema namespaces are associated with each filter.

Figure 39.1. Filter Map



Filter behavior

When a header is filtered, the filter mechanism peeks at the header to discover the header's XML namespace. The filter then looks up the XML namespace in the filter map to find the corresponding filter implementation. This filter is then applied to the header.

PAYLOAD mode

In PAYLOAD mode, both in-band and out-of-band messages pass through the installed filters.

POJO mode

In POJO mode, only out-of-band messages pass through the installed filters. In-band messages bypass the filters and are propagated by default.

39.3. IMPLEMENTING A CUSTOM FILTER

Overview

You can implement your own customer message header filters by implementing the **MessageHeaderFilter** Java interface. You must associate a filter with one or more XML schema namespaces (representing the header's namespace) and it is possible to differentiate between request message headers and response message headers.

MessageHeaderFilter interface

The `MessageHeaderFilter` interface is defined in the `org.apache.camel.component.cxf.common.header` package, as follows:

```
// Java
package org.apache.camel.component.cxf.common.header;

import java.util.List;

import org.apache.camel.spi.HeaderFilterStrategy.Direction;
import org.apache.cxf.headers.Header;

public interface MessageHeaderFilter {
    List<String> getActivationNamespaces();

    void filter(Direction direction, List<Header> headers);
}
```

Implementing the filter() method

The `MessageHeaderFilter.filter()` method is responsible for applying header filtering. Filtering is applied both before and after an operation is invoked on an endpoint. Hence, there are two directions to which filtering is applied, as follows:

Direction.OUT

When the `direction` parameter equals `Direction.OUT`, the filter is being applied to a request either leaving a consumer endpoint or entering a producer endpoint (that is, it applies to a WS request message propagating through a route).

Direction.IN

When the `direction` parameter equals `Direction.IN`, the filter is being applied to a response either leaving a producer endpoint or entering a consumer endpoint (that is, it applies to a WS response message being sent back).

Filtering can be applied by removing elements from the list of headers, `headers`. Any headers left in the list are propagated.

Binding filters to XML namespaces

It is possible to register multiple header filters against a given CXF endpoint. The CXF endpoint selects the appropriate filter to use based on the XML namespace of the WSDL binding protocol (for example, the namespace for the SOAP 1.1 binding or for the SOAP 1.2 binding). If a header's namespace is unknown, the header is propagated by default.

To bind a filter to one or more namespaces, implement the `getActivationNamespaces()` method, which returns the list of bound XML namespaces.

Identifying the namespace to bind to

[Example 39.1, "Sample Binding Namespaces"](#) illustrates how to identify the namespaces to which you can bind a filter. This example shows the WSDL file for a Bank server that exposes SOAP endpoints.

Example 39.1. Sample Binding Namespaces

```

<wsdl:definitions
targetNamespace="http://cxf.apache.org/schemas/cxf/idl/bank"
  xmlns:tns="http://cxf.apache.org/schemas/cxf/idl/bank"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
  ...
  <wsdl:binding name="BankSOAPBinding" type="tns:Bank">
    <soap:binding style="document"
transport="http://schemas.xmlsoap.org/soap/http" />
    <wsdl:operation name="getAccount">
      ...
    </wsdl:operation>
  </wsdl:binding>
  ...
</wsdl>

```

From the **soap:binding** tag, you can infer that namespace associated with the SOAP binding is **http://schemas.xmlsoap.org/wsdl/soap/**.

Implementing a custom filter

If you want to implement your own custom filter, define a class that inherits from the **MessageHeaderFilter** interface and implement its methods as described in this section. For example, [Example 39.2, “Sample Header Filter Implementation”](#) shows an example of a custom filter, **CustomHeaderFilter**, that binds to the namespace, **http://cxf.apache.org/bindings/custom**, and relays all of the headers that pass through it.

Example 39.2. Sample Header Filter Implementation

```

// Java
package org.apache.camel.component.cxf.soap.headers;

import java.util.Arrays;
import java.util.List;

import org.apache.camel.component.cxf.common.header.MessageHeaderFilter;
import org.apache.camel.spi.HeaderFilterStrategy.Direction;
import org.apache.cxf.headers.Header;

public class CustomHeaderFilter implements MessageHeaderFilter {

    public static final String ACTIVATION_NAMESPACE =
"http://cxf.apache.org/bindings/custom";
    public static final List<String> ACTIVATION_NAMESPACES =
Arrays.asList(ACTIVATION_NAMESPACE);

    public List<String> getActivationNamespaces() {
        return ACTIVATION_NAMESPACES;
    }
}

```

```

    public void filter(Direction direction, List<Header> headers) {
    }
}

```

39.4. INSTALLING FILTERS

Overview

To install message header filters, set the **messageHeaderFilters** property of the **CxfHeaderFilterStrategy** object. When you initialize this property with a list of message header filters, the header filter strategy combines the specified filters to make a filter map.

The **messageHeaderFilters** property is of type, **List<MessageHeaderFilter>**.

Installing filters in XML

The following example shows how to create a **CxfHeaderFilterStrategy** instance, specifying a customized list of header filters in the **messageHeaderFilters** property. There are two header filters in this example: **SoapMessageHeaderFilter** and **CustomHeaderFilter**.

```

<bean id="customMessageFilterStrategy"
class="org.apache.camel.component.cxf.common.header.CxfHeaderFilterStrategy">
  <property name="messageHeaderFilters">
    <list>
      <!-- SoapMessageHeaderFilter is the built in filter. It can
be removed by omitting it. -->
      <bean
class="org.apache.camel.component.cxf.common.header.SoapMessageHeaderFilter"/>

      <!-- Add custom filter here -->
      <bean
class="org.apache.camel.component.cxf.soap.headers.CustomHeaderFilter"/>
    </list>
  </property>
  <!-- The 'relayHeaders' property is 'true' by default -->
</bean>

```

SoapMessageHeaderFilter

The first header filter in the preceding example is the **SoapMessageHeaderFilter** filter, which is the default header filter. This filter is designed to filter standard SOAP headers and is bound to the following XML namespaces:

```

http://schemas.xmlsoap.org/soap/
http://schemas.xmlsoap.org/wsdl/soap/
http://schemas.xmlsoap.org/wsdl/soap12/

```

This filter peeks at the header element, in order to decide whether or not to block a particular header. If the **soap:actor** attribute (SOAP 1.1) or the **soap:role** attribute (SOAP

1.2) is present and has the value *next*, the header is removed from the message. Otherwise, the header is propagated.

Namespace clashes

Normally, each namespace should be bound to just a single header filter. If a namespace is bound to more than one header filter, this normally causes an error. It is possible, however, to override this policy by setting the **allowFilterNamespaceClash** property to **true** in the **CxfHeaderFilterStrategy** instance. When this policy is set to **true**, the nearest to last filter is selected, in the event of a namespace clash.

PART IV. PROGRAMMING EIP COMPONENTS

Abstract

This guide describes how to use the Apache Camel API.

CHAPTER 40. UNDERSTANDING MESSAGE FORMATS

Abstract

Before you can begin programming with Apache Camel, you should have a clear understanding of how messages and message exchanges are modelled. Because Apache Camel can process many message formats, the basic message type is designed to have an abstract format. Apache Camel provides the APIs needed to access and transform the data formats that underly message bodies and message headers.

40.1. EXCHANGES

Overview

An *exchange object* is a wrapper that encapsulates a received message and stores its associated metadata (including the *exchange properties*). In addition, if the current message is dispatched to a producer endpoint, the exchange provides a temporary slot to hold the reply (the *Out* message).

An important feature of exchanges in Apache Camel is that they support lazy creation of messages. This can provide a significant optimization in the case of routes that do not require explicit access to messages.

Figure 40.1. Exchange Object Passing through a Route




Figure 40.1, “Exchange Object Passing through a Route” shows an exchange object passing through a route. In the context of a route, an exchange object gets passed as the argument of the `Processor.process()` method. This means that the exchange object is directly accessible to the source endpoint, the target endpoint, and all of the processors in between.

The Exchange interface

The `org.apache.camel.Exchange` interface defines methods to access *In* and *Out* messages, as shown in [Example 40.1, “Exchange Methods”](#).

Example 40.1. Exchange Methods

```
// Access the In message
Message getIn();
void    setIn(Message in);
```

```
// Access the Out message (if any)
Message getOut();
void    setOut(Message out);
boolean hasOut();

// Access the exchange ID
String  getExchangeId();
void    setExchangeId(String id);
```

For a complete description of the methods in the **Exchange** interface, see [Section 49.1, “The Exchange Interface”](#).

Lazy creation of messages

Apache Camel supports lazy creation of *In*, *Out*, and *Fault* messages. This means that message instances are not created until you try to access them (for example, by calling **getIn()** or **getOut()**). The lazy message creation semantics are implemented by the **org.apache.camel.impl.DefaultExchange** class.

If you call one of the no-argument accessors (**getIn()** or **getOut()**), or if you call an accessor with the boolean argument equal to **true** (that is, **getIn(true)** or **getOut(true)**), the default method implementation creates a new message instance, if one does not already exist.

If you call an accessor with the boolean argument equal to **false** (that is, **getIn(false)** or **getOut(false)**), the default method implementation returns the current message value^[2]

Lazy creation of exchange IDs

Apache Camel supports lazy creation of exchange IDs. You can call **getExchangeId()** on any exchange to obtain a unique ID for that exchange instance, but the ID is generated only when you actually call the method. The **DefaultExchange.getExchangeId()** implementation of this method delegates ID generation to the UUID generator that is registered with the **CamelContext**.

For details of how to register UUID generators with the **CamelContext**, see [Section 40.4, “Built-In UUID Generators”](#).

40.2. MESSAGES

Overview

Message objects represent messages using the following abstract model:

- *Message body*
- *Message headers*
- *Message attachments*

The message body and the message headers can be of arbitrary type (they are declared as type **Object**) and the message attachments are declared to be of type **javax.activation.DataHandler**, which can contain arbitrary MIME types. If you need to

obtain a concrete representation of the message contents, you can convert the body and headers to another type using the type converter mechanism and, possibly, using the marshalling and unmarshalling mechanism.

One important feature of Apache Camel messages is that they support *lazy creation* of message bodies and headers. In some cases, this means that a message can pass through a route without needing to be parsed at all.

The Message interface

The `org.apache.camel.Message` interface defines methods to access the message body, message headers and message attachments, as shown in [Example 40.2, “Message Interface”](#).

Example 40.2. Message Interface

```
// Access the message body
Object getBody();
<T> T getBody(Class<T> type);
void setBody(Object body);
<T> void setBody(Object body, Class<T> type);

// Access message headers
Object getHeader(String name);
<T> T getHeader(String name, Class<T> type);
void setHeader(String name, Object value);
Object removeHeader(String name);
Map<String, Object> getHeaders();
void setHeaders(Map<String, Object> headers);

// Access message attachments
javax.activation.DataHandler getAttachment(String id);
java.util.Map<String, javax.activation.DataHandler> getAttachments();
java.util.Set<String> getAttachmentNames();
void addAttachment(String id, javax.activation.DataHandler content)

// Access the message ID
String getMessageId();
void setMessageId(String messageId);
```

For a complete description of the methods in the `Message` interface, see [Section 50.1, “The Message Interface”](#).

Lazy creation of bodies, headers, and attachments

Apache Camel supports lazy creation of bodies, headers, and attachments. This means that the objects that represent a message body, a message header, or a message attachment are not created until they are needed.

For example, consider the following route that accesses the `foo` message header from the `In` message:

```
from("SourceURL")
```

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

In this route, if we assume that the component referenced by *SourceURL* supports lazy creation, the *In* message headers are not actually parsed until the `header("foo")` call is executed. At that point, the underlying message implementation parses the headers and populates the header map. The message *body* is not parsed until you reach the end of the route, at the `to("TargetURL")` call. At that point, the body is converted into the format required for writing it to the target endpoint, *TargetURL*.

By waiting until the last possible moment before populating the bodies, headers, and attachments, you can ensure that unnecessary type conversions are avoided. In some cases, you can completely avoid parsing. For example, if a route contains no explicit references to message headers, a message could traverse the route without ever parsing the headers.

Whether or not lazy creation is implemented in practice depends on the underlying component implementation. In general, lazy creation is valuable for those cases where creating a message body, a message header, or a message attachment is expensive. For details about implementing a message type that supports lazy creation, see [Section 50.2, “Implementing the Message Interface”](#).

Lazy creation of message IDs

Apache Camel supports lazy creation of message IDs. That is, a message ID is generated only when you actually call the `getMessageId()` method. The `DefaultExchange.getExchangeId()` implementation of this method delegates ID generation to the UUID generator that is registered with the `CamelContext`.

Some endpoint implementations would call the `getMessageId()` method implicitly, if the endpoint implements a protocol that requires a unique message ID. In particular, JMS messages normally include a header containing unique message ID, so the JMS component automatically calls `getMessageId()` to obtain the message ID (this is controlled by the `messageIdEnabled` option on the JMS endpoint).

For details of how to register UUID generators with the `CamelContext`, see [Section 40.4, “Built-In UUID Generators”](#).

Initial message format

The initial format of an *In* message is determined by the source endpoint, and the initial format of an *Out* message is determined by the target endpoint. If lazy creation is supported by the underlying component, the message remains unparsed until it is accessed explicitly by the application. Most Apache Camel components create the message body in a relatively raw form—for example, representing it using types such as `byte[]`, `ByteBuffer`, `InputStream`, or `OutputStream`. This ensures that the overhead required for creating the initial message is minimal. Where more elaborate message formats are required components usually rely on *type converters* or *marshalling processors*.

Type converters

It does not matter what the initial format of the message is, because you can easily convert a message from one format to another using the built-in type converters (see [Section 40.3, “Built-In Type Converters”](#)). There are various methods in the Apache Camel API that

expose type conversion functionality. For example, the **convertBodyTo(Class type)** method can be inserted into a route to convert the body of an *In* message, as follows:

```
from("SourceURL").convertBodyTo(String.class).to("TargetURL");
```

Where the body of the *In* message is converted to a **java.lang.String**. The following example shows how to append a string to the end of the *In* message body:

```
from("SourceURL").setBody(bodyAs(String.class).append("My Special Signature")).to("TargetURL");
```

Where the message body is converted to a string format before appending a string to the end. It is not necessary to convert the message body explicitly in this example. You can also use:

```
from("SourceURL").setBody(body().append("My Special Signature")).to("TargetURL");
```

Where the **append()** method automatically converts the message body to a string before appending its argument.

Type conversion methods in Message

The **org.apache.camel.Message** interface exposes some methods that perform type conversion explicitly:

- **getBody(Class<T> type)**—Returns the message body as type, T.
- **getHeader(String name, Class<T> type)**—Returns the named header value as type, T.

For the complete list of supported conversion types, see [Section 40.3, “Built-In Type Converters”](#).

Converting to XML

In addition to supporting conversion between simple types (such as **byte[]**, **ByteBuffer**, **String**, and so on), the built-in type converter also supports conversion to XML formats. For example, you can convert a message body to the **org.w3c.dom.Document** type. This conversion is more expensive than the simple conversions, because it involves parsing the entire message and then creating a tree of nodes to represent the XML document structure. You can convert to the following XML document types:

- **org.w3c.dom.Document**
- **javax.xml.transform.sax.SAXSource**

XML type conversions have narrower applicability than the simpler conversions. Because not every message body conforms to an XML structure, you have to remember that this type conversion might fail. On the other hand, there are many scenarios where a router deals exclusively with XML message types.

Marshalling and unmarshalling

Marshalling involves converting a high-level format to a low-level format, and *unmarshalling* involves converting a low-level format to a high-level format. The following two processors are used to perform marshalling or unmarshalling in a route:

- `marshal()`
- `unmarshal()`

For example, to read a serialized Java object from a file and unmarshal it into a Java object, you could use the route definition shown in [Example 40.3, “Unmarshalling a Java Object”](#).

Example 40.3. Unmarshalling a Java Object

```
from("file://tmp/appfiles/serialized")
    .unmarshal()
    .serialization()
    .<FurtherProcessing>
    .to("TargetURL");
```

Final message format

When an *In* message reaches the end of a route, the target endpoint must be able to convert the message body into a format that can be written to the physical endpoint. The same rule applies to *Out* messages that arrive back at the source endpoint. This conversion is usually performed implicitly, using the Apache Camel type converter. Typically, this involves converting from a low-level format to another low-level format, such as converting from a `byte[]` array to an `InputStream` type.

40.3. BUILT-IN TYPE CONVERTERS

Overview

This section describes the conversions supported by the master type converter. These conversions are built into the Apache Camel core.

Usually, the type converter is called through convenience functions, such as `Message.getBody(Class<T> type)` or `Message.getHeader(String name, Class<T> type)`. It is also possible to invoke the master type converter directly. For example, if you have an exchange object, `exchange`, you could convert a given value to a `String` as shown in [Example 40.4, “Converting a Value to a String”](#).

Example 40.4. Converting a Value to a String

```
org.apache.camel.TypeConverter tc =
exchange.getContext().getTypeConverter();
String str_value = tc.convertTo(String.class, value);
```

Basic type converters

Apache Camel provides built-in type converters that perform conversions to and from the following basic types:

- `java.io.File`
- `String`
- `byte[]` and `java.nio.ByteBuffer`
- `java.io.InputStream` and `java.io.OutputStream`
- `java.io.Reader` and `java.io.Writer`
- `java.io.BufferedReader` and `java.io.BufferedWriter`
- `java.io.StringReader`

However, not all of these types are inter-convertible. The built-in converter is mainly focused on providing conversions from the **File** and **String** types. The **File** type can be converted to any of the preceding types, except **Reader**, **Writer**, and **StringReader**. The **String** type can be converted to **File**, **byte[]**, **ByteBuffer**, **InputStream**, or **StringReader**. The conversion from **String** to **File** works by interpreting the string as a file name. The trio of **String**, **byte[]**, and **ByteBuffer** are completely inter-convertible.



NOTE

You can explicitly specify which character encoding to use for conversion from **byte[]** to **String** and from **String** to **byte[]** by setting the **Exchange.CHARSET_NAME** exchange property in the current exchange. For example, to perform conversions using the UTF-8 character encoding, call `exchange.setProperty("Exchange.CHARSET_NAME", "UTF-8")`. The supported character sets are described in the [java.nio.charset.Charset](#) class.

Collection type converters

Apache Camel provides built-in type converters that perform conversions to and from the following collection types:

- `Object[]`
- `java.util.Set`
- `java.util.List`

All permutations of conversions between the preceding collection types are supported.

Map type converters

Apache Camel provides built-in type converters that perform conversions to and from the following map types:

- `java.util.Map`
- `java.util.HashMap`
- `java.util.Hashtable`
- `java.util.Properties`

The preceding map types can also be converted into a set, of `java.util.Set` type, where the set elements are of the `MapEntry<K,V>` type.

DOM type converters

You can perform type conversions to the following Document Object Model (DOM) types:

- `org.w3c.dom.Document`—convertible from `byte[]`, `String`, `java.io.File`, and `java.io.InputStream`.
- `org.w3c.dom.Node`
- `javax.xml.transform.dom.DOMSource`—convertible from `String`.
- `javax.xml.transform.Source`—convertible from `byte[]` and `String`.

All permutations of conversions between the preceding DOM types are supported.

SAX type converters

You can also perform conversions to the `javax.xml.transform.sax.SAXSource` type, which supports the SAX event-driven XML parser (see the [SAX Web site](#) for details). You can convert to `SAXSource` from the following types:

- `String`
- `InputStream`
- `Source`
- `StreamSource`
- `DOMSource`

Custom type converters

Apache Camel also enables you to implement your own custom type converters. For details on how to implement a custom type converter, see [Chapter 42, Type Converters](#).

40.4. BUILT-IN UUID GENERATORS

Overview

Apache Camel enables you to register a UUID generator in the `CamelContext`. This UUID generator is then used whenever Apache Camel needs to generate a unique ID—in particular, the registered UUID generator is called to generate the IDs returned by the `Exchange.getExchangeId()` and the `Message.getMessageId()` methods.

For example, you might prefer to replace the default UUID generator, if part of your application does not support IDs with a length of 36 characters (like Websphere MQ). Also, it can be convenient to generate IDs using a simple counter (see the `SimpleUuidGenerator`) for testing purposes.

Provided UUID generators

You can configure Apache Camel to use one of the following UUID generators, which are provided in the core:

- **org.apache.camel.impl.ActiveMQUuidGenerator**—(*Default*) generates the same style of ID as is used by Apache ActiveMQ. This implementation might not be suitable for all applications, because it uses some JDK APIs that are forbidden in the context of cloud computing (such as the Google App Engine).
- **org.apache.camel.impl.SimpleUuidGenerator**—implements a simple counter ID, starting at **1**. The underlying implementation uses the **java.util.concurrent.atomic.AtomicLong** type, so that it is thread-safe.
- **org.apache.camel.impl.JavaUuidGenerator**—implements an ID based on the **java.util.UUID** type. Because **java.util.UUID** is synchronized, this might affect performance on some highly concurrent systems.

Custom UUID generator

To implement a custom UUID generator, implement the **org.apache.camel.spi.UuidGenerator** interface, which is shown in [Example 40.5](#), “[UuidGenerator Interface](#)”. The **generateUuid()** must be implemented to return a unique ID string.

Example 40.5. UuidGenerator Interface


```
// Java
package org.apache.camel.spi;

/**
 * Generator to generate UUID strings.
 */
public interface UuidGenerator {
    String generateUuid();
}
```

Specifying the UUID generator using Java

To replace the default UUID generator using Java, call the **setUuidGenerator()** method on the current **CamelContext** object. For example, you can register a **SimpleUuidGenerator** instance with the current **CamelContext**, as follows:

```
// Java
getContext().setUuidGenerator(new
org.apache.camel.impl.SimpleUuidGenerator());
```



NOTE

The **setUuidGenerator()** method should be called during startup, *before* any routes are activated.

Specifying the UUID generator using Spring

To replace the default UUID generator using Spring, all you need to do is to create an instance of a UUID generator using the Spring **bean** element. When a **camelContext** instance is created, it automatically looks up the Spring registry, searching for a bean that implements **org.apache.camel.spi.UuidGenerator**. For example, you can register a **SimpleUuidGenerator** instance with the **CamelContext** as follows:

```
<beans ...>
  <bean id="simpleUuidGenerator"
        class="org.apache.camel.impl.SimpleUuidGenerator" />

  <camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
    ...
  </camelContext>
  ...
</beans>
```

[2] If there is no active method the returned value will be **null**.

CHAPTER 41. IMPLEMENTING A PROCESSOR

Abstract

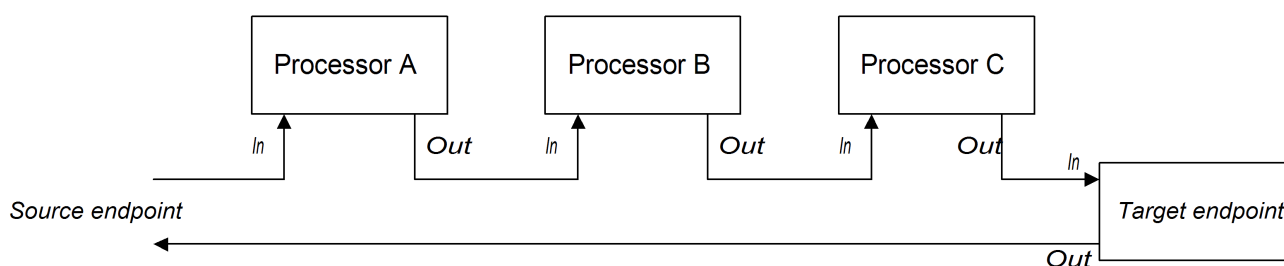
Apache Camel allows you to implement a custom processor. You can then insert the custom processor into a route to perform operations on exchange objects as they pass through the route.

41.1. PROCESSING MODEL

Pipelining model

The *pipelining model* describes the way in which processors are arranged in [Section 4.4, “Pipes and Filters”](#). Pipelining is the most common way to process a sequence of endpoints (a producer endpoint is just a special type of processor). When the processors are arranged in this way, the exchange's *In* and *Out* messages are processed as shown in [Figure 41.1, “Pipelining Model”](#).

Figure 41.1. Pipelining Model



The processors in the pipeline look like services, where the *In* message is analogous to a request, and the *Out* message is analogous to a reply. In fact, in a realistic pipeline, the nodes in the pipeline are often implemented by Web service endpoints, such as the CXF component.

For example, [Example 41.1, “Java DSL Pipeline”](#) shows a Java DSL pipeline constructed from a sequence of two processors, **ProcessorA**, **ProcessorB**, and a producer endpoint, *TargetURI*.

Example 41.1. Java DSL Pipeline

```
from(SourceURI).pipeline(ProcessorA, ProcessorB, TargetURI);
```

41.2. IMPLEMENTING A SIMPLE PROCESSOR

Overview

This section describes how to implement a simple processor that executes message processing logic before delegating the exchange to the next processor in the route.

Processor interface

Simple processors are created by implementing the `org.apache.camel.Processor` interface. As shown in [Example 41.2, “Processor Interface”](#), the interface defines a single method, `process()`, which processes an exchange object.

Example 41.2. Processor Interface

```
package org.apache.camel;

public interface Processor {
    void process(Exchange exchange) throws Exception;
}
```

Implementing the Processor interface

To create a simple processor you must implement the `Processor` interface and provide the logic for the `process()` method. [Example 41.3, “Simple Processor Implementation”](#) shows the outline of a simple processor implementation.

Example 41.3. Simple Processor Implementation

```
import org.apache.camel.Processor;

public class MyProcessor implements Processor {
    public MyProcessor() { }

    public void process(Exchange exchange) throws Exception
    {
        // Insert code that gets executed *before* delegating
        // to the next processor in the chain.
        ...
    }
}
```

All of the code in the `process()` method gets executed *before* the exchange object is delegated to the next processor in the chain.

For examples of how to access the message body and header values inside a simple processor, see [Section 41.3, “Accessing Message Content”](#).

Inserting the simple processor into a route

Use the `process()` DSL command to insert a simple processor into a route. Create an instance of your custom processor and then pass this instance as an argument to the `process()` method, as follows:

```
org.apache.camel.Processor myProc = new MyProcessor();

from("SourceURL").process(myProc).to("TargetURL");
```

41.3. ACCESSING MESSAGE CONTENT

Accessing message headers

Message headers typically contain the most useful message content from the perspective of a router, because headers are often intended to be processed in a router service. To access header data, you must first get the message from the exchange object (for example, using `Exchange.getIn()`), and then use the `Message` interface to retrieve the individual headers (for example, using `Message.getHeader()`).

[Example 41.4, “Accessing an Authorization Header”](#) shows an example of a custom processor that accesses the value of a header named `Authorization`. This example uses the `ExchangeHelper.getMandatoryHeader()` method, which eliminates the need to test for a null header value.

Example 41.4. Accessing an Authorization Header

```
import org.apache.camel.*;
import org.apache.camel.util.ExchangeHelper;

public class MyProcessor implements Processor {
    public void process(Exchange exchange) {
        String auth = ExchangeHelper.getMandatoryHeader(
            exchange,
            "Authorization",
            String.class
        );
        // process the authorization string...
        // ...
    }
}
```

For full details of the `Message` interface, see [Section 40.2, “Messages”](#).

Accessing the message body

You can also access the message body. For example, to append a string to the end of the `In` message, you can use the processor shown in [Example 41.5, “Accessing the Message Body”](#).

Example 41.5. Accessing the Message Body

```
import org.apache.camel.*;
import org.apache.camel.util.ExchangeHelper;

public class MyProcessor implements Processor {
    public void process(Exchange exchange) {
        Message in = exchange.getIn();
        in.setBody(in.getBody(String.class) + " World!");
    }
}
```

Accessing message attachments

You can access a message's attachments using either the `Message.getAttachment()` method or the `Message.getAttachments()` method. See [Example 40.2, "Message Interface"](#) for more details.

41.4. THE EXCHANGEHELPER CLASS

Overview

The `org.apache.camel.util.ExchangeHelper` class is a Apache Camel utility class that provides methods that are useful when implementing a processor.

Resolve an endpoint

The static `resolveEndpoint()` method is one of the most useful methods in the `ExchangeHelper` class. You use it inside a processor to create new `Endpoint` instances on the fly.

Example 41.6. The `resolveEndpoint()` Method

```
public final class ExchangeHelper {
    ...
    @SuppressWarnings({"unchecked" })
    public static Endpoint
    resolveEndpoint(Exchange exchange, Object value)
        throws NoSuchEndpointException { ... }
    ...
}
```

The first argument to `resolveEndpoint()` is an exchange instance, and the second argument is usually an endpoint URI string. [Example 41.7, "Creating a File Endpoint"](#) shows how to create a new file endpoint from an exchange instance `exchange`

Example 41.7. Creating a File Endpoint

```
Endpoint file_endp = ExchangeHelper.resolveEndpoint(exchange,
"file://tmp/messages/in.xml");
```

Wrapping the exchange accessors

The `ExchangeHelper` class provides several static methods of the form `getMandatoryBeanProperty()`, which wrap the corresponding `getBeanProperty()` methods on the `Exchange` class. The difference between them is that the original `getBeanProperty()` accessors return `null`, if the corresponding property is unavailable, and the `getMandatoryBeanProperty()` wrapper methods throw a Java exception. The following wrapper methods are implemented in the `ExchangeHelper` class:

```
public final class ExchangeHelper {
    ...
    public static <T> T getMandatoryProperty(Exchange exchange, String
propertyName, Class<T> type)
```

```

        throws NoSuchPropertyException { ... }

    public static <T> T getMandatoryHeader(Exchange exchange, String
propertyName, Class<T> type)
        throws NoSuchHeaderException { ... }

    public static Object getMandatoryInBody(Exchange exchange)
        throws InvalidPayloadException { ... }

    public static <T> T getMandatoryInBody(Exchange exchange, Class<T>
type)
        throws InvalidPayloadException { ... }

    public static Object getMandatoryOutBody(Exchange exchange)
        throws InvalidPayloadException { ... }

    public static <T> T getMandatoryOutBody(Exchange exchange, Class<T>
type)
        throws InvalidPayloadException { ... }
    ...
}

```

Testing the exchange pattern

Several different exchange patterns are compatible with holding an *In* message. Several different exchange patterns are also compatible with holding an *Out* message. To provide a quick way of checking whether or not an exchange object is capable of holding an *In* message or an *Out* message, the **ExchangeHelper** class provides the following methods:

```

public final class ExchangeHelper {
    ...
    public static boolean isInCapable(Exchange exchange) { ... }

    public static boolean isOutCapable(Exchange exchange) { ... }
    ...
}

```

Get the *In* message's MIME content type

If you want to find out the MIME content type of the exchange's *In* message, you can access it by calling the **ExchangeHelper.getContentTypes(exchange)** method. To implement this, the **ExchangeHelper** object looks up the value of the *In* message's **Content-Type** header—this method relies on the underlying component to populate the header value).

CHAPTER 42. TYPE CONVERTERS

Abstract

Apache Camel has a built-in type conversion mechanism, which is used to convert message bodies and message headers to different types. This chapter explains how to extend the type conversion mechanism by adding your own custom converter methods.

42.1. TYPE CONVERTER ARCHITECTURE

Overview

This section describes the overall architecture of the type converter mechanism, which you must understand, if you want to write custom type converters. If you only need to use the built-in type converters, see [Chapter 40, Understanding Message Formats](#).

Type converter interface

[Example 42.1, “TypeConverter Interface”](#) shows the definition of the `org.apache.camel.TypeConverter` interface, which all type converters must implement.

Example 42.1. TypeConverter Interface

```
package org.apache.camel;

public interface TypeConverter {
    <T> T convertTo(Class<T> type, Object value);
}
```

Master type converter

The Apache Camel type converter mechanism follows a master/slave pattern. There are many *slave* type converters, which are each capable of performing a limited number of type conversions, and a single *master* type converter, which aggregates the type conversions performed by the slaves. The master type converter acts as a front-end for the slave type converters. When you request the master to perform a type conversion, it selects the appropriate slave and delegates the conversion task to that slave.

For users of the type conversion mechanism, the master type converter is the most important because it provides the entry point for accessing the conversion mechanism. During start up, Apache Camel automatically associates a master type converter instance with the `CamelContext` object. To obtain a reference to the master type converter, you call the `CamelContext.getTypeConverter()` method. For example, if you have an exchange object, `exchange`, you can obtain a reference to the master type converter as shown in [Example 42.2, “Getting a Master Type Converter”](#).

Example 42.2. Getting a Master Type Converter

```
org.apache.camel.TypeConverter tc =
exchange.getContext().getTypeConverter();
```

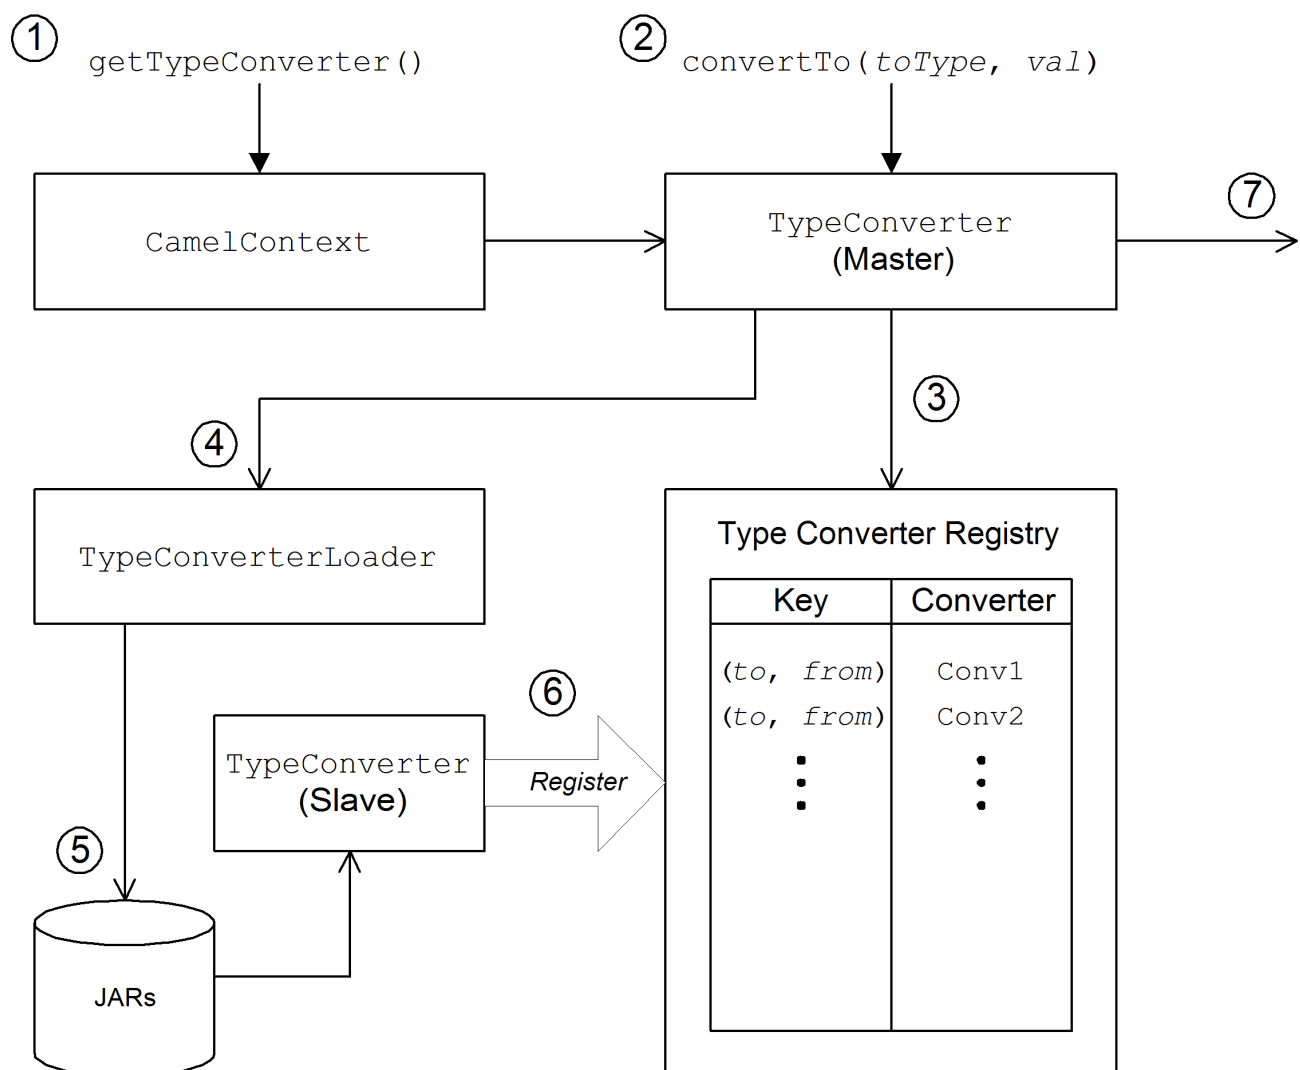

Type converter loader

The master type converter uses a *type converter loader* to populate the registry of slave type converters. A type converter loader is any class that implements the **TypeConverterLoader** interface. Apache Camel currently uses only one kind of type converter loader—the *annotation type converter loader* (of **AnnotationTypeConverterLoader** type).

Type conversion process

Figure 42.1, “Type Conversion Process” gives an overview of the type conversion process, showing the steps involved in converting a given data value, **value**, to a specified type, **toType**.

Figure 42.1. Type Conversion Process



The type conversion mechanism proceeds as follows:

1. The **CamelContext** object holds a reference to the master **TypeConverter** instance. The first step in the conversion process is to retrieve the master type converter by calling **CamelContext.getTypeConverter()**.

2. Type conversion is initiated by calling the **convertTo()** method on the master type converter. This method instructs the type converter to convert the data object, **value**, from its original type to the type specified by the **toType** argument.
3. Because the master type converter is a front end for many different slave type converters, it looks up the appropriate slave type converter by checking a registry of type mappings. The registry of type converters is keyed by a type mapping pair (**toType**, **fromType**). If a suitable type converter is found in the registry, the master type converter calls the slave's **convertTo()** method and returns the result.
4. If a suitable type converter *cannot* be found in the registry, the master type converter loads a new type converter, using the type converter loader.
5. The type converter loader searches the available JAR libraries on the classpath to find a suitable type converter. Currently, the loader strategy that is used is implemented by the annotation type converter loader, which attempts to load a class annotated by the **org.apache.camel.Converter** annotation. See [the section called "Create a TypeConverter file"](#).
6. If the type converter loader is successful, a new slave type converter is loaded and entered into the type converter registry. This type converter is then used to convert the **value** argument to the **toType** type.
7. If the data is successfully converted, the converted data value is returned. If the conversion does not succeed, **null** is returned.

42.2. IMPLEMENTING TYPE CONVERTER USING ANNOTATIONS

Overview

The type conversion mechanism can easily be customized by adding a new slave type converter. This section describes how to implement a slave type converter and how to integrate it with Apache Camel, so that it is automatically loaded by the annotation type converter loader.

How to implement a type converter

To implement a custom type converter, perform the following steps:

1. [Implement an annotated converter class](#)
2. [Create a TypeConverter file.](#)
3. [Package the type converter.](#)

Implement an annotated converter class

You can implement a custom type converter class using the **@Converter** annotation. You must annotate the class itself and each of the **static** methods intended to perform type conversion. Each converter method takes an argument that defines the *from* type, optionally takes a second **Exchange** argument, and has a non-void return value that defines the *to* type. The type converter loader uses Java reflection to find the annotated methods and integrate them into the type converter mechanism. [Example 42.3, "Example of an Annotated Converter Class"](#) shows an example of an annotated converter class that defines

a converter method for converting from `java.io.File` to `java.io.InputStream` and another converter method (with an **Exchange** argument) for converting from `byte[]` to **String**.

Example 42.3. Example of an Annotated Converter Class

```
package com.YourDomain.YourPackageName;

import org.apache.camel.Converter;

import java.io.*;

@Converter
public class IOConverter {
    private IOConverter() {
    }

    @Converter
    public static InputStream toInputStream(File file) throws
    FileNotFoundException {
        return new BufferedInputStream(new FileInputStream(file));
    }

    @Converter
    public static String toString(byte[] data, Exchange exchange) {
        if (exchange != null) {
            String charsetName =
            exchange.getProperty(Exchange.CHARSET_NAME, String.class);
            if (charsetName != null) {
                try {
                    return new String(data, charsetName);
                } catch (UnsupportedEncodingException e) {
                    LOG.warn("Can't convert the byte to String with the
                    charset " + charsetName, e);
                }
            }
        }
        return new String(data);
    }
}
```

The `toInputStream()` method is responsible for performing the conversion from the **File** type to the **InputStream** type and the `toString()` method is responsible for performing the conversion from the `byte[]` type to the **String** type.



NOTE

The method name is unimportant, and can be anything you choose. What is important are the argument type, the return type, and the presence of the `@Converter` annotation.

Create a TypeConverter file

To enable the discovery mechanism (which is implemented by the *annotation type*

converter loader) for your custom converter, create a **TypeConverter** file at the following location:

```
META-INF/services/org/apache/camel/TypeConverter
```

The **TypeConverter** file must contain a comma-separated list of package names identifying the packages that contain type converter classes. For example, if you want the type converter loader to search the **com.YourDomain.YourPackageName** package for annotated converter classes, the **TypeConverter** file would have the following contents:

```
com.YourDomain.YourPackageName
```

Package the type converter

The type converter is packaged as a JAR file containing the compiled classes of your custom type converters and the **META-INF** directory. Put this JAR file on your classpath to make it available to your Apache Camel application.

Fallback converter method

In addition to defining regular converter methods using the **@Converter** annotation, you can optionally define a fallback converter method using the **@FallbackConverter** annotation. The fallback converter method will only be tried, if the master type converter fails to find a regular converter method in the type registry.

The essential difference between a regular converter method and a fallback converter method is that whereas a regular converter is defined to perform conversion between a specific pair of types (for example, from **byte[]** to **String**), a fallback converter can potentially perform conversion between *any* pair of types. It is up to the code in the body of the fallback converter method to figure out which conversions it is able to perform. At run time, if a conversion cannot be performed by a regular converter, the master type converter iterates through every available fallback converter until it finds one that can perform the conversion.

The method signature of a fallback converter can have either of the following forms:

```
// 1. Non-generic form of signature
@FallbackConverter
public static Object MethodName(
    Class type,
    Exchange exchange,
    Object value,
    TypeConverterRegistry registry
)

// 2. Templating form of signature
@FallbackConverter
public static <T> T MethodName(
    Class<T> type,
    Exchange exchange,
    Object value,
    TypeConverterRegistry registry
)
```

Where *MethodName* is an arbitrary method name for the fallback converter.

For example, the following code extract (taken from the implementation of the File component) shows a fallback converter that can convert the body of a **GenericFile** object, exploiting the type converters already available in the type converter registry:

```
package org.apache.camel.component.file;

import org.apache.camel.Converter;
import org.apache.camel.FallbackConverter;
import org.apache.camel.Exchange;
import org.apache.camel.TypeConverter;
import org.apache.camel.spi.TypeConverterRegistry;

@Converter
public final class GenericFileConverter {

    private GenericFileConverter() {
        // Helper Class
    }

    @FallbackConverter
    public static <T> T convertTo(Class<T> type, Exchange exchange, Object
value, TypeConverterRegistry registry) {
        // use a fallback type converter so we can convert the embedded
body if the value is GenericFile
        if (GenericFile.class.isAssignableFrom(value.getClass())) {
            GenericFile file = (GenericFile) value;
            Class from = file.getBody().getClass();
            TypeConverter tc = registry.lookup(type, from);
            if (tc != null) {
                Object body = file.getBody();
                return tc.convertTo(type, exchange, body);
            }
        }

        return null;
    }
    ...
}
```

42.3. IMPLEMENTING A TYPE CONVERTER DIRECTLY

Overview

Generally, the recommended way to implement a type converter is to use an annotated class, as described in the previous section, [Section 42.2, “Implementing Type Converter Using Annotations”](#). But if you want to have complete control over the registration of your type converter, you can implement a custom slave type converter and add it directly to the type converter registry, as described here.

Implement the `TypeConverter` interface

To implement your own type converter class, define a class that implements the

TypeConverter interface. For example, the following **MyOrderTypeConverter** class converts an integer value to a **MyOrder** object, where the integer value is used to initialize the order ID in the **MyOrder** object.

```
import org.apache.camel.TypeConverter

private class MyOrderTypeConverter implements TypeConverter {

    public <T> T convertTo(Class<T> type, Object value) {
        // converter from value to the MyOrder bean
        MyOrder order = new MyOrder();
        order.setId(Integer.parseInt(value.toString()));
        return (T) order;
    }

    public <T> T convertTo(Class<T> type, Exchange exchange, Object value)
{
    // this method with the Exchange parameter will be preferred by
    Camel to invoke
    // this allows you to fetch information from the exchange during
    conversions
    // such as an encoding parameter or the likes
    return convertTo(type, value);
}

    public <T> T mandatoryConvertTo(Class<T> type, Object value) {
        return convertTo(type, value);
    }

    public <T> T mandatoryConvertTo(Class<T> type, Exchange exchange,
    Object value) {
        return convertTo(type, value);
    }
}
```

Add the type converter to the registry

You can add the custom type converter *directly* to the type converter registry using code like the following:

```
// Add the custom type converter to the type converter registry
context.getTypeConverterRegistry().addTypeConverter(MyOrder.class,
String.class, new MyOrderTypeConverter());
```

Where **context** is the current **org.apache.camel.CamelContext** instance. The **addTypeConverter()** method registers the **MyOrderTypeConverter** class against the specific type conversion, from **String.class** to **MyOrder.class**.

CHAPTER 43. PRODUCER AND CONSUMER TEMPLATES

Abstract

The producer and consumer templates in Apache Camel are modelled after a feature of the Spring container API, whereby access to a resource is provided through a simplified, easy-to-use API known as a *template*. In the case of Apache Camel, the producer template and consumer template provide simplified interfaces for sending messages to and receiving messages from producer endpoints and consumer endpoints.

43.1. USING THE PRODUCER TEMPLATE

43.1.1. Introduction to the Producer Template

Overview

The producer template supports a variety of different approaches to invoking producer endpoints. There are methods that support different formats for the request message (as an **Exchange** object, as a message body, as a message body with a single header setting, and so on) and there are methods to support both the synchronous and the asynchronous style of invocation. Overall, producer template methods can be grouped into the following categories:

- [the section called “Synchronous invocation”](#).
- [the section called “Synchronous invocation with a processor”](#).
- [the section called “Asynchronous invocation”](#).
- [the section called “Asynchronous invocation with a callback”](#).

Synchronous invocation

The methods for invoking endpoints synchronously have names of the form **sendSuffix()** and **requestSuffix()**. For example, the methods for invoking an endpoint using either the default message exchange pattern (MEP) or an explicitly specified MEP are named **send()**, **sendBody()**, and **sendBodyAndHeader()** (where these methods respectively send an **Exchange** object, a message body, or a message body and header value). If you want to force the MEP to be *InOut* (request/reply semantics), you can call the **request()**, **requestBody()**, and **requestBodyAndHeader()** methods instead.

The following example shows how to create a **ProducerTemplate** instance and use it to send a message body to the **activemq:MyQueue** endpoint. The example also shows how to send a message body and header value using **sendBodyAndHeader()**.

```
import org.apache.camel.ProducerTemplate
import org.apache.camel.impl.DefaultProducerTemplate
...
ProducerTemplate template = context.createProducerTemplate();

// Send to a specific queue
```

```

template.sendBody("activemq:MyQueue", "<hello>world!</hello>");

// Send with a body and header
template.sendBodyAndHeader(
    "activemq:MyQueue",
    "<hello>world!</hello>",
    "CustomerRating", "Gold" );

```

Synchronous invocation with a processor

A special case of synchronous invocation is where you provide the `send()` method with a **Processor** argument instead of an **Exchange** argument. In this case, the producer template implicitly asks the specified endpoint to create an **Exchange** instance (typically, but not always having the *InOnly* MEP by default). This default exchange is then passed to the processor, which initializes the contents of the exchange object.

The following example shows how to send an exchange initialized by the **MyProcessor** processor to the `activemq:MyQueue` endpoint.

```

import org.apache.camel.ProducerTemplate
import org.apache.camel.impl.DefaultProducerTemplate
...
ProducerTemplate template = context.createProducerTemplate();

// Send to a specific queue, using a processor to initialize
template.send("activemq:MyQueue", new MyProcessor());

```

The **MyProcessor** class is implemented as shown in the following example. In addition to setting the *In* message body (as shown here), you could also initialize message headers and exchange properties.

```

import org.apache.camel.Processor;
import org.apache.camel.Exchange;
...
public class MyProcessor implements Processor {
    public MyProcessor() { }

    public void process(Exchange ex) {
        ex.getIn().setBody("<hello>world!</hello>");
    }
}

```

Asynchronous invocation

The methods for invoking endpoints *asynchronously* have names of the form **asyncSendSuffix()** and **asyncRequestSuffix()**. For example, the methods for invoking an endpoint using either the default message exchange pattern (MEP) or an explicitly specified MEP are named **asyncSend()** and **asyncSendBody()** (where these methods respectively send an **Exchange** object or a message body). If you want to force the MEP to be *InOut* (request/reply semantics), you can call the **asyncRequestBody()**, **asyncRequestBodyAndHeader()**, and **asyncRequestBodyAndHeaders()** methods instead.

The following example shows how to send an exchange asynchronously to the `direct:start` endpoint. The **asyncSend()** method returns a

`java.util.concurrent.Future` object, which is used to retrieve the invocation result at a later time.

```
import java.util.concurrent.Future;

import org.apache.camel.Exchange;
import org.apache.camel.impl.DefaultExchange;
...
Exchange exchange = new DefaultExchange(context);
exchange.getIn().setBody("Hello");

Future<Exchange> future = template.asyncSend("direct:start", exchange);

// You can do other things, whilst waiting for the invocation to complete
...
// Now, retrieve the resulting exchange from the Future
Exchange result = future.get();
```

The producer template also provides methods to send a message body asynchronously (for example, using `asyncSendBody()` or `asyncRequestBody()`). In this case, you can use one of the following helper methods to extract the returned message body from the **Future** object:

```
<T> T extractFutureBody(Future future, Class<T> type);
<T> T extractFutureBody(Future future, long timeout, TimeUnit unit,
Class<T> type) throws TimeoutException;
```

The first version of the `extractFutureBody()` method blocks until the invocation completes and the reply message is available. The second version of the `extractFutureBody()` method allows you to specify a timeout. Both methods have a type argument, **type**, which casts the returned message body to the specified type using a built-in type converter.

The following example shows how to use the `asyncRequestBody()` method to send a message body to the `direct:start` endpoint. The blocking `extractFutureBody()` method is then used to retrieve the reply message body from the **Future** object.

```
Future<Object> future = template.asyncRequestBody("direct:start",
"Hello");

// You can do other things, whilst waiting for the invocation to complete
...
// Now, retrieve the reply message body as a String type
String result = template.extractFutureBody(future, String.class);
```

Asynchronous invocation with a callback

In the preceding asynchronous examples, the request message is dispatched in a sub-thread, while the reply is retrieved and processed by the main thread. The producer template also gives you the option, however, of processing replies in the sub-thread, using one of the `asyncCallback()`, `asyncCallbackSendBody()`, or `asyncCallbackRequestBody()` methods. In this case, you supply a callback object (of `org.apache.camel.impl.SynchronizationAdapter` type), which automatically gets invoked in the sub-thread as soon as a reply message arrives.

The **Synchronization** callback interface is defined as follows:

```
package org.apache.camel.spi;

import org.apache.camel.Exchange;

public interface Synchronization {
    void onComplete(Exchange exchange);
    void onFailure(Exchange exchange);
}
```

Where the **onComplete()** method is called on receipt of a normal reply and the **onFailure()** method is called on receipt of a fault message reply. Only one of these methods gets called back, so you must override both of them to ensure that all types of reply are processed.

The following example shows how to send an exchange to the **direct:start** endpoint, where the reply message is processed in the sub-thread by the **SynchronizationAdapter** callback object.

```
import java.util.concurrent.Future;
import java.util.concurrent.TimeUnit;

import org.apache.camel.Exchange;
import org.apache.camel.impl.DefaultExchange;
import org.apache.camel.impl.SynchronizationAdapter;
...
Exchange exchange = context.getEndpoint("direct:start").createExchange();
exchange.getIn().setBody("Hello");

Future<Exchange> future = template.asyncCallback("direct:start", exchange,
new SynchronizationAdapter() {
    @Override
    public void onComplete(Exchange exchange) {
        assertEquals("Hello World", exchange.getIn().getBody());
    }
});
```

Where the **SynchronizationAdapter** class is a default implementation of the **Synchronization** interface, which you can override to provide your own definitions of the **onComplete()** and **onFailure()** callback methods.

You still have the option of accessing the reply from the main thread, because the **asyncCallback()** method also returns a **Future** object—for example:

```
// Retrieve the reply from the main thread, specifying a timeout
Exchange reply = future.get(10, TimeUnit.SECONDS);
```

43.1.2. Synchronous Send

Overview

The *synchronous send* methods are a collection of methods that you can use to invoke a producer endpoint, where the current thread blocks until the method invocation is

complete and the reply (if any) has been received. These methods are compatible with any kind of message exchange protocol.

Send an exchange

The basic `send()` method is a general-purpose method that sends the contents of an **Exchange** object to an endpoint, using the message exchange pattern (MEP) of the exchange. The return value is the exchange that you get after it has been processed by the producer endpoint (possibly containing an *Out* message, depending on the MEP).

There are three varieties of `send()` method for sending an exchange that let you specify the target endpoint in one of the following ways: as the default endpoint, as an endpoint URI, or as an **Endpoint** object.

```
Exchange send(Exchange exchange);
Exchange send(String endpointUri, Exchange exchange);
Exchange send(Endpoint endpoint, Exchange exchange);
```

Send an exchange populated by a processor

A simple variation of the general `send()` method is to use a processor to populate a default exchange, instead of supplying the exchange object explicitly (see [the section called “Synchronous invocation with a processor”](#) for details).

The `send()` methods for sending an exchange populated by a processor let you specify the target endpoint in one of the following ways: as the default endpoint, as an endpoint URI, or as an **Endpoint** object. In addition, you can optionally specify the exchange's MEP by supplying the **pattern** argument, instead of accepting the default.

```
Exchange send(Processor processor);
Exchange send(String endpointUri, Processor processor);
Exchange send(Endpoint endpoint, Processor processor);
Exchange send(
    String endpointUri,
    ExchangePattern pattern,
    Processor processor
);
Exchange send(
    Endpoint endpoint,
    ExchangePattern pattern,
    Processor processor
);
```

Send a message body

If you are only concerned with the contents of the message body that you want to send, you can use the `sendBody()` methods to provide the message body as an argument and let the producer template take care of inserting the body into a default exchange object.

The `sendBody()` methods let you specify the target endpoint in one of the following ways: as the default endpoint, as an endpoint URI, or as an **Endpoint** object. In addition, you can optionally specify the exchange's MEP by supplying the **pattern** argument, instead of accepting the default. The methods *without* a **pattern** argument return **void** (even though

the invocation might give rise to a reply in some cases); and the methods *with* a **pattern** argument return either the body of the *Out* message (if there is one) or the body of the *In* message (otherwise).

```
void sendBody(Object body);
void sendBody(String endpointUri, Object body);
void sendBody(Endpoint endpoint, Object body);
Object sendBody(
    String endpointUri,
    ExchangePattern pattern,
    Object body
);
Object sendBody(
    Endpoint endpoint,
    ExchangePattern pattern,
    Object body
);
```

Send a message body and header(s)

For testing purposes, it is often interesting to try out the effect of a *single* header setting and the **sendBodyAndHeader()** methods are useful for this kind of header testing. You supply the message body and header setting as arguments to **sendBodyAndHeader()** and let the producer template take care of inserting the body and header setting into a default exchange object.

The **sendBodyAndHeader()** methods let you specify the target endpoint in one of the following ways: as the default endpoint, as an endpoint URI, or as an **Endpoint** object. In addition, you can optionally specify the exchange's MEP by supplying the **pattern** argument, instead of accepting the default. The methods *without* a **pattern** argument return **void** (even though the invocation might give rise to a reply in some cases); and the methods *with* a **pattern** argument return either the body of the *Out* message (if there is one) or the body of the *In* message (otherwise).

```
void sendBodyAndHeader(
    Object body,
    String header,
    Object headerValue
);
void sendBodyAndHeader(
    String endpointUri,
    Object body,
    String header,
    Object headerValue
);
void sendBodyAndHeader(
    Endpoint endpoint,
    Object body,
    String header,
    Object headerValue
);
Object sendBodyAndHeader(
    String endpointUri,
    ExchangePattern pattern,
    Object body,
```

```

        String header,
        Object headerValue
    );
    Object sendBodyAndHeader(
        Endpoint endpoint,
        ExchangePattern pattern,
        Object body,
        String header,
        Object headerValue
    );

```

The **sendBodyAndHeaders()** methods are similar to the **sendBodyAndHeader()** methods, except that instead of supplying just a single header setting, these methods allow you to specify a complete hash map of header settings.

```

void sendBodyAndHeaders(
    Object body,
    Map<String, Object> headers
);
void sendBodyAndHeaders(
    String endpointUri,
    Object body,
    Map<String, Object> headers
);
void sendBodyAndHeaders(
    Endpoint endpoint,
    Object body,
    Map<String, Object> headers
);
Object sendBodyAndHeaders(
    String endpointUri,
    ExchangePattern pattern,
    Object body,
    Map<String, Object> headers
);
Object sendBodyAndHeaders(
    Endpoint endpoint,
    ExchangePattern pattern,
    Object body,
    Map<String, Object> headers
);

```

Send a message body and exchange property

You can try out the effect of setting a single exchange property using the **sendBodyAndProperty()** methods. You supply the message body and property setting as arguments to **sendBodyAndProperty()** and let the producer template take care of inserting the body and exchange property into a default exchange object.

The **sendBodyAndProperty()** methods let you specify the target endpoint in one of the following ways: as the default endpoint, as an endpoint URI, or as an **Endpoint** object. In addition, you can optionally specify the exchange's MEP by supplying the **pattern** argument, instead of accepting the default. The methods *without* a **pattern** argument

return **void** (even though the invocation might give rise to a reply in some cases); and the methods *with* a **pattern** argument return either the body of the *Out* message (if there is one) or the body of the *In* message (otherwise).

```
void sendBodyAndProperty(
    Object body,
    String property,
    Object propertyValue
);
void sendBodyAndProperty(
    String endpointUri,
    Object body,
    String property,
    Object propertyValue
);
void sendBodyAndProperty(
    Endpoint endpoint,
    Object body,
    String property,
    Object propertyValue
);
Object sendBodyAndProperty(
    String endpoint,
    ExchangePattern pattern,
    Object body,
    String property,
    Object propertyValue
);
Object sendBodyAndProperty(
    Endpoint endpoint,
    ExchangePattern pattern,
    Object body,
    String property,
    Object propertyValue
);
```

43.1.3. Synchronous Request with InOut Pattern

Overview

The *synchronous request* methods are similar to the synchronous send methods, except that the request methods force the message exchange pattern to be *InOut* (conforming to request/reply semantics). Hence, it is generally convenient to use a synchronous request method, if you expect to receive a reply from the producer endpoint.

Request an exchange populated by a processor

The basic **request()** method is a general-purpose method that uses a processor to populate a default exchange and forces the message exchange pattern to be *InOut* (so that the invocation obeys request/reply semantics). The return value is the exchange that you get after it has been processed by the producer endpoint, where the *Out* message contains the reply message.

The **request()** methods for sending an exchange populated by a processor let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint**

object.

```
Exchange request(String endpointUri, Processor processor);
Exchange request(Endpoint endpoint, Processor processor);
```

Request a message body

If you are only concerned with the contents of the message body in the request and in the reply, you can use the **requestBody()** methods to provide the request message body as an argument and let the producer template take care of inserting the body into a default exchange object.

The **requestBody()** methods let you specify the target endpoint in one of the following ways: as the default endpoint, as an endpoint URI, or as an **Endpoint** object. The return value is the body of the reply message (*Out* message body), which can either be returned as plain **Object** or converted to a specific type, **T**, using the built-in type converters (see [Section 40.3, “Built-In Type Converters”](#)).

```
Object requestBody(Object body);
<T> T requestBody(Object body, Class<T> type);
Object requestBody(
    String endpointUri,
    Object body
);
<T> T requestBody(
    String endpointUri,
    Object body,
    Class<T> type
);
Object requestBody(
    Endpoint endpoint,
    Object body
);
<T> T requestBody(
    Endpoint endpoint,
    Object body,
    Class<T> type
);
```

Request a message body and header(s)

You can try out the effect of setting a single header value using the **requestBodyAndHeader()** methods. You supply the message body and header setting as arguments to **requestBodyAndHeader()** and let the producer template take care of inserting the body and exchange property into a default exchange object.

The **requestBodyAndHeader()** methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object. The return value is the body of the reply message (*Out* message body), which can either be returned as plain **Object** or converted to a specific type, **T**, using the built-in type converters (see [Section 40.3, “Built-In Type Converters”](#)).

```
Object requestBodyAndHeader(
    String endpointUri,
```

```

    Object body,
    String header,
    Object headerValue
);
<T> T requestBodyAndHeader(
    String endpointUri,
    Object body,
    String header,
    Object headerValue,
    Class<T> type
);
Object requestBodyAndHeader(
    Endpoint endpoint,
    Object body,
    String header,
    Object headerValue
);
<T> T requestBodyAndHeader(
    Endpoint endpoint,
    Object body,
    String header,
    Object headerValue,
    Class<T> type
);

```

The **requestBodyAndHeaders()** methods are similar to the **requestBodyAndHeader()** methods, except that instead of supplying just a single header setting, these methods allow you to specify a complete hash map of header settings.

```

Object requestBodyAndHeaders(
    String endpointUri,
    Object body,
    Map<String, Object> headers
);
<T> T requestBodyAndHeaders(
    String endpointUri,
    Object body,
    Map<String, Object> headers,
    Class<T> type
);
Object requestBodyAndHeaders(
    Endpoint endpoint,
    Object body,
    Map<String, Object> headers
);
<T> T requestBodyAndHeaders(
    Endpoint endpoint,
    Object body,
    Map<String, Object> headers,
    Class<T> type
);

```

43.1.4. Asynchronous Send

Overview

The producer template provides a variety of methods for invoking a producer endpoint asynchronously, so that the main thread does not block while waiting for the invocation to complete and the reply message can be retrieved at a later time. The asynchronous send methods described in this section are compatible with any kind of message exchange protocol.

Send an exchange

The basic `asyncSend()` method takes an **Exchange** argument and invokes an endpoint asynchronously, using the message exchange pattern (MEP) of the specified exchange. The return value is a `java.util.concurrent.Future` object, which is a ticket you can use to collect the reply message at a later time—for details of how to obtain the return value from the **Future** object, see [the section called “Asynchronous invocation”](#).

The following `asyncSend()` methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object.

```
Future<Exchange> asyncSend(String endpointUri, Exchange exchange);
Future<Exchange> asyncSend(Endpoint endpoint, Exchange exchange);
```

Send an exchange populated by a processor

A simple variation of the general `asyncSend()` method is to use a processor to populate a default exchange, instead of supplying the exchange object explicitly.

The following `asyncSend()` methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object.

```
Future<Exchange> asyncSend(String endpointUri, Processor processor);
Future<Exchange> asyncSend(Endpoint endpoint, Processor processor);
```

Send a message body

If you are only concerned with the contents of the message body that you want to send, you can use the `asyncSendBody()` methods to send a message body asynchronously and let the producer template take care of inserting the body into a default exchange object.

The `asyncSendBody()` methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object.

```
Future<Object> asyncSendBody(String endpointUri, Object body);
Future<Object> asyncSendBody(Endpoint endpoint, Object body);
```

43.1.5. Asynchronous Request with InOut Pattern

Overview

The *asynchronous request* methods are similar to the asynchronous send methods, except that the request methods force the message exchange pattern to be *InOut* (conforming to request/reply semantics). Hence, it is generally convenient to use an asynchronous request method, if you expect to receive a reply from the producer endpoint.

Request a message body

If you are only concerned with the contents of the message body in the request and in the reply, you can use the **requestBody()** methods to provide the request message body as an argument and let the producer template take care of inserting the body into a default exchange object.

The **asyncRequestBody()** methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object. The return value that is retrievable from the **Future** object is the body of the reply message (*Out* message body), which can be returned either as a plain **Object** or converted to a specific type, **T**, using a built-in type converter (see [the section called “Asynchronous invocation”](#)).

```
Future<Object> asyncRequestBody(
    String endpointUri,
    Object body
);
<T> Future<T> asyncRequestBody(
    String endpointUri,
    Object body,
    Class<T> type
);
Future<Object> asyncRequestBody(
    Endpoint endpoint,
    Object body
);
<T> Future<T> asyncRequestBody(
    Endpoint endpoint,
    Object body,
    Class<T> type
);
```

Request a message body and header(s)

You can try out the effect of setting a single header value using the **asyncRequestBodyAndHeader()** methods. You supply the message body and header setting as arguments to **asyncRequestBodyAndHeader()** and let the producer template take care of inserting the body and exchange property into a default exchange object.

The **asyncRequestBodyAndHeader()** methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object. The return value that is retrievable from the **Future** object is the body of the reply message (*Out* message body), which can be returned either as a plain **Object** or converted to a specific type, **T**, using a built-in type converter (see [the section called “Asynchronous invocation”](#)).

```
Future<Object> asyncRequestBodyAndHeader(
    String endpointUri,
    Object body,
    String header,
    Object headerValue
);
<T> Future<T> asyncRequestBodyAndHeader(
    String endpointUri,
    Object body,
    String header,
```

```

        Object headerValue,
        Class<T> type
    );
    Future<Object> asyncRequestBodyAndHeader(
        Endpoint endpoint,
        Object body,
        String header,
        Object headerValue
    );
    <T> Future<T> asyncRequestBodyAndHeader(
        Endpoint endpoint,
        Object body,
        String header,
        Object headerValue,
        Class<T> type
    );

```

The **asyncRequestBodyAndHeaders()** methods are similar to the **asyncRequestBodyAndHeader()** methods, except that instead of supplying just a single header setting, these methods allow you to specify a complete hash map of header settings.

```

    Future<Object> asyncRequestBodyAndHeaders(
        String endpointUri,
        Object body,
        Map<String, Object> headers
    );
    <T> Future<T> asyncRequestBodyAndHeaders(
        String endpointUri,
        Object body,
        Map<String, Object> headers,
        Class<T> type
    );
    Future<Object> asyncRequestBodyAndHeaders(
        Endpoint endpoint,
        Object body,
        Map<String, Object> headers
    );
    <T> Future<T> asyncRequestBodyAndHeaders(
        Endpoint endpoint,
        Object body,
        Map<String, Object> headers,
        Class<T> type
    );

```

43.1.6. Asynchronous Send with Callback

Overview

The producer template also provides the option of processing the reply message in the same sub-thread that is used to invoke the producer endpoint. In this case, you provide a callback object, which automatically gets invoked in the sub-thread as soon as the reply message is received. In other words, the *asynchronous send with callback* methods enable

you to initiate an invocation in your main thread and then have all of the associated processing—invocation of the producer endpoint, waiting for a reply and processing the reply—occur asynchronously in a sub-thread.

Send an exchange

The basic `asyncCallback()` method takes an **Exchange** argument and invokes an endpoint asynchronously, using the message exchange pattern (MEP) of the specified exchange. This method is similar to the `asyncSend()` method for exchanges, except that it takes an additional `org.apache.camel.spi.Synchronization` argument, which is a callback interface with two methods: `onComplete()` and `onFailure()`. For details of how to use the **Synchronization** callback, see [the section called “Asynchronous invocation with a callback”](#).

The following `asyncCallback()` methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object.

```
Future<Exchange> asyncCallback(
    String endpointUri,
    Exchange exchange,
    Synchronization onCompletion
);
Future<Exchange> asyncCallback(
    Endpoint endpoint,
    Exchange exchange,
    Synchronization onCompletion
);
```

Send an exchange populated by a processor

The `asyncCallback()` method for processors calls a processor to populate a default exchange and forces the message exchange pattern to be *InOut* (so that the invocation obeys request/reply semantics).

The following `asyncCallback()` methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object.

```
Future<Exchange> asyncCallback(
    String endpointUri,
    Processor processor,
    Synchronization onCompletion
);
Future<Exchange> asyncCallback(
    Endpoint endpoint,
    Processor processor,
    Synchronization onCompletion
);
```

Send a message body

If you are only concerned with the contents of the message body that you want to send, you can use the `asyncCallbackSendBody()` methods to send a message body asynchronously and let the producer template take care of inserting the body into a default exchange object.

The **asyncCallbackSendBody()** methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object.

```
Future<Object> asyncCallbackSendBody(
    String endpointUri,
    Object body,
    Synchronization onCompletion
);
Future<Object> asyncCallbackSendBody(
    Endpoint endpoint,
    Object body,
    Synchronization onCompletion
);
```

Request a message body

If you are only concerned with the contents of the message body in the request and in the reply, you can use the **asyncCallbackRequestBody()** methods to provide the request message body as an argument and let the producer template take care of inserting the body into a default exchange object.

The **asyncCallbackRequestBody()** methods let you specify the target endpoint in one of the following ways: as an endpoint URI, or as an **Endpoint** object.

```
Future<Object> asyncCallbackRequestBody(
    String endpointUri,
    Object body,
    Synchronization onCompletion
);
Future<Object> asyncCallbackRequestBody(
    Endpoint endpoint,
    Object body,
    Synchronization onCompletion
);
```

43.2. USING THE CONSUMER TEMPLATE

Overview

The consumer template provides methods for polling a consumer endpoint in order to receive incoming messages. You can choose to receive the incoming message either in the form of an exchange object or in the form of a message body (where the message body can be cast to a particular type using a built-in type converter).

Example of polling exchanges

You can use a consumer template to poll a consumer endpoint for exchanges using one of the following polling methods: blocking **receive()**; **receive()** with a timeout; or **receiveNoWait()**, which returns immediately. Because a consumer endpoint represents a service, it is also essential to start the service thread by calling **start()** before you attempt to poll for exchanges.

The following example shows how to poll an exchange from the **seda:foo** consumer endpoint using the blocking **receive()** method:

```
import org.apache.camel.ProducerTemplate;
import org.apache.camel.ConsumerTemplate;
import org.apache.camel.Exchange;
...
ProducerTemplate template = context.createProducerTemplate();
ConsumerTemplate consumer = context.createConsumerTemplate();

// Start the consumer service
consumer.start();
...
template.sendBody("seda:foo", "Hello");
Exchange out = consumer.receive("seda:foo");
...
// Stop the consumer service
consumer.stop();
```

Where the consumer template instance, **consumer**, is instantiated using the **CamelContext.createConsumerTemplate()** method and the consumer service thread is started by calling **ConsumerTemplate.start()**.

Example of polling message bodies

You can also poll a consumer endpoint for incoming message bodies using one of the following methods: blocking **receiveBody()**; **receiveBody()** with a timeout; or **receiveBodyNowait()**, which returns immediately. As in the previous example, it is also essential to start the service thread by calling **start()** before you attempt to poll for exchanges.

The following example shows how to poll an incoming message body from the **seda:foo** consumer endpoint using the blocking **receiveBody()** method:

```
import org.apache.camel.ProducerTemplate;
import org.apache.camel.ConsumerTemplate;
...
ProducerTemplate template = context.createProducerTemplate();
ConsumerTemplate consumer = context.createConsumerTemplate();

// Start the consumer service
consumer.start();
...
template.sendBody("seda:foo", "Hello");
Object body = consumer.receiveBody("seda:foo");
...
// Stop the consumer service
consumer.stop();
```

Methods for polling exchanges

There are three basic methods for polling *exchanges* from a consumer endpoint: **receive()** without a timeout blocks indefinitely; **receive()** with a timeout blocks for the specified period of milliseconds; and **receiveNowait()** is non-blocking. You can specify the consumer endpoint either as an endpoint URI or as an **Endpoint** instance.

```

Exchange receive(String endpointUri);
Exchange receive(String endpointUri, long timeout);
Exchange receiveNoWait(String endpointUri);

Exchange receive(Endpoint endpoint);
Exchange receive(Endpoint endpoint, long timeout);
Exchange receiveNoWait(Endpoint endpoint);

```

Methods for polling message bodies

There are three basic methods for polling *message bodies* from a consumer endpoint: **receiveBody()** without a timeout blocks indefinitely; **receiveBody()** with a timeout blocks for the specified period of milliseconds; and **receiveBodyNoWait()** is non-blocking. You can specify the consumer endpoint either as an endpoint URI or as an **Endpoint** instance. Moreover, by calling the templating forms of these methods, you can convert the returned body to a particular type, **T**, using a built-in type converter.

```

Object receiveBody(String endpointUri);
Object receiveBody(String endpointUri, long timeout);
Object receiveBodyNoWait(String endpointUri);

Object receiveBody(Endpoint endpoint);
Object receiveBody(Endpoint endpoint, long timeout);
Object receiveBodyNoWait(Endpoint endpoint);

<T> T receiveBody(String endpointUri, Class<T> type);
<T> T receiveBody(String endpointUri, long timeout, Class<T> type);
<T> T receiveBodyNoWait(String endpointUri, Class<T> type);

<T> T receiveBody(Endpoint endpoint, Class<T> type);
<T> T receiveBody(Endpoint endpoint, long timeout, Class<T> type);
<T> T receiveBodyNoWait(Endpoint endpoint, Class<T> type);

```

CHAPTER 44. IMPLEMENTING A COMPONENT

Abstract

This chapter provides a general overview of the approaches can be used to implement a Apache Camel component.

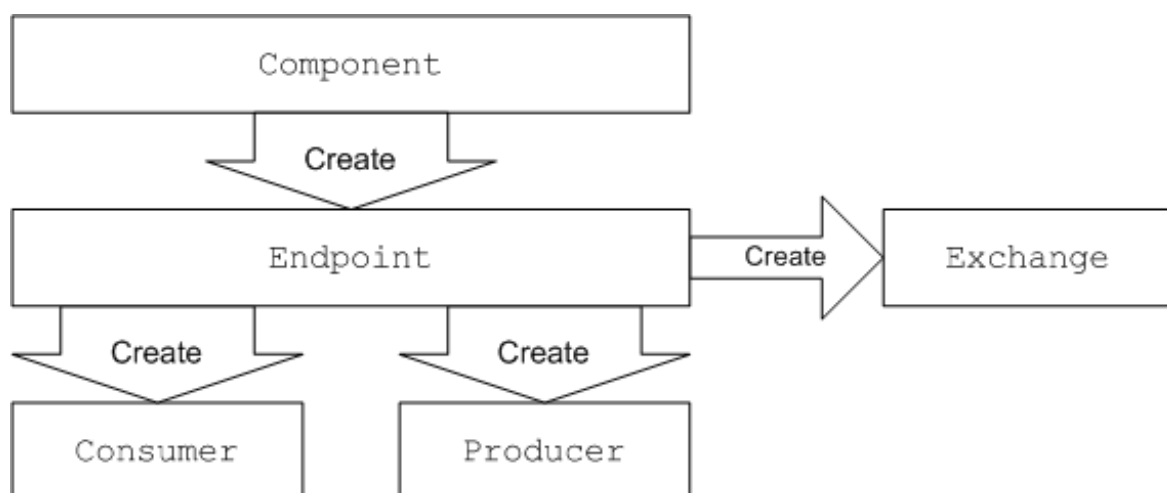
44.1. COMPONENT ARCHITECTURE

44.1.1. Factory Patterns for a Component

Overview

A Apache Camel component consists of a set of classes that are related to each other through a factory pattern. The primary entry point to a component is the **Component** object itself (an instance of `org.apache.camel.Component` type). You can use the **Component** object as a factory to create **Endpoint** objects, which in turn act as factories for creating **Consumer**, **Producer**, and **Exchange** objects. These relationships are summarized in [Figure 44.1, “Component Factory Patterns”](#)

Figure 44.1. Component Factory Patterns



Component

A component implementation is an endpoint factory. The main task of a component implementor is to implement the `Component.createEndpoint()` method, which is responsible for creating new endpoints on demand.

Each kind of component must be associated with a *component prefix* that appears in an endpoint URI. For example, the file component is usually associated with the file prefix, which can be used in an endpoint URI like `file://tmp/messages/input`. When you install a new component in Apache Camel, you must define the association between a particular component prefix and the name of the class that implements the component.

Endpoint

Each endpoint instance encapsulates a particular endpoint URI. Every time Apache Camel encounters a new endpoint URI, it creates a new endpoint instance. An endpoint object is also a factory for creating consumer endpoints and producer endpoints.

Endpoints must implement the **org.apache.camel.Endpoint** interface. The **Endpoint** interface defines the following factory methods:

- **createConsumer()** and **createPollingConsumer()**—Creates a consumer endpoint, which represents the source endpoint at the beginning of a route.
- **createProducer()**—Creates a producer endpoint, which represents the target endpoint at the end of a route.
- **createExchange()**—Creates an exchange object, which encapsulates the messages passed up and down the route.

Consumer

Consumer endpoints *consume* requests. They always appear at the start of a route and they encapsulate the code responsible for receiving incoming requests and dispatching outgoing replies. From a service-oriented perspective a consumer represents a *service*.

Consumers must implement the **org.apache.camel.Consumer** interface. There are a number of different patterns you can follow when implementing a consumer. These patterns are described in [Section 44.1.3, “Consumer Patterns and Threading”](#).

Producer

Producer endpoints *produce* requests. They always appears at the end of a route and they encapsulate the code responsible for dispatching outgoing requests and receiving incoming replies. From a service-oriented perspective a producer represents a *service consumer*.

Producers must implement the **org.apache.camel.Producer** interface. You can optionally implement the producer to support an asynchronous style of processing. See [Section 44.1.4, “Asynchronous Processing”](#) for details.

Exchange

Exchange objects encapsulate a related set of messages. For example, one kind of message exchange is a synchronous invocation, which consists of a request message and its related reply.

Exchanges must implement the **org.apache.camel.Exchange** interface. The default implementation, **DefaultExchange**, is sufficient for many component implementations. However, if you want to associated extra data with the exchanges or have the exchanges preform additional processing, it can be useful to customize the exchange implementation.

Message

There are two different message slots in an **Exchange** object:

- *In* message—holds the current message.
- *Out* message—temporarily holds a reply message.


All of the message types are represented by the same Java object, `org.apache.camel.Message`. It is not always necessary to customize the message implementation—the default implementation, `DefaultMessage`, is usually adequate.

44.1.2. Using a Component in a Route

Overview

A Apache Camel route is essentially a pipeline of processors, of `org.apache.camel.Processor` type. Messages are encapsulated in an exchange object, `E`, which gets passed from node to node by invoking the `process()` method. The architecture of the processor pipeline is illustrated in [Figure 44.2, “Consumer and Producer Instances in a Route”](#).

Figure 44.2. Consumer and Producer Instances in a Route



Source endpoint

At the start of the route, you have the source endpoint, which is represented by an `org.apache.camel.Consumer` object. The source endpoint is responsible for accepting incoming request messages and dispatching replies. When constructing the route, Apache Camel creates the appropriate `Consumer` type based on the component prefix from the endpoint URI, as described in [Section 44.1.1, “Factory Patterns for a Component”](#).

Processors

Each intermediate node in the pipeline is represented by a processor object (implementing the `org.apache.camel.Processor` interface). You can insert either standard processors (for example, `filter`, `throttler`, or `delayer`) or insert your own custom processor implementations.

Target endpoint

At the end of the route is the target endpoint, which is represented by an `org.apache.camel.Producer` object. Because it comes at the end of a processor pipeline, the producer is also a processor object (implementing the `org.apache.camel.Processor` interface). The target endpoint is responsible for sending outgoing request messages and receiving incoming replies. When constructing the route, Apache Camel creates the appropriate `Producer` type based on the component prefix from the endpoint URI.

44.1.3. Consumer Patterns and Threading

Overview

The pattern used to implement the consumer determines the threading model used in processing the incoming exchanges. Consumers can be implemented using one of the following patterns:


- **Event-driven pattern**—The consumer is driven by an external thread.
- **Scheduled poll pattern**—The consumer is driven by a dedicated thread pool.
- **Polling pattern**—The threading model is left undefined.

Event-driven pattern

In the event-driven pattern, the processing of an incoming request is initiated when another part of the application (typically a third-party library) calls a method implemented by the consumer. A good example of an event-driven consumer is the Apache Camel JMX component, where events are initiated by the JMX library. The JMX library calls the **handleNotification()** method to initiate request processing—see [Example 47.4, “JMXConsumer Implementation”](#) for details.

[Figure 44.3, “Event-Driven Consumer”](#) shows an outline of the event-driven consumer pattern. In this example, it is assumed that processing is triggered by a call to the **notify()** method.

Figure 44.3. Event-Driven Consumer



The event-driven consumer processes incoming requests as follows:

1. The consumer must implement a method to receive the incoming event (in [Figure 44.3, “Event-Driven Consumer”](#) this is represented by the **notify()** method). The thread that calls **notify()** is normally a separate part of the application, so the consumer's threading policy is externally driven.

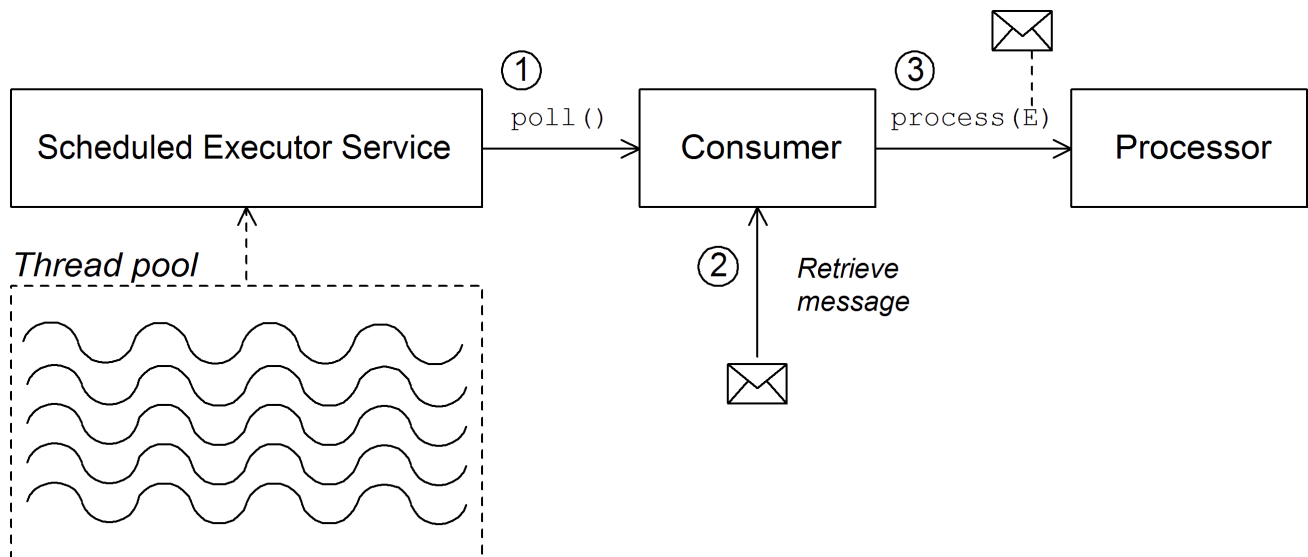
For example, in the case of the JMX consumer implementation, the consumer implements the **NotificationListener.handleNotification()** method to receive notifications from JMX. The threads that drive the consumer processing are created within the JMX layer.

2. In the body of the **notify()** method, the consumer first converts the incoming event into an exchange object, **E**, and then calls **process()** on the next processor in the route, passing the exchange object as its argument.

Scheduled poll pattern

In the scheduled poll pattern, the consumer retrieves incoming requests by checking at regular time intervals whether or not a request has arrived. Checking for requests is scheduled automatically by a built-in timer class, the *scheduled executor service*, which is a standard pattern provided by the `java.util.concurrent` library. The scheduled executor service executes a particular task at timed intervals and it also manages a pool of threads, which are used to run the task instances.

[Figure 44.4, “Scheduled Poll Consumer”](#) shows an outline of the scheduled poll consumer pattern.

Figure 44.4. Scheduled Poll Consumer

The scheduled poll consumer processes incoming requests as follows:

1. The scheduled executor service has a pool of threads at its disposal, that can be used to initiate consumer processing. After each scheduled time interval has elapsed, the scheduled executor service attempts to grab a free thread from its pool (there are five threads in the pool by default). If a free thread is available, it uses that thread to call the `poll()` method on the consumer.
2. The consumer's `poll()` method is intended to trigger processing of an incoming request. In the body of the `poll()` method, the consumer attempts to retrieve an incoming message. If no request is available, the `poll()` method returns immediately.
3. If a request message is available, the consumer inserts it into an exchange object and then calls `process()` on the next processor in the route, passing the exchange object as its argument.

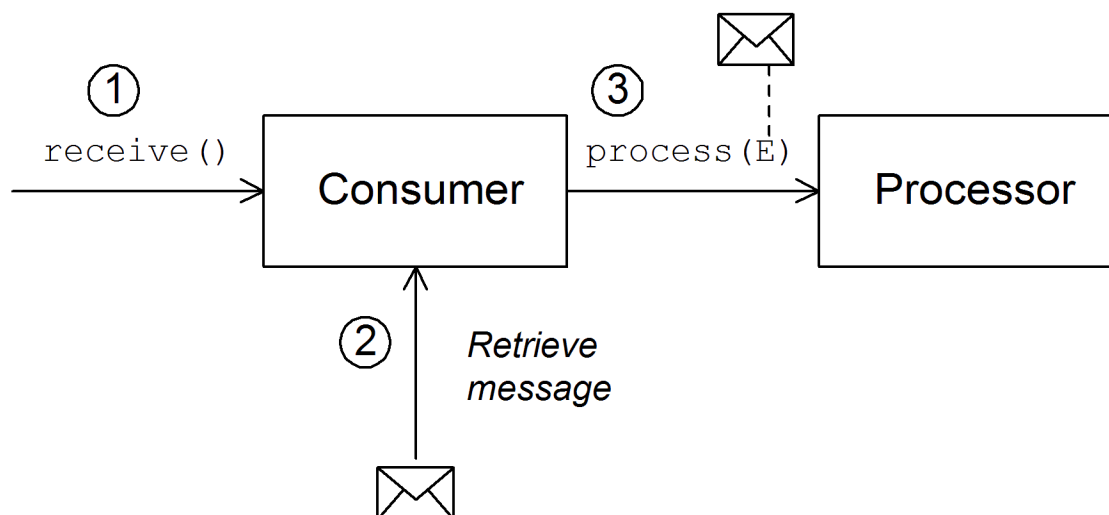
Polling pattern

In the polling pattern, processing of an incoming request is initiated when a third-party calls one of the consumer's polling methods:

- `receive()`
- `receiveNoWait()`
- `receive(long timeout)`

It is up to the component implementation to define the precise mechanism for initiating calls on the polling methods. This mechanism is not specified by the polling pattern.

Figure 44.5, “Polling Consumer” shows an outline of the polling consumer pattern.

Figure 44.5. Polling Consumer

The polling consumer processes incoming requests as follows:

1. Processing of an incoming request is initiated whenever one of the consumer's polling methods is called. The mechanism for calling these polling methods is implementation defined.
2. In the body of the **receive()** method, the consumer attempts to retrieve an incoming request message. If no message is currently available, the behavior depends on which receive method was called.
 - **receiveNowait()** returns immediately
 - **receive(long timeout)** waits for the specified timeout interval^[3] before returning
 - **receive()** waits until a message is received
3. If a request message is available, the consumer inserts it into an exchange object and then calls **process()** on the next processor in the route, passing the exchange object as its argument.

44.1.4. Asynchronous Processing

Overview

Producer endpoints normally follow a *synchronous* pattern when processing an exchange. When the preceding processor in a pipeline calls **process()** on a producer, the **process()** method blocks until a reply is received. In this case, the processor's thread remains blocked until the producer has completed the cycle of sending the request and receiving the reply.

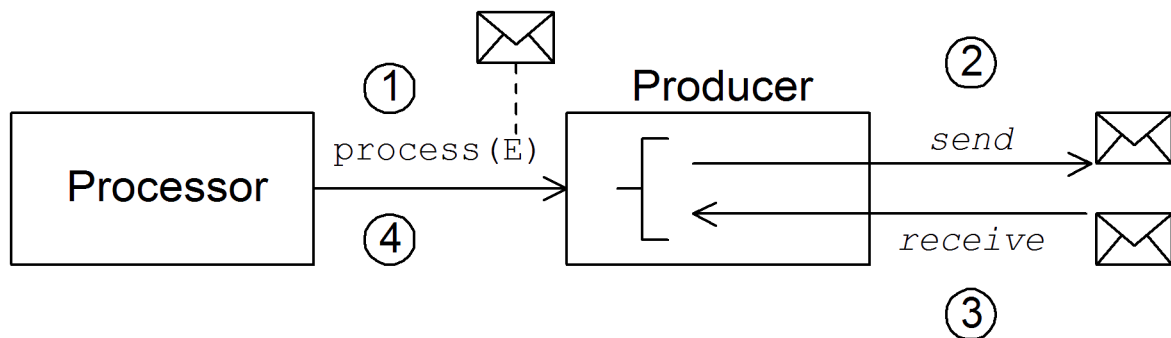
Sometimes, however, you might prefer to decouple the preceding processor from the producer, so that the processor's thread is released immediately and the **process()** call does *not* block. In this case, you should implement the producer using an *asynchronous* pattern, which gives the preceding processor the option of invoking a non-blocking version of the **process()** method.

To give you an overview of the different implementation options, this section describes both the synchronous and the asynchronous patterns for implementing a producer endpoint.

Synchronous producer

Figure 44.6, “Synchronous Producer” shows an outline of a synchronous producer, where the preceding processor blocks until the producer has finished processing the exchange.

Figure 44.6. Synchronous Producer



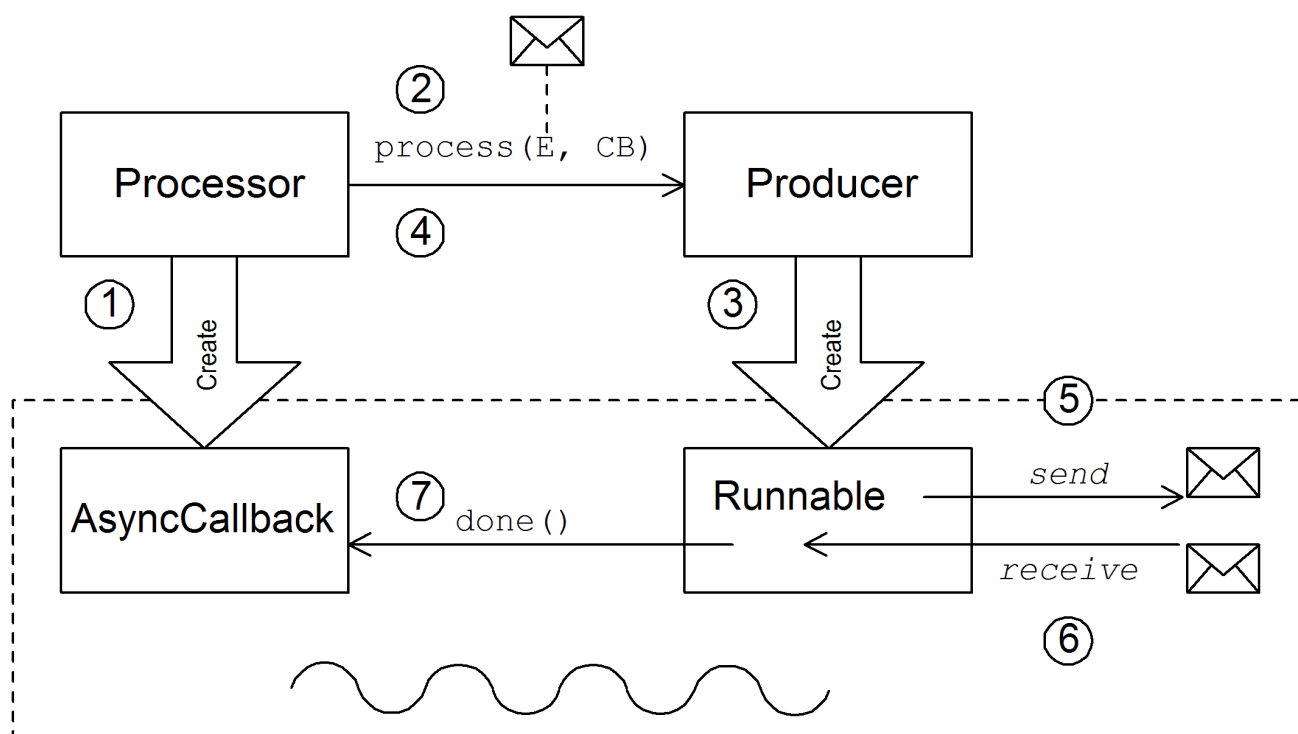
The synchronous producer processes an exchange as follows:

1. The preceding processor in the pipeline calls the synchronous `process()` method on the producer to initiate synchronous processing. The synchronous `process()` method takes a single exchange argument.
2. In the body of the `process()` method, the producer sends the request (`n` message) to the endpoint.
3. If required by the exchange pattern, the producer waits for the reply (`Out` message) to arrive from the endpoint. This step can cause the `process()` method to block indefinitely. However, if the exchange pattern does not mandate a reply, the `process()` method can return immediately after sending the request.
4. When the `process()` method returns, the exchange object contains the reply from the synchronous call (an `Out` message).

Asynchronous producer

Figure 44.7, “Asynchronous Producer” shows an outline of an asynchronous producer, where the producer processes the exchange in a sub-thread, and the preceding processor is not blocked for any significant length of time.

Figure 44.7. Asynchronous Producer



The asynchronous producer processes an exchange as follows:

1. Before the processor can call the asynchronous **process()** method, it must create an *asynchronous callback* object, which is responsible for processing the exchange on the return portion of the route. For the asynchronous callback, the processor must implement a class that inherits from the **AsyncCallback** interface.
2. The processor calls the asynchronous **process()** method on the producer to initiate asynchronous processing. The asynchronous **process()** method takes two arguments:
 - an exchange object
 - a synchronous callback object
3. In the body of the **process()** method, the producer creates a **Runnable** object that encapsulates the processing code. The producer then delegates the execution of this **Runnable** object to a sub-thread.
4. The asynchronous **process()** method returns, thereby freeing up the processor's thread. The exchange processing continues in a separate sub-thread.
5. The **Runnable** object sends the *In* message to the endpoint.
6. If required by the exchange pattern, the **Runnable** object waits for the reply (*Out* or *Fault* message) to arrive from the endpoint. The **Runnable** object remains blocked until the reply is received.
7. After the reply arrives, the **Runnable** object inserts the reply (*Out* message) into the exchange object and then calls **done()** on the asynchronous callback object. The asynchronous callback is then responsible for processing the reply message (executed in the sub-thread).

44.2. HOW TO IMPLEMENT A COMPONENT

Overview

This section gives a brief overview of the steps required to implement a custom Apache Camel component.

Which interfaces do you need to implement?

When implementing a component, it is usually necessary to implement the following Java interfaces:

- `org.apache.camel.Component`
- `org.apache.camel.Endpoint`
- `org.apache.camel.Consumer`
- `org.apache.camel.Producer`

In addition, it can also be necessary to implement the following Java interfaces:

- `org.apache.camel.Exchange`
- `org.apache.camel.Message`

Implementation steps

You typically implement a custom component as follows:

1. *Implement the **Component** interface*—A component object acts as an endpoint factory. You extend the **DefaultComponent** class and implement the **createEndpoint()** method.

See [Chapter 45, Component Interface](#).

2. *Implement the **Endpoint** interface*—An endpoint represents a resource identified by a specific URI. The approach taken when implementing an endpoint depends on whether the consumers follow an *event-driven* pattern, a *scheduled poll* pattern, or a *polling* pattern.

For an event-driven pattern, implement the endpoint by extending the **DefaultEndpoint** class and implementing the following methods:

- `createProducer()`
- `createConsumer()`

For a scheduled poll pattern, implement the endpoint by extending the **ScheduledPollEndpoint** class and implementing the following methods:

- `createProducer()`
- `createConsumer()`

For a polling pattern, implement the endpoint by extending the **DefaultPollingEndpoint** class and implementing the following methods:

- `createProducer()`
- `createPollConsumer()`

See [Chapter 46, Endpoint Interface](#).

3. *Implement the **Consumer** interface*—There are several different approaches you can take to implementing a consumer, depending on which pattern you need to implement (event-driven, scheduled poll, or polling). The consumer implementation is also crucially important for determining the threading model used for processing a message exchange.

See [Section 47.2, “Implementing the Consumer Interface”](#).

4. *Implement the **Producer** interface*—To implement a producer, you extend the **DefaultProducer** class and implement the `process()` method.

See [Chapter 48, Producer Interface](#).

5. *Optionally implement the **Exchange** or the **Message** interface*—The default implementations of **Exchange** and **Message** can be used directly, but occasionally, you might find it necessary to customize these types.

See [Chapter 49, Exchange Interface](#) and [Chapter 50, Message Interface](#).

Installing and configuring the component

You can install a custom component in one of the following ways:

- *Add the component directly to the CamelContext*—The `CamelContext.addComponent()` method adds a component programatically.
- *Add the component using Spring configuration*—The standard Spring **bean** element creates a component instance. The bean's `id` attribute implicitly defines the component prefix. For details, see [Section 44.3.2, “Configuring a Component”](#).
- *Configure Apache Camel to auto-discover the component*—Auto-discovery, ensures that Apache Camel automatically loads the component on demand. For details, see [Section 44.3.1, “Setting Up Auto-Discovery”](#).

44.3. AUTO-DISCOVERY AND CONFIGURATION

44.3.1. Setting Up Auto-Discovery

Overview

Auto-discovery is a mechanism that enables you to dynamically add components to your Apache Camel application. The component URI prefix is used as a key to load components on demand. For example, if Apache Camel encounters the endpoint URI, `activemq://MyQName`, and the ActiveMQ endpoint is not yet loaded, Apache Camel searches for the component identified by the `activemq` prefix and dynamically loads the component.

Availability of component classes

Before configuring auto-discovery, you must ensure that your custom component classes are accessible from your current classpath. Typically, you bundle the custom component classes into a JAR file, and add the JAR file to your classpath.

Configuring auto-discovery

To enable auto-discovery of your component, create a Java properties file named after the component prefix, *component-prefix*, and store that file in the following location:

```
/META-INF/services/org/apache/camel/component/component-prefix
```

The *component-prefix* properties file must contain the following property setting:

```
class=component-class-name
```

Where *component-class-name* is the fully-qualified name of your custom component class. You can also define additional system property settings in this file.

Example

For example, you can enable auto-discovery for the Apache Camel FTP component by creating the following Java properties file:

```
/META-INF/services/org/apache/camel/component/ftp
```

Which contains the following Java property setting:

```
class=org.apache.camel.component.file.remote.RemoteFileComponent
```



NOTE

The Java properties file for the FTP component is already defined in the JAR file, **camel-ftp-Version.jar**.

44.3.2. Configuring a Component

Overview

You can add a component by configuring it in the Apache Camel Spring configuration file, **META-INF/spring/camel-context.xml**. To find the component, the component's URI prefix is matched against the ID attribute of a **bean** element in the Spring configuration. If the component prefix matches a bean element ID, Apache Camel instantiates the referenced class and injects the properties specified in the Spring configuration.



NOTE

This mechanism has priority over auto-discovery. If the CamelContext finds a Spring bean with the requisite ID, it will not attempt to find the component using auto-discovery.

Define bean properties on your component class

If there are any properties that you want to inject into your component class, define them as bean properties. For example:

```
public class CustomComponent extends
    DefaultComponent<CustomExchange> {
    ...
    PropType getProperty() { ... }
    void setProperty(PropType v) { ... }
}
```

The `getProperty()` method and the `setProperty()` method access the value of *property*.

Configure the component in Spring

To configure a component in Spring, edit the configuration file, `META-INF/spring/camel-context.xml`, as shown in [Example 44.1, “Configuring a Component in Spring”](#).

Example 44.1. Configuring a Component in Spring

```
<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xsi:schemaLocation="
         http://www.springframework.org/schema/beans
         http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
         http://camel.apache.org/schema/spring
         http://camel.apache.org/schema/spring/camel-spring.xsd">

  <camelContext id="camel"
    xmlns="http://camel.apache.org/schema/spring">
    <package>RouteBuilderPackage</package>
  </camelContext>

  <bean id="component-prefix" class="component-class-name">
    <property name="property" value="propertyValue"/>
  </bean>
</beans>
```

The `bean` element with ID `component-prefix` configures the `component-class-name` component. You can inject properties into the component instance using `property` elements. For example, the `property` element in the preceding example would inject the value, `propertyValue`, into the `property` property by calling `setProperty()` on the component.

Examples

[Example 44.2, “JMS Component Spring Configuration”](#) shows an example of how to configure the Apache Camel's JMS component by defining a bean element with ID equal to `jms`. These settings are added to the Spring configuration file `camel-context.xml`.

Example 44.2. JMS Component Spring Configuration

```

<?xml version="1.0" encoding="UTF-8"?>
<beans xmlns="http://www.springframework.org/schema/beans"
       xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
       xsi:schemaLocation="
         http://www.springframework.org/schema/beans
         http://www.springframework.org/schema/beans/spring-beans-2.0.xsd
         http://camel.apache.org/schema/spring
         http://camel.apache.org/schema/spring/camel-spring.xsd">

  <camelContext id="camel"
    xmlns="http://camel.apache.org/schema/spring">
    ❶ <package>org.apache.camel.example.spring</package>
    </camelContext>

    <bean id="jms" class="org.apache.camel.component.jms.JmsComponent">
    ❷❸ <property name="connectionFactory">
      <bean
        class="org.apache.activemq.ActiveMQConnectionFactory">
        <property name="brokerURL"
          value="vm://localhost?
    ❹ broker.persistent=false&broker.useJmx=false"/>
      </bean>
    </property>
    </bean>
  </beans>

```

- ❶ The **CamelContext** automatically instantiates any **RouteBuilder** classes that it finds in the specified Java package, `org.apache.camel.example.spring`.
- ❷ The bean element with ID, **jms**, configures the JMS component. The bean ID corresponds to the component's URI prefix. For example, if a route specifies an endpoint with the URI, `jms://MyQName`, Apache Camel automatically loads the JMS component using the settings from the **jms** bean element.
- ❸ JMS is just a wrapper for a messaging service. You must specify the concrete implementation of the messaging system by setting the **connectionFactory** property on the **JmsComponent** class.
- ❹ In this example, the concrete implementation of the JMS messaging service is Apache ActiveMQ. The **brokerURL** property initializes a connection to an ActiveMQ broker instance, where the message broker is embedded in the local Java virtual machine (JVM). If a broker is not already present in the JVM, ActiveMQ will instantiate it with the options **broker.persistent=false** (the broker does not persist messages) and **broker.useJmx=false** (the broker does not open a JMX port).

[3] The timeout interval is typically specified in milliseconds.

CHAPTER 45. COMPONENT INTERFACE

Abstract


This chapter describes how to implement the **Component** interface.

45.1. THE COMPONENT INTERFACE

Overview

To implement a Apache Camel component, you must implement the **org.apache.camel.Component** interface. An instance of **Component** type provides the entry point into a custom component. That is, all of the other objects in a component are ultimately accessible through the **Component** instance. [Figure 45.1, “Component Inheritance Hierarchy”](#) shows the relevant Java interfaces and classes that make up the **Component** inheritance hierarchy.

Figure 45.1. Component Inheritance Hierarchy



The Component interface

[Example 45.1, “Component Interface”](#) shows the definition of the **org.apache.camel.Component** interface.

Example 45.1. Component Interface

```

package org.apache.camel;

public interface Component {
    CamelContext getCamelContext();
    void setCamelContext(CamelContext context);

    Endpoint createEndpoint(String uri) throws Exception;
}
  
```

Component methods

The **Component** interface defines the following methods:

- **getCamelContext()** and **setCamelContext()**—References the **CamelContext** to which this **Component** belongs. The **setCamelContext()** method is automatically called when you add the component to a **CamelContext**.
- **createEndpoint()**—The factory method that gets called to create **Endpoint** instances for this component. The **uri** parameter is the endpoint URI, which contains the details required to create the endpoint.

45.2. IMPLEMENTING THE COMPONENT INTERFACE

The `DefaultComponent` class

You implement a new component by extending the **org.apache.camel.impl.DefaultComponent** class, which provides some standard functionality and default implementations for some of the methods. In particular, the **DefaultComponent** class provides support for URI parsing and for creating *æcheduled executor* (which is used for the scheduled poll pattern).

URI parsing

The **createEndpoint(String uri)** method defined in the base **Component** interface takes a complete, unparsed endpoint URI as its sole argument. The **DefaultComponent** class, on the other hand, defines a three-argument version of the **createEndpoint()** method with the following signature:

```
protected abstract Endpoint createEndpoint(
    String uri,
    String remaining,
    Map parameters
)
throws Exception;
```

uri is the original, unparsed URI; **remaining** is the part of the URI that remains after stripping off the component prefix at the start and cutting off the query options at the end; and **parameters** contains the parsed query options. It is this version of the **createEndpoint()** method that you must override when inheriting from **DefaultComponent**. This has the advantage that the endpoint URI is already parsed for you.

The following sample endpoint URI for the **file** component shows how URI parsing works in practice:

```
file:///tmp/messages/foo?delete=true&moveNamePostfix=.old
```

For this URI, the following arguments are passed to the three-argument version of **createEndpoint()**:

Argument	Sample Value
uri	file:///tmp/messages/foo?delete=true&moveNamePostfix=.old
remaining	/tmp/messages/foo

Argument	Sample Value
parameters	Two entries are set in <code>java.util.Map</code> : <ul style="list-style-type: none"> parameter delete is boolean true parameter moveNamePostfix has the string value, .old.

Parameter injection

By default, the parameters extracted from the URI query options are injected on the endpoint's bean properties. The **DefaultComponent** class automatically injects the parameters for you.

For example, if you want to define a custom endpoint that supports two URI query options: **delete** and **moveNamePostfix**. All you must do is define the corresponding bean methods (getters and setters) in the endpoint class:

```
public class FileEndpoint extends ScheduledPollEndpoint {
    ...
    public boolean isDelete() {
        return delete;
    }
    public void setDelete(boolean delete) {
        this.delete = delete;
    }
    ...
    public String getMoveNamePostfix() {
        return moveNamePostfix;
    }
    public void setMoveNamePostfix(String moveNamePostfix) {
        this.moveNamePostfix = moveNamePostfix;
    }
}
```

It is also possible to inject URI query options into *consumer* parameters. For details, see [the section called “Consumer parameter injection”](#).

Disabling endpoint parameter injection

If there are no parameters defined on your **Endpoint** class, you can optimize the process of endpoint creation by disabling endpoint parameter injection. To disable parameter injection on endpoints, override the **useIntrospectionOnEndpoint()** method and implement it to return **false**, as follows:

```
protected boolean useIntrospectionOnEndpoint() {
    return false;
}
```

**NOTE**

The `useIntrospectionOnEndpoint()` method does *not* affect the parameter injection that might be performed on a **Consumer** class. Parameter injection at that level is controlled by the `Endpoint.configureProperties()` method (see [Section 46.2, “Implementing the Endpoint Interface”](#)).

Scheduled executor service

The scheduled executor is used in the scheduled poll pattern, where it is responsible for driving the periodic polling of a consumer endpoint (a scheduled executor is effectively a thread pool implementation).

To instantiate a scheduled executor service, use the **ExecutorServiceStrategy** object that is returned by the `CamelContext.getExecutorServiceStrategy()` method. For details of the Apache Camel threading model, see [Section 2.9, “Threading Model”](#).

**NOTE**

Prior to Apache Camel 2.3, the **DefaultComponent** class provided a `getExecutorService()` method for creating thread pool instances. Since 2.3, however, the creation of thread pools is now managed centrally by the **ExecutorServiceStrategy** object.

Validating the URI

If you want to validate the URI before creating an endpoint instance, you can override the `validateURI()` method from the **DefaultComponent** class, which has the following signature:

```
protected void validateURI(String uri,
                           String path,
                           Map parameters)
    throws ResolveEndpointFailedException;
```

If the supplied URI does not have the required format, the implementation of `validateURI()` should throw the `org.apache.camel.ResolveEndpointFailedException` exception.

Creating an endpoint

[Example 45.2, “Implementation of `createEndpoint\(\)`”](#) outlines how to implement the **DefaultComponent.createEndpoint()** method, which is responsible for creating endpoint instances on demand.

Example 45.2. Implementation of `createEndpoint()`

```
1 public class CustomComponent extends DefaultComponent {
    ...
    2 protected Endpoint createEndpoint(String uri, String remaining, Map
    3 parameters) throws Exception {
        CustomEndpoint result = new CustomEndpoint(uri, this);
        // ...
        return result;
    }
}
```


- 1 The *CustomComponent* is the name of your custom component class, which is defined by extending the **DefaultComponent** class.
- 2 When extending **DefaultComponent**, you must implement the **createEndpoint()** method with three arguments (see [the section called “URI parsing”](#)).
- 3 Create an instance of your custom endpoint type, *CustomEndpoint*, by calling its constructor. At a minimum, this constructor takes a copy of the original URI string, **uri**, and a reference to this component instance, **this**.

Example

[Example 45.3, “FileComponent Implementation”](#) shows a sample implementation of a **FileComponent** class.

Example 45.3. FileComponent Implementation

```
package org.apache.camel.component.file;

import org.apache.camel.CamelContext;
import org.apache.camel.Endpoint;
import org.apache.camel.impl.DefaultComponent;

import java.io.File;
import java.util.Map;

public class FileComponent extends DefaultComponent {
    public static final String HEADER_FILE_NAME =
"org.apache.camel.file.name";

    1 public FileComponent() {
    }

    2 public FileComponent(CamelContext context) {
        super(context);
    }

    3 protected Endpoint createEndpoint(String uri, String remaining, Map
parameters) throws Exception {
        File file = new File(remaining);
        FileEndpoint result = new FileEndpoint(file, uri, this);
        return result;
    }
}
```

- 1 Always define a no-argument constructor for the component class in order to facilitate automatic instantiation of the class.
- 2 A constructor that takes the parent **CamelContext** instance as an argument is convenient when creating a component instance by programming.

- 3** The implementation of the **FileComponent.createEndpoint()** method follows the pattern described in [Example 45.2, “Implementation of createEndpoint\(\)”](#). The implementation creates a **FileEndpoint** object.

CHAPTER 46. ENDPOINT INTERFACE

Abstract

This chapter describes how to implement the **Endpoint** interface, which is an essential step in the implementation of a Apache Camel component.

46.1. THE ENDPOINT INTERFACE


Overview

An instance of `org.apache.camel.Endpoint` type encapsulates an endpoint URI, and it also serves as a factory for **Consumer**, **Producer**, and **Exchange** objects. There are three different approaches to implementing an endpoint:

- Event-driven
- scheduled poll
- polling

These endpoint implementation patterns complement the corresponding patterns for implementing a consumer—see [Section 47.2, “Implementing the Consumer Interface”](#).

[Figure 46.1, “Endpoint Inheritance Hierarchy”](#) shows the relevant Java interfaces and classes that make up the **Endpoint** inheritance hierarchy.

Figure 46.1. Endpoint Inheritance Hierarchy

The Endpoint interface

Example 46.1, “Endpoint Interface” shows the definition of the `org.apache.camel.Endpoint` interface.

Example 46.1. Endpoint Interface

```

package org.apache.camel;

public interface Endpoint {
    boolean isSingleton();

    String getEndpointUri();

    String getEndpointKey();

    CamelContext getCamelContext();
    void setCamelContext(CamelContext context);

    void configureProperties(Map options);

    boolean isLenientProperties();
  }

```

```

Exchange createExchange();
Exchange createExchange(ExchangePattern pattern);
Exchange createExchange(Exchange exchange);

Producer createProducer() throws Exception;

Consumer createConsumer(Processor processor) throws Exception;
PollingConsumer createPollingConsumer() throws Exception;
}

```

Endpoint methods

The **Endpoint** interface defines the following methods:

- **isSingleton()**—Returns **true**, if you want to ensure that each URI maps to a single endpoint within a **CamelContext**. When this property is **true**, multiple references to the identical URI within your routes always refer to a *single* endpoint instance. When this property is **false**, on the other hand, multiple references to the same URI within your routes refer to *distinct* endpoint instances. Each time you refer to the URI in a route, a new endpoint instance is created.
- **getEndpointUri()**—Returns the endpoint URI of this endpoint.
- **getEndpointKey()**—Used by **org.apache.camel.spi.LifecycleStrategy** when registering the endpoint.
- **getCamelContext()**—return a reference to the **CamelContext** instance to which this endpoint belongs.
- **setCamelContext()**—Sets the **CamelContext** instance to which this endpoint belongs.
- **configureProperties()**—Stores a copy of the parameter map that is used to inject parameters when creating a new **Consumer** instance.
- **isLenientProperties()**—Returns **true** to indicate that the URI is allowed to contain unknown parameters (that is, parameters that cannot be injected on the **Endpoint** or the **Consumer** class). Normally, this method should be implemented to return **false**.
- **createExchange()**—An overloaded method with the following variants:
 - **Exchange createExchange()**—Creates a new exchange instance with a default exchange pattern setting.
 - **Exchange createExchange(ExchangePattern pattern)**—Creates a new exchange instance with the specified exchange pattern.
 - **Exchange createExchange(Exchange exchange)**—Converts the given **exchange** argument to the type of exchange needed for this endpoint. If the given exchange is not already of the correct type, this method copies it into a new instance of the correct type. A default implementation of this method is provided in the **DefaultEndpoint** class.
- **createProducer()**—Factory method used to create new **Producer** instances.

- **createConsumer()**—Factory method to create new event-driven consumer instances. The **processor** argument is a reference to the first processor in the route.
- **createPollingConsumer()**—Factory method to create new polling consumer instances.

Endpoint singletons

In order to avoid unnecessary overhead, it is a good idea to create a *single* endpoint instance for all endpoints that have the same URI (within a CamelContext). You can enforce this condition by implementing **isSingleton()** to return **true**.



NOTE

In this context, *same URI* means that two URIs are the same when compared using string equality. In principle, it is possible to have two URIs that are equivalent, though represented by different strings. In that case, the URIs would not be treated as the same.

46.2. IMPLEMENTING THE ENDPOINT INTERFACE

Alternative ways of implementing an endpoint

The following alternative endpoint implementation patterns are supported:

- [Event-driven endpoint implementation](#)
- [Scheduled poll endpoint implementation](#)
- [Polling endpoint implementation](#)

Event-driven endpoint implementation

If your custom endpoint conforms to the event-driven pattern (see [Section 44.1.3, “Consumer Patterns and Threading”](#)), it is implemented by extending the abstract class, **org.apache.camel.impl.DefaultEndpoint**, as shown in [Example 46.2, “Implementing DefaultEndpoint”](#).

Example 46.2. Implementing DefaultEndpoint

```
import java.util.Map;
import java.util.concurrent.BlockingQueue;

import org.apache.camel.Component;
import org.apache.camel.Consumer;
import org.apache.camel.Exchange;
import org.apache.camel.Processor;
import org.apache.camel.Producer;
import org.apache.camel.impl.DefaultEndpoint;
import org.apache.camel.impl.DefaultExchange;

1 public class CustomEndpoint extends DefaultEndpoint {
2     public CustomEndpoint(String endpointUri, Component component) {
```

```

        super(endpointUri, component);
        // Do any other initialization...
    }

    3 public Producer createProducer() throws Exception {
        return new CustomProducer(this);
    }

    public Consumer createConsumer(Processor processor) throws Exception
    4 {
        return new CustomConsumer(this, processor);
    }

    public boolean isSingleton() {
        return true;
    }

    // Implement the following methods, only if you need to set exchange
    // properties.
    //
    5 public Exchange createExchange() {
        return this.createExchange(getExchangePattern());
    }

    public Exchange createExchange(ExchangePattern pattern) {
        Exchange result = new DefaultExchange(getCamelContext(),
        pattern);
        // Set exchange properties
        ...
        return result;
    }
}

```

- 1 Implement an event-driven custom endpoint, *CustomEndpoint*, by extending the **DefaultEndpoint** class.
- 2 You must have at least one constructor that takes the endpoint URI, **endpointUri**, and the parent component reference, **component**, as arguments.
- 3 Implement the **createProducer()** factory method to create producer endpoints.
- 4 Implement the **createConsumer()** factory method to create event-driven consumer instances.



IMPORTANT

Do *not* override the **createPollingConsumer()** method.

- 5 In general, it is *not* necessary to override the **createExchange()** methods. The implementations inherited from **DefaultEndpoint** create a **DefaultExchange** object by default, which can be used in any Apache Camel component. If you need to initialize some exchange properties in the **DefaultExchange** object, however, it is appropriate to override the **createExchange()** methods here in order to add the exchange property settings.

The **DefaultEndpoint** class provides default implementations of the following methods, which you might find useful when writing your custom endpoint code:

- **getEndpointUri()**—Returns the endpoint URI.
- **getCamelContext()**—Returns a reference to the **CamelContext**.
- **getComponent()**—Returns a reference to the parent component.
- **createPollingConsumer()**—Creates a polling consumer. The created polling consumer's functionality is based on the event-driven consumer. If you override the event-driven consumer method, **createConsumer()**, you get a polling consumer implementation for free.
- **createExchange(Exchange e)**—Converts the given exchange object, **e**, to the type required for this endpoint. This method creates a new endpoint using the overridden **createExchange()** endpoints. This ensures that the method also works for custom exchange types.

Scheduled poll endpoint implementation

If your custom endpoint conforms to the scheduled poll pattern (see [Section 44.1.3, “Consumer Patterns and Threading”](#)) it is implemented by inheriting from the abstract class, **org.apache.camel.impl.ScheduledPollEndpoint**, as shown in [Example 46.3, “ScheduledPollEndpoint Implementation”](#).

Example 46.3. ScheduledPollEndpoint Implementation

```
import org.apache.camel.Consumer;
import org.apache.camel.Processor;
import org.apache.camel.Producer;
import org.apache.camel.ExchangePattern;
import org.apache.camel.Message;
import org.apache.camel.impl.ScheduledPollEndpoint;

1 public class CustomEndpoint extends ScheduledPollEndpoint {
    protected CustomEndpoint(String endpointUri, CustomComponent
2 component) {
        super(endpointUri, component);
        // Do any other initialization...
    }

3     public Producer createProducer() throws Exception {
        Producer result = new CustomProducer(this);
        return result;
    }

    public Consumer createConsumer(Processor processor) throws Exception
4 {
5     Consumer result = new CustomConsumer(this, processor);
        configureConsumer(result);
        return result;
    }

    public boolean isSingleton() {
```



```


        return true;
    }

    // Implement the following methods, only if you need to set exchange
    // properties.
    //
    6 public Exchange createExchange() {
        return this.createExchange(getExchangePattern());
    }

    public Exchange createExchange(ExchangePattern pattern) {
        Exchange result = new DefaultExchange(getCamelContext(),
        pattern);
        // Set exchange properties
        ...
        return result;
    }
}

```

- 1 Implement a scheduled poll custom endpoint, *CustomEndpoint*, by extending the **ScheduledPollEndpoint** class.
- 2 You must to have at least one constructor that takes the endpoint URI, **endpointUri**, and the parent component reference, **component**, as arguments.
- 3 Implement the **createProducer()** factory method to create a producer endpoint.
- 4 Implement the **createConsumer()** factory method to create a scheduled poll consumer instance.



IMPORTANT

Do *not* override the **createPollingConsumer()** method.

- 5 The **configureConsumer()** method, defined in the **ScheduledPollEndpoint** base class, is responsible for injecting consumer query options into the consumer. See [the section called “Consumer parameter injection”](#).
- 6 In general, it is *not* necessary to override the **createExchange()** methods. The implementations inherited from **DefaultEndpoint** create a **DefaultExchange** object by default, which can be used in any Apache Camel component. If you need to initialize some exchange properties in the **DefaultExchange** object, however, it is appropriate to override the **createExchange()** methods here in order to add the exchange property settings.

Polling endpoint implementation

If your custom endpoint conforms to the polling consumer pattern (see [Section 44.1.3, “Consumer Patterns and Threading”](#)), it is implemented by inheriting from the abstract class, **org.apache.camel.impl.DefaultPollingEndpoint**, as shown in [Example 46.4, “DefaultPollingEndpoint Implementation”](#).

Example 46.4. DefaultPollingEndpoint Implementation

```

import org.apache.camel.Consumer;
import org.apache.camel.Processor;
import org.apache.camel.Producer;
import org.apache.camel.ExchangePattern;
import org.apache.camel.Message;
import org.apache.camel.impl.DefaultPollingEndpoint;

public class CustomEndpoint extends DefaultPollingEndpoint {
    ...
    public PollingConsumer createPollingConsumer() throws Exception {
        PollingConsumer result = new CustomConsumer(this);
        configureConsumer(result);
        return result;
    }

    // Do NOT implement createConsumer(). It is already implemented in
    // DefaultPollingEndpoint.
    ...
}

```

Because this *CustomEndpoint* class is a polling endpoint, you must implement the **createPollingConsumer()** method instead of the **createConsumer()** method. The consumer instance returned from **createPollingConsumer()** must inherit from the **PollingConsumer** interface. For details of how to implement a polling consumer, see [the section called “Polling consumer implementation”](#).

Apart from the implementation of the **createPollingConsumer()** method, the steps for implementing a **DefaultPollingEndpoint** are similar to the steps for implementing a **ScheduledPollEndpoint**. See [Example 46.3, “ScheduledPollEndpoint Implementation”](#) for details.

Implementing the **BrowsableEndpoint** interface

If you want to expose the list of exchange instances that are pending in the current endpoint, you can implement the **org.apache.camel.spi.BrowsableEndpoint** interface, as shown in [Example 46.5, “BrowsableEndpoint Interface”](#). It makes sense to implement this interface if the endpoint performs some sort of buffering of incoming events. For example, the Apache Camel SEDA endpoint implements the **BrowsableEndpoint** interface—see [Example 46.6, “SedaEndpoint Implementation”](#).

Example 46.5. BrowsableEndpoint Interface

```

package org.apache.camel.spi;

import java.util.List;

import org.apache.camel.Endpoint;
import org.apache.camel.Exchange;

public interface BrowsableEndpoint extends Endpoint {
    List<Exchange> getExchanges();
}

```

Example

Example 46.6, “SedaEndpoint Implementation” shows a sample implementation of **SedaEndpoint**. The SEDA endpoint is an example of an *event-driven endpoint*. Incoming events are stored in a FIFO queue (an instance of **java.util.concurrent.BlockingQueue**) and a SEDA consumer starts up a thread to read and process the events. The events themselves are represented by **org.apache.camel.Exchange** objects.

Example 46.6. SedaEndpoint Implementation

```

package org.apache.camel.component.seda;

import java.util.ArrayList;
import java.util.List;
import java.util.Map;
import java.util.concurrent.BlockingQueue;

import org.apache.camel.Component;
import org.apache.camel.Consumer;
import org.apache.camel.Exchange;
import org.apache.camel.Processor;
import org.apache.camel.Producer;
import org.apache.camel.impl.DefaultEndpoint;
import org.apache.camel.spi.BrowsableEndpoint;

public class SedaEndpoint extends DefaultEndpoint implements
  ❶ BrowsableEndpoint {
    private BlockingQueue<Exchange> queue;

    public SedaEndpoint(String endpointUri, Component component,
  ❷ BlockingQueue<Exchange> queue) {
        super(endpointUri, component);
        this.queue = queue;
    }

    public SedaEndpoint(String uri, SedaComponent component, Map
  ❸ parameters) {
        this(uri, component, component.createQueue(uri, parameters));
    }

    ❹ public Producer createProducer() throws Exception {
        return new CollectionProducer(this, getQueue());
    }

    public Consumer createConsumer(Processor processor) throws Exception
  ❺ {
        return new SedaConsumer(this, processor);
    }

    ❻ public BlockingQueue<Exchange> getQueue() {
        return queue;
    }

    ❼ public boolean isSingleton() {
        return true;
    }
}

```

```
8 public List<Exchange> getExchanges() {  
    return new ArrayList<Exchange>(getQueue());  
}  
}
```

- 1 The **SedaEndpoint** class follows the pattern for implementing an event-driven endpoint by extending the **DefaultEndpoint** class. The **SedaEndpoint** class also implements the **BrowsableEndpoint** interface, which provides access to the list of exchange objects in the queue.
- 2 Following the usual pattern for an event-driven consumer, **SedaEndpoint** defines a constructor that takes an endpoint argument, **endpointUri**, and a component reference argument, **component**.
- 3 Another constructor is provided, which delegates queue creation to the parent component instance.
- 4 The **createProducer()** factory method creates an instance of **CollectionProducer**, which is a producer implementation that adds events to the queue.
- 5 The **createConsumer()** factory method creates an instance of **SedaConsumer**, which is responsible for pulling events off the queue and processing them.
- 6 The **getQueue()** method returns a reference to the queue.
- 7 The **isSingleton()** method returns **true**, indicating that a single endpoint instance should be created for each unique URI string.
- 8 The **getExchanges()** method implements the corresponding abstract method from **BrowsableEndpoint**.

CHAPTER 47. CONSUMER INTERFACE

Abstract

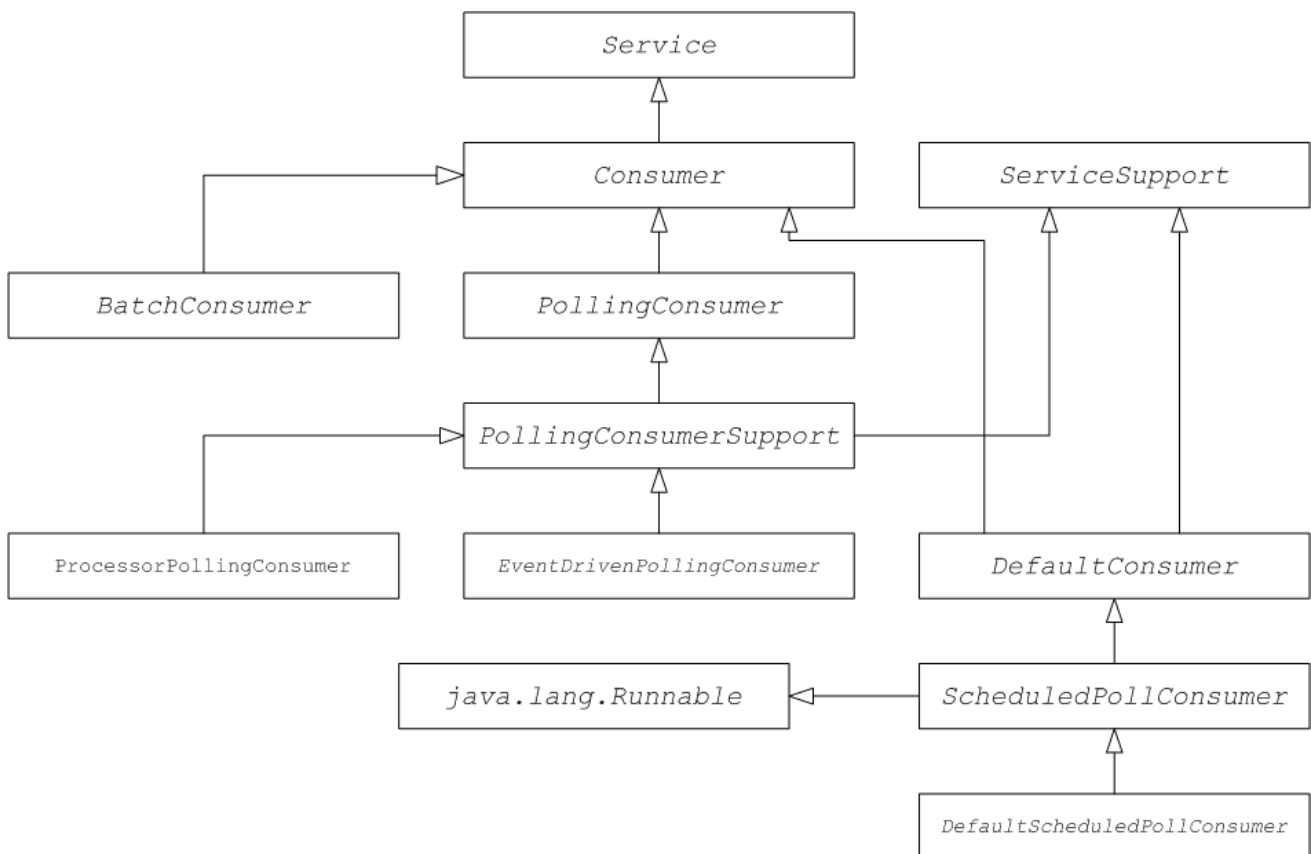
This chapter describes how to implement the **Consumer** interface, which is an essential step in the implementation of a Apache Camel component.

47.1. THE CONSUMER INTERFACE

Overview

An instance of `org.apache.camel.Consumer` type represents a source endpoint in a route. There are several different ways of implementing a consumer (see [Section 44.1.3, “Consumer Patterns and Threading”](#)), and this degree of flexibility is reflected in the inheritance hierarchy (see [Figure 47.1, “Consumer Inheritance Hierarchy”](#)), which includes several different base classes for implementing a consumer.

Figure 47.1. Consumer Inheritance Hierarchy



Consumer parameter injection

For consumers that follow the scheduled poll pattern (see [the section called “Scheduled poll pattern”](#)), Apache Camel provides support for injecting parameters into consumer instances. For example, consider the following endpoint URI for a component identified by the **custom** prefix:

```
custom:destination?consumer.myConsumerParam
```

Apache Camel provides support for automatically injecting query options of the form **consumer.***. For the **consumer.myConsumerParam** parameter, you need to define corresponding setter and getter methods on the **Consumer** implementation class as follows:

```
public class CustomConsumer extends ScheduledPollConsumer {
    ...
    String getMyConsumerParam() { ... }
    void setMyConsumerParam(String s) { ... }
    ...
}
```

Where the getter and setter methods follow the usual Java bean conventions (including capitalizing the first letter of the property name).

In addition to defining the bean methods in your Consumer implementation, you must also remember to call the **configureConsumer()** method in the implementation of **Endpoint.createConsumer()**. See the section called “[Scheduled poll endpoint implementation](#)”. [Example 47.1, “FileEndpoint createConsumer\(\) Implementation](#)” shows an example of a **createConsumer()** method implementation, taken from the **FileEndpoint** class in the file component:

Example 47.1. FileEndpoint createConsumer() Implementation

```
...
public class FileEndpoint extends ScheduledPollEndpoint {
    ...
    public Consumer createConsumer(Processor processor) throws
    Exception {
        Consumer result = new FileConsumer(this, processor);
        configureConsumer(result);
        return result;
    }
    ...
}
```

At run time, consumer parameter injection works as follows:

1. When the endpoint is created, the default implementation of **DefaultComponent.createEndpoint(String uri)** parses the URI to extract the consumer parameters, and stores them in the endpoint instance by calling **ScheduledPollEndpoint.configureProperties()**.
2. When **createConsumer()** is called, the method implementation calls **configureConsumer()** to inject the consumer parameters (see [Example 47.1, “FileEndpoint createConsumer\(\) Implementation](#)”).
3. The **configureConsumer()** method uses Java reflection to call the setter methods whose names match the relevant options after the **consumer.** prefix has been stripped off.

Scheduled poll parameters

A consumer that follows the scheduled poll pattern automatically supports the consumer parameters shown in [Table 47.1, “Scheduled Poll Parameters”](#) (which can appear as query options in the endpoint URI).

Table 47.1. Scheduled Poll Parameters

Name	Default	Description
initialDelay	1000	Delay, in milliseconds, before the first poll.
delay	500	Depends on the value of the useFixedDelay flag (time unit is milliseconds).
useFixedDelay	false	<p>If false, the delay parameter is interpreted as the polling period. Polls will occur at initialDelay, initialDelay+delay, initialDelay+2*delay, and so on.</p> <p>If true, the delay parameter is interpreted as the time elapsed between the previous execution and the next execution. Polls will occur at initialDelay, initialDelay+[ProcessingTime]+delay, and so on. Where <i>ProcessingTime</i> is the time taken to process an exchange object in the current thread.</p>

Converting between event-driven and polling consumers

Apache Camel provides two special consumer implementations which can be used to convert back and forth between an event-driven consumer and a polling consumer. The following conversion classes are provided:

- **org.apache.camel.impl.EventDrivenPollingConsumer**—Converts an event-driven consumer into a polling consumer instance.
- **org.apache.camel.impl.DefaultScheduledPollConsumer**—Converts a polling consumer into an event-driven consumer instance.

In practice, these classes are used to simplify the task of implementing an **Endpoint** type. The **Endpoint** interface defines the following two methods for creating a consumer instance:

```
package org.apache.camel;

public interface Endpoint {
    ...
    Consumer createConsumer(Processor processor) throws Exception;
    PollingConsumer createPollingConsumer() throws Exception;
}
```

createConsumer() returns an event-driven consumer and **createPollingConsumer()** returns a polling consumer. You would only implement one these methods. For example, if you are following the event-driven pattern for your consumer, you would implement the **createConsumer()** method provide a method implementation for **createPollingConsumer()** that simply raises an exception. With the help of the conversion classes, however, Apache Camel is able to provide a more useful default implementation.

For example, if you want to implement your consumer according to the event-driven pattern, you implement the endpoint by extending **DefaultEndpoint** and implementing the **createConsumer()** method. The implementation of **createPollingConsumer()** is inherited from **DefaultEndpoint**, where it is defined as follows:

```
public PollingConsumer<E> createPollingConsumer() throws Exception {
    return new EventDrivenPollingConsumer<E>(this);
}
```

The **EventDrivenPollingConsumer** constructor takes a reference to the event-driven consumer, **this**, effectively wrapping it and converting it into a polling consumer. To implement the conversion, the **EventDrivenPollingConsumer** instance buffers incoming events and makes them available on demand through the **receive()**, the **receive(long timeout)**, and the **receiveNowait()** methods.

Analogously, if you are implementing your consumer according to the polling pattern, you implement the endpoint by extending **DefaultPollingEndpoint** and implementing the **createPollingConsumer()** method. In this case, the implementation of the **createConsumer()** method is inherited from **DefaultPollingEndpoint**, and the default implementation returns a **DefaultScheduledPollConsumer** instance (which converts the polling consumer into an event-driven consumer).

ShutdownPrepared interface

Consumer classes can optionally implement the **org.apache.camel.spi.ShutdownPrepared** interface, which enables your custom consumer endpoint to receive shutdown notifications.

[Example 47.2, “ShutdownPrepared Interface”](#) shows the definition of the **ShutdownPrepared** interface.

Example 47.2. ShutdownPrepared Interface

```
package org.apache.camel.spi;

public interface ShutdownPrepared {

    void prepareShutdown(boolean forced);

}
```

The **ShutdownPrepared** interface defines the following methods:

prepareShutdown

Receives notifications to shut down the consumer endpoint in one or two phases, as follows:

1. *Graceful shutdown*—where the **forced** argument has the value **false**. Attempt to clean up resources gracefully. For example, by stopping threads gracefully.
2. *Forced shutdown*—where the **forced** argument has the value **true**. This means that the shutdown has timed out, so you must clean up resources more aggressively. This is the last chance to clean up resources before the process

exits.

ShutdownAware interface

Consumer classes can optionally implement the `org.apache.camel.spi.ShutdownAware` interface, which interacts with the graceful shutdown mechanism, enabling a consumer to ask for extra time to shut down. This is typically needed for components such as SEDA, which can have pending exchanges stored in an internal queue. Normally, you would want to process all of the exchanges in the queue before shutting down the SEDA consumer.

[Example 47.3, “ShutdownAware Interface”](#) shows the definition of the `ShutdownAware` interface.

Example 47.3. ShutdownAware Interface

```
// Java
package org.apache.camel.spi;

import org.apache.camel.ShutdownRunningTask;

public interface ShutdownAware extends ShutdownPrepared {

    boolean deferShutdown(ShutdownRunningTask shutdownRunningTask);

    int getPendingExchangesSize();
}
```

The `ShutdownAware` interface defines the following methods:

deferShutdown

Return `true` from this method, if you want to delay shutdown of the consumer. The `shutdownRunningTask` argument is an `enum` which can take either of the following values:

- `ShutdownRunningTask.CompleteCurrentTaskOnly`—finish processing the exchanges that are currently being processed by the consumer's thread pool, but do not attempt to process any more exchanges than that.
- `ShutdownRunningTask.CompleteAllTasks`—process *all* of the pending exchanges. For example, in the case of the SEDA component, the consumer would process all of the exchanges from its incoming queue.

getPendingExchangesSize

Indicates how many exchanges remain to be processed by the consumer. A zero value indicates that processing is finished and the consumer can be shut down.

For an example of how to define the `ShutdownAware` methods, see [Example 47.7, “Custom Threading Implementation”](#).

47.2. IMPLEMENTING THE CONSUMER INTERFACE

Alternative ways of implementing a consumer

You can implement a consumer in one of the following ways:

- [Event-driven consumer implementation](#)
- [Scheduled poll consumer implementation](#)
- [Polling consumer implementation](#)
- [Custom threading implementation](#)

Event-driven consumer implementation

In an event-driven consumer, processing is driven explicitly by external events. The events are received through an event-listener interface, where the listener interface is specific to the particular event source.

[Example 47.4, “JMXConsumer Implementation”](#) shows the implementation of the **JMXConsumer** class, which is taken from the Apache Camel JMX component implementation. The **JMXConsumer** class is an example of an event-driven consumer, which is implemented by inheriting from the **org.apache.camel.impl.DefaultConsumer** class. In the case of the **JMXConsumer** example, events are represented by calls on the **NotificationListener.handleNotification()** method, which is a standard way of receiving JMX events. In order to receive these JMX events, it is necessary to implement the **NotificationListener** interface and override the **handleNotification()** method, as shown in [Example 47.4, “JMXConsumer Implementation”](#).

Example 47.4. JMXConsumer Implementation

```
package org.apache.camel.component.jmx;

import javax.management.Notification;
import javax.management.NotificationListener;
import org.apache.camel.Processor;
import org.apache.camel.impl.DefaultConsumer;

public class JMXConsumer extends DefaultConsumer implements
1 NotificationListener {
    JMXEndpoint jmxEndpoint;

2    public JMXConsumer(JMXEndpoint endpoint, Processor processor) {
        super(endpoint, processor);
        this.jmxEndpoint = endpoint;
    }


    public void handleNotification(Notification notification, Object
3 handback) {
        try {

4 getProcessor().process(jmxEndpoint.createExchange(notification));
        } catch (Throwable e) {

5             handleException(e);
        }
    }
}
```

```
| | }
```

- 1 The **JMXConsumer** pattern follows the usual pattern for event-driven consumers by extending the **DefaultConsumer** class. Additionally, because this consumer is designed to receive events from JMX (which are represented by JMX notifications), it is necessary to implement the **NotificationListener** interface.
- 2 You must implement at least one constructor that takes a reference to the parent endpoint, **endpoint**, and a reference to the next processor in the chain, **processor**, as arguments.
- 3 The **handleNotification()** method (which is defined in **NotificationListener**) is automatically invoked by JMX whenever a JMX notification arrives. The body of this method should contain the code that performs the consumer's event processing. Because the **handleNotification()** call originates from the JMX layer, the consumer's threading model is implicitly controlled by the JMX layer, not by the **JMXConsumer** class.



NOTE

The **handleNotification()** method is specific to the JMX example. When implementing your own event-driven consumer, you must identify an analogous event listener method to implement in your custom consumer.

- 4 This line of code combines two steps. First, the JMX notification object is converted into an exchange object, which is the generic representation of an event in Apache Camel. Then the newly created exchange object is passed to the next processor in the route (invoked synchronously).
- 5 The **handleException()** method is implemented by the **DefaultConsumer** base class. By default, it handles exceptions using the **org.apache.camel.impl.LoggingExceptionHandler** class.

Scheduled poll consumer implementation

In a scheduled poll consumer, polling events are automatically generated by a timer class, **java.util.concurrent.ScheduledExecutorService**. To receive the generated polling events, you must implement the **ScheduledPollConsumer.poll()** method (see [Section 44.1.3, “Consumer Patterns and Threading”](#)).

[Example 47.5, “ScheduledPollConsumer Implementation”](#) shows how to implement a consumer that follows the scheduled poll pattern, which is implemented by extending the **ScheduledPollConsumer** class.

Example 47.5. ScheduledPollConsumer Implementation

```
import java.util.concurrent.ScheduledExecutorService;

import org.apache.camel.Consumer;
import org.apache.camel.Endpoint;
import org.apache.camel.Exchange;
import org.apache.camel.Message;
import org.apache.camel.PollingConsumer;
import org.apache.camel.Processor;
```

```

import org.apache.camel.impl.ScheduledPollConsumer;

1 public class CustomConsumer extends ScheduledPollConsumer {
    private final CustomEndpoint endpoint;

    2 public CustomConsumer(CustomEndpoint endpoint, Processor processor)
    {
        super(endpoint, processor);
        this.endpoint = endpoint;
    }

    3 protected void poll() throws Exception {
        Exchange exchange = /* Receive exchange object ... */;

        4 // Example of a synchronous processor.
        getProcessor().process(exchange);
    }

    @Override
    5 protected void doStart() throws Exception {
        // Pre-Start:
        // Place code here to execute just before start of processing.
        super.doStart();
        // Post-Start:
        // Place code here to execute just after start of processing.
    }

    @Override
    6 protected void doStop() throws Exception {
        // Pre-Stop:
        // Place code here to execute just before processing stops.
        super.doStop();
        // Post-Stop:
        // Place code here to execute just after processing stops.
    }
}

```

- 1 Implement a scheduled poll consumer class, *CustomConsumer*, by extending the `org.apache.camel.impl.ScheduledPollConsumer` class.
- 2 You must implement at least one constructor that takes a reference to the parent endpoint, **endpoint**, and a reference to the next processor in the chain, **processor**, as arguments.
- 3 Override the `poll()` method to receive the scheduled polling events. This is where you should put the code that retrieves and processes incoming events (represented by exchange objects).
- 4 In this example, the event is processed synchronously. If you want to process events asynchronously, you should use a reference to an asynchronous processor instead, by calling `getAsyncProcessor()`. For details of how to process events asynchronously, see [Section 44.1.4, “Asynchronous Processing”](#).

5

(Optional) If you want some lines of code to execute as the consumer is starting up, override the `doStart()` method as shown.

- 6 (Optional) If you want some lines of code to execute as the consumer is stopping, override the `doStop()` method as shown.

Polling consumer implementation

Example 47.6, “[PollingConsumerSupport Implementation](#)” outlines how to implement a consumer that follows the polling pattern, which is implemented by extending the `PollingConsumerSupport` class.

Example 47.6. `PollingConsumerSupport` Implementation

```
import org.apache.camel.Exchange;
import org.apache.camel.RuntimeCamelException;
import org.apache.camel.impl.PollingConsumerSupport;

1 public class CustomConsumer extends PollingConsumerSupport {
    private final CustomEndpoint endpoint;

2     public CustomConsumer(CustomEndpoint endpoint) {
        super(endpoint);
        this.endpoint = endpoint;
    }

3     public Exchange receiveNowait() {
        Exchange exchange = /* Obtain an exchange object. */;
        // Further processing ...
        return exchange;
    }

4     public Exchange receive() {
        // Blocking poll ...
    }

5     public Exchange receive(long timeout) {
        // Poll with timeout ...
    }

6     protected void doStart() throws Exception {
        // Code to execute whilst starting up.
    }

    protected void doStop() throws Exception {
        // Code to execute whilst shutting down.
    }
}
```

- 1 Implement your polling consumer class, `CustomConsumer`, by extending the `org.apache.camel.impl.PollingConsumerSupport` class.

2

You must implement at least one constructor that takes a reference to the parent endpoint, **endpoint**, as an argument. A polling consumer does not need a reference to

- 3 The **receiveNowait()** method should implement a non-blocking algorithm for retrieving an event (exchange object). If no event is available, it should return **null**.
- 4 The **receive()** method should implement a blocking algorithm for retrieving an event. This method can block indefinitely, if events remain unavailable.
- 5 The **receive(long timeout)** method implements an algorithm that can block for as long as the specified timeout (typically specified in units of milliseconds).
- 6 If you want to insert code that executes while a consumer is starting up or shutting down, implement the **doStart()** method and the **doStop()** method, respectively.

Custom threading implementation

If the standard consumer patterns are not suitable for your consumer implementation, you can implement the **Consumer** interface directly and write the threading code yourself. When writing the threading code, however, it is important that you comply with the standard Apache Camel threading model, as described in [Section 2.9, “Threading Model”](#).

For example, the SEDA component from **camel-core** implements its own consumer threading, which is consistent with the Apache Camel threading model. [Example 47.7, “Custom Threading Implementation”](#) shows an outline of how the **SedaConsumer** class implements its threading.

Example 47.7. Custom Threading Implementation

```
package org.apache.camel.component.seda;

import java.util.ArrayList;
import java.util.List;
import java.util.concurrent.BlockingQueue;
import java.util.concurrent.ExecutorService;
import java.util.concurrent.TimeUnit;

import org.apache.camel.Consumer;
import org.apache.camel.Endpoint;
import org.apache.camel.Exchange;
import org.apache.camel.Processor;
import org.apache.camel.ShutdownRunningTask;
import org.apache.camel.impl.LoggingExceptionHandler;
import org.apache.camel.impl.ServiceSupport;
import org.apache.camel.util.ServiceHelper;
...
import org.apache.commons.logging.Log;
import org.apache.commons.logging.LogFactory;

/**
 * A Consumer for the SEDA component.
 *
 * @version $Revision: 922485 $
 */
public class SedaConsumer extends ServiceSupport implements Consumer,
```

```

1 Runnable, ShutdownAware {
    private static final transient Log LOG =
LogFactory.getLog(SedaConsumer.class);

    private SedaEndpoint endpoint;
    private Processor processor;
    private ExecutorService executor;
    ...
    public SedaConsumer(SedaEndpoint endpoint, Processor processor) {
        this.endpoint = endpoint;
        this.processor = processor;
    }
    ...

2    public void run() {
        BlockingQueue<Exchange> queue = endpoint.getQueue();
        // Poll the queue and process exchanges
        ...
    }

    ...

3    protected void doStart() throws Exception {
        int poolSize = endpoint.getConcurrentConsumers();
        executor =
endpoint.getCamelContext().getExecutorServiceStrategy()
        .newFixedThreadPool(this, endpoint.getEndpointUri(),
4 poolSize);
5        for (int i = 0; i < poolSize; i++) {
            executor.execute(this);
        }
        endpoint.onStarted(this);
    }

6    protected void doStop() throws Exception {
        endpoint.onStopped(this);
        // must shutdown executor on stop to avoid overhead of having
them running

endpoint.getCamelContext().getExecutorServiceStrategy().shutdownNow(exec
7 utor);
        executor = null;

        if (multicast != null) {
            ServiceHelper.stopServices(multicast);
        }
    }
    ...
//-----
// Implementation of ShutdownAware interface

    public boolean deferShutdown(ShutdownRunningTask
shutdownRunningTask) {
        // deny stopping on shutdown as we want seda consumers to run
in case some other queues
        // depend on this consumer to run, so it can complete its
exchanges

```

```
        return true;
    }

    public int getPendingExchangesSize() {
        // number of pending messages on the queue
        return endpoint.getQueue().size();
    }
}
```

- 1 The **SedaConsumer** class is implemented by extending the **org.apache.camel.impl.ServiceSupport** class and implementing the **Consumer**, **Runnable**, and **ShutdownAware** interfaces.
- 2 Implement the **Runnable.run()** method to define what the consumer does while it is running in a thread. In this case, the consumer runs in a loop, polling the queue for new exchanges and then processing the exchanges in the latter part of the queue.
- 3 The **doStart()** method is inherited from **ServiceSupport**. You override this method in order to define what the consumer does when it starts up.
- 4 Instead of creating threads directly, you should create a thread pool using the **ExecutorServiceStrategy** object that is registered with the **CamelContext**. This is important, because it enables Apache Camel to implement centralized management of threads and support such features as graceful shutdown.

For details, see [Section 2.9, “Threading Model”](#).

- 5 Kick off the threads by calling the **ExecutorService.execute()** method **poolSize** times.
- 6 The **doStop()** method is inherited from **ServiceSupport**. You override this method in order to define what the consumer does when it shuts down.
- 7 Shut down the thread pool, which is represented by the **executor** instance.

CHAPTER 48. PRODUCER INTERFACE

Abstract


This chapter describes how to implement the **Producer** interface, which is an essential step in the implementation of a Apache Camel component.

48.1. THE PRODUCER INTERFACE

Overview

An instance of `org.apache.camel.Producer` type represents a target endpoint in a route. The role of the producer is to send requests (*In* messages) to a specific physical endpoint and to receive the corresponding response (*Out* or *Fault* message). A **Producer** object is essentially a special kind of **Processor** that appears at the end of a processor chain (equivalent to a route). [Figure 48.1, “Producer Inheritance Hierarchy”](#) shows the inheritance hierarchy for producers.

Figure 48.1. Producer Inheritance Hierarchy



The Producer interface

[Example 48.1, “Producer Interface”](#) shows the definition of the `org.apache.camel.Producer` interface.

Example 48.1. Producer Interface

```

package org.apache.camel;

public interface Producer extends Processor, Service, IsSingleton {

    Endpoint<E> getEndpoint();

    Exchange createExchange();

    Exchange createExchange(ExchangePattern pattern);

    Exchange createExchange(E exchange);
}

```

Producer methods

The **Producer** interface defines the following methods:

- **process()** (*inherited from Processor*)—The most important method. A producer is essentially a special type of processor that sends a request to an endpoint, instead of forwarding the exchange object to another processor. By overriding the **process()** method, you define how the producer sends and receives messages to and from the relevant endpoint.
- **getEndpoint()**—Returns a reference to the parent endpoint instance.
- **createExchange()**—These overloaded methods are analogous to the corresponding methods defined in the **Endpoint** interface. Normally, these methods delegate to the corresponding methods defined on the parent **Endpoint** instance (this is what the **DefaultEndpoint** class does by default). Occasionally, you might need to override these methods.

Asynchronous processing

Processing an exchange object in a producer—which usually involves sending a message to a remote destination and waiting for a reply—can potentially block for a significant length of time. If you want to avoid blocking the current thread, you can opt to implement the producer as an *asynchronous processor*. The asynchronous processing pattern decouples the preceding processor from the producer, so that the **process()** method returns without delay. See [Section 44.1.4, “Asynchronous Processing”](#).

When implementing a producer, you can support the asynchronous processing model by implementing the **org.apache.camel.AsyncProcessor** interface. On its own, this is not enough to ensure that the asynchronous processing model will be used: it is also necessary for the preceding processor in the chain to call the asynchronous version of the **process()** method. The definition of the **AsyncProcessor** interface is shown in [Example 48.2, “AsyncProcessor Interface”](#).

Example 48.2. AsyncProcessor Interface

```
package org.apache.camel;

public interface AsyncProcessor extends Processor {
    boolean process(Exchange exchange, AsyncCallback callback);
}
```

The asynchronous version of the **process()** method takes an extra argument, **callback**, of **org.apache.camel.AsyncCallback** type. The corresponding **AsyncCallback** interface is defined as shown in [Example 48.3, “AsyncCallback Interface”](#).

Example 48.3. AsyncCallback Interface

```
package org.apache.camel;

public interface AsyncCallback {
```

```

    void done(boolean doneSynchronously);
}

```

The caller of `AsyncProcessor.process()` must provide an implementation of `AsyncCallback` to receive the notification that processing has finished. The `AsyncCallback.done()` method takes a boolean argument that indicates whether the processing was performed synchronously or not. Normally, the flag would be `false`, to indicate asynchronous processing. In some cases, however, it can make sense for the producer *not* to process asynchronously (in spite of being asked to do so). For example, if the producer knows that the processing of the exchange will complete rapidly, it could optimise the processing by doing it synchronously. In this case, the `doneSynchronously` flag should be set to `true`.

ExchangeHelper class

When implementing a producer, you might find it helpful to call some of the methods in the `org.apache.camel.util.ExchangeHelper` utility class. For full details of the `ExchangeHelper` class, see [Section 41.4, “The ExchangeHelper Class”](#).

48.2. IMPLEMENTING THE PRODUCER INTERFACE

Alternative ways of implementing a producer

You can implement a producer in one of the following ways:

- [How to implement a synchronous producer](#).
- [How to implement an asynchronous producer](#).

How to implement a synchronous producer

[Example 48.4, “DefaultProducer Implementation”](#) outlines how to implement a synchronous producer. In this case, call to `Producer.process()` blocks until a reply is received.

Example 48.4. DefaultProducer Implementation

```

import org.apache.camel.Endpoint;
import org.apache.camel.Exchange;
import org.apache.camel.Producer;
import org.apache.camel.impl.DefaultProducer;

1 public class CustomProducer extends DefaultProducer {
2     public CustomProducer(Endpoint endpoint) {
        super(endpoint);
        // Perform other initialization tasks...
    }
3     public void process(Exchange exchange) throws Exception {
        // Process exchange synchronously.
        // ...
    }
}

```

- 1 Implement a custom synchronous producer class, *CustomProducer*, by extending the `org.apache.camel.impl.DefaultProducer` class.
- 2 Implement a constructor that takes a reference to the parent endpoint.
- 3 The `process()` method implementation represents the core of the producer code. The implementation of the `process()` method is entirely dependent on the type of component that you are implementing. In outline, the `process()` method is normally implemented as follows:
 - If the exchange contains an *In* message, and if this is consistent with the specified exchange pattern, then send the *In* message to the designated endpoint.
 - If the exchange pattern anticipates the receipt of an *Out* message, then wait until the *Out* message has been received. This typically causes the `process()` method to block for a significant length of time.
 - When a reply is received, call `exchange.setOut()` to attach the reply to the exchange object. If the reply contains a fault message, set the fault flag on the *Out* message using `Message.setFault(true)`.

How to implement an asynchronous producer

[Example 48.5, “CollectionProducer Implementation”](#) outlines how to implement an asynchronous producer. In this case, you must implement both a synchronous `process()` method and an asynchronous `process()` method (which takes an additional `AsyncCallback` argument).

Example 48.5. CollectionProducer Implementation

```
import org.apache.camel.AsyncCallback;
import org.apache.camel.AsyncProcessor;
import org.apache.camel.Endpoint;
import org.apache.camel.Exchange;
import org.apache.camel.Producer;
import org.apache.camel.impl.DefaultProducer;

public class CustomProducer extends DefaultProducer implements
1 AsyncProcessor {
2     public CustomProducer(Endpoint endpoint) {
        super(endpoint);
        // ...
    }
3     public void process(Exchange exchange) throws Exception {
        // Process exchange synchronously.
        // ...
    }

    public boolean process(Exchange exchange, AsyncCallback callback) {
```

```

4         // Process exchange asynchronously.
           CustomProducerTask task = new CustomProducerTask(exchange,
callback);
           // Process 'task' in a separate thread...
           // ...
5         return false;
       }
}

6 public class CustomProducerTask implements Runnable {
    private Exchange exchange;
    private AsyncCallback callback;

    public CustomProducerTask(Exchange exchange, AsyncCallback
callback) {
        this.exchange = exchange;
        this.callback = callback;
    }

7     public void run() {
        // Process exchange.
        // ...
        callback.done(false);
    }
}

```

- 1 Implement a custom asynchronous producer class, *CustomProducer*, by extending the `org.apache.camel.impl.DefaultProducer` class, and implementing the `AsyncProcessor` interface.
- 2 Implement a constructor that takes a reference to the parent endpoint.
- 3 Implement the synchronous `process()` method.
- 4 Implement the asynchronous `process()` method. You can implement the asynchronous method in several ways. The approach shown here is to create a `java.lang.Runnable` instance, `task`, that represents the code that runs in a sub-thread. You then use the Java threading API to run the task in a sub-thread (for example, by creating a new thread or by allocating the task to an existing thread pool).
- 5 Normally, you return `false` from the asynchronous `process()` method, to indicate that the exchange was processed asynchronously.
- 6 The `CustomProducerTask` class encapsulates the processing code that runs in a sub-thread. This class must store a copy of the `Exchange` object, `exchange`, and the `AsyncCallback` object, `callback`, as private member variables.
- 7 The `run()` method contains the code that sends the `in` message to the producer endpoint and waits to receive the reply, if any. After receiving the reply (`Out` message or `Fault` message) and inserting it into the exchange object, you must call `callback.done()` to notify the caller that processing is complete.

CHAPTER 49. EXCHANGE INTERFACE

Abstract

This chapter describes the **Exchange** interface. Since the refactoring of the camel-core module performed in Apache Camel 2.0, there is no longer any necessity to define custom exchange types. The **DefaultExchange** implementation can now be used in all cases.


49.1. THE EXCHANGE INTERFACE

Overview

An instance of `org.apache.camel.Exchange` type encapsulates the current message passing through a route, with additional metadata encoded as exchange properties.

Figure 49.1, “Exchange Inheritance Hierarchy” shows the inheritance hierarchy for the exchange type. The default implementation, **DefaultExchange**, is always used.

Figure 49.1. Exchange Inheritance Hierarchy



The Exchange interface

Example 49.1, “Exchange Interface” shows the definition of the `org.apache.camel.Exchange` interface.

Example 49.1. Exchange Interface

```

package org.apache.camel;

import java.util.Map;

import org.apache.camel.spi.Synchronization;
import org.apache.camel.spi.UnitOfWork;

public interface Exchange {
    // Exchange property names (string constants)
    // (Not shown here)
    ...

    ExchangePattern getPattern();
    void setPattern(ExchangePattern pattern);
  
```

```

Object getProperty(String name);
Object getProperty(String name, Object defaultValue);
<T> T getProperty(String name, Class<T> type);
<T> T getProperty(String name, Object defaultValue, Class<T>
type);
void setProperty(String name, Object value);
Object removeProperty(String name);
Map<String, Object> getProperties();
boolean hasProperties();

Message getIn();
<T> T getIn(Class<T> type);
void setIn(Message in);

Message getOut();
<T> T getOut(Class<T> type);
void setOut(Message out);
boolean hasOut();

Throwable getException();
<T> T getException(Class<T> type);
void setException(Throwable e);

boolean isFailed();

boolean isTransacted();

boolean isRollbackOnly();

CamelContext getContext();

Exchange copy();

Endpoint getFromEndpoint();
void setFromEndpoint(Endpoint fromEndpoint);

String getFromRouteId();
void setFromRouteId(String fromRouteId);

UnitOfWork getUnitOfWork();
void setUnitOfWork(UnitOfWork unitOfWork);

String getExchangeId();
void setExchangeId(String id);

void addOnCompletion(Synchronization onCompletion);
void handoverCompletions(Exchange target);
}

```

Exchange methods

The **Exchange** interface defines the following methods:

- **getPattern(), setPattern()**—The exchange pattern can be one of the values enumerated in `org.apache.camel.ExchangePattern`. The following exchange pattern values are supported:
 - **InOnly**
 - **RobustInOnly**
 - **InOut**
 - **InOptionalOut**
 - **OutOnly**
 - **RobustOutOnly**
 - **OutIn**
 - **OutOptionalIn**
- **setProperty(), getProperty(), getProperties(), removeProperty(), hasProperties()**—Use the property setter and getter methods to associate named properties with the exchange instance. The properties consist of miscellaneous metadata that you might need for your component implementation.
- **setIn(), getIn()**—Setter and getter methods for the *In* message.

The `getIn()` implementation provided by the `DefaultExchange` class implements lazy creation semantics: if the *In* message is null when `getIn()` is called, the `DefaultExchange` class creates a default *In* message.

- **setOut(), getOut(), hasOut()**—Setter and getter methods for the *Out* message.

The `getOut()` method implicitly supports lazy creation of an *Out* message. That is, if the current *Out* message is `null`, a new message instance is automatically created.
- **setException(), getException()**—Getter and setter methods for an exception object (of `Throwable` type).
- **isFailed()**—Returns `true`, if the exchange failed either due to an exception or due to a fault.
- **isTransacted()**—Returns `true`, if the exchange is transacted.
- **isRollback()**—Returns `true`, if the exchange is marked for rollback.
- **getContext()**—Returns a reference to the associated `CamelContext` instance.
- **copy()**—Creates a new, identical (apart from the exchange ID) copy of the current custom exchange object. The body and headers of the *In* message, the *Out* message (if any), and the *Fault* message (if any) are also copied by this operation.
- **setFromEndpoint(), getFromEndpoint()**—Getter and setter methods for the consumer endpoint that originated this message (which is typically the endpoint appearing in the `from()` DSL command at the start of a route).

- **setFromRouteId()**, **getFromRouteId()**—Getters and setters for the route ID that originated this exchange. The **getFromRouteId()** method should only be called internally.
- **setUnitOfWork()**, **getUnitOfWork()**—Getter and setter methods for the **org.apache.camel.spi.UnitOfWork** bean property. This property is only required for exchanges that can participate in a transaction.
- **setExchangeId()**, **getExchangeId()**—Getter and setter methods for the exchange ID. Whether or not a custom component uses and exchange ID is an implementation detail.
- **addOnCompletion()**—Adds an **org.apache.camel.spi.Synchronization** callback object, which gets called when processing of the exchange has completed.
- **handoverCompletions()**—Hands over all of the *OnCompletion* callback objects to the specified exchange object.

CHAPTER 50. MESSAGE INTERFACE

Abstract

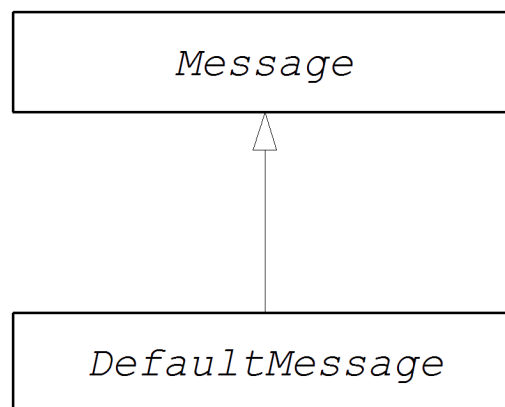
This chapter describes how to implement the **Message** interface, which is an optional step in the implementation of a Apache Camel component.

50.1. THE MESSAGE INTERFACE

Overview

An instance of **org.apache.camel.Message** type can represent any kind of message (*In* or *Out*). [Figure 50.1, “Message Inheritance Hierarchy”](#) shows the inheritance hierarchy for the message type. You do not always need to implement a custom message type for a component. In many cases, the default implementation, **DefaultMessage**, is adequate.

Figure 50.1. Message Inheritance Hierarchy



The Message interface

[Example 50.1, “Message Interface”](#) shows the definition of the **org.apache.camel.Message** interface.

Example 50.1. Message Interface

```
package org.apache.camel;

import java.util.Map;
import java.util.Set;

import javax.activation.DataHandler;

public interface Message {

    String getMessageId();
    void setMessageId(String messageId);

    Exchange getExchange();

    boolean isFault();
}
```

```

void    setFault(boolean fault);

Object getHeader(String name);
Object getHeader(String name, Object defaultValue);
<T> T getHeader(String name, Class<T> type);
<T> T getHeader(String name, Object defaultValue, Class<T> type);
Map<String, Object> getHeaders();
void setHeader(String name, Object value);
void setHeaders(Map<String, Object> headers);
Object removeHeader(String name);
boolean removeHeaders(String pattern);
boolean hasHeaders();

Object getBody();
Object getMandatoryBody() throws InvalidPayloadException;
<T> T getBody(Class<T> type);
<T> T getMandatoryBody(Class<T> type) throws
InvalidPayloadException;
void    setBody(Object body);
<T> void setBody(Object body, Class<T> type);

DataHandler getAttachment(String id);
Map<String, DataHandler> getAttachments();
Set<String> getAttachmentNames();
void removeAttachment(String id);
void addAttachment(String id, DataHandler content);
void setAttachments(Map<String, DataHandler> attachments);
boolean hasAttachments();

Message copy();

void copyFrom(Message message);

String createExchangeId();
}

```

Message methods

The **Message** interface defines the following methods:

- **setMessageId()**, **getMessageId()**—Getter and setter methods for the message ID. Whether or not you need to use a message ID in your custom component is an implementation detail.
- **getExchange()**—Returns a reference to the parent exchange object.
- **isFault()**, **setFault()**—Getter and setter methods for the fault flag, which indicates whether or not this message is a fault message.
- **getHeader()**, **getHeaders()**, **setHeader()**, **setHeaders()**, **removeHeader()**, **hasHeaders()**—Getter and setter methods for the message headers. In general, these message headers can be used either to store actual header data, or to store miscellaneous metadata.

- **getBody()**, **getMandatoryBody()**, **setBody()**—Getter and setter methods for the message body. The `getMandatoryBody()` accessor guarantees that the returned body is non-null, otherwise the **InvalidPayloadException** exception is thrown.
- **getAttachment()**, **getAttachments()**, **getAttachmentNames()**, **removeAttachment()**, **addAttachment()**, **setAttachments()**, **hasAttachments()**—Methods to get, set, add, and remove attachments.
- **copy()**—Creates a new, identical (including the message ID) copy of the current custom message object.
- **copyFrom()**—Copies the complete contents (including the message ID) of the specified generic message object, **message**, into the current message instance. Because this method must be able to copy from *any* message type, it copies the generic message properties, but not the custom properties.
- **createExchangeId()**—Returns the unique ID for this exchange, if the message implementation is capable of providing an ID; otherwise, return **null**.

50.2. IMPLEMENTING THE MESSAGE INTERFACE

How to implement a custom message

[Example 50.2, “Custom Message Implementation”](#) outlines how to implement a message by extending the **DefaultMessage** class.

Example 50.2. Custom Message Implementation

```
import org.apache.camel.Exchange;
import org.apache.camel.impl.DefaultMessage;

1 public class CustomMessage extends DefaultMessage {
2     public CustomMessage() {
        // Create message with default properties...
    }

    @Override
3     public String toString() {
        // Return a stringified message...
    }

    @Override
4     public CustomMessage newInstance() {
        return new CustomMessage( ... );
    }

    @Override
5     protected Object createBody() {
        // Return message body (lazy creation).
    }

    @Override
6     protected void populateInitialHeaders(Map<String, Object> map) {
        // Initialize headers from underlying message (lazy
```

```

creation).
    }

    @Override
    protected void populateInitialAttachments(Map<String, DataHandler>
7 map) {
        // Initialize attachments from underlying message (lazy
creation).
    }
}

```

- 1 Implements a custom message class, *CustomMessage*, by extending the `org.apache.camel.impl.DefaultMessage` class.
- 2 Typically, you need a default constructor that creates a message with default properties.
- 3 Override the `toString()` method to customize message stringification.
- 4 The `newInstance()` method is called from inside the `MessageSupport.copy()` method. Customization of the `newInstance()` method should focus on copying all of the *custom* properties of the current message instance into the new message instance. The `MessageSupport.copy()` method copies the generic message properties by calling `copyFrom()`.
- 5 The `createBody()` method works in conjunction with the `MessageSupport.getBody()` method to implement lazy access to the message body. By default, the message body is `null`. It is only when the application code tries to access the body (by calling `getBody()`), that the body should be created. The `MessageSupport.getBody()` automatically calls `createBody()`, when the message body is accessed for the first time.
- 6 The `populateInitialHeaders()` method works in conjunction with the header getter and setter methods to implement lazy access to the message headers. This method parses the message to extract any message headers and inserts them into the hash map, `map`. The `populateInitialHeaders()` method is automatically called when a user attempts to access a header (or headers) for the first time (by calling `getHeader()`, `getHeaders()`, `setHeader()`, or `setHeaders()`).
- 7 The `populateInitialAttachments()` method works in conjunction with the attachment getter and setter methods to implement lazy access to the attachments. This method extracts the message attachments and inserts them into the hash map, `map`. The `populateInitialAttachments()` method is automatically called when a user attempts to access an attachment (or attachments) for the first time by calling `getAttachment()`, `getAttachments()`, `getAttachmentNames()`, or `addAttachment()`.

INDEX

Symbols

@Converter, [Implement an annotated converter class](#)

A

AsyncCallback, [Asynchronous processing](#)

asynchronous producer

implementing, [How to implement an asynchronous producer](#)

AsyncProcessor, [Asynchronous processing](#)

auto-discovery

configuration, [Configuring auto-discovery](#)

C

Component

createEndpoint(), [URI parsing](#)

definition, [The Component interface](#)

methods, [Component methods](#)

component prefix, [Component](#)

components, [Component](#)

bean properties, [Define bean properties on your component class](#)

configuring, [Installing and configuring the component](#)

implementation steps, [Implementation steps](#)

installing, [Installing and configuring the component](#)

interfaces to implement, [Which interfaces do you need to implement?](#)

parameter injection, [Parameter injection](#)

Spring configuration, [Configure the component in Spring](#)

Consumer, [Consumer](#)

consumers, [Consumer](#)

event-driven, [Event-driven pattern](#), [Implementation steps](#)

polling, [Polling pattern](#), [Implementation steps](#)

scheduled, [Scheduled poll pattern](#), [Implementation steps](#)

threading, [Overview](#)

D

DefaultComponent

createEndpoint(), [URI parsing](#)

DefaultEndpoint, [Event-driven endpoint implementation](#)

createExchange(), [Event-driven endpoint implementation](#)

createPollingConsumer(), Event-driven endpoint implementation
getCamelConext(), Event-driven endpoint implementation
getComponent(), Event-driven endpoint implementation
getEndpointUri(), Event-driven endpoint implementation

E

Endpoint, [Endpoint](#)

createConsumer(), Endpoint methods
createExchange(), Endpoint methods
createPollingConsumer(), Endpoint methods
createProducer(), Endpoint methods
getCamelContext(), Endpoint methods
getEndpointURI(), Endpoint methods
interface definition, The Endpoint interface
isLenientProperties(), Endpoint methods
isSingleton(), Endpoint methods
setCamelContext(), Endpoint methods

endpoint

event-driven, Event-driven endpoint implementation
scheduled, Scheduled poll endpoint implementation

endpoints, [Endpoint](#)

Exchange, [Exchange](#), [The Exchange interface](#)

copy(), Exchange methods
getExchangeId(), Exchange methods
getIn(), Accessing message headers, Exchange methods
getOut(), Exchange methods
getPattern(), Exchange methods
getProperties(), Exchange methods
getProperty(), Exchange methods
getUnitOfWork(), Exchange methods
removeProperty(), Exchange methods
setExchangeId(), Exchange methods
setIn(), Exchange methods

setOut(), [Exchange methods](#)

setProperty(), [Exchange methods](#)

setUnitOfWork(), [Exchange methods](#)

exchange

in capable, [Testing the exchange pattern](#)

out capable, [Testing the exchange pattern](#)

exchange properties

accessing, [Wrapping the exchange accessors](#)

ExchangeHelper, [The ExchangeHelper Class](#)

getContentType(), [Get the In message's MIME content type](#)

getMandatoryHeader(), [Accessing message headers](#), [Wrapping the exchange accessors](#)

getMandatoryInBody(), [Wrapping the exchange accessors](#)

getMandatoryOutBody(), [Wrapping the exchange accessors](#)

getMandatoryProperty(), [Wrapping the exchange accessors](#)

isInCapable(), [Testing the exchange pattern](#)

isOutCapable(), [Testing the exchange pattern](#)

resolveEndpoint(), [Resolve an endpoint](#)

exchanges, [Exchange](#)

I

in message

MIME type, [Get the In message's MIME content type](#)

M

Message, [Message](#)

getHeader(), [Accessing message headers](#)

message headers

accessing, [Accessing message headers](#)

messages, [Message](#)

P

performer, [Overview](#)

pipeline, [Pipelining model](#)

Processor, [Processor interface](#)

implementing, [Implementing the Processor interface](#)

producer, [Producer](#)

Producer, [Producer](#)

createExchange(), [Producer methods](#)

getEndpoint(), [Producer methods](#)

process(), [Producer methods](#)

producers

asynchronous, [Asynchronous producer](#)

synchronous, [Synchronous producer](#)

S

ScheduledPollEndpoint, [Scheduled poll endpoint implementation](#)

simple processor

implementing, [Implementing the Processor interface](#)

synchronous producer

implementing, [How to implement a synchronous producer](#)

T

type conversion

runtime process, [Type conversion process](#)

type converter

annotating the implementation, [Implement an annotated converter class](#)

discovery file, [Create a TypeConverter file](#)

implementation steps, [How to implement a type converter](#)

master, [Master type converter](#)

packaging, [Package the type converter](#)

slave, [Master type converter](#)

TypeConverter, [Type converter interface](#)

TypeConverterLoader, [Type converter loader](#)

U

`useIntrospectionOnEndpoint()`, [Disabling endpoint parameter injection](#)

W

wire tap pattern, [System Management](#)