

# MicroProfile Reactive Messaging Specification

James Roper, Clement Escoffier, Gordon Hutchison

1.0-RC2, June 26, 2019

# Table of Contents

MicroProfile Reactive Messaging .....	2
Rationale .....	3
Reactive Systems .....	3
On JMS and Message Driven Beans .....	4
Use cases .....	4
Architecture .....	6
Concepts .....	6
Overall architecture .....	6
Channel .....	6
Message .....	7
Message consumption with @Incoming .....	7
Message production with @Outgoing .....	8
Method consuming and producing .....	8
Connectors .....	9
Message stream operation .....	9
Supported CDI scopes .....	10
Supported method signatures .....	11
Methods producing data .....	11
Methods consuming data .....	13
Methods processing data .....	15
Examples of simple method streams .....	20
Examples of methods using Reactive Streams or MicroProfile Reactive Streams Operators type .....	20
Message acknowledgement .....	21
Acknowledgement Examples .....	25
Connector .....	27
Connector concepts .....	28
Configuration .....	29
Acknowledgement .....	30

Specification: MicroProfile Reactive Messaging Specification

Version: 1.0-RC2

Status: Draft

Release: June 26, 2019

Copyright (c) 2018-2019 Contributors to the Eclipse Foundation

Licensed under the Apache License, Version 2.0 (the "License");  
you may not use this file except in compliance with the License.  
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software  
distributed under the License is distributed on an "AS IS" BASIS,  
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.  
See the License for the specific language governing permissions and  
limitations under the License.

# MicroProfile Reactive Messaging

# Rationale

State-of-the-art systems must be able to adapt themselves to emerging needs and requirements, such as market change and user expectations but also fluctuating load and inevitable failures. Leading-edge applications are capable of dynamic and adaptive capabilities aiming to provide *responsive* systems. While microservices aim to offer this agility, HTTP-based connecting tissue tends to fail to provide the required runtime adaptations, especially when facing failures.

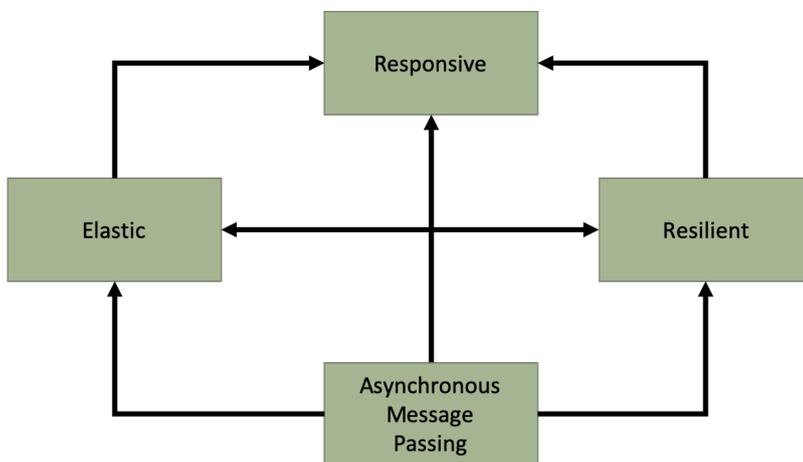
Asynchronous communication allows temporal decoupling of services in a microservice based architecture. This temporal decoupling is necessary if communication is to be enabled to occur regardless of when the parties involved in the communication are running, whether they are loaded or overloaded, and whether they are successfully processing messages or failing.

In contrast, synchronous communication couples services together, binding their uptime, failure, and handling of the load to each other. In a chain of synchronous interactions, the entire conversation can only be successful if all parties in the chain are responsive - if they are all running, processing messages successfully, and not overloaded. If just one party has a problem, all effectively exhibit the same problem. Therefore, systems of microservices relying on synchronous HTTP or relying on synchronous protocols tend to be *fragile*, and failures limit their availability. Indeed, in a microservice-based architecture, temporal coupling results in a fragile system, with resilience and scaling properties that are worse than a monolith, hence, it is essential for microservice based architectures to embrace asynchronous communication as much as possible.

The role of the MicroProfile Reactive Messaging specification is to deliver a way to build systems of microservices promoting both location transparency and temporal decoupling, enforcing asynchronous communication between the different parts of the system.

## Reactive Systems

[Reactive Systems](#) provide an architecture style to deliver *responsive* systems. By infusing asynchronous messaging passing at the core of the system, applications enforcing the reactive system's characteristics are inherently resilient and become more elastic by scaling up and down the number of message consumers.



Microservices as part of reactive systems interact using *messages*. The location and temporal decoupling, promoted by this interaction mechanism, enable numerous benefits such as:

- Better failure handling as the temporal decoupling enables message brokers to resend or reroute messages in the case of remote service failures.
- Improved elasticity as under fluctuating load the system can decide to scale up and down some of the microservices.
- The ability to introduce new features more easily as components are more loosely coupled by receiving and publishing messages.

The MicroProfile Reactive Messaging specification aims to deliver applications embracing the characteristics of reactive systems.

## On JMS and Message Driven Beans

Java EE offers JMS and Message Driven Beans for handling asynchronous communication; however, there are some problems with these specifications:

- Both are designed for a technology landscape where messaging was typically on the edge of the system to hand control of a transaction from one system to another; consequently, these technologies can appear heavyweight when used between microservices.
- It is assumed in their design that consistency is handled using distributed transactions. However, many message brokers, popular in microservice deployments, such as Apache Kafka, Amazon Kinesis and Azure Event Hubs, do not support XA transactions, rather, message acknowledgment is handled using offsets with at least once delivery guarantees.
- They do not have support for asynchronous IO; it is assumed that message processing is done on a single thread, however, many modern specifications are moving to asynchronous IO.

Hence a lighter weight, reactive solution to messaging is desirable for MicroProfile to ensure microservices written using MicroProfile are able to meet the demands required by the architecture.

## Use cases

MicroProfile Reactive Messaging aims to provide a way to connect event-driven microservices. The key characteristics of the specification make it versatile and suitable for building different types of architecture and applications.

First, asynchronous interactions with different services and resources can be implemented using Reactive Messaging. Typically, asynchronous database drivers can be used in conjunction with Reactive Messaging to read and write into a data store in a non-blocking and asynchronous manner.

When building microservices, the CQRS and event-sourcing patterns provide an answer to the data sharing between microservices. Reactive Messaging can also be used as the foundation to CQRS and Event-Sourcing mechanism, as these patterns embrace message-passing as core communication pattern.

IOT applications, dealing with events from various devices, and data streaming applications can also be implemented using Reactive Messaging. The application receives events or messages,

process them, transform them, and may forward them to another microservices. It allows for more fluid architecture for building data-centric applications.

# Architecture

The Reactive Messaging specification defines a development model for declaring CDI *beans* producing, consuming and processing messages. The communication between these components uses Reactive Streams.

This specification relies on [Eclipse MicroProfile Reactive Streams Operators](#) and [CDI](#).

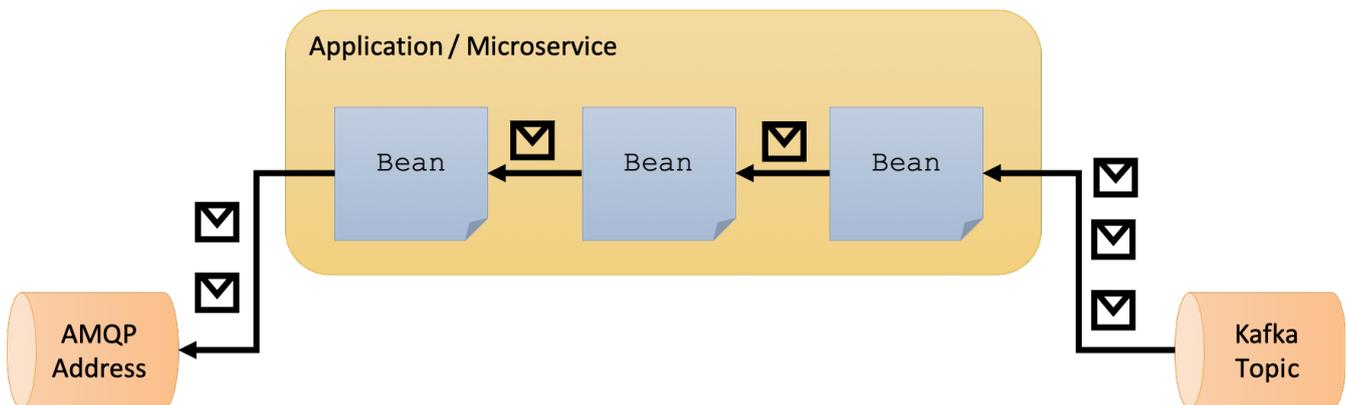
## Concepts

This section describes the different concepts introduced by the Reactive Messaging specification

### Overall architecture

An application using Reactive Messaging is composed of CDI beans consuming, producing and processing *messages*.

These messages can be wholly *internal* to the application or can be sent and received via different message brokers.



Application's beans contain methods annotated with `@Incoming` and `@Outgoing` annotations. A method with an `@Incoming` annotation consumes messages from a *channel*. A method with an `@Outgoing` annotation publishes messages to a *channel*. A method with both an `@Incoming` and an `@Outgoing` annotation is a message processor, it consumes messages from a *channel*, does some transformation to them, and publishes messages to another *channel*.

### Channel

A *channel* is a name indicating which source or destination of messages is used. Channels are opaque *Strings*.

There are two types of channel:

- Internal channels are local to the application. They allow implementing multi-step processing where several *beans* from the same application form a chain of processing.
- Channels can be *connected* to remote brokers or various message transport layers such as Apache Kafka or to an AMQP broker. These channels are managed by *connectors*.

## Message

At the core of the Reactive Messaging specification is the concept of *message*. A *message* is an envelope wrapping a *payload*. A message is sent to a specific channel and, when received and processed successfully, *acknowledged*.

Reactive Messaging application components are addressable recipients which await the arrival of messages on a channel and react to them, otherwise lying dormant.

Messages are represented by the `org.eclipse.microprofile.reactive.messaging.Message` class. This interface is intentionally kept minimal. The aim is that *connectors* will provide their own implementations with additional metadata that is relevant to that connector. For instance, a `KafkaMessage` would provide access to the *topic* and *partition*.

The `org.eclipse.microprofile.reactive.messaging.Message#getPayload` method retrieves the wrapped payload. The `org.eclipse.microprofile.reactive.messaging.Message#ack` method acknowledges the message. Note that the `ack` method is asynchronous as acknowledgement is generally an asynchronous process.

*Plain* messages are created using:

- `org.eclipse.microprofile.reactive.messaging.Message#of(T)` - wraps the given payload, no acknowledgement
- `org.eclipse.microprofile.reactive.messaging.Message#of(T, java.util.function.Supplier<java.util.concurrent.CompletionStage<java.lang.Void>>)` - wraps the given payload and provides the acknowledgment logic

## Message consumption with `@Incoming`

The `org.eclipse.microprofile.reactive.messaging.Incoming` annotation is used on a method from a CDI bean to indicate that the method consumes messages from the specified channel:

```
@Incoming("my-channel")           ①
public CompletionStage<Void> consume(Message<String> message) {       ②
    return message.ack();
}
```

1. `my-channel` is the channel
2. the method is called for every message sent to the `my-channel` channel

Reactive Messaging supports various forms of method signatures. This is detailed in the next section.

Remember that Reactive Messaging interactions are assembled from Reactive Streams. A method annotated with `@Incoming` is a Reactive Streams *subscriber* and so consumes messages that fit with the message signature and its annotations. Note that the handling of the Reactive Streams protocol, such as subscriptions and back pressure, is managed by the Reactive Messaging implementation. The MicroProfile Reactive Streams specification used as a foundation for this version of Reactive

Messaging is a single subscriber model where a stream **Publisher** is connected to a single **Subscriber** which controls back pressure. This implies that a Reactive Messaging channel should appear in a single **@Incoming** annotation. The annotation of more than one **@Incoming** method to be associated with the same channel is not supported and will cause an error during deployment.

From the user perspective, whether the incoming messages comes from co-located beans or a remote message broker is transparent. However, the user may decide to consume a specific subclass of **Message** (e.g. **KafkaMessage** in the following example) if the user is aware of this characteristic:

```
@Incoming("my-kafka-topic")
public CompletionStage<Void> consume(KafkaMessage<String> message) { ①
    return message.ack();
}
```

1. Explicit consumption of a **KafkaMessage**

## Message production with @Outgoing

The **org.eclipse.microprofile.reactive.messaging.Outgoing** annotation is used to annotate a method from a CDI bean to indicate that the method publishes messages to a specified channel:

```
@Outgoing("my-channel") ①
public Message<String> publish() { ②
    return Message.of("hello"); ③
}
```

1. **my-channel** is the targeted channel
2. the method is called for every consumer *request*
3. you can create a *plain* **org.eclipse.microprofile.reactive.messaging.Message** using **org.eclipse.microprofile.reactive.messaging.Message#of(T)**

Reactive Messaging supports various forms of method signatures. This is detailed in the next section.

A method annotated with **@Outgoing** is a Reactive Streams *publisher* and so publishes messages according to the requests it receives. The downstream **@Incoming** method or outgoing connector with a matching channel name will be linked to this publisher. Only a single method can be annotated with **@Outgoing** for a particular channel name. Having the same channel name in more than one **@Outgoing** annotated method is not supported and will result in an error during deployment.

## Method consuming and producing

A method can combine the **@Incoming** and **@Outgoing** annotation and will then act as a Reactive Streams processor:

```
@Incoming("my-incoming-channel") ①
@Outgoing("my-outgoing-channel") ②
public Message<String> process(Message<String> message) {
    return Message.of(message.getPayload().toUpperCase());
}
```

1. The incoming channel
2. The outgoing channel

Having the same channel appear in the `@Outgoing` and `@Incoming` annotations of a processor is not supported and will result in an error during deployment.

## Connectors

The application can receive and forward messages from various message brokers or transport layers. For instance, an application can be connected to a Kafka cluster, an AMQP broker or an MQTT server.

Reactive Messaging *Connectors* are extensions managing the communication with a specific transport technology. They are responsible for mapping a specific *channel* to remote sink or source of messages. This mapping is configured in the application configuration. Note that an implementation may provide various ways to configure the mapping, but support for MicroProfile Config as a configuration source is mandatory.

Connector implementations are associated with a name corresponding to a messaging transport, such as Apache Kafka, Amazon Kinesis, RabbitMQ or Apache ActiveMQ. For instance, an hypothetical Kafka connector could be associated with the following name: `acme.kafka`. This name is indicated using a qualifier on the connector implementation.

The user can associate a channel with this connector using the associated name:

```
mp.messaging.incoming.my-kafka-topic.connector=acme.kafka ①
```

1. the name associated with the connector.

The configuration format is detailed later in this document.

The Reactive Messaging implementation is responsible for finding the connector implementation associated with the given name in the user configuration. If the connector cannot be found, the deployment of the application must be failed.

The Reactive Messaging specification provides an SPI to implement connectors.

## Message stream operation

Message stream operation occurs according to the principles of reactive programming. The back pressure mechanism of reactive streams means that a publisher will not send data to a subscriber

unless there are outstanding subscriber requests. This implies that data flow along the stream is enabled by the first request for data received by the publisher. For methods that are annotated with `@Incoming` and `@Outgoing` this data flow control is handled automatically by the underlying system which will call the `@Incoming` and `@Outgoing` methods as appropriate.

#### NOTE

Although `@Incoming` and `@Outgoing` methods remain callable from Java code, calling them directly will not affect the reactive streams they are associated with. For example, calling an `@Outgoing` annotated method from user code will not post a message on a message queue and calling an `@Incoming` method cannot be used to read a message. Enabling this would bypass the automatic back pressure mechanism that is one of the benefits of the specification. The `@Incoming` and `@Outgoing` method annotations are used to declaratively define the stream which is then run by the implementation of MicroProfile Reactive Messaging without the user's code needing to handle concerns such as subscriptions or flow control within the stream.

## Supported CDI scopes

Implementations of the Reactive Messaging specification must support at least the following CDI scopes:

- `@ApplicationScoped` beans
- `@Dependent` beans

The following code gives an example of a bean annotated with `@ApplicationScoped`:

```
@ApplicationScoped
public class ApplicationScopeBeans {

    @Outgoing("source")
    public Publisher<Integer> source() {
        return ReactiveStreams.of(id).buildRs();
    }

    @Incoming("source")
    @Outgoing("output")
    public int process(int i) {
        return i + 1;
    }

    @Incoming("output")
    public void sink(int v) {
        System.out.println(v);
    }
}
```

Implementations can provide support for other scopes. However the behavior is not defined.

# Supported method signatures

The signature of message stream methods can have a number of different distinct types, offering differing levels of power and simplicity to application developers. Different shapes are supported depending on whether the method is a publisher, subscriber or processor, for example, a publishing stream supports returning MicroProfile Reactive Streams `PublisherBuilder`, but not `SubscriberBuilder`, the inverse is true for a subscribing stream.

This section lists the methods signatures that must be supported by the Reactive Messaging implementation. Implementations must validate that the stream shape matches the `@Outgoing` and `@Incoming` annotations, if they don't, a CDI definition exception should be raised to the CDI container during initialization.

It's important to remember that users must not call these methods directly. They are invoked by the Reactive Messaging implementation following the Reactive Streams protocol.

Also the method must be implemented in a non-blocking fashion. For blocking transformations, asynchronous variants can be used.

**NOTE** *assembly* time is when the Reactive Messaging implementation initializes itself and creates the different bean instances and connects them together.

**NOTE** In the following lists, `Message` can be an implementation of the `Message` interface.

## Methods producing data

Signature	Behavior	Invocation
<code>@Outgoing("name") Publisher&lt;Message&lt;0&gt;&gt; method()</code>	Returns a stream of <code>Message</code> associated with the channel <code>name</code> .	Method called once at assembly time.
<code>@Outgoing("channel") Publisher&lt;0&gt; method()</code>	Returns a stream of <i>payload</i> of type <code>0</code> associated with the channel <code>channel</code> . Produced payloads are mapped to <code>Message&lt;0&gt;</code> by the Reactive Messaging implementation.	Method called once at assembly time.
<code>@Outgoing("channel") PublisherBuilder&lt;Message&lt;0&gt;&gt; method()</code>	Returns a stream of <code>Message</code> associated with the channel <code>channel</code> .	Method called once at assembly time.

Signature	Behavior	Invocation
<pre>@Outgoing("channel") PublisherBuilder&lt;O&gt; method()</pre>	<p>Returns a stream of <i>payload</i> associated with the channel <code>channel</code>. Produced payloads are mapped to <code>Message&lt;O&gt;</code> by the Reactive Messaging implementation.</p>	<p>Method called once at subscription time.</p>
<pre>@Outgoing("channel") Message&lt;O&gt; method()</pre>	<p>Produces an infinite stream of <code>Message</code> associated with the channel <code>channel</code>.</p>	<p>This method is called for each <i>request</i> made by the subscriber.</p>
<pre>@Outgoing("channel") O method()</pre>	<p>Produces an infinite stream of <i>payload</i> associated with the channel <code>channel</code>. Produced payloads are mapped to <code>Message&lt;O&gt;</code> by the Reactive Messaging implementation.</p>	<p>This method is called for each <i>request</i> made by the subscriber.</p>
<pre>@Outgoing("channel") CompletionStage&lt;Message&lt;O&gt;&gt; method()</pre>	<p>Produces an infinite stream of <code>Message</code> associated with the channel <code>channel</code>. The result is a <code>CompletionStage</code>. The method should not be called by the reactive messaging implementation until the <code>CompletionStage</code> returned previously is completed.</p>	<p>This method is called for each <i>request</i> made by the subscriber.</p>

Signature	Behavior	Invocation
<pre>@Outgoing("channel") CompletionStage&lt;O&gt; method()</pre>	<p>Produces an infinite stream of <i>payload</i> associated with the channel <code>channel</code>. Produced payloads are mapped to <code>Message&lt;O&gt;</code> by the Reactive Messaging implementation. The result is a <code>CompletionStage</code>. The method should not be called by the reactive messaging implementation until the <code>CompletionStage</code> returned previously is completed.</p>	<p>This method is called for each <i>request</i> made by the subscriber.</p>

## Methods consuming data

Signature	Behavior	Invocation
<pre>@Incoming("channel") Subscriber&lt;Message&lt;I&gt;&gt; method()</pre>	<p>Returns a <code>Subscriber</code> that receives the <code>Message</code> objects transiting on the channel <code>channel</code>.</p>	<p>The method is called only once to retrieve the <code>Subscriber</code> object at assembly time. This subscriber is connected to the matching channel.</p>
<pre>@Incoming("channel") Subscriber&lt;I&gt; method()</pre>	<p>Returns a <code>Subscriber</code> that receives the <i>payload</i> objects transiting on the channel <code>channel</code>. The payload is automatically extracted from the inflight messages using <code>Message.getPayload()</code>.</p>	<p>The method is called only once to retrieve the <code>Subscriber</code> object at assembly time. This subscriber is connected to the matching channel.</p>
<pre>@Incoming("channel") SubscriberBuilder&lt;Message&lt;I&gt;&gt; method()</pre>	<p>Returns a <code>SubscriberBuilder</code> that receives the <code>Message</code> objects transiting on the channel <code>channel</code>.</p>	<p>The method is called only once at assembly time to retrieve a <code>SubscriberBuilder</code> that is used to build a <code>CompletionSubscriber</code> that is subscribed to the matching channel.</p>

Signature	Behavior	Invocation
<pre data-bbox="140 159 783 293">@Incoming("channel") SubscriberBuilder&lt;I&gt; method()</pre>	<p data-bbox="802 159 1121 566">Returns a <code>SubscriberBuilder</code> that is used to build a <code>CompletionSubscriber&lt;I&gt;</code> that receives the <i>payload</i> of each <code>Message</code>. The payload is automatically extracted from the inflight messages using <code>Message.getPayload()</code>.</p>	<p data-bbox="1137 159 1458 461">The method is called only once at assembly time to retrieve a <code>SubscriberBuilder</code> that is used to build a <code>CompletionSubscriber</code> that is subscribed to the matching channel.</p>
<pre data-bbox="140 589 783 723">@Incoming("channel") void method(I payload)</pre>	<p data-bbox="802 589 1121 775">Consumes the <i>payload</i>. The method can return <code>void</code> or any object or <code>null</code>. The returned value is ignored.</p>	<p data-bbox="1137 589 1458 1149">This method is called for every <code>Message&lt;I&gt;</code> instance transiting on the channel <code>channel</code>. The payload is automatically extracted from the inflight messages using <code>Message.getPayload()</code>. The user method is never called concurrently and so must return before being called with the next payload.</p>
<pre data-bbox="140 1171 783 1305">@Incoming("channel") CompletionStage&lt;?&gt; method(Message&lt;I&gt; msg)</pre>	<p data-bbox="802 1171 1121 1205">Consumes the <code>Message</code></p>	<p data-bbox="1137 1171 1458 1984">This method is called for every <code>Message&lt;I&gt;</code> instance transiting on the channel <code>channel</code>. The user method is never called concurrently. The reactive messaging implementation must wait until the completion of the previously returned <code>CompletionStage</code> before calling the method again with the next <code>Message</code>. Note that <code>@Incoming("channel") void method(Message&lt;I&gt; msg)</code> is not allowed as message acknowledgement is asynchronous.</p>

Signature	Behavior	Invocation
<pre>@Incoming("channel") CompletionStage&lt;?&gt; method(I payload)</pre>	Consumes the <i>payload</i> asynchronously	This method is called for every <code>Message&lt;I&gt;</code> instance transiting on the channel <code>channel</code> . The payload is automatically extracted from the inflight messages using <code>Message.getPayload()</code> . The user method is never called concurrently. The reactive messaging implementation must wait until the completion of the previously returned <code>CompletionStage</code> before calling the method again with the next <i>payload</i> .

## Methods processing data

Signature	Behavior	Invocation
<pre>@Incoming("in") @Outgoing("out") Processor&lt;Message&lt;I&gt;, Message&lt;O&gt;&gt; method()</pre>	Returns a Reactive Streams processor consuming incoming <code>Message</code> instances and produces <code>Message</code> instances.	This method is called once; at assembly time.
<pre>@Incoming("in") @Outgoing("out") Processor&lt;I, O&gt; method();</pre>	Returns a Reactive Streams processor consuming incoming <i>payload</i> instances and produces <i>payload</i> instances.	This method is called once; at assembly time.
<pre>@Incoming("in") @Outgoing("out") ProcessorBuilder&lt;Message&lt;I&gt;, Message&lt;O&gt;&gt; method();</pre>	Returns a <code>ProcessorBuilder</code> consuming incoming <code>Message</code> instances and produces <code>Message</code> instances.	This method is called once; at assembly time.

Signature	Behavior	Invocation
<pre data-bbox="129 152 794 398"> @Incoming("in") @Outgoing("out") ProcessorBuilder&lt;I, O&gt; method(); </pre>	<p data-bbox="794 152 1129 398">Returns a Reactive Streams processor that consuming incoming <i>payload</i> instances and produces <i>payload</i> instances.</p>	<p data-bbox="1129 152 1465 398">This method is called once; at assembly time.</p>
<pre data-bbox="129 398 794 1164"> @Incoming("in") @Outgoing("out") Publisher&lt;Message&lt;O&gt;&gt; method(Message&lt;I&gt; msg) </pre>	<p data-bbox="794 398 1129 1164">Returns a Reactive Streams <b>Publisher</b> for each incoming <b>Message</b>. The returned <b>Publisher</b> can be empty or emits multiple <b>Message</b> instances. If the returned <b>Publisher</b> emits several elements, these elements are <i>flattened</i> in the outgoing stream as a concatenation of elements. The <i>flattening</i> follows the same semantics as the <b>flatMap</b> operator from the MicroProfile Reactive Streams specification.</p>	<p data-bbox="1129 398 1465 1164">This method is called for every incoming message. Implementations must not call the method subsequently until the stream from the previously returned <b>Publisher</b> is completed.</p>

Signature	Behavior	Invocation
<pre data-bbox="135 159 788 331"> @Incoming("in") @Outgoing("out") Publisher&lt;O&gt; method(I payload) </pre>	<p data-bbox="801 159 1123 1216">Returns a Reactive Streams <b>Publisher</b> for each incoming <i>payload</i>. The returned <b>Publisher</b> can be empty or emits multiple <i>payload</i> instances. If the returned <b>Publisher</b> emits several elements, these elements are <i>flattened</i> in the outgoing stream as a concatenation of elements. The <i>flattening</i> follows the same semantics as the <b>flatMap</b> operator from the MicroProfile Reactive Streams specification. The Reactive Messaging implementation must create new <b>Message</b> instances for each emitted <i>payload</i> as well as extracting the payload for each incoming <b>Message</b>.</p>	<p data-bbox="1136 159 1458 499">This method is called for every incoming message. Implementations must not call the method subsequently until the stream from the previously returned <b>Publisher</b> is completed.</p>
<pre data-bbox="135 1238 788 1411"> @Incoming("in") @Outgoing("out") PublisherBuilder&lt;Message&lt;O&gt;&gt; method (Message&lt;I&gt; msg) </pre>	<p data-bbox="801 1238 1123 2016">Returns a <b>PublisherBuilder</b> for each incoming <b>Message</b>. The stream resulting from the built <b>Publisher</b> can be empty or emits multiple <b>Message</b> instances. If the stream emitted from the built <b>Publisher</b> emits several elements, these elements are <i>flattened</i> in the outgoing stream as a concatenation of elements. The <i>flattening</i> follows the same semantics as the <b>flatMap</b> operator from the MicroProfile Reactive Streams specification.</p>	<p data-bbox="1136 1238 1458 1615">This method is called for every incoming message. Implementations must not call the method subsequently until the stream built from the previously returned <b>PublisherBuilder</b> is completed.</p>

Signature	Behavior	Invocation
<pre data-bbox="140 159 783 331"> @Incoming("in") @Outgoing("out") PublisherBuilder&lt;O&gt; method(I payload) </pre>	<p data-bbox="799 159 1125 1261">Returns a <code>PublisherBuilder</code> for each incoming <code>payload</code>. The stream resulting from the built <code>Publisher</code> can be empty or emits multiple <code>payload</code> instances. If the stream emitted from the built <code>Publisher</code> emits several elements, these elements are <i>flattened</i> in the outgoing stream as a concatenation of elements. The <i>flattening</i> follows the same semantics as the <code>flatMap</code> operator from the MicroProfile Reactive Streams specification. The Reactive Messaging implementation must create new <code>Message</code> instances for each emitted <code>payload</code> as well as extracting the payload for each incoming <code>Message</code>.</p>	<p data-bbox="1134 159 1460 539">This method is called for every incoming message. Implementations must not call the method subsequently until the stream built from the previously returned <code>PublisherBuilder</code> is completed.</p>
<pre data-bbox="140 1279 783 1451"> @Incoming("in") @Outgoing("out") Message&lt;O&gt; method(Message&lt;I&gt; msg) </pre>	<p data-bbox="799 1279 1125 1579">Returns a <code>Message</code> for each incoming <code>Message</code>.</p>	<p data-bbox="1134 1279 1460 1579">This method is called for every incoming message. Implementations must not call the method subsequently until the previous call must have returned.</p>
<pre data-bbox="140 1597 783 1769"> @Incoming("in") @Outgoing("out") O method(I payload) </pre>	<p data-bbox="799 1597 1125 1975">Returns a <code>payload</code> for each incoming <code>payload</code>. <i>The Reactive Messaging implementation is responsible for unwrapping the <code>_payload</code> from the incoming <code>Message</code> and creating a <code>Message</code> from the returned <code>payload</code>.</i></p>	<p data-bbox="1134 1597 1460 1899">This method is called for every incoming message. Implementations must not call the method subsequently until the previous call must have returned.</p>

Signature	Behavior	Invocation
<pre data-bbox="129 147 794 622"> @Incoming("in") @Outgoing("out") CompletionStage&lt;Message&lt;O&gt;&gt; method (Message&lt;I&gt; msg) </pre>	<p data-bbox="794 147 1129 622">Produces a <b>Message</b> for each incoming <b>Message</b>. This method returns a <b>CompletionStage</b> that can redeem the <b>Message</b> instance asynchronously. The returned <b>CompletionStage</b> must not be completed with <b>null</b>.</p>	<p data-bbox="1129 147 1465 622">This method is called for every incoming message. Never concurrently. The implementations must wait until the completion of the previously returned <b>CompletionStage</b> before calling the method again with the next <b>Message</b>.</p>
<pre data-bbox="129 622 794 1093"> @Incoming("in") @Outgoing("out") CompletionStage&lt;O&gt; method(I payload) </pre>	<p data-bbox="794 622 1129 1093">Produces a <i>payload</i> for each incoming <i>payload</i>. This method returns a <b>CompletionStage</b> that can redeem the <i>payload</i> instance asynchronously. The returned <b>CompletionStage</b> must not be completed with <b>null</b>.</p>	<p data-bbox="1129 622 1465 1093">This method is called for every incoming <i>payload</i>. Never concurrently. The implementations must wait until the completion of the previously returned <b>CompletionStage</b> before calling the method again with the next <i>payload</i>.</p>
<pre data-bbox="129 1093 794 1413"> @Incoming("in") @Outgoing("out") Publisher&lt;Message&lt;O&gt;&gt; method(Publisher &lt;Message&lt;I&gt;&gt; pub) </pre>	<p data-bbox="794 1093 1129 1413">Applies a transformation to the incoming stream of <b>Message</b>. This method is used to manipulate streams and apply stream transformations.</p>	<p data-bbox="1129 1093 1465 1413">This method is called once, at <i>assembly</i> time.</p>
<pre data-bbox="129 1413 794 1771"> @Incoming("in") @Outgoing("out") PublisherBuilder&lt;Message&lt;O&gt;&gt; method (PublisherBuilder&lt;Message&lt;I&gt;&gt; pub) </pre>	<p data-bbox="794 1413 1129 1771">Applies a transformation to the stream represented by the <b>PublisherBuilder</b> of <b>Message</b>. This method is used to manipulate streams and apply stream transformations.</p>	<p data-bbox="1129 1413 1465 1771">This method is called once, at <i>assembly</i> time.</p>
<pre data-bbox="129 1771 794 2085"> @Incoming("in") @Outgoing("out") Publisher&lt;O&gt; method(Publisher&lt;I&gt; pub) </pre>	<p data-bbox="794 1771 1129 2085">Applies a transformation to the incoming streams of <i>payloads</i>. This method is used to manipulate streams and apply stream transformations.</p>	<p data-bbox="1129 1771 1465 2085">This method is called once, at <i>assembly</i> time.</p>

Signature	Behavior	Invocation
<pre>@Incoming("in") @Outgoing("out") PublisherBuilder&lt;O&gt; method (PublisherBuilder&lt;I&gt; pub)</pre>	<p>Applies a transformation to the stream represented by the <code>PublisherBuilder</code> of <i>payloads</i>. This method is used to manipulate streams and apply stream transformations.</p>	<p>This method is called once, at <i>assembly</i> time.</p>

## Examples of simple method streams

The simplest shape that an application may use is a simple method. This is a method that accepts an incoming message, and possibly publishes an outgoing message:

```
@Incoming("in")
@Outgoing("out")
public Message<O> process(Message<I> msg) {
    return convert(msg);
}
```

In the above example, the stream is both a publishing and subscribing stream, with a 1:1 mapping of incoming to outgoing messages. Asynchronous processing may also be used, by returning a `CompletionStage`:

```
@Incoming("in")
@Outgoing("out")
public CompletionStage<Message<O>> process(Message<I> msg) {
    return asyncConvert(msg);
}
```

If the method is not `@Outgoing` annotated, then the returned value is ignored - however, note that for asynchronous methods, the returned `CompletionStage` is still important for determining when message processing has completed successfully, for the purposes of message acknowledgement. When there is no `@Outgoing` annotation, `void` may also be returned.

In addition to `Message`, implementations must allow:

- payloads (the content wrapped in a `Message`)
- implementation of the `Message` interface

## Examples of methods using Reactive Streams or MicroProfile Reactive Streams Operators types

For more power, developers may use Reactive Streams instances. Reactive Streams shaped methods accept no parameters, and return one of the following:

- `org.eclipse.microprofile.reactive.streams.PublisherBuilder`
- `org.eclipse.microprofile.reactive.streams.SubscriberBuilder`
- `org.eclipse.microprofile.reactive.streams.ProcessorBuilder`
- `org.reactivestreams.Publisher`
- `org.reactivestreams.Subscriber`
- `org.reactivestreams.Processor`

Implementations may optionally support other types, such as JDK9 Flow publishers, subscribers and processors, or other representations of Reactive Streams. Application developers are recommended to use the MicroProfile Reactive Streams Operators *builders* in order to allow for the highest level of portability.

For example, here's a message processor:

```
@Incoming("in")
@Outgoing("out")
public PublisherBuilder<Message<I>, Message<O>> process() {
    return ReactiveStreams.<Message<I>>builder()
        .map(this::convert);
}
```

**NOTE** | Implementations must support implementations of the `Message` interface.

## Message acknowledgement

Acknowledgement is an important part of message processing. Messages are either acknowledged explicitly, or implicitly by the implementation.

Acknowledgement for the `@Incoming` messages is controlled by the `org.eclipse.microprofile.reactive.messaging.Acknowledgment` annotation. The annotation allows configuring the acknowledgement strategy among:

- **NONE** - no acknowledgment is performed
- **MANUAL** - the user is responsible for the acknowledgement, by calling the `Message#ack()` method, so the Reactive Messaging implementation does not apply implicit acknowledgement
- **PRE\_PROCESSING** - the Reactive Messaging implementation acknowledges the message before the annotated method or processing is executed
- **POST\_PROCESSING** - the Reactive Messaging implementation acknowledges the message once:
  1. the method or processing completes if the method does not emit data
  2. when the emitted data is acknowledged

Each method signature type has different defaults and can implement different acknowledgement policies. If the `Acknowledgment` annotation is not set, the default policy is applied.

**IMPORTANT**

Method only annotated with `@Outgoing` do not support acknowledgement as they don't receive an input `Message`.

When a method annotated with `@Incoming` defines its acknowledgement policy to be `PRE_PROCESSING` or `POST_PROCESSING`, the Reactive Messaging implementation is responsible for the acknowledgement of the message. When the `POST_PROCESSING` policy is used, the incoming message is acknowledged when the outgoing message is acknowledged. Thus, it creates a chain of acknowledgements, making sure that the messages produced by an `IncomingConnectorFactory` are only acknowledged when the dispatching of the messages has been completed successfully.

The `NONE` strategy indicates that the incoming message is not acknowledged. The `MANUAL` strategy indicates that the incoming message acknowledgement is managed by the user code. The `MANUAL` strategy is often used to acknowledge incoming messages when the produced messages are acknowledged. For example, in the next snippet, the received `KafkaMessage` is acknowledged when the produced message is acknowledged.

```
@Incoming("data")
@Outgoing("sink")
@Acknowledgment(Acknowledgment.Strategy.MANUAL)
public Message<Integer> process(KafkaMessage<String, Integer> input) {
    return Message.of(processThePayload(input.getPayload(), () -> input.ack()));
}
```

The following table indicates the defaults and supported acknowledgement for each supported signature:

Signature	Default Acknowledgement Strategy	Supported Strategies
<code>@Incoming("channel")</code> Subscriber<Message<I>> method()	Post-Processing	None, Pre-Processing, Post-Processing (when the <code>onNext</code> method returns), Manual
<code>@Incoming("channel")</code> Subscriber<I> method()	Post-Processing	None, Pre-Processing, Post-Processing (when the <code>onNext</code> method returns)
<code>@Incoming("channel")</code> SubscriberBuilder<Message<I>> method()	Post-Processing	None, Pre-Processing, Post-Processing (when the <code>onNext</code> method returns), Manual
<code>@Incoming("channel")</code> SubscriberBuilder<I> method()	Post-Processing	None, Pre-Processing, Post-Processing (when the <code>onNext</code> method returns)

Signature	Default Acknowledgement Strategy	Supported Strategies
<pre>@Incoming("channel") void method(I payload)</pre>	Post-Processing	None, Pre-Processing, Post-Processing (when the method returns)
<pre>@Incoming("channel") CompletionStage&lt;?&gt; method(Message&lt;I&gt; msg)</pre>	Post-Processing	None, Pre-Processing, Post-Processing (when the returned <b>CompletionStage</b> is completed), Manual
<pre>@Incoming("channel") CompletionStage&lt;?&gt; method(I payload)</pre>	Post-Processing	None, Pre-Processing, Post-Processing (when the returned <b>CompletionStage</b> is completed)
<pre>@Incoming("in") @Outgoing("out") Processor&lt;Message&lt;I&gt;, Message&lt;O&gt;&gt; method()</pre>	Pre-Processing	None, Pre-Processing, Manual
<pre>@Incoming("in") @Outgoing("out") Processor&lt;I, O&gt; method();</pre>	Pre-Processing	None, Pre-Processing Post-Processing can be optionally supported by implementations, however it requires a 1:1 mapping between the incoming element and the outgoing element.
<pre>@Incoming("in") @Outgoing("out") ProcessorBuilder&lt;Message&lt;I&gt;, Message&lt;O&gt;&gt; method();</pre>	Pre-Processing	None, Pre-Processing, Manual
<pre>@Incoming("in") @Outgoing("out") ProcessorBuilder&lt;I, O&gt; method();</pre>	Pre-Processing	None, Pre-Processing Post-Processing can be optionally supported by implementations, however it requires a 1:1 mapping the incoming element and the outgoing element.

Signature	Default Acknowledgement Strategy	Supported Strategies
<pre>@Incoming("in") @Outgoing("out") Publisher&lt;Message&lt;0&gt;&gt; method(Message&lt;I&gt; msg)</pre>	Pre-Processing	None, Manual, Pre-Processing
<pre>@Incoming("in") @Outgoing("out") Publisher&lt;0&gt; method(I payload)</pre>	Pre-Processing	None, Pre-Processing
<pre>@Incoming("in") @Outgoing("out") PublisherBuilder&lt;Message&lt;0&gt;&gt; method (Message&lt;I&gt; msg)</pre>	Pre-Processing	None, Manual, Pre-Processing
<pre>@Incoming("in") @Outgoing("out") PublisherBuilder&lt;0&gt; method(I payload)</pre>	Pre-Processing	None, Pre-Processing
<pre>@Incoming("in") @Outgoing("out") Message&lt;0&gt; method(Message&lt;I&gt; msg)</pre>	Pre-Processing	None, Manual, Pre-Processing
<pre>@Incoming("in") @Outgoing("out") 0 method(I payload)</pre>	Post-Processing	None, Pre-Processing, Post-Processing (when the message wrapping the produced payload is acknowledged)
<pre>@Incoming("in") @Outgoing("out") CompletionStage&lt;Message&lt;0&gt;&gt; method (Message&lt;I&gt; msg)</pre>	Pre-Processing	None, Manual, Pre-Processing

Signature	Default Acknowledgement Strategy	Supported Strategies
<pre>@Incoming("in") @Outgoing("out") CompletionStage&lt;O&gt; method(I payload)</pre>	Post-Processing	None, Pre-Processing, Post-Processing (when the message wrapping the produced payload is acknowledged)
<pre>@Incoming("in") @Outgoing("out") Publisher&lt;Message&lt;O&gt;&gt; method(Publisher&lt;Message&lt;I&gt;&gt; pub)</pre>	Pre-Processing	None, Manual, Pre-Processing
<pre>@Incoming("in") @Outgoing("out") PublisherBuilder&lt;Message&lt;O&gt;&gt; method (PublisherBuilder&lt;Message&lt;I&gt;&gt; pub)</pre>	Pre-Processing	None, Manual, Pre-Processing
<pre>@Incoming("in") @Outgoing("out") Publisher&lt;O&gt; method(Publisher&lt;I&gt; pub)</pre>	Pre-Processing	None, Pre-Processing
<pre>@Incoming("in") @Outgoing("out") PublisherBuilder&lt;O&gt; method (PublisherBuilder&lt;I&gt; pub)</pre>	Pre-Processing	None, Pre-Processing

Invalid acknowledgement policies must be detected and a `DeploymentException` raised when the application is deployed.

## Acknowledgement Examples

Transiting data may be wrapped in a `Message`, which can be used to supply metadata, and also allows messages to be acknowledged. The contract for acknowledging messages is anything that accepts a `Message` is required to acknowledge it. So, if the application receives an incoming message wrapped in `Message`, it is responsible for invoking `Message.ack()`, and if the application publish an outgoing message wrapped in `Message`, then the spec implementation is responsible for invoking `Message.ack()`.

For example, the following application code is incorrect, since it accepts a message wrapped in `Message`, but does not acknowledge the messages:

```

@Incoming("in")
@Acknowledgment(Acknowledgment.Strategy.MANUAL)
public void process(Message<I> msg) {
    System.out.println("Got message " + msg.getPayload());
}

```

Here is a correct implementation:

```

@Incoming("in")
@Acknowledgment(Acknowledgment.Strategy.MANUAL)
public CompletionStage<Void> process(Message<I> msg) {
    System.out.println("Got message " + msg.getPayload());
    return msg.ack();
}

```

This implementation is also correct, since the application receives a payload wrapped in a message. It's the implementations responsibility to invoke `ack()` on the incoming message:

```

@Incoming("in")
public void process(I payload) {
    System.out.println("Got payload " + payload);
}

```

When dealing with payloads, the `POST_PROCESSING` strategy is the default strategy. In the following snippet, the incoming payload is transported into a message and unwrapped before calling the method. The produced result is wrapped into another `Message`. Following the `POST_PROCESSING` strategy, the incoming message must only be acknowledged when the output message is acknowledged. The implementation is responsible to chain the acknowledgements.

```

@Incoming("in")
@Outgoing("out")
public O process(I payload) {
    ...
}

```

The acknowledgment strategy can be changed. For instance, using the `PRE_PROCESSING` strategy, the incoming message is acknowledged before the method is called. It also means that the acknowledgment of the outgoing message would not acknowledge the incoming message anymore, as it's already acknowledged.

```

@Incoming("in")
@Outgoing("out")
@Acknowledgment(Acknowledgment.Strategy.PRE_PROCESSING)
public 0 process(I payload) {
    ...
}

```

The **NONE** strategy indicates that the incoming message is not acknowledged and the acknowledgment of the outgoing message would not acknowledge the incoming message anymore. The **NONE** strategy may be used for protocols that do not support acknowledgment.

```

@Incoming("in")
@Outgoing("out")
@Acknowledgment(Acknowledgment.Strategy.NONE)
public 0 process(I payload) {
    ...
}

```

The **MANUAL** strategy indicates that the acknowledgment is managed by the user code. The following snippet is particularly useful for processing messages that are also being sent to a destination, as the implementation must not invoke **ack** until after the outgoing message has been sent to the destination:

```

@Incoming("in")
@Outgoing("out")
@Acknowledgment(Acknowledgment.Strategy.MANUAL)
public Message<O> process(Message<I> msg) {
    return Message.of(convert(msg.getPayload()), msg::ack);
}

```

The implementation is responsible for enforcing the acknowledgement strategy defined by the user when the **@Acknowledgement** policy is used. If the annotation is not used, the default policy must be enforced.

## Connector

Reactive Messaging connects matching **@Incoming** and **@Outgoing** stream elements running inside the same application. Additionally, it maps specific *channels* to *external* technologies such as Apache Kafka, MQTT, Web Sockets, AMQP, or JMS. This means that Reactive Messaging can receive messages from virtually any messaging technology and dispatch messages to any messaging technology. This bridging to an external messaging technology is done using a reactive messaging *connector*.

## Connector concepts

Each *connector* is responsible for a specific technology. A connector can:

- act as a *Publisher*, meaning it retrieves or receives messages from an external messaging technology and publishes them to a reactive stream. The messages will then be sent to a method annotated with `@Incoming`.
- act as a *Subscriber*, meaning it subscribes to a reactive stream and dispatches messages to an external messaging technology. The messages are received from a method annotated with `@Outgoing`.
- handle both directions.

It's essential that connectors implement the back-pressure protocol defined by the Reactive Streams specification.

A *connector* is implemented as a CDI Bean, generally *application* scoped implementing:

- the `org.eclipse.microprofile.reactive.messaging.connector.IncomingConnectorFactory` interface to receive messages from an external source;
- the `org.eclipse.microprofile.reactive.messaging.connector.OutgoingConnectorFactory` interface to dispatch messages to an external sink

### NOTE

Depending on the integrated technology, the *connector* can implement one of the interface or both.

The bean is a factory called by the Reactive Messaging implementation to create `PublisherBuilder` or `SubscriberBuilder` objects. These objects are then connected to methods annotated with `@Incoming` or `@Outgoing`.

Beans implementing the `IncomingConnectorFactory` or `OutgoingConnectorFactory` must use the `org.eclipse.microprofile.reactive.messaging.spi.Connector` qualifier. This qualifier defined the name associated with the connector.

The `@Connector` qualifier is used as follows:

```
package org.eclipse.reactive.sample.kafka;

import org.eclipse.microprofile.reactive.messaging.spi.*;

@ApplicationScoped
@Connector("acme.kafka")
public class KafkaConnector implements IncomingConnectorFactory,
    OutgoingConnectorFactory {
    // ...
}
```

Once defined, the user can, in the configuration, refer to this connector using the given name (`acme.kafka` in this example). When the Reactive Messaging implementation processes the

configuration, it determines the connector to be used based on the `connector` attribute.

## Configuration

Reactive Messaging connectors are configured using MicroProfile Config. The implementation processes the global configuration and determines:

- which channels are defined
- which connectors are used (using the `connector`) attribute
- the configuration for each channel

The builder methods defined in the `IncomingConnectorFactory` and `OutgoingConnectorFactory` receive a `org.eclipse.microprofile.config.Config` as parameter. The `Config` object contains key-value pairs to configure the connector. The configuration is specific to the connector. For example, a Kafka connector expects a `bootstrap.servers` entry as well as a `topic` entry.

The Reactive Messaging implementation reads the global application configuration and must support the following format:

- `mp.messaging.incoming.[channel-name].[attribute]=[value]`
- `mp.messaging.outgoing.[channel-name].[attribute]=[value]`
- `mp.messaging.connector.[connector-name].[attribute]=[value]`

For each extracted `channel-name`:

1. The `connector` attribute of the channel is read, and the connector implementation identified. If no loadable connector implementation matches, the deployment must be failed with a `DeploymentException`;
2. Relevant attributes are those matching either the `channel-name` or the resolved `connector-name`.
3. Relevant attributes are processed to generate a `Config` object containing only `attribute=value` entries. It is valid to have an attribute specified at a connector level and also for a specific channel. If an attribute appears for both a channel and its relevant connector, the channel specific value will be used. In the example below, the `acme.kafka` default value for `bootstrap.servers` is overridden for `my-channel` to be `9096`.

The following snippet gives an example for a hypothetical Kafka connector:

```
mp.messaging.incoming.my-channel.connector=acme.kafka
mp.messaging.incoming.my-channel.bootstrap.servers=localhost:9096
mp.messaging.incoming.my-channel.topic=my-topic
mp.messaging.connector.acme.kafka.bootstrap.servers=localhost:9092
```

For properties that have a `mp.messaging.incoming.` or `mp.messaging.outgoing` prefix, this prefix is stripped off the property name and the remainder of the property name up to the first occurrence of `.` is treated as the channel name. Channel names may not include the `.` character.

For properties that have a `mp.messaging.connector.` prefix, this prefix is stripped off the property

name and the longest remaining prefix that matches any configured `connector` is treated as a connector name. The remainder of the property name, minus the expected initial `.` separator, is taken as the name of an attribute for this connector. For example `bootstrap.servers` appears as a default attribute for all channels that use the `acme.kafka` connector.

The Reactive Messaging implementation:

1. Reads the configuration
2. Identifies that a `my-channel` source needs to be managed
3. Searches for the `connector` attribute and finds `acme.kafka`
4. Looks for a bean implementing the `IncomingConnectorFactory` interface qualified with `@Connector("acme.kafka")`. If the configuration had contained a `mp.messaging.outgoing.my-channel...` entry, a bean implementing the `OutgoingConnectorFactory` interface would have been searched for.
5. Creates a new `Config` object with just the relevant `key=value` pairs:

```
bootstrap.servers=localhost:9096
topic=my-topic
```

6. Calls the `PublisherBuilder<? extends Message> getPublisherBuilder(Config config)` method with the created `Config` object. If the configuration is invalid, the connector can throw:
  - a `NoSuchElementException` if a mandatory attribute is missing in the configuration
  - an `IllegalArgumentException` if the initialization of the connector fails for any other reasons.

The Reactive Messaging implementation catches these exceptions and wraps them into a `DeploymentException`, failing the deployment of the application.

7. The built `PublisherBuilder` is connected to a method using the `@Incoming("my-stream")` annotation. The implementation of the connector must map every received message to an `org.eclipse.microprofile.reactive.messaging.Message`. Optionally, it can provide its own implementation of `org.eclipse.microprofile.reactive.messaging.Message` providing additional metadata.

The configuration passed to the `IncomingConnectorFactory` and `OutgoingConnectorFactory` contains at least the: `* channel-name` attribute indicating the name of the channel being configured, `* connector` attribute indicating the fully qualified name of the connector.

## Acknowledgement

The connector is responsible for the acknowledgment of the incoming and outgoing messages:

- An incoming connector must only acknowledge the received message when the produced `org.eclipse.microprofile.reactive.messaging.Message` is acknowledged.
- An outgoing connector must acknowledge the incoming `org.eclipse.microprofile.reactive.messaging.Message` once it has successfully dispatched the message.