

Spring for Apache Pulsar

Table of Contents

1. Introduction	2
1.1. Minimum Supported Versions	2
1.2. Building the Project	2
2. Reference	3
2.1. Using Spring for Apache Pulsar	3
2.1.1. Quick Tour	3
Dependencies	3
Application Code	3
2.1.2. Pulsar Client	4
Authentication	5
2.1.3. Message Production	7
Pulsar Template	7
Simple API	8
Fluent API	8
Message customization	8
Producer customization	8
Specifying Schema Information	9
Pulsar Producer Factory	9
Pulsar Producer Caching	9
Intercept Messages on the Producer	10
2.1.4. Message Consumption	10
Pulsar Listener	10
Specifying Schema Information	14
Accessing the Pulsar Consumer Object	14
Pulsar Message Listener Container	14
DefaultPulsarMessageListenerContainer	15
ConcurrentPulsarMessageListenerContainer	16
Message Consumption	17
Pulsar Headers	18
Accessing in Single Record based Consumer	18
Accessing in Batch Record based Consumer	19
Message Acknowledgment	19
Message ACK modes	20
Automatic Message Ack in Single Record Mode	20
Manual Message Ack in Single Record Mode	21
Automatic Message Ack in Batch Consumption	23
Manual Message Ack in Batch Consumption	23
Message Redelivery and Error Handling	24

Specifying Acknowledgment Timeout for Message Redelivery	24
Specifying Negative Acknowledgment Redelivery	26
Using Dead Letter Topic from Apache Pulsar for Message Redelivery and Error Handling	26
Native Error Handling in Spring for Apache Pulsar	28
Batch listener with PulsarConsumerErrorHandler	31
2.1.5. Publishing and Consuming Partitioned Topics	32
2.2. Reactive Support	35
2.2.1. Quick Tour	35
Dependencies	35
Application Code	36
2.2.2. Design	37
Apache Pulsar Reactive	37
Additive Auto-Configuration	37
2.2.3. Reactive Pulsar Client	37
Authentication	37
2.2.4. Message Production	37
ReactivePulsarTemplate	37
Fluent API	38
Message customization	38
Sender customization	38
Specifying Schema Information	38
ReactivePulsarSenderFactory	39
Producer Caching	39
2.2.5. Message Consumption	39
@ReactivePulsarListener	39
Streaming	41
Configuration - Application Properties	42
Consumer Customization	42
Specifying Schema Information	43
Message Listener Container Infrastructure	43
ReactivePulsarMessageListenerContainer	43
ReactiveMessagePipeline	43
ReactivePulsarMessageHandler	43
Concurrency	44
Pulsar Headers	44
Accessing In OneByOne Listener	44
Accessing In Streaming Listener	45
Message Acknowledgment	45
OneByOne Listener	45
Streaming Listener	45

Message Redelivery and Error Handling	45
Acknowledgment Timeout	45
Negative Acknowledgment Redelivery Delay	46
Dead Letter Topic	46
Pulsar Reader Support	48
2.3. Pulsar Administration	48
2.3.1. Pulsar Admin Client	48
Authentication	48
2.3.2. Automatic Topic Creation	48
2.4. Observability	49
2.4.1. Micrometer Observations	49
Custom tags	49
Observability - Metrics	49
Listener Observation	49
Template Observation	50
Observability - Spans	51
Listener Observation Span	51
Template Observation Span	51
Manual Configuration without Spring Boot	52
Auto-Configuration with Spring Boot	52
Example Configuration	52
Other Resources	54
Appendices	55
Appendix A: Application Properties	55
Pulsar Client Properties	55
Pulsar Producer Properties	59
Pulsar Consumer Properties	61
Pulsar Administration Properties	64
Pulsar Reactive Sender Properties	65
Pulsar Reactive Consumer Properties	67
Pulsar Reactive Reader Properties	70
Appendix B: Non-GA Versions	71
Appendix C: GraalVM Native Image Support	72

Soby Chacko; Chris Bono; Alexander Preuß; Jay Bryant; Christophe Bornet
(v0.1.0)

© 2022 VMware, Inc.

Copies of this document may be made for your own use and for distribution to others, provided that you do not charge any fee for such copies and further provided that each copy contains this Copyright Notice, whether distributed in print or electronically.

Chapter 1. Introduction

This project provides a basic Spring-friendly API for developing [Apache Pulsar](#) applications.

On a very high level, Spring for Apache Pulsar provides a `PulsarTemplate` for publishing to a Pulsar topic and a `PulsarListener` annotation for consuming from a Pulsar topic. In addition, it also provides various convenience APIs for Spring developers to ramp up their development journey into Apache Pulsar.

1.1. Minimum Supported Versions

The minimum supported versions for the underlying libraries required by the framework are as follows:

Library	Version
Java	17
Apache Pulsar	2.10.0
Spring Boot	3.0.0
Spring Framework	6.0.0
Gradle	7.5

1.2. Building the Project

If you have cloned the project locally, follow these steps to build the project from the source code.

Spring for Apache Pulsar uses Gradle as its build tool. Run the following command to do a full build of the project:

```
./gradlew clean build
```

You can build without running tests by using the following command:

```
./gradlew clean build -x test
```

Chapter 2. Reference

This part of the reference documentation goes through the details of the various components in Spring for Apache Pulsar.

2.1. Using Spring for Apache Pulsar

2.1.1. Quick Tour

We will take a quick tour of Spring for Apache Pulsar by showing a sample Spring Boot application that produces and consumes. This is a complete application and does not require any additional configuration, as long as you have a Pulsar cluster running on the default location - `localhost:6650`.



We recommend using a Spring-Boot-First approach for Spring for Apache Pulsar-based application, as that simplifies things tremendously. To do so, you can add the `spring-pulsar-spring-boot-starter` module as a dependency.

Dependencies

Spring Boot applications need only the `spring-pulsar-spring-boot-starter` dependency. The following listings show how to define the dependency for Maven and Gradle, respectively:

Maven

```
<dependencies>
  <dependency>
    <groupId>org.springframework.pulsar</groupId>
    <artifactId>spring-pulsar-spring-boot-starter</artifactId>
    <version>0.1.0</version>
  </dependency>
</dependencies>
```

Gradle

```
dependencies {
    implementation 'org.springframework.pulsar:spring-pulsar-spring-boot-
starter:0.1.0'
}
```

Application Code

The following listing shows the Spring Boot application case for the example:

```

@SpringBootApplication
public class PulsarBootHelloWorld {

    public static void main(String[] args) {
        SpringApplication.run(PulsarBootHelloWorld.class, args);
    }

    @Bean
    ApplicationRunner runner(PulsarTemplate<String> pulsarTemplate) {
        return (args) -> pulsarTemplate.send("hello-pulsar-topic", "Hello Pulsar
World!");
    }

    @PulsarListener(subscriptionName = "hello-pulsar-sub", topics = "hello-pulsar-
topic")
    void listen(String message) {
        System.out.println("Message Received: " + message);
    }
}

```

Let us quickly go through the higher-level details of this application. Later in the documentation we see these components in much more detail.

In the preceding sample, we heavily rely on Spring Boot auto-configuration. Spring Boot auto-configures several components for our application. It automatically provides a `PulsarClient`, which is used by both the producer and the consumer, for the application.

Spring Boot also auto-configures `PulsarTemplate`, which we inject in the application and start sending records to a Pulsar topic. The application sends messages to a topic named `hello-pulsar`. Note that the application does not specify any schema information, because Spring for Apache Pulsar library automatically infers the schema type from the type of the data that you send.

We use the `PulsarListener` annotation to consume from the `hello-pulsar` topic where we publish the data. `PulsarListener` is a convenience annotation that wraps the message listener container infrastructure in Spring for Apache Pulsar. Behind the scenes, it creates a message listener container to create and manage the Pulsar consumer. As with a regular Pulsar consumer, the default subscription type when using `PulsarListener` is the `Exclusive` mode. As records are published in to the `hello-pulsar` topic, the `PulsarListener` consumes them and prints them on the console. The framework also infers the schema type used from the data type that the `PulsarListener` method uses as the payload — `String`, in this case.

2.1.2. Pulsar Client

When you use the Pulsar Spring Boot Starter, you get the `PulsarClient` auto-configured. This is done through a factory bean called `PulsarClientFactoryBean`, which takes a configuration object called `PulsarClientConfiguration`. By default, the application tries to connect to a local Pulsar instance at `pulsar://localhost:6650`. However, there are many application properties available to configure the client. See the [Appendix](#) for more detail.

Authentication

To connect to a Pulsar cluster that requires authentication, you need to set the `authPluginClassName` and any parameters required by the authentication plugin. You can set the parameters as a single JSON-encoded string or as map of parameter names to parameter values. The following listings show both approaches:

Map

```
spring:
  pulsar:
    client:
      auth-plugin-class-name:
org.apache.pulsar.client.impl.auth.oauth2.AuthenticationOAuth2
      authentication:
        issuer-url: https://auth.server.cloud/
        private-key: file:///Users/some-key.json
        audience: urn:sn:acme:dev:my-instance
```

JSON encoded string

```
spring:
  pulsar:
    client:
      auth-plugin-class-name:
org.apache.pulsar.client.impl.auth.oauth2.AuthenticationOAuth2
      auth-params: "{\"privateKey\":\"file:///Users/some-
key.json\", \"issuerUrl\":\"https://auth.server.cloud/\",
\"audience\":\"urn:sn:acme:dev:my-instance\"}"
```



Using a map is the recommended approach as it is less error-prone and easier to read.

The following listings show how to configure each of the supported authentication mechanisms.

Click [here](#) for Athenz

```
spring:
  pulsar:
    client:
      auth-plugin-class-name:
org.apache.pulsar.client.impl.auth.AuthenticationAthenz
      authentication:
        tenant-domain: ...
        tenant-service: ...
        provider-domain: ...
        private-key: ...
        key-id: ...
      enable-tls: true
      tls-trust-certs-file: /path/to/cacert.pem
```

Click [here](#) for Basic

```
spring:
  pulsar:
    client:
      auth-plugin-class-name:
org.apache.pulsar.client.impl.auth.AuthenticationBasic
      authentication:
        user-id: ...
        password: ...
```

Click [here](#) for OAuth2

```
spring:
  pulsar:
    client:
      auth-plugin-class-name:
org.apache.pulsar.client.impl.auth.oauth2.AuthenticationFactoryOAuth2
      authentication:
        issuer-url: ...
        private-key: ...
        audience: ...
        scope: ...
```

[Click here](#) for Sasl

```
spring:
  pulsar:
    client:
      auth-plugin-class-name:
org.apache.pulsar.client.impl.auth.AuthenticationSasl
      authentication:
        sasl-jaas-client-section-name: ...
        server-type: ...
```

[Click here](#) for Tls

```
spring:
  pulsar:
    client:
      auth-plugin-class-name: org.apache.pulsar.client.impl.auth.AuthenticationTls
      authentication:
        tls-cert-file: /path/to/my-role.cert.pem
        tls-key-file: /path/to/my-role.key-pk8.pem
      enable-tls: true
      tls-trust-certs-file: /path/to/cacert.pem
```

[Click here](#) for Token

```
spring:
  pulsar:
    client:
      auth-plugin-class-name:
org.apache.pulsar.client.impl.auth.AuthenticationToken
      authentication:
        token: some-token-goes-here
```



You can find more information on each of the schemes and their required properties in the official [Pulsar security](#) documentation.

2.1.3. Message Production

Pulsar Template

On the Pulsar producer side, Spring Boot auto-configuration provides a `PulsarTemplate` for publishing records. The template implements an interface called `PulsarOperations` and provides methods to publish records through its contract.

There are two categories of these send API methods: `send` and `sendAsync`. The `send` methods block calls by using the synchronous sending capabilities on the Pulsar producer. They return the `MessageId` of the message that was published once the message is persisted on the broker. The `sendAsync` method calls are asynchronous calls that are non-blocking. They return a `CompletableFuture`, which you can use to asynchronously receive the message ID once the messages are published.

Simple API

The template provides a handful of methods (prefixed with `'send'`) for simple send requests that contain only a message or a destination topic. For more complicated send requests, a fluent API lets you configure more options.



Both `send` and `sendAsync` methods have a variety that allows publishing with only the message. When you do that, the application must provide the topic name by setting the property `spring.pulsar.producer.topic-name`.

Fluent API

The template provides a `fluent builder` to handle more complicated send requests.

Message customization

You can specify a `TypedMessageBuilderCustomizer` to configure the outgoing message. For example, the following code shows how to send a keyed message:

```
template.newMessage(msg)
    .withMessageCustomizer((mb) -> mb.key("foo-msg-key"))
    .send();
```

Producer customization

You can specify a `ProducerBuilderCustomizer` to configure the underlying Pulsar producer builder that ultimately constructs the producer used to send the outgoing message.



Use with caution as this gives full access to the producer builder and invoking some of its methods (such as `create`) may have unintended side effects.

For example, the following code shows how to disable batching and enable chunking:

```
template.newMessage(msg)
    .withProducerCustomizer((pb) -> pb.enableChunking(true).enableBatching(false))
    .send();
```

This other example shows how to use custom routing when publishing records to partitioned topics. Specify your custom `MessageRouter` implementation on the `Producer` builder such as:

```
template.newMessage(msg)
    .withProducerCustomizer((pb) -> pb.messageRouter(messageRouter))
    .send();
```



Note that, when using a `MessageRouter`, the only valid setting for `spring.pulsar.producer.message-routing-mode` is `custom`.

This other example shows how to add a `ProducerInterceptor` that will intercept and mutate messages received by the producer before being published to the brokers:

```
template.newMessage(msg)
    .withProducerCustomizer((pb) -> pb.intercept(interceptor))
    .send();
```

Specifying Schema Information

If you use Java primitive types, the framework auto-detects the schema for you, and you need not specify any schema types for publishing the data. However, if you use any complex types (such as `JSON`, `AVRO`, `PROTOBUF`, and others), you need to set the proper schema type on the `PulsarTemplate` before invoking any send operations, as the following example shows for `JSON`:

```
pulsarTemplate.setSchema(Schema.JSON(Foo.class));
```



Complex Schema types that are currently supported are `JSON`, `AVRO`, `PROTOBUF`, and `KEY_VALUE` w/ `INLINE` encoding.

See the [Appendix](#) for Pulsar producer properties.

Pulsar Producer Factory

The `PulsarTemplate` relies on a `PulsarProducerFactory` to actually create the underlying producer. Spring Boot auto-configuration also provides this producer factory. Additionally, you can configure the factory by specifying any of the available producer-centric application properties. See the [Appendix](#).

Pulsar Producer Caching

Each underlying Pulsar producer consumes resources. To improve performance and avoid

continual creation of producers, the producer factory caches the producers that it creates. They are cached in an LRU fashion and evicted when they have not been used within a configured time period. The [cache key](#) is composed of just enough information to ensure that callers are returned the same producer on subsequent creation requests.

Additionally, you can configure the cache settings by specifying any of the `spring.pulsar.producer.cache` prefixed application properties. See the [Appendix](#).

Intercept Messages on the Producer

Adding a `ProducerInterceptor` lets you intercept and mutate messages received by the producer before they are published to the brokers. To do so, you can pass a list of interceptors into the `PulsarTemplate` constructor. When using multiple interceptors, the order they are applied in is the order in which they appear in the list.

If you use Spring Boot auto-configuration, you can specify the interceptors as Beans. They are passed automatically to the `PulsarTemplate`. Ordering of the interceptors is achieved by using the `@Order` annotation as follows:

```
@Bean
@Order(100)
ProducerInterceptor firstInterceptor() {
    ...
}

@Bean
@Order(200)
ProducerInterceptor secondInterceptor() {
    ...
}
```

2.1.4. Message Consumption

Pulsar Listener

When it comes to Pulsar consumers, we recommend that end-user applications use the `PulsarListener` annotation. To use `PulsarListener`, you need to use the `@EnablePulsar` annotation. When you use Spring Boot support, it automatically enables this annotation and configures all the components necessary for `PulsarListener`, such as the message listener infrastructure (which is responsible for creating the Pulsar consumer). `PulsarMessageListenerContainer` uses a `PulsarConsumerFactory` to create and manage the Pulsar consumer.

This consumer factory is also auto-configured through Spring Boot. See the [Appendix](#) for Pulsar consumer properties.

Let us revisit the `PulsarListener` code snippet we saw in the quick-tour section:

```
@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-pulsar")
public void listen(String message) {
    System.out.println("Message Received: " + message);
}
```

You can further simplify this method:

```
@PulsarListener
public void listen(String message) {
    System.out.println("Message Received: " + message);
}
```

In this most basic form, you must provide the following two properties with their corresponding values:

```
spring.pulsar.consumer:
  topic-names: hello-pulsar
  subscription-name: hello-pulsar-subscription
```

In the `PulsarListener` method shown earlier, we receive the data as `String`, but we do not specify any schema types. Internally, the framework relies on Pulsar's schema mechanism to convert the data to the required type. The framework detects that you expect the `String` type and then infers the schema type based on that information. Then it provides that schema to the consumer. For all the primitive types in Java, the framework does this inference. For any complex types (such as JSON, AVRO, and others), the framework cannot do this inference and the user needs to provide the schema type on the annotation using the `schemaType` property.

The following example shows another `PulsarListener` method, which takes an `Integer`:

```
@PulsarListener(subscriptionName = "my-subscription-1", topics = "my-topic-1")
public void listen(Integer message) {
    System.out.println(message);
}
```

The following `PulsarListener` method shows how we can consume complex types from a topic:

```
@PulsarListener(subscriptionName = "my-subscription-2", topics = "my-topic-2",
    schemaType = SchemaType.JSON)
public void listen(Foo message) {
    System.out.println(message);
}
```

Note the addition of a `schemaType` property on `PulsarListener`. That is because the library is not capable of inferring the schema type from the provided type: `Foo`. We must tell the framework what schema to use.

Let us look at a few more ways.

You can consume the Pulsar message directly:

```
@PulsarListener(subscriptionName = "my-subscription", topics = "my-topic")
public void listen(org.apache.pulsar.client.api.Message<String> message) {
    System.out.println(message.getValue());
}
```

The following example consumes the record by using the Spring messaging envelope:

```
@PulsarListener(subscriptionName = "my-subscription", topics = "my-topic")
public void listen(org.springframework.messaging.Message<String> message) {
    System.out.println(message.getPayload());
}
```

Now let us see how we can consume records in batches. The following example uses `PulsarListener` to consume records in batches as POJOs:

```
@PulsarListener(subscriptionName = "hello-batch-subscription", topics = "hello-batch",
    schemaType = SchemaType.JSON, batch = true)
public void listen(List<Foo> messages) {
    System.out.println("records received : " + messages.size());
    messages.forEach((message) -> System.out.println("record : " + message));
}
```

Note that, in this example, we receive the records as a collection (`List`) of objects. In addition, to enable batch consumption at the `PulsarListener` level, you need to set the `batch` property on the annotation to `true`.

Based on the actual type that the `List` holds, the framework tries to infer the schema to use. If the `List` contains a complex type, you still need to provide the `schemaType` on `PulsarListener`.

The following uses the `Message` envelope provided by the Pulsar Java client:

```
@PulsarListener(subscriptionName = "hello-batch-subscription", topics = "hello-batch", schemaType = SchemaType.JSON, batch = true)
public void listen(List<Message<Foo>> messages) {
    System.out.println("records received :" + messages.size());
    messages.forEach((message) -> System.out.println("record : " +
message.getValue()));
}
```

The following example consumes batch records with an envelope of the Spring messaging `Message` type:

```
@PulsarListener(subscriptionName = "hello-batch-subscription", topics = "hello-batch", schemaType = SchemaType.JSON, batch = true)
public void listen(List<org.springframework.messaging.Message<Foo>> messages) {
    System.out.println("records received :" + messages.size());
    messages.forEach((message) -> System.out.println("record : " +
message.getPayload()));
}
```

Finally, you can also use the `Messages` holder object from Pulsar for the batch listener:

```
@PulsarListener(subscriptionName = "hello-batch-subscription", topics = "hello-batch", schemaType = SchemaType.JSON, batch = true)
public void listen(org.apache.pulsar.client.api.Messages<Foo>> messages) {
    System.out.println("records received :" + messages.size());
    messages.forEach((message) -> System.out.println("record : " +
message.getValue()));
}
```

When you use `PulsarListener`, you can provide Pulsar consumer properties directly on the annotation itself. This is convenient if you do not want to use the Boot configuration properties mentioned earlier or have multiple `PulsarListener` methods.

The following example uses Pulsar consumer properties directly on `PulsarListener`:

```
@PulsarListener(properties = { "subscriptionName=subscription-1", "topicNames=foo-1", "receiverQueueSize=5000" })
void listen(String message) {
}
```



The properties used are direct Pulsar consumer properties, not the `spring.pulsar.consumer` application configuration properties

Specifying Schema Information

As indicated earlier, for Java primitives, the Spring Pulsar framework can infer the proper Schema to use on the `PulsarListener`. However, for more complex types (such as JSON or AVRO), you need to specify the schema type on the annotation.



Complex Schema types that are currently supported are JSON, AVRO, PROTOBUF, and KEY_VALUE w/ INLINE encoding.

Accessing the Pulsar Consumer Object

Sometimes, you need direct access to the Pulsar Consumer object. The following example shows how to get it:

```
@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-pulsar")
public void listen(String message, org.apache.pulsar.client.api.Consumer<String> consumer) {
    System.out.println("Message Received: " + message);
    ConsumerStats stats = consumer.getStats();
    ...
}
```



When accessing the `Consumer` object this way, do NOT invoke any operations that would change the Consumer's cursor position by invoking any receive methods. All such operations must be done by the container.

Pulsar Message Listener Container

Now that we saw the basic interactions on the consumer side through `PulsarListener`. Let us now dive into the inner workings of how `PulsarListener` interacts with the underlying Pulsar consumer. Keep in mind that, for end-user applications, in most scenarios, we recommend using the `PulsarListener` annotation directly for consuming from a Pulsar topic when using Spring for Apache Pulsar, as that model covers a broad set of application use cases. However, it is important to understand how `PulsarListener` works internally. This section goes through those details.

As briefly mentioned earlier, the message listener container is at the heart of message consumption when you use Spring for Apache Pulsar. `PulsarListener` uses the message listener container infrastructure behind the scenes to create and manage the Pulsar consumer. Spring for Apache Pulsar provides the contract for this message listener container through `PulsarMessageListenerContainer`. The default implementation for this message listener container is provided through `DefaultPulsarMessageListenerContainer`. As its name indicates, `PulsarMessageListenerContainer` contains the message listener. The container creates the Pulsar consumer and then runs a separate thread to receive and handle the data. The data is handled by the provided message listener implementation.

The message listener container consumes the data in batch by using the consumer's `batchReceive` method. Once data is received, it is handed over to the selected message listener implementation.

The following message listener types are available when you use Spring for Apache Pulsar.

- [PulsarRecordMessageListener](#)
- [PulsarAcknowledgingMessageListener](#)
- [PulsarBatchMessageListener](#)
- [PulsarBatchAcknowledgingMessageListener](#)

We see the details about these various message listeners in the following sections.

Before doing so, however, let us take a closer look at the container itself.

DefaultPulsarMessageListenerContainer

This is a single consumer-based message listener container. The following listing shows its constructor:

```
public DefaultPulsarMessageListenerContainer(PulsarConsumerFactory<? super T>
pulsarConsumerFactory,
        PulsarContainerProperties pulsarContainerProperties)
}
```

It receives a `PulsarConsumerFactory` (which it uses to create the consumer) and a `PulsarContainerProperties` object (which contains information about the container properties). `PulsarContainerProperties` has the following constructors:

```
public PulsarContainerProperties(String... topics)

public PulsarContainerProperties(Pattern topicPattern)
```

You can provide the topic information through `PulsarContainerProperties` or as a consumer property that is provided to the consumer factory. The following example uses the

DefaultPulsarMessageListenerContainer:

```
Map<String, Object> config = new HashMap<>();
config.put("topics", "my-topic");
PulsarConsumerFactory<String> pulsarConsumerFactory =
DefaultPulsarConsumerFactory<>(pulsarClient, config);

PulsarContainerProperties pulsarContainerProperties = new
PulsarContainerProperties();

pulsarContainerProperties.setMessageListener((PulsarRecordMessageListener<?>)
(consumer, msg) -> {
    });

DefaultPulsarMessageListenerContainer<String> pulsarListenerContainer = new
DefaultPulsarMessageListenerContainer(pulsarConsumerFactory,
    pulsarContainerProperties);

return pulsarListenerContainer;
```

DefaultPulsarMessageListenerContainer creates only a single consumer. If you want to have multiple consumers managed through multiple threads, you need to use **ConcurrentPulsarMessageListenerContainer**.

ConcurrentPulsarMessageListenerContainer

ConcurrentPulsarMessageListenerContainer has the following constructor:

```
public ConcurrentPulsarMessageListenerContainer(PulsarConsumerFactory<? super T>
pulsarConsumerFactory,
    PulsarContainerProperties pulsarContainerProperties)
```

ConcurrentPulsarMessageListenerContainer lets you specify a **concurrency** property through a setter. Concurrency of more than 1 is allowed only on non-exclusive subscriptions (**failover**, **shared**, and **key-shared**). You can only have the default 1 for concurrency when you have an exclusive subscription mode.

The following example enables **concurrency** through the **PulsarListener** annotation for a **failover** subscription.

```
@PulsarListener(topics = "my-topic", subscriptionName = "subscription-1",
                subscriptionType = SubscriptionType.Failover, concurrency = "3")
void listen(String message, Consumer<String> consumer) {
    ...
    System.out.println("Current Thread: " + Thread.currentThread().getName());
    System.out.println("Current Consumer: " + consumer.getConsumerName());
}
```

In the preceding listener, it is assumed that the topic `my-topic` has three partitions. If it is a non-partitioned topic, having concurrency set to `3` does nothing. You get two idle consumers in addition to the main active one. If the topic has more than three partitions, messages are load-balanced across the consumers that the container creates. If you run this `PulsarListener`, you see that messages from different partitions are consumed through different consumers, as implied by the thread name and consumer names printouts in the preceding example.



When you use the `Failover` subscription this way on partitioned topics, Pulsar guarantees message ordering.

The following listing shows another example of `PulsarListener`, but with `Shared` subscription and `concurrency` enabled.

```
@PulsarListener(topics = "my-topic", subscriptionName = "subscription-1",
                subscriptionType = SubscriptionType.Shared, concurrency = "5")
void listen(String message) {
    ...
}
```

In the preceding example, the `PulsarListener` creates five different consumers (this time, we assume that the topic has five partitions).



In this version, there is no message ordering, as `Shared` subscriptions do not guarantee any message ordering in Pulsar.

If you need message ordering and still want a shared subscription types, you need to use the `Key_Shared` subscription type.

Message Consumption

Let us take a look at how the message listener container enables both single-record and batch-based message consumption.

Single Record Consumption

Let us revisit our basic `PulsarListener` for the sake of this discussion:

```
@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-pulsar")
public void listen(String message) {
    System.out.println("Message Received: " + message);
}
```

With this `PulsarListener` method, we essentially ask Spring for Apache Pulsar to invoke the listener method with a single record each time. We mentioned that the message listener container consumes the data in batches using the `batchReceive` method on the consumer. The framework detects that the `PulsarListener`, in this case, receives a single record. This means that, on each invocation of the method, it needs a single record. Although the records are consumed by the message listener container in batches, it iterates through the received batch and invokes the listener method through an adapter for `PulsarRecordMessageListener`. As you can see in the previous section, `PulsarRecordMessageListener` extends from the `MessageListener` provided by the Pulsar Java client, and it supports the basic `received` method.

Batch Consumption

The following example shows the `PulsarListener` consuming records in batches:

```
@PulsarListener(subscriptionName = "hello-batch-subscription", topics = "hello-batch", schemaType = SchemaType.JSON, batch = true)
public void listen4(List<Foo> messages) {
    System.out.println("records received : " + messages.size());
    messages.forEach((message) -> System.out.println("record : " + message));
}
```

When you use this type of `PulsarListener`, the framework detects that you are in batch mode. Since it already received the data in batches by using the Consumer's `batchReceive` method, it hands off the entire batch to the listener method through an adapter for `PulsarBatchMessageListener`.

Pulsar Headers

The Pulsar message metadata can be consumed as Spring message headers. The list of available headers can be found in [PulsarHeaders.java](#).

Accessing in Single Record based Consumer

The following example shows how you can access the various Pulsar Headers in an application that uses the single record mode of consuming:

```
@PulsarListener(topics = "simpleListenerWithHeaders")
void simpleListenerWithHeaders(String data, @Header(PulsarHeaders.MESSAGE_ID)
MessageId messageId,
                                @Header(PulsarHeaders.RAW_DATA) byte[] rawData,
                                @Header("foo") String foo) {

}
```

In the preceding example, we access the values for the `messageId` and `rawData` message metadata as well as a custom message property named `foo`. The Spring `@Header` annotation is used for each header field.

You can also use Pulsar's `Message` as the envelope to carry the payload. When doing so, the user can directly call the corresponding methods on the Pulsar message for retrieving the metadata. However, as a convenience, you can also retrieve it by using the `Header` annotation. Note that you can also use the Spring messaging `Message` envelope to carry the payload and then retrieve the Pulsar headers by using `@Header`.

Accessing in Batch Record based Consumer

In this section, we see how to access the various Pulsar Headers in an application that uses a batch consumer:

```
@PulsarListener(topics = "simpleBatchListenerWithHeaders", batch = true)
void simpleBatchListenerWithHeaders(List<String> data,
                                @Header(PulsarHeaders.MESSAGE_ID) List<MessageId> messageIds,
                                @Header(PulsarHeaders.TOPIC_NAME) List<String> topicNames,
                                @Header("foo") List<String> fooValues) {

}
```

In the preceding example, we consume the data as a `List<String>`. When extracting the various headers, we do so as a `List<>` as well. Spring Pulsar ensures that the headers list corresponds to the data list.

You can also extract headers in the same manner when you use the batch listener and receive payloads as `List<org.apache.pulsar.client.api.Message<?>, org.apache.pulsar.client.api.Messages<?>, or org.springframework.messaging.Message<?>`.

Message Acknowledgment

When you use Spring for Apache Pulsar, the message acknowledgment is handled by the framework, unless opted out by the application. In this section, we go through the details of how the framework takes care of message acknowledgment.

Message ACK modes

Spring for Apache Pulsar provides the following modes for acknowledging messages:

- **BATCH**
- **RECORD**
- **MANUAL**

BATCH acknowledgment mode is the default, but you can change it on the message listener container. In the following sections, we see how acknowledgment works when you use both single and batch versions of **PulsarListener** and how they translate to the backing message listener container (and, ultimately, to the Pulsar consumer).

Automatic Message Ack in Single Record Mode

Let us revisit our basic single message based **PulsarListener**:

```
@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-pulsar")
public void listen(String message) {
    System.out.println("Message Received: " + message);
}
```

It is natural to wonder, how acknowledgment works when you use **PulsarListener**, especially if you are familiar with using Pulsar consumer directly. The answer comes down to the message listener container, as that is the central place in Spring for Apache Pulsar that coordinates all the consumer related activities.

Assuming you are not overriding the default behavior, this is what happens behind the scenes when you use the preceding **PulsarListener**:

1. First, the listener container receives messages as batches from the Pulsar consumer.
2. The received messages are handed down to **PulsarListener** one message at a time.
3. When all the records are handed down to the listener method and successfully processed, the container acknowledges all the messages from the original batch.

This is the normal flow. If any records from the original batch throw an exception, Spring for Apache Pulsar track those records separately. When all the records from the batch are processed, Spring for Apache Pulsar acknowledges all the successful messages and negatively acknowledges (nack) all the failed messages. In other words, when consuming single records by using **PulsarRecordMessageListener** and the default ack mode of **BATCH** is used, the framework waits for all the records received from the **batchReceive** call to process successfully and then calls the **acknowledge** method on the Pulsar consumer. If any particular record throws an exception when invoking the handler method, Spring for Apache Pulsar tracks those records and separately calls **negativeAcknowledge** on those records after the entire batch is processed.

If the application wants the acknowledgment or negative acknowledgment to occur per record, the **RECORD** ack mode can be enabled. In that case, after handling each record, the message is acknowledged if no error and negatively acknowledged if there was an error. The following example enables **RECORD** ack mode on the Pulsar listener:

```
@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-pulsar", ackMode = AckMode.RECORD)
public void listen(String message) {
    System.out.println("Message Received: " + message);
}
```

You can also set the listener property, **spring.pulsar.listener.ack-mode**, to set the ack mode application-wide. When doing this, you need not set this on the **PulsarListener** annotation. In that case, all the **PulsarListener** methods in the application acquire that property.

Manual Message Ack in Single Record Mode

You might not always want the framework to send acknowledgments but, rather, do that directly from the application itself. Spring for Apache Pulsar provides a couple of ways to enable manual message acknowledgments. The following example shows one of them:

```
@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-pulsar", ackMode = AckMode.MANUAL)
public void listen(Message<String> message, Acknowledgment acknowledgment) {
    System.out.println("Message Received: " + message.getValue());
    acknowledgment.acknowledge();
}
```

A few things merit explanation here. First, we enable manual ack mode by setting **ackMode** on **PulsarListener**. When enabling manual ack mode, Spring for Apache Pulsar lets the application inject an **Acknowledgment** object. The framework achieves this by selecting a compatible message listener container: **PulsarAcknowledgingMessageListener** for single record based consumption, which gives you access to an **Acknowledgment** object.

The **Acknowledgment** object provides the following API methods:

```

void acknowledge();

void acknowledge(MessageId messageId);

void acknowledge(List<MessageId> messageIds);

void nack();

void nack(MessageId messageId);

```

You can inject this `Acknowledgment` object into your `PulsarListener` while using `MANUAL` ack mode and then call one of the corresponding methods.

In the preceding `PulsarListener` example, we call a parameter-less `acknowledge` method. This is because the framework knows which `Message` it is currently operating under. When calling `acknowledge()`, you need not receive the payload with the `Message` envelope but, rather, use the target type—`String`, in this example. You can also call a different variant of `acknowledge` by providing the message ID: `acknowledge.acknowledge(message.getMessageId());` When you use `acknowledge(messageId)`, you must receive the payload by using the `Message<?>` envelope.

Similar to what is possible for acknowledging, the `Acknowledgment` API also provides options for negatively acknowledging. See the `nack` methods shown earlier.

You can also call `acknowledge` directly on the Pulsar consumer:

```

@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-
pulsar", ackMode = AckMode.MANUAL)
public void listen(Message<String> message, Consumer<String> consumer) {
    System.out.println("Message Received: " + message.getValue());
    try {
        consumer.acknowledge(message);
    }
    catch (Exception e) {
        ....
    }
}

```

When calling `acknowledge` directly on the underlying consumer, you need to do error handling by yourself. Using the `Acknowledgment` does not require that, as the framework can do that for you. Therefore, you should use the `Acknowledgment` object approach when using manual acknowledgment.



When using manual acknowledgment, it is important to understand that the framework completely stays from any acknowledgment at all. Hence, it is extremely important to think through the right acknowledgment strategies when designing applications.

Automatic Message Ack in Batch Consumption

When you consume records in batches (see “[Message ACK modes](#)”) and you use the default ack mode of **BATCH** is used, when the entire batch is processed successfully, the entire batch is acknowledged. If any records throw an exception, the entire batch is negatively acknowledged. Note that this may not be the same batch that was batched on the producer side. Rather, this is the batch that returned from calling `batchReceive` on the consumer

Consider the following batch listener:

```
@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-pulsar", batch = true)
public void batchListen(List<Foo> messages) {
    for (Foo foo : messages) {
        ...
    }
}
```

When all the messages in the incoming collection (`messages` in this example) are processed, the framework acknowledges all of them.

When consuming in batch mode, **RECORD** is not an allowed ack mode. This might cause an issue, as an application may not want the entire batch to be re-delivered again. In such situations, you need to use the **MANUAL** acknowledgement mode.

Manual Message Ack in Batch Consumption

As seen in the previous section, when **MANUAL** ack mode is set on the message listener container, the framework does not do any acknowledgment, positive or negative. It is entirely up to the application to take care of such concerns. When **MANUAL** ack mode is set, Spring for Apache Pulsar selects a compatible message listener container: `PulsarBatchAcknowledgingMessageListener` for batch consumption, which gives you access to an `Acknowledgment` object. The following are the methods available in the `Acknowledgment` API:

```

void acknowledge();

void acknowledge(MessageId messageId);

void acknowledge(List<MessageId> messageIds);

void nack();

void nack(MessageId messageId);

```

You can inject this `Acknowledgment` object into your `PulsarListener` while using `MANUAL` ack mode. The following listing shows a basic example for a batch based listener:

```

@PulsarListener(subscriptionName = "hello-pulsar-subscription", topics = "hello-
pulsar")
public void listen(List<Message<String>> messages, Acknowledgment acknowledgment)
{
    for (Message<String> message : messages) {
        try {
            ...
            acknowledgment.acknowledge(message.getMessageId());
        }
        catch (Exception e) {
            acknowledgment.nack(message.getMessageId());
        }
    }
}

```

When you use a batch listener, the message listener container cannot know which record it is currently operating upon. Therefore, to manually acknowledge, you need to use one of the overloaded `acknowledge` method that takes a `MessageId` or a `List<MessageId>`. You can also negatively acknowledge with the `MessageId` for the batch listener.

Message Redelivery and Error Handling

Now that we have seen both `PulsarListener` and the message listener container infrastructure and its various functions, let us now try to understand message redelivery and error handling. Apache Pulsar provides various native strategies for message redelivery and error handling. We take a look at them and see how we can use them through Spring for Apache Pulsar.

Specifying Acknowledgment Timeout for Message Redelivery

By default, Pulsar consumers does not redeliver messages unless the consumer crashes, but you can change this behavior by setting an ack timeout on the Pulsar consumer. When you use Spring for Apache Pulsar, you can enable this property by setting the `spring.pulsar.consumer.ack-timeout` Boot

property. If this property has a value above zero and if the Pulsar consumer does not acknowledge a message within that timeout period, the message is redelivered.

You can also specify this property directly as a Pulsar consumer property on the `PulsarListener` itself:

```
@PulsarListener(subscriptionName = "subscription-1", topics = "topic-1"
                 properties = {"ackTimeout=60s"})
public void listen(String s) {
    ...
}
```

When you specify `ackTimeout` (as seen in the preceding `PulsarListener` method), if the consumer does not send an acknowledgement within 60 seconds, the message is redelivered by Pulsar to the consumer.

If you want to specify some advanced backoff options for ack timeout with different delays, you can do the following:

```
@EnablePulsar
@Configuration
class AckTimeoutRedeliveryConfig {

    @PulsarListener(subscriptionName =
"withAckTimeoutRedeliveryBackoffSubscription",
                    topics = "withAckTimeoutRedeliveryBackoff-test-topic",
                    ackTimeoutRedeliveryBackoff = "ackTimeoutRedeliveryBackoff",
                    properties = { "ackTimeout=60s" })
    void listen(String msg) {
        // some long-running process that may cause an ack timeout
    }

    @Bean
    RedeliveryBackoff ackTimeoutRedeliveryBackoff() {
        return
MultiplierRedeliveryBackoff.builder().minDelayMs(1000).maxDelayMs(10 *
1000).multiplier(2)
                                .build();
    }
}
```

In the preceding example, we specify a bean for Pulsar's `RedeliveryBackoff` with a minimum delay of 1 second, a maximum delay of 10 seconds, and a backoff multiplier of 2. After the initial ack timeout occurs, the message redeliveries are controlled through this backoff bean. We provide the

backoff bean to the `PulsarListener` annotation by setting the `ackTimeoutRedeliveryBackoff` property to the actual bean name — `ackTimeoutRedeliveryBackoff`, in this case.

Specifying Negative Acknowledgment Redelivery

When acknowledging negatively, Pulsar consumer lets you specify how the application wants the message to be re-delivered. The default is to redeliver the message in one minute, but you can change it by setting `spring.pulsar.consumer.negative-ack-redelivery-delay`. You can also set it as a consumer property directly on `PulsarListener`, as follows:

```
@PulsarListener(subscriptionName = "subscription-1", topics = "topic-1"
                 properties = {"negativeAckRedeliveryDelay=10ms"})
public void listen(String s) {
    ...
}
```

You can also specify different delays and backoff mechanisms with a multiplier by providing a `RedeliveryBackoff` bean and providing the bean name as the `negativeAckRedeliveryBackoff` property on the `PulsarProducer`, as follows:

```
@EnablePulsar
@Configuration
class NegativeAckRedeliveryConfig {

    @PulsarListener(subscriptionName = "withNegRedeliveryBackoffSubscription",
                    topics = "withNegRedeliveryBackoff-test-topic",
                    negativeAckRedeliveryBackoff = "redeliveryBackoff",
                    subscriptionType = SubscriptionType.Shared)
    void listen(String msg) {
        throw new RuntimeException("fail " + msg);
    }

    @Bean
    RedeliveryBackoff redeliveryBackoff() {
        return
        MultiplierRedeliveryBackoff.builder().minDelayMs(1000).maxDelayMs(10 *
        1000).multiplier(2)
            .build();
    }
}
```

Using Dead Letter Topic from Apache Pulsar for Message Redelivery and Error Handling

Apache Pulsar lets applications use a dead letter topic on consumers with a `Shared` subscription

type. For the **Exclusive** and **Failover** subscription types, this feature is not available. The basic idea is that, if a message is retried a certain number of times (maybe due to an ack timeout or nack redelivery), once the number of retries are exhausted, the message can be sent to a special topic called the dead letter queue (DLQ). Let us see some details around this feature in action by inspecting some code snippets:

```
@EnablePulsar
@Configuration
class DeadLetterPolicyConfig {

    @PulsarListener(id = "deadLetterPolicyListener", subscriptionName =
"deadLetterPolicySubscription",
        topics = "topic-with-dlp", deadLetterPolicy = "deadLetterPolicy",
        subscriptionType = SubscriptionType.Shared, properties = {
"ackTimeout=1s" })
    void listen(String msg) {
        throw new RuntimeException("fail " + msg);
    }

    @PulsarListener(id = "dlqListener", topics = "my-dlq-topic")
    void listenDlq(String msg) {
        System.out.println("From DLQ: " + msg);
    }

    @Bean
    DeadLetterPolicy deadLetterPolicy() {
        return
DeadLetterPolicy.builder().maxRedeliverCount(10).deadLetterTopic("my-dlq-
topic").build();
    }

}
```

First, we have a special bean for **DeadLetterPolicy**, and it is named as **deadLetterPolicy** (it can be any name as you wish). This bean specifies a number of things, such as the max delivery (10, in this case) and the name of the dead letter topic—**my-dlq-topic**, in this case. If you do not specify a DLQ topic name, it defaults to **<topicname>-<subscriptionname>-DLQ** in Pulsar. Next, we provide this bean name to **PulsarListener** by setting the **deadLetterPolicy** property. Note that the **PulsarListener** has a subscription type of **Shared**, as the DLQ feature only works with shared subscriptions. This code is primarily for demonstration purposes, so we provide an **ackTimeout** value of 1 second. The idea is that the code throws the exception and, if Pulsar does not receive an ack within 1 second, it does a retry. If that cycle continues ten times (as that is our max redelivery count in the **DeadLetterPolicy**), the Pulsar consumer publishes the messages to the DLQ topic. We have another **PulsarListener** that listens on the DLQ topic to receive data as it is published to the DLQ topic.

Special note on DLQ topics when using partitioned topics

If the main topic is partitioned, behind the scenes, each partition is treated as a separate topic by Pulsar. Pulsar appends `partition-<n>`, where `n` stands for the partition number to the main topic name. The problem is that, if you do not specify a DLQ topic (as opposed to what we did above), Pulsar publishes to a default topic name that has this `'partition-<n>` info in it—for example: `topic-with-dlp-partition-0-deadLetterPolicySubscription-DLQ`. The easy way to solve this is to provide a DLQ topic name always.

Native Error Handling in Spring for Apache Pulsar

As we noted earlier, the DLQ feature in Apache Pulsar works only for shared subscriptions. What does an application do if it needs to use some similar feature for non-shared subscriptions? The main reason Pulsar does not support DLQ on exclusive and failover subscriptions is because those subscription types are order-guaranteed. Allowing redeliveries, DLQ, and so on effectively receives messages out of order. However, what if an application are okay with that but, more importantly, needs this DLQ feature for non-shared subscriptions? For that, Spring for Apache Pulsar provides a `PulsarConsumerErrorHandler`, which you can use across any subscription types in Pulsar: `Exclusive`, `Failover`, `Shared`, or `Key_Shared`.

When you use `PulsarConsumerErrorHandler` from Spring for Apache Pulsar, make sure not to set the ack timeout properties on the listener.

Let us see some details by examining a few code snippets:

```

@EnablePulsar
@Configuration
class PulsarConsumerErrorHandlerConfig {

    @Bean
    PulsarConsumerErrorHandler<String> pulsarConsumerErrorHandler(
        PulsarTemplate<String> pulsarTemplate) {
        return new DefaultPulsarConsumerErrorHandler<>(
            new PulsarDeadLetterPublishingRecoverer<>(pulsarTemplate, (c, m)
-> "my-foo-dlt"), new FixedBackOff(100, 10));
    }

    @PulsarListener(id = "pulsarConsumerErrorHandler-id", subscriptionName =
"pulsarConsumerErrorHandler-subscription",
        topics = "pulsarConsumerErrorHandler-topic",
        pulsarConsumerErrorHandler = "pulsarConsumerErrorHandler")
    void listen(String msg) {
        throw new RuntimeException("fail " + msg);
    }

    @PulsarListener(id = "pceh-dltListener", topics = "my-foo-dlt")
    void listenDlt(String msg) {
        System.out.println("From DLT: " + msg);
    }
}

```

Consider the `pulsarConsumerErrorHandler` bean. This creates a bean of type `PulsarConsumerErrorHandler` and uses the default implementation provided out of the box by Spring for Apache Pulsar: `DefaultPulsarConsumerErrorHandler`. `DefaultPulsarConsumerErrorHandler` has a constructor that takes a `PulsarMessageRecovererFactory` and a `org.springframework.util.backoff.Backoff`. `PulsarMessageRecovererFactory` is a functional interface with the following API:

```

@FunctionalInterface
public interface PulsarMessageRecovererFactory<T> {

    /**
     * Provides a message recoverer {@link PulsarMessageRecoverer}.
     * @param consumer Pulsar consumer
     * @return {@link PulsarMessageRecoverer}.
     */
    PulsarMessageRecoverer<T> recovererForConsumer(Consumer<T> consumer);
}

```

The `recovererForConsumer` method takes a Pulsar consumer and returns a `PulsarMessageRecoverer`, which is another functional interface. Here is the API of `PulsarMessageRecoverer`:

```
public interface PulsarMessageRecoverer<T> {

    /**
     * Recover a failed message, for e.g. send the message to a DLT.
     * @param message Pulsar message
     * @param exception exception from failed message
     */
    void recoverMessage(Message<T> message, Exception exception);

}
```

Spring for Apache Pulsar provides an implementation for `PulsarMessageRecovererFactory` called `PulsarDeadLetterPublishingRecoverer` that provides a default implementation that can recover the message by sending it to a Dead Letter Topic (DLT). We provide this implementation to the constructor for the preceding `DefaultPulsarConsumerErrorHandler`. As the second argument, we provide a `FixedBackOff`. You can also provide the `ExponentialBackoff` from Spring for advanced backoff features. Then we provide this bean name for the `PulsarConsumerErrorHandler` as a property to the `PulsarListener`. The property is called `pulsarConsumerErrorHandler`. Each time the `PulsarListener` method fails for a message, it gets retried. The number of retries are controlled by the `Backoff` provided implementation values. In our example, we do 10 retries (11 total tries — the first one and then the 10 retries). Once all the retries are exhausted, the message is sent to the DLT topic.

The `PulsarDeadLetterPublishingRecoverer` implementation we provide uses a `PulsarTemplate` that is used for publishing the message to the DLT. In most cases, the same auto-configured `PulsarTemplate` from Spring Boot is sufficient with the caveat for partitioned topics. When using partitioned topics and using custom message routing for the main topic, you must use a different `PulsarTemplate` that does not take the auto-configured `PulsarProducerFactory` that is populated with a value of `custompartition` for `message-routing-mode`. You can use a `PulsarConsumerErrorHandler` with the following blueprint:

```

@Bean
PulsarConsumerErrorHandler<Integer> pulsarConsumerErrorHandler(PulsarClient
pulsarClient) {
    PulsarProducerFactory<Integer> pulsarProducerFactory = new
DefaultPulsarProducerFactory<>(pulsarClient, Map.of());
    PulsarTemplate<Integer> pulsarTemplate = new
PulsarTemplate<>(pulsarProducerFactory);

    BiFunction<Consumer<?>, Message<?>, String> destinationResolver =
        (c, m) -> "my-foo-dlt";

    PulsarDeadLetterPublishingRecoverer<Integer>
pulsarDeadLetterPublishingRecoverer =
        new PulsarDeadLetterPublishingRecoverer<>(pulsarTemplate,
destinationResolver);

    return new
DefaultPulsarConsumerErrorHandler<>(pulsarDeadLetterPublishingRecoverer,
        new FixedBackOff(100, 5));
}

```

Note that we are provide a destination resolver to the `PulsarDeadLetterPublishingRecoverer` as the second constructor argument. If not provided, `PulsarDeadLetterPublishingRecoverer` uses `<subscription-name>-<topic-name>-DLT` as the DLT topic name. When using this feature, you should use a proper destination name by setting the destination resolver rather than using the default.

When using a single record message listener, as we did with `PulsarConsumerErrorHnadler`, and if you use manual acknowledgement, make sure to not negatively acknowledge the message when an exception is thrown. Rather, re-throw the exception back to the container. Otherwise, the container thinks the message is handled separately, and the error handling is not triggered.

Finally, we have a second `PulsarListener` that receives messages from the DLT topic.

In the examples provided in this section so far, we only saw how to use `PulsarConsumerErrorHandler` with a single record message listener. Next, we look at how you can use this on batch listeners.

Batch listener with PulsarConsumerErrorHandler

First, let us look at a batch `PulsarListener` method:

```

@PulsarListener(subscriptionName = "batch-demo-5-sub", topics = "batch-demo-4",
batch = true, concurrency = "3",
    subscriptionType = SubscriptionType.Failover,
    pulsarConsumerErrorHandler = "pulsarConsumerErrorHandler", ackMode =
AckMode.MANUAL)
void listen(List<Message<Integer>> data, Consumer<Integer> consumer,
Acknowledgment acknowledgment) {
    for (Message<Integer> datum : data) {
        if (datum.getValue() == 5) {
            throw new PulsarBatchListenerFailedException("failed", datum);
        }
        acknowledgment.acknowledge(datum.getMessageId());
    }
}

@Bean
PulsarConsumerErrorHandler<String> pulsarConsumerErrorHandler(
    PulsarTemplate<String> pulsarTemplate) {
    return new DefaultPulsarConsumerErrorHandler<>(
        new PulsarDeadLetterPublishingRecoverer<>(pulsarTemplate, (c, m) ->
"my-foo-dlt"), new FixedBackOff(100, 10));
}

@PulsarListener(subscriptionName = "my-dlt-subscription", topics = "my-foo-dlt")
void dltReceiver(Message<Integer> message) {
    System.out.println("DLT - RECEIVED: " + message.getValue());
}

```

Once again, we provide the `pulsarConsumerErrorHandler` property with the `PulsarConsumerErrorHandler` bean name. When you use a batch listener (as shown in the preceding example) and want to use the `PulsarConsumerErrorHandler` from Spring for Apache Pulsar, you need to use manual acknowledgment. This way, you can acknowledge all the successful individual messages. For the ones that fail, you must throw a `PulsarBatchListenerFailedException` with the message on which it fails. Without this exception, the framework does not know what to do with the failure. On retry, the container sends a new batch of messages, starting with the failed message to the listener. If it fails again, it is retried, until the retries are exhausted, at which point the message is sent to the DLT. At that point, the message is acknowledged by the container, and the listener is handed over with the subsequent messages in the original batch.

2.1.5. Publishing and Consuming Partitioned Topics

In the following example, we publish to a topic called `hello-pulsar-partitioned`. It is a topic that is partitioned, and, for this sample, we assume that the topic is already created with three partitions.

```

@SpringBootApplication
public class PulsarBootPartitioned {

    public static void main(String[] args) {
        SpringApplication.run(PulsarBootPartitioned.class, "--spring.pulsar.producer.message-routing-mode=CustomPartition");
    }

    @Bean
    public ApplicationRunner runner(PulsarTemplate<String> pulsarTemplate) {
        pulsarTemplate.setDefaultTopicName("hello-pulsar-partitioned");
        return args -> {
            for (int i = 0; i < 10; i++) {
                pulsarTemplate.sendAsync("hello john doe 0 ", new FooRouter());
                pulsarTemplate.sendAsync("hello alice doe 1", new BarRouter());
                pulsarTemplate.sendAsync("hello buzz doe 2", new BuzzRouter());
            }
        };
    }

    @PulsarListener(subscriptionName = "hello-pulsar-partitioned-subscription",
        topics = "hello-pulsar-partitioned")
    public void listen(String message) {
        System.out.println("Message Received: " + message);
    }

    static class FooRouter implements MessageRouter {

        @Override
        public int choosePartition(Message<?> msg, TopicMetadata metadata) {
            return 0;
        }
    }

    static class BarRouter implements MessageRouter {

        @Override
        public int choosePartition(Message<?> msg, TopicMetadata metadata) {
            return 1;
        }
    }

    static class BuzzRouter implements MessageRouter {

        @Override
        public int choosePartition(Message<?> msg, TopicMetadata metadata) {
            return 2;
        }
    }
}

```

```
}
```

In the preceding example, we publish to a partitioned topic, and we would like to publish some data segment to a specific partition. If you leave it to Pulsar's default, it follows a round-robin mode of partition assignments, and we would like to override that. To do so, we provide a message router object with the `send` method. Consider the three message routers implemented. `FooRouter` always sends data to partition 0, `BarRouter` sends to partition 1, and `BuzzRouter` sends to partition 2. Also note that we now use the `sendAsync` method of `PulsarTemplate` that returns a `CompletableFuture`. When running the application, we also need to set the `messageRoutingMode` on the producer to `CustomPartition` (`spring.pulsar.producer.message-routing-mode`).

On the consumer side, we use a `PulsarListener` with the exclusive subscription type. This means that data from all the partitions ends up in the same consumer and there is no ordering guarantee.

What can we do if we want each partition to be consumed by a single distinct consumer? We can switch to the `failover` subscription mode and add three separate consumers:

```
@PulsarListener(subscriptionName = "hello-pulsar-partitioned-subscription", topics = "hello-pulsar-partitioned", subscriptionType = SubscriptionType.Failover)
public void listen1(String foo) {
    System.out.println("Message Received 1: " + foo);
}

@PulsarListener(subscriptionName = "hello-pulsar-partitioned-subscription", topics = "hello-pulsar-partitioned", subscriptionType = SubscriptionType.Failover)
public void listen2(String foo) {
    System.out.println("Message Received 2: " + foo);
}

@PulsarListener(subscriptionName = "hello-pulsar-partitioned-subscription", topics = "hello-pulsar-partitioned", subscriptionType = SubscriptionType.Failover)
public void listen3(String foo) {
    System.out.println("Message Received 3: " + foo);
}
```

When you follow this approach, a single partition always gets consumed by a dedicated consumer.

In a similar vein, if you want to use Pulsar's shared consumer type, you can use the `shared` subscription type. However, when you use the `shared` mode, you lose any ordering guarantees, as a single consumer may receive messages from all the partitions before another consumer gets a chance.

Consider the following example:

```

@PulsarListener(subscriptionName = "hello-pulsar-shared-subscription", topics =
"hello-pulsar-partitioned", subscriptionType = SubscriptionType.Shared)
public void listen1(String foo) {
    System.out.println("Message Received 1: " + foo);
}

@PulsarListener(subscriptionName = "hello-pulsar-shared-subscription", topics =
"hello-pulsar-partitioned", subscriptionType = SubscriptionType.Shared)
public void listen2(String foo) {
    System.out.println("Message Received 2: " + foo);
}

```

2.2. Reactive Support

The framework provides a Reactive counterpart for almost all supported features.



If you put the word **Reactive** in front of a provided imperative component, you will likely find its Reactive counterpart.

- **PulsarTemplate** → **ReactivePulsarTemplate**
- **PulsarListener** → **ReactivePulsarListener**
- **PulsarConsumerFactory** → **ReactivePulsarConsumerFactory**
- etc..

However, the following is not yet supported:

- Error Handling in non-shared subscriptions
- Accessing Pulsar headers via **@Header** in streaming mode
- Observations

2.2.1. Quick Tour

We will take a quick tour of the Reactive support in Spring for Apache Pulsar by showing a sample Spring Boot application that produces and consumes in a Reactive fashion. This is a complete application and does not require any additional configuration, as long as you have a Pulsar cluster running on the default location - **localhost:6650**.



We recommend using a Spring-Boot-First approach for Spring for Apache Pulsar-based applications, as that simplifies things tremendously. To do so, you can add the **spring-pulsar-reactive-spring-boot-starter** module as a dependency.

Dependencies

Spring Boot applications need only the **spring-pulsar-reactive-spring-boot-starter** dependency.

The following listings show how to define the dependency for Maven and Gradle, respectively:

Maven

```
<dependencies>
  <dependency>
    <groupId>org.springframework.pulsar</groupId>
    <artifactId>spring-pulsar-reactive-spring-boot-starter</artifactId>
    <version>0.1.0</version>
  </dependency>
</dependencies>
```

Gradle

```
dependencies {
    implementation 'org.springframework.pulsar:spring-pulsar-reactive-spring-boot-
starter:0.1.0'
}
```

Application Code

Here is the application source code:

```
@SpringBootApplication
public class ReactiveSpringPulsarHelloWorld {

    public static void main(String[] args) {
        SpringApplication.run(ReactiveSpringPulsarHelloWorld.class, args);
    }

    @Bean
    ApplicationRunner runner(ReactivePulsarTemplate<String> pulsarTemplate) {
        return (args) -> pulsarTemplate.send("hello-pulsar-topic", "Hello Reactive
Pulsar World!").subscribe();
    }

    @ReactivePulsarListener(subscriptionName = "hello-pulsar-sub", topics = "hello-
pulsar-topic")
    Mono<Void> listen(String message) {
        System.out.println("Reactive listener received: " + message);
        return Mono.empty();
    }
}
```

That is it, with just a few lines of code we have a working Spring Boot app that is producing and consuming messages from a Pulsar topic in a Reactive fashion.

Once started, the application uses a `ReactivePulsarTemplate` to send messages to the `hello-pulsar-topic`. It then consumes from the `hello-pulsar-topic` using a `@ReactivePulsarListener`.



One of the key ingredients to the simplicity is the Spring Boot starter which auto-configures and provides the required components to the application

2.2.2. Design

Here are a few key design points to keep in mind.

Apache Pulsar Reactive

The reactive support is ultimately provided by the [Apache Pulsar Reactive client](#) whose current implementation is a fully non-blocking adapter around the regular Pulsar client's asynchronous API. This implies that the Reactive client requires the regular client.

Additive Auto-Configuration

Due to the dependence on the regular (imperative) client, the Reactive auto-configuration provided by the framework is additive to the imperative auto-configuration. In other words, The imperative starter only includes the imperative components but the reactive starter includes both imperative and reactive components.

2.2.3. Reactive Pulsar Client

When you use the Reactive Pulsar Spring Boot Starter, you get the `ReactivePulsarClient` auto-configured. By default, the application tries to connect to a local Pulsar instance at `pulsar://localhost:6650`. However, there are many application properties (inherited from the adapted imperative client) available to configure.

See the [Appendix](#) for properties prefixed with `spring.pulsar.client`.

Authentication

To connect to a Pulsar cluster that requires authentication, follow [the same steps](#) as the imperative client. Again, this is because the reactive client adapts the imperative client which handles all security configuration.

2.2.4. Message Production

ReactivePulsarTemplate

On the Pulsar producer side, Spring Boot auto-configuration provides a `ReactivePulsarTemplate` for publishing records. The template implements an interface called `ReactivePulsarOperations` and provides methods to publish records through its contract.

The template provides send methods that accept a single message and return a `Mono<MessageId>`. It also provides send methods that accept multiple messages (in the form of the `ReactiveStreams Publisher` type) and return a `Flux<MessageId>`.



The send methods that do not have a topic input parameter require the topic name to be provided via the property `spring.pulsar.reactive.sender.topic-name`.

Fluent API

The template provides a [fluent builder](#) to handle more complicated send requests.

Message customization

You can specify a [MessageSpecBuilderCustomizer](#) to configure the outgoing message. For example, the following code shows how to send a keyed message:

```
template.newMessage(msg)
    .withMessageCustomizer((mc) -> mc.key("foo-msg-key"))
    .send();
```

Sender customization

You can specify a [ReactiveMessageSenderBuilderCustomizer](#) to configure the underlying Pulsar sender builder that ultimately constructs the sender used to send the outgoing message.



Use with caution as this gives full access to the sender builder and invoking some of its methods (such as [create](#)) may have unintended side effects.

For example, the following code shows how to disable batching and enable chunking:

```
template.newMessage(msg)
    .withSenderCustomizer((sc) -> sc.enableChunking(true).enableBatching(false))
    .send();
```

This other example shows how to use custom routing when publishing records to partitioned topics. Specify your custom [MessageRouter](#) implementation on the sender builder such as:

```
template.newMessage(msg)
    .withSenderCustomizer((sc) -> sc.messageRouter(messageRouter))
    .send();
```



Note that, when using a [MessageRouter](#), the only valid setting for `spring.pulsar.reactive.sender.message-routing-mode` is `custom`.

Specifying Schema Information

If you use Java primitive types, the framework auto-detects the schema for you, and you need not specify any schema types for publishing the data. However, if you use any complex types (such as [JSON](#), [AVRO](#), [PROTOBUF](#), and others), you need to set the proper schema type on the

`ReactivePulsarTemplate` before invoking any send operations, as the following example shows for JSON:

```
template.setSchema(Schema.JSON(Foo.class));
```



Complex Schema types that are currently supported are JSON, AVRO, PROTOBUF, and KEY_VALUE w/ INLINE encoding.

ReactivePulsarSenderFactory

The `ReactivePulsarTemplate` relies on a `ReactivePulsarSenderFactory` to actually create the underlying sender.

Spring Boot provides this sender factory which can be configured with any of the `spring.pulsar.reactive.sender` prefixed application properties.

Producer Caching

Each underlying Pulsar producer consumes resources. To improve performance and avoid continual creation of producers, the `ReactiveMessageSenderCache` in the underlying Apache Pulsar Reactive client caches the producers that it creates. They are cached in an LRU fashion and evicted when they have not been used within a configured time period.

You can configure the cache settings by specifying any of the `spring.pulsar.reactive.sender.cache` prefixed application properties.

2.2.5. Message Consumption

@ReactivePulsarListener

When it comes to Pulsar consumers, we recommend that end-user applications use the `ReactivePulsarListener` annotation. To use `ReactivePulsarListener`, you need to use the `@EnableReactivePulsar` annotation. When you use Spring Boot support, it automatically enables this annotation and configures all necessary components, such as the message listener infrastructure (which is responsible for creating the underlying Pulsar consumer).

Let us revisit the `ReactivePulsarListener` code snippet we saw in the quick-tour section:

```
@ReactivePulsarListener(subscriptionName = "hello-pulsar-sub", topics = "hello-pulsar-topic")
Mono<Void> listen(String message) {
    System.out.println(message);
    return Mono.empty();
}
```



The listener method returns a `Mono<Void>` to signal whether the message was successfully processed. `Mono.empty()` indicates success (acknowledgment) and `Mono.error()` indicates failure (negative acknowledgment).

You can also further simplify this method:

```
@ReactivePulsarListener
Mono<Void> listen(String message) {
    System.out.println(message);
    return Mono.empty();
}
```

In this most basic form, you must still provide the topic name by setting the following property:

```
spring.pulsar.reactive.consumer:
  topic-names: hello-pulsar-topic
```



If `subscription-name` is not provided an auto-generated subscription name will be used.

In the `ReactivePulsarListener` method shown earlier, we receive the data as `String`, but we do not specify any schema types. Internally, the framework relies on Pulsar's schema mechanism to convert the data to the required type. The framework detects that you expect the `String` type and then infers the schema type based on that information. Then it provides that schema to the consumer. For all the primitive types in Java, the framework does this inference. For any complex types (such as JSON, AVRO, and others), the framework cannot do this inference and the user needs to provide the schema type on the annotation using the `schemaType` property.

This example shows how we can consume complex types from a topic:

```
@ReactivePulsarListener(topics = "my-topic-2", schemaType = SchemaType.JSON)
Mono<Void> listen(Foo message) {
    System.out.println(message);
    return Mono.empty();
}
```

Note the addition of a `schemaType` property on `ReactivePulsarListener`. That is because the library is not capable of inferring the schema type from the provided type: `Foo`. We must tell the framework what schema to use.

Let us look at a few more ways we can consume.

This example consumes the Pulsar message directly:

```
@ReactivePulsarListener(topics = "my-topic")
Mono<Void> listen(org.apache.pulsar.client.api.Message<String> message) {
    System.out.println(message.getValue());
    return Mono.empty();
}
```

This example consumes the record wrapped in a Spring messaging envelope:

```
@ReactivePulsarListener(topics = "my-topic")
Mono<Void> listen(org.springframework.messaging.Message<String> message) {
    System.out.println(message.getPayload());
    return Mono.empty();
}
```

Streaming

All of the above are examples of consuming a single record one-by-one. However, one of the compelling reasons to use Reactive is for the streaming capability with backpressure support.

The following example uses `ReactivePulsarListener` to consume a stream of POJOs:

```
@ReactivePulsarListener(topics = "streaming-1", stream = true)
Flux<MessageResult<Void>> listen(Flux<Message<String>> messages) {
    return messages
        .doOnNext((msg) -> System.out.println("Received: " + msg.getValue()))
        .map(MessageResult::acknowledge);
}
```

Here we receive the records as a `Flux` of messages. In addition, to enable stream consumption at the `ReactivePulsarListener` level, you need to set the `stream` property on the annotation to `true`.



The listener method returns a `Flux<MessageResult<Void>>` where each element represents a processed message and holds the message id, value and whether it was acknowledged. The `MessageResult` has a set of static factory methods that can be used to create the appropriate `MessageResult` instance.

Based on the actual type of the messages in the `Flux`, the framework tries to infer the schema to use. If it contains a complex type, you still need to provide the `schemaType` on `ReactivePulsarListener`.

The following listener uses the Spring messaging `Message` envelope with a complex type :

```
@ReactivePulsarListener(topics = "streaming-2", stream = true, schemaType =
SchemaType.JSON)
Flux<MessageResult<Void>> listen2(Flux<org.springframework.messaging.Message<Foo>>
messages) {
    return messages
        .doOnNext((msg) -> System.out.println("Received: " + msg.getPayload()))
        .map(MessageResult::acknowledge);
}
```

Configuration - Application Properties

The listener ultimately relies on `ReactivePulsarConsumerFactory` to create and manage the underlying Pulsar consumer.

Spring Boot provides this consumer factory which can be configured with any of the `spring.pulsar.reactive.consumer` prefixed application properties.

Consumer Customization

You can specify a `ReactiveMessageConsumerBuilderCustomizer` to configure the underlying Pulsar consumer builder that ultimately constructs the consumer used by the listener to receive the messages.



Use with caution as this gives full access to the consumer builder and invoking some of its methods (such as `create`) may have unintended side effects.

For example, the following code shows how to set the initial position of the subscription to the earliest message on the topic.

```
@ReactivePulsarListener(topics = "hello-pulsar-topic", consumerCustomizer =
"myConsumerCustomizer")
Mono<Void> listen(String message) {
    System.out.println(message);
    return Mono.empty();
}

@Bean
ReactiveMessageConsumerBuilderCustomizer<String> myConsumerCustomizer() {
    return b ->
        b.subscriptionInitialPosition(SubscriptionInitialPosition.Earliest);
}
```

You can also use the customizer to provide direct Pulsar consumer properties to the consumer builder. This is convenient if you do not want to use the Boot configuration properties mentioned earlier or have multiple `ReactivePulsarListener` methods whose configuration varies.

The following customizer example uses direct Pulsar consumer properties:

```
@Bean
ReactiveMessageConsumerBuilderCustomizer<String> directConsumerPropsCustomizer() {
    return b -> b.property("subscriptionName", "subscription-
1").property("topicNames", "foo-1");
}
```



The properties used are direct Pulsar consumer properties, not the `spring.pulsar.reactive.consumer` Spring Boot configuration properties

Specifying Schema Information

As indicated earlier, for Java primitives, the Spring Pulsar framework can infer the proper Schema to use on the `ReactivePulsarListener`. However, for more complex types (such as JSON or AVRO), you need to specify the schema type on the annotation.



Complex Schema types that are currently supported are JSON, AVRO, PROTOBUF, and KEY_VALUE w/ INLINE encoding.

Message Listener Container Infrastructure

In most scenarios, we recommend using the `ReactivePulsarListener` annotation directly for consuming from a Pulsar topic as that model covers a broad set of application use cases. However, it is important to understand how `ReactivePulsarListener` works internally.

The message listener container is at the heart of message consumption when you use Spring for Apache Pulsar. The `ReactivePulsarListener` uses the message listener container infrastructure behind the scenes to create and manage the underlying Pulsar consumer.

ReactivePulsarMessageListenerContainer

The contract for this message listener container is provided through `ReactivePulsarMessageListenerContainer` whose default implementation creates a reactive Pulsar consumer and wires up a reactive message pipeline that uses the created consumer.

ReactiveMessagePipeline

The pipeline is a feature of the underlying Apache Pulsar Reactive client which does the heavy lifting of receiving the data in a reactive manner and then handing it over to the provided message handler. The reactive message listener container implementation is much simpler because the pipeline handles the majority of the work.

ReactivePulsarMessageHandler

The "listener" aspect is provided by the `ReactivePulsarMessageHandler` of which there are two provided implementations:

- `ReactivePulsarOneByOneMessageHandler` - handles a single message one-by-one
- `ReactivePulsarStreamingHandler` - handles multiple messages via a `Flux`

Concurrency

When consuming records in streaming mode (`stream = true`) concurrency comes naturally via the underlying Reactive support in the client implementation.

However, when handling messages one-by-one, concurrency can be specified to increase processing throughput. Simply set the `concurrency` property on `@ReactivePulsarListener`. Additionally, when `concurrency > 1` you can ensure messages are ordered by key and therefore sent to the same handler by setting `useKeyOrderedProcessing = "true"` on the annotation.

Again, the `ReactiveMessagePipeline` does the heavy lifting, we simply set the properties on it.

Reactive vs Imperative

Concurrency in the reactive container is different from its imperative counterpart. The latter creates multiple threads (each with a Pulsar consumer) whereas the former dispatches the messages to multiple handler instances concurrently on the Reactive parallel scheduler.

One advantage of reactive concurrency is that it can be used with `Exclusive` and `Failover` subscriptions to increase processing throughput if strict ordering is not required. In contrast to imperative concurrency that can not currently be used with `Exclusive` and does not provide more processing power with `Failover`.

Pulsar Headers

The Pulsar message metadata can be consumed as Spring message headers. The list of available headers can be found in [PulsarHeaders.java](#).

Accessing In OneByOne Listener

The following example shows how you can access Pulsar Headers when using a one-by-one message listener:

```
@ReactivePulsarListener(topics = "some-topic")
Mono<Void> listen(String data,
    @Header(PulsarHeaders.MESSAGE_ID) MessageId messageId,
    @Header("foo") String foo) {
    System.out.println("Received " + data + " w/ id=" + messageId + " w/ foo=" +
foo);
    return Mono.empty();
}
```

In the preceding example, we access the values for the `messageId` message metadata as well as a

custom message property named `foo`. The Spring `@Header` annotation is used for each header field.

You can also use Pulsar's `Message` as the envelope to carry the payload. When doing so, the user can directly call the corresponding methods on the Pulsar message for retrieving the metadata. However, as a convenience, you can also retrieve it by using the `Header` annotation. Note that you can also use the Spring messaging `Message` envelope to carry the payload and then retrieve the Pulsar headers by using `@Header`.

Accessing In Streaming Listener

When using a streaming message listener the header support is limited. Only when the `Flux` contains Spring `org.springframework.messaging.Message` elements will the headers be populated. Additionally, the Spring `@Header` annotation can not be used to retrieve the data. You must directly call the corresponding methods on the Spring message to retrieve the data.

Message Acknowledgment

The framework automatically handles message acknowledgement. However, the listener method must send a signal indicating whether the message was successfully processed. The container implementation then uses that signal to perform the ack or nack operation. This is a slightly different from its imperative counterpart where the signal is implied as positive unless the method throws an exception.

OneByOne Listener

The single message (aka OneByOne) message listener method returns a `Mono<Void>` to signal whether the message was successfully processed. `Mono.empty()` indicates success (acknowledgment) and `Mono.error()` indicates failure (negative acknowledgment).

Streaming Listener

The streaming listener method returns a `Flux<MessageResult<Void>>` where each `MessageResult` element represents a processed message and holds the message id, value and whether it was acknowledged. The `MessageResult` has a set of `acknowledge` and `negativeAcknowledge` static factory methods that can be used to create the appropriate `MessageResult` instance.

Message Redelivery and Error Handling

Apache Pulsar provides various native strategies for message redelivery and error handling. We will take a look at them and see how to use them through Spring for Apache Pulsar.

Acknowledgment Timeout

By default, Pulsar consumers do not redeliver messages unless the consumer crashes, but you can change this behavior by setting an ack timeout on the Pulsar consumer. When you use Spring for Apache Pulsar, you can enable this property by setting the `spring.pulsar.reactive.consumer.ack-timeout` Boot property. If this property has a value above zero and if the Pulsar consumer does not acknowledge a message within that timeout period, the message is redelivered.

You can also specify this property directly as a Pulsar consumer property via a `consumer customizer` such as:

```
@Bean
ReactiveMessageConsumerBuilderCustomizer<String> consumerCustomizer() {
    return b -> b.property("ackTimeout", "60s");
}
```

Negative Acknowledgment Redelivery Delay

When acknowledging negatively, Pulsar consumer lets you specify how the application wants the message to be re-delivered. The default is to redeliver the message in one minute, but you can change it by setting `spring.pulsar.reactive.consumer.negative-ack-redelivery-delay`.

You can also set it directly as a Pulsar consumer property via a `consumer customizer` such as:

```
@Bean
ReactiveMessageConsumerBuilderCustomizer<String> consumerCustomizer() {
    return b -> b.property("negativeAckRedeliveryDelay", "10ms");
}
```

Dead Letter Topic

Apache Pulsar lets applications use a dead letter topic on consumers with a `Shared` subscription type. For the `Exclusive` and `Failover` subscription types, this feature is not available. The basic idea is that, if a message is retried a certain number of times (maybe due to an ack timeout or nack redelivery), once the number of retries are exhausted, the message can be sent to a special topic called the dead letter queue (DLQ). Let us see some details around this feature in action by inspecting some code snippets:

```

@Configuration(proxyBeanMethods = false)
class DeadLetterPolicyConfig {

    @ReactivePulsarListener(
        topics = "topic-with-dlp",
        subscriptionType = SubscriptionType.Shared,
        deadLetterPolicy = "myDeadLetterPolicy",
        consumerCustomizer = "ackTimeoutCustomizer" )
    void listen(String msg) {
        throw new RuntimeException("fail " + msg);
    }

    @ReactivePulsarListener(topics = "my-dlq-topic")
    void listenDlq(String msg) {
        System.out.println("From DLQ: " + msg);
    }

    @Bean
    DeadLetterPolicy myDeadLetterPolicy() {
        return
        DeadLetterPolicy.builder().maxRedeliverCount(10).deadLetterTopic("my-dlq-
        topic").build();
    }

    @Bean
    ReactiveMessageConsumerBuilderCustomizer<String> ackTimeoutCustomizer() {
        return b -> b.property("ackTimeout", "1s");
    }
}

```

First, we have a special bean for `DeadLetterPolicy`, and it is named as `deadLetterPolicy` (it can be any name as you wish). This bean specifies a number of things, such as the max delivery (10, in this case) and the name of the dead letter topic — `my-dlq-topic`, in this case. If you do not specify a DLQ topic name, it defaults to `<topicname>-<subscriptionname>-DLQ` in Pulsar. Next, we provide this bean name to `ReactivePulsarListener` by setting the `deadLetterPolicy` property. Note that the `ReactivePulsarListener` has a subscription type of `Shared`, as the DLQ feature only works with shared subscriptions. This code is primarily for demonstration purposes, so we provide an `ackTimeout` value of 1 second. The idea is that the code throws the exception and, if Pulsar does not receive an ack within 1 second, it does a retry. If that cycle continues ten times (as that is our max redelivery count in the `DeadLetterPolicy`), the Pulsar consumer publishes the messages to the DLQ topic. We have another `ReactivePulsarListener` that listens on the DLQ topic to receive data as it is published to the DLQ topic.

Special note on DLQ topics when using partitioned topics

If the main topic is partitioned, behind the scenes, each partition is treated as a separate topic by Pulsar. Pulsar appends `partition-<n>`, where `n` stands for the partition number to the main topic name. The problem is that, if you do not specify a DLQ topic (as opposed to what we did above), Pulsar publishes to a default topic name that has this `'partition-<n>` info in it—for example: `topic-with-dlp-partition-0-deadLetterPolicySubscription-DLQ`. The easy way to solve this is to provide a DLQ topic name always.

Pulsar Reader Support

The framework provides support for using [Pulsar Reader](#) in a Reactive fashion via the `ReactivePulsarReaderFactory`.

Spring Boot provides this reader factory which can be configured with any of the `spring.pulsar.reactive.reader` prefixed application properties.

2.3. Pulsar Administration

2.3.1. Pulsar Admin Client

On the Pulsar administration side, Spring Boot auto-configuration provides a `PulsarAdministration` to manage Pulsar clusters. The administration implements an interface called `PulsarAdminOperations` and provides a `createOrModify` method to handle topic administration through its contract.

When you use the Pulsar Spring Boot starter, you get the `PulsarAdministration` auto-configured. By default, the application tries to connect to a local Pulsar instance at `localhost:8080`. However, there are many application properties available to help you configure the client. See the [Appendix](#) for application properties prefixed with `spring.pulsar.administration`.

Authentication

When accessing a Pulsar cluster that requires authentication, the admin client requires the same security configuration as the regular Pulsar client. You can use the aforementioned [security configuration](#) by replacing `spring.pulsar.client` with `spring.pulsar.administration`.

2.3.2. Automatic Topic Creation

On initialization, the `PulsarAdministration` checks if there are any `PulsarTopic` beans in the application context. For all such beans, the `PulsarAdministration` either creates the corresponding topic or, if necessary, modifies the number of partitions.

The following example shows how to add `PulsarTopic` beans to let the `PulsarAdministration` auto-create topics for you:

```

@Bean
PulsarTopic simpleTopic {
    // This will create a non-partitioned topic in the public/default namespace
    return PulsarTopic.builder("simple-topic").build();
}

@Bean
PulsarTopic partitionedTopic {
    // This will create a partitioned topic with 3 partitions in the provided
    tenant and namespace
    return PulsarTopic.builder("persistent://my-tenant/my-namespace/partitioned-
topic", 3).build();
}

```

2.4. Observability

Spring for Apache Pulsar includes a way to manage observability through [Micrometer](#).



Observability has not been added to the Reactive components yet

2.4.1. Micrometer Observations

The `PulsarTemplate` and `PulsarListener` are instrumented with the Micrometer observations API. When a Micrometer `ObservationRegistry` bean is provided, send and receive operations are traced and timed.

Custom tags

The default implementation adds the `bean.name` tag for template observations and `listener.id` tag for listener observations. To add other tags to timers and traces, configure a custom `PulsarTemplateObservationConvention` or `PulsarListenerObservationConvention` to the template or listener container, respectively.



You can subclass either `DefaultPulsarTemplateObservationConvention` or `DefaultPulsarListenerObservationConvention` or provide completely new implementations.

Observability - Metrics

Below you can find a list of all metrics declared by this project.

Listener Observation

Observation created when a Pulsar listener receives a message.

Metric	name	spring.pulsar.listener	(defined	by	convention	class
--------	------	------------------------	----------	----	------------	-------

`org.springframework.pulsar.observation.DefaultPulsarListenerObservationConvention`). **Type** `timer`.

Metric name `spring.pulsar.listener.active` (defined by convention class `org.springframework.pulsar.observation.DefaultPulsarListenerObservationConvention`). **Type** `long task timer`.



KeyValues that are added after starting the Observation might be missing from the `*.active metrics`.



Micrometer internally uses `nanoseconds` for the baseunit. However, each backend determines the actual baseunit. (i.e. Prometheus uses seconds)

Fully qualified name of the enclosing class `org.springframework.pulsar.observation.PulsarListenerObservation`.



All tags must be prefixed with `spring.pulsar.listener` prefix!

Table 1. Low cardinality Keys

Name	Description
<code>spring.pulsar.listener.id</code> (required)	Id of the listener container that received the message.

Template Observation

Observation created when a Pulsar template sends a message.

Metric name `spring.pulsar.template` (defined by convention class `org.springframework.pulsar.observation.DefaultPulsarTemplateObservationConvention`). **Type** `timer`.

Metric name `spring.pulsar.template.active` (defined by convention class `org.springframework.pulsar.observation.DefaultPulsarTemplateObservationConvention`). **Type** `long task timer`.



KeyValues that are added after starting the Observation might be missing from the `*.active metrics`.



Micrometer internally uses `nanoseconds` for the baseunit. However, each backend determines the actual baseunit. (i.e. Prometheus uses seconds)

Fully qualified name of the enclosing class `org.springframework.pulsar.observation.PulsarTemplateObservation`.



All tags must be prefixed with `spring.pulsar.template` prefix!

Table 2. Low cardinality Keys

Name	Description
<code>spring.pulsar.template.name</code> <i>(required)</i>	Bean name of the template that sent the message.

Observability - Spans

Below you can find a list of all spans declared by this project.

Listener Observation Span

Observation created when a Pulsar listener receives a message.

Span name `spring.pulsar.listener` (defined by convention class `org.springframework.pulsar.observation.DefaultPulsarListenerObservationConvention`).

Fully qualified name of the enclosing class `org.springframework.pulsar.observation.PulsarListenerObservation`.



All tags must be prefixed with `spring.pulsar.listener` prefix!

Table 3. Tag Keys

Name	Description
<code>spring.pulsar.listener.id</code> <i>(required)</i>	Id of the listener container that received the message.

Template Observation Span

Observation created when a Pulsar template sends a message.

Span name `spring.pulsar.template` (defined by convention class `org.springframework.pulsar.observation.DefaultPulsarTemplateObservationConvention`).

Fully qualified name of the enclosing class `org.springframework.pulsar.observation.PulsarTemplateObservation`.



All tags must be prefixed with `spring.pulsar.template` prefix!

Table 4. Tag Keys

Name	Description
<code>spring.pulsar.template.name</code> <i>(required)</i>	Bean name of the template that sent the message.

See [Micrometer Tracing](#) for more information.

Manual Configuration without Spring Boot

If you do not use Spring Boot, you need to configure and provide an `ObservationRegistry` as well as Micrometer Tracing. See [Micrometer Tracing](#) for more information.

Auto-Configuration with Spring Boot

If you use Spring Boot, the Spring Boot Actuator auto-configures an instance of `ObservationRegistry` for you. If `micrometer-core` is on the classpath, every stopped observation leads to a timer.

Spring Boot also auto-configures Micrometer Tracing for you. This includes support for Brave OpenTelemetry, Zipkin, and Wavefront. When using the Micrometer Observation API, finishing observations leads to spans reported to Zipkin or Wavefront. You can control tracing by setting properties under `management.tracing`. You can use Zipkin with `management.zipkin.tracing`, while Wavefront uses `management.wavefront`.

Example Configuration

The following example shows the steps to configure your Spring Boot application to use Zipkin with Brave.

1. Add the required dependencies to your application (in Maven or Gradle, respectively):

Maven

```
<dependencies>
  <dependency>
    <groupId>org.springframework.boot</groupId>
    <artifactId>spring-boot-starter-actuator</artifactId>
  </dependency>
  <dependency>
    <groupId>io.micrometer</groupId>
    <artifactId>micrometer-tracing-bridge-brave</artifactId>
  </dependency>
  <dependency>
    <groupId>io.zipkin.reporter2</groupId>
    <artifactId>zipkin-reporter-brave</artifactId>
  </dependency>
  <dependency>
    <groupId>io.zipkin.reporter2</groupId>
    <artifactId>zipkin-sender-urlconnection</artifactId>
  </dependency>
</dependencies>
```

Gradle

```
dependencies {  
    implementation 'org.springframework.boot:spring-boot-starter-actuator'  
    implementation 'io.micrometer:micrometer-tracing-bridge-brave'  
    implementation 'io.zipkin.reporter2:zipkin-reporter-brave'  
    implementation 'io.zipkin.reporter2:zipkin-sender-urlconnection'  
}
```

NOTE

You need the `'io.zipkin.reporter2:zipkin-sender-urlconnection'` dependency only if your application does not have a configured WebClient or RestTemplate.

1. Add the required properties to your application:

```
management:  
    tracing.enabled: true  
zipkin:  
    tracing.endpoint: "http://localhost:9411/api/v2/spans"
```

The `tracing.endpoint` above expects Zipkin is running locally as described [here](#).

At this point, your application should record traces when you send and receive Pulsar messages. You should be able to view them in the Zipkin UI (at [localhost:9411](#), when running locally).



You can also see the preceding configuration on the [Spring Pulsar Sample Apps](#).

The steps are very similar to configuring any of the other supported Tracing environments.

Other Resources

In addition to this reference documentation, we recommend a number of other resources that may help you learn about Spring and Apache Pulsar.

- [Spring for Apache Pulsar GitHub Repository](#)
- [Apache Pulsar Project Home Page](#)
- [Apache Pulsar Java Client](#)
- [Apache Pulsar GitHub Repository](#)
- [Apache Pulsar Reactive Client GitHub Repository](#)

Appendices

Appendix A: Application Properties

You can specify various properties inside your `application.properties` file, inside your `application.yml` file, or as command line switches. This appendix provides a list of Spring Pulsar properties and references to the underlying classes that consume them.



Spring Boot provides various conversion mechanisms with advanced value formatting. See [the properties conversion section](#) for more detail.

Pulsar Client Properties

Name	Description	Default Value
<code>spring.pulsar.client.auth-params</code>	Authentication parameter(s) as a JSON encoded string.	
<code>spring.pulsar.client.auth-plugin-class-name</code>	Fully qualified class name of the authentication plugin.	
<code>spring.pulsar.client.authentication.*</code>	Authentication parameter(s) as a map of parameter names to parameter values.	
<code>spring.pulsar.client.connection-timeout</code>	Duration to wait for a connection to a broker to be established.	<code>10s</code>
<code>spring.pulsar.client.dns-lookup-bind-address</code>	DNS lookup bind address.	
<code>spring.pulsar.client.dns-lookup-bind-port</code>	DNS lookup bind port.	<code>0</code>
<code>spring.pulsar.client.enable-busy-wait</code>	Enables spin-waiting on executors and IO threads in order to reduce latency during context switches.	<code>false</code>
<code>spring.pulsar.client.enable-transaction</code>	Enables transactions. To use this, start the <code>transactionCoordinatorClient</code> with the pulsar client.	<code>false</code>
<code>spring.pulsar.client.initial-backoff-interval</code>	Initial backoff interval.	<code>100ms</code>
<code>spring.pulsar.client.keep-alive-interval</code>	Keep alive interval for broker-client connection.	<code>30s</code>

Name	Description	Default Value
<code>spring.pulsar.client.listener-name</code>	Listener name for lookup. Clients can use <code>listenerName</code> to choose one of the listeners as the service URL to create a connection to the broker. To use this, "advertisedListeners" must be enabled on the broker.	
<code>spring.pulsar.client.lookup-timeout</code>	Client lookup timeout.	-1ms
<code>spring.pulsar.client.max-backoff-interval</code>	Maximum backoff interval.	30s
<code>spring.pulsar.client.max-concurrent-lookup-request</code>	Number of concurrent lookup-requests allowed to send on each broker-connection to prevent overload on broker.	5000
<code>spring.pulsar.client.max-lookup-redirects</code>	Maximum number of times a lookup-request to a broker will be redirected.	20
<code>spring.pulsar.client.max-lookup-request</code>	Number of max lookup-requests allowed on each broker-connection to prevent overload on broker.	50000
<code>spring.pulsar.client.max-number-of-rejected-request-per-connection</code>	Maximum number of broker-rejected requests in a certain timeframe, after which the current connection is closed and a new connection is created by the client.	50
<code>spring.pulsar.client.memory-limit</code>	Limit of direct memory that will be allocated by the client.	64MB
<code>spring.pulsar.client.num-connections-per-broker</code>	Maximum number of connections that the client will open to a single broker.	1
<code>spring.pulsar.client.num-io-threads</code>	Number of threads to be used for handling connections to brokers.	1

Name	Description	Default Value
<code>spring.pulsar.client.num-listener-threads</code>	Number of threads to be used for message listeners. The listener thread pool is shared across all the consumers and readers that are using a "listener" model to get messages. For a given consumer, the listener will always be invoked from the same thread, to ensure ordering.	1
<code>spring.pulsar.client.operation-timeout</code>	Client operation timeout.	30s
<code>spring.pulsar.client.proxy-protocol</code>	Protocol of proxy service. proxyServiceUrl and proxyProtocol must be mutually inclusive.	
<code>spring.pulsar.client.proxy-service-url</code>	URL of proxy service. proxyServiceUrl and proxyProtocol must be mutually inclusive.	
<code>spring.pulsar.client.request-timeout</code>	Maximum duration for completing a request.	1m
<code>spring.pulsar.client.service-url</code>	Pulsar cluster URL to connect to a broker.	
<code>spring.pulsar.client.socks5-proxy-address</code>	SOCKS5 proxy address.	
<code>spring.pulsar.client.socks5-proxy-password</code>	SOCKS5 proxy password.	
<code>spring.pulsar.client.socks5-proxy-username</code>	SOCKS5 proxy username.	
<code>spring.pulsar.client.ssl-provider</code>	Name of the security provider used for SSL connections.	
<code>spring.pulsar.client.stats-interval</code>	Interval between each stat info.	60s
<code>spring.pulsar.client.tls-allow-insecure-connection</code>	Whether the client accepts untrusted TLS certificates from the broker.	false

Name	Description	Default Value
<code>spring.pulsar.client.tls-ciphers</code>	Comma-separated list of cipher suites. This is a named combination of authentication, encryption, MAC and key exchange algorithm used to negotiate the security settings for a network connection using TLS or SSL network protocol. By default, all the available cipher suites are supported.	
<code>spring.pulsar.client.tls-hostname-verification-enable</code>	Whether the hostname is validated when the proxy creates a TLS connection with brokers.	<code>false</code>
<code>spring.pulsar.client.tls-protocols</code>	Comma-separated list of SSL protocols used to generate the SSLContext. Allowed values in recent JVMs are TLS, TLSv1.3, TLSv1.2 and TLSv1.1.	
<code>spring.pulsar.client.tls-trust-certs-file-path</code>	Path to the trusted TLS certificate file.	
<code>spring.pulsar.client.tls-trust-store-password</code>	Store password for the key store file.	
<code>spring.pulsar.client.tls-trust-store-path</code>	Location of the trust store file.	
<code>spring.pulsar.client.tls-trust-store-type</code>	File format of the trust store file.	
<code>spring.pulsar.client.use-key-store-tls</code>	Enable KeyStore instead of PEM type configuration if TLS is enabled.	<code>false</code>
<code>spring.pulsar.client.use-tcp-no-delay</code>	Whether to use TCP no-delay flag on the connection, to disable Nagle algorithm.	<code>true</code>
<code>spring.pulsar.client.use-tls</code>	Whether to use TLS encryption on the connection.	<code>false</code>

Pulsar Producer Properties

Name	Description	Default Value
<code>spring.pulsar.producer.auto-update-partitions</code>	Whether partitioned producer automatically discover new partitions at runtime.	<code>true</code>
<code>spring.pulsar.producer.auto-update-partitions-interval</code>	Interval of partitions discovery updates.	<code>1m</code>
<code>spring.pulsar.producer.batching-enabled</code>	Whether to automatically batch messages.	<code>true</code>
<code>spring.pulsar.producer.batching-max-bytes</code>	Maximum number of bytes permitted in a batch.	<code>128KB</code>
<code>spring.pulsar.producer.batching-max-messages</code>	Maximum number of messages to be batched.	<code>1000</code>
<code>spring.pulsar.producer.batching-max-publish-delay</code>	Time period within which the messages sent will be batched.	<code>1ms</code>
<code>spring.pulsar.producer.batching-partition-switch-frequency-by-publish-delay</code>	Partition switch frequency while batching of messages is enabled and using round-robin routing mode for non-keyed message.	<code>10</code>
<code>spring.pulsar.producer.block-if-queue-full</code>	Whether the "send" and "sendAsync" methods should block if the outgoing message queue is full.	<code>false</code>
<code>spring.pulsar.producer.cache.expire-after-access</code>	Time period to expire unused entries in the cache.	<code>1m</code>
<code>spring.pulsar.producer.cache.initial-capacity</code>	Initial size of cache.	<code>50</code>
<code>spring.pulsar.producer.cache.maximum-size</code>	Maximum size of cache (entries).	<code>1000</code>
<code>spring.pulsar.producer.chunking-enabled</code>	Whether to split large-size messages into multiple chunks.	<code>false</code>
<code>spring.pulsar.producer.compression-type</code>	Message compression type.	
<code>spring.pulsar.producer.crypto-failure-action</code>	Action the producer will take in case of encryption failure.	

Name	Description	Default Value
<code>spring.pulsar.producer.encryption-keys</code>	Names of the public encryption keys to use when encrypting data.	
<code>spring.pulsar.producer.hashing-scheme</code>	Message hashing scheme to choose the partition to which the message is published.	
<code>spring.pulsar.producer.initial-sequence-id</code>	Baseline for the sequence ids for messages published by the producer.	
<code>spring.pulsar.producer.lazy-start-partitioned-producers</code>	Whether producers in Shared mode register and connect immediately to the owner broker of each partition or start lazily on demand.	<code>false</code>
<code>spring.pulsar.producer.max-pending-messages</code>	Maximum number of pending messages for the producer.	<code>1000</code>
<code>spring.pulsar.producer.max-pending-messages-across-partitions</code>	Maximum number of pending messages across all the partitions.	<code>50000</code>
<code>spring.pulsar.producer.message-routing-mode</code>	Message routing mode for a partitioned producer.	
<code>spring.pulsar.producer.multi-schema</code>	Whether the multiple schema mode is enabled.	<code>true</code>
<code>spring.pulsar.producer.producer-access-mode</code>	Type of access to the topic the producer requires.	
<code>spring.pulsar.producer.producer-name</code>	Name for the producer. If not assigned, a unique name is generated.	
<code>spring.pulsar.producer.properties.*</code>	Map of properties to add to the producer.	
<code>spring.pulsar.producer.send-timeout</code>	Time before a message has to be acknowledged by the broker.	<code>30s</code>
<code>spring.pulsar.producer.topic-name</code>	Topic the producer will publish to.	

Name	Description	Default Value
<code>spring.pulsar.template.observations-enabled</code>	Whether to record observations for send operations when the Observations API is available.	<code>true</code>

Pulsar Consumer Properties

Name	Description	Default Value
<code>spring.pulsar.consumer.ack-receipt-enabled</code>	Whether an acknowledgement receipt is enabled.	<code>false</code>
<code>spring.pulsar.consumer.ack-timeout</code>	Timeout for unacked messages to be redelivered.	<code>0</code>
<code>spring.pulsar.consumer.acknowledgements-group-time</code>	Time to group acknowledgements before sending them to the broker.	<code>100ms</code>
<code>spring.pulsar.consumer.auto-ack-oldest-chunked-message-on-queue-full</code>	Whether to automatically drop outstanding un-acked messages if the queue is full.	<code>true</code>
<code>spring.pulsar.consumer.auto-update-partitions</code>	Whether the consumer auto-subscribes for partition increase. This is only for partitioned consumers.	<code>true</code>
<code>spring.pulsar.consumer.auto-update-partitions-interval</code>	Interval of partitions discovery updates.	<code>1m</code>
<code>spring.pulsar.consumer.batch-index-ack-enabled</code>	Whether the batch index acknowledgment is enabled.	<code>false</code>
<code>spring.pulsar.consumer.consumer-name</code>	Consumer name to identify a particular consumer from the topic stats.	
<code>spring.pulsar.consumer.crypto-failure-action</code>	Action the consumer will take in case of decryption failure.	
<code>spring.pulsar.consumer.dead-letter-policy.dead-letter-topic</code>		
<code>spring.pulsar.consumer.dead-letter-policy.initial-subscription-name</code>		
<code>spring.pulsar.consumer.dead-letter-policy.max-redeliver-count</code>		

Name	Description	Default Value
<code>spring.pulsar.consumer.dead-letter-policy.retry-letter-topic</code>		
<code>spring.pulsar.consumer.expire-time-of-incomplete-chunked-message</code>	Time to expire incomplete chunks if the consumer won't be able to receive all chunks before.	<code>1m</code>
<code>spring.pulsar.consumer.max-pending-chunked-message</code>	Maximum number of chunked messages to be kept in memory.	<code>10</code>
<code>spring.pulsar.consumer.max-total-receiver-queue-size-across-partitions</code>	Maximum number of messages that a consumer can be pushed at once from a broker across all partitions.	<code>50000</code>
<code>spring.pulsar.consumer.negative-ack-redelivery-delay</code>	Delay before re-delivering messages that have failed to be processed.	<code>1m</code>
<code>spring.pulsar.consumer.pattern-auto-discovery-period</code>	Auto-discovery period for topics when topic pattern is used in minutes.	<code>1</code>
<code>spring.pulsar.consumer.pool-messages</code>	Whether pooling of messages and the underlying data buffers is enabled.	<code>false</code>
<code>spring.pulsar.consumer.priority-level</code>	Priority level for shared subscription consumers.	<code>0</code>
<code>spring.pulsar.consumer.properties</code>	Map of properties to add to the consumer.	
<code>spring.pulsar.consumer.read-compacted</code>	Whether to read messages from the compacted topic rather than the full message backlog.	<code>false</code>
<code>spring.pulsar.consumer.receiver-queue-size</code>	Number of messages that can be accumulated before the consumer calls "receive".	<code>1000</code>
<code>spring.pulsar.consumer.regex-subscription-mode</code>	Determines which topics the consumer should be subscribed to when using pattern subscriptions.	
<code>spring.pulsar.consumer.replicate-subscription-state</code>	Whether to replicate subscription state.	<code>false</code>

Name	Description	Default Value
<code>spring.pulsar.consumer.reset-include-head</code>	Whether to include the given position of any reset operation like {@link org.apache.pulsar.client.api.Consumer#seek(long)} or {@link Consumer#seek(MessageId)}.	false
<code>spring.pulsar.consumer.retry-enable</code>	Whether to auto retry messages.	false
<code>spring.pulsar.consumer.start-paused</code>	Whether to start the consumer in a paused state.	false
<code>spring.pulsar.consumer.subscription-initial-position</code>	Position where to initialize a newly created subscription.	
<code>spring.pulsar.consumer.subscription-mode</code>	Subscription mode to be used when subscribing to the topic.	
<code>spring.pulsar.consumer.subscription-name</code>	Subscription name for the consumer.	
<code>spring.pulsar.consumer.subscription-properties.*</code>	Map of properties to add to the subscription.	
<code>spring.pulsar.consumer.subscription-type</code>	Subscription type to be used when subscribing to a topic.	
<code>spring.pulsar.consumer.tick-duration</code>	Precision for the ack timeout messages tracker.	1s
<code>spring.pulsar.consumer.topics</code>	Comma-separated list of topics the consumer subscribes to.	
<code>spring.pulsar.consumer.topics-pattern</code>	Pattern for topics the consumer subscribes to.	
<code>spring.pulsar.listener.ack-mode</code>	AckMode for acknowledgements. Allowed values are RECORD, BATCH, MANUAL.	
<code>spring.pulsar.listener.batch-timeout</code>	Duration to wait for enough message to fill a batch request before timing out.	100ms
<code>spring.pulsar.listener.max-num-bytes</code>	Max size in a single batch request.	10MB

Name	Description	Default Value
<code>spring.pulsar.listener.max-num-messages</code>	Max number of messages in a single batch request.	<code>-1</code>
<code>spring.pulsar.listener.observations-enabled</code>	Whether to record observations for receive operations when the Observations API is available.	<code>true</code>
<code>spring.pulsar.listener.schema-type</code>	SchemaType of the consumed messages.	

Pulsar Administration Properties

Name	Description	Default Value
<code>spring.pulsar.administration.auth-params</code>	Authentication parameter(s) as a JSON encoded string.	
<code>spring.pulsar.administration.auth-plugin-class-name</code>	Fully qualified class name of the authentication plugin.	
<code>spring.pulsar.administration.authentication.*</code>	Authentication parameter(s) as a map of parameter names to parameter values.	
<code>spring.pulsar.administration.auto-cert-refresh-time</code>	Certificates auto refresh time if Pulsar admin uses tls authentication.	<code>5m</code>
<code>spring.pulsar.administration.connection-timeout</code>	Duration to wait for a connection to server to be established.	<code>1m</code>
<code>spring.pulsar.administration.read-timeout</code>	Server response read time out for any request.	<code>1m</code>
<code>spring.pulsar.administration.request-timeout</code>	Server request time out for any request.	<code>5m</code>
<code>spring.pulsar.administration.service-url</code>	Pulsar service URL for the admin endpoint.	
<code>spring.pulsar.administration.ssl-provider</code>	Name of the security provider used for SSL connections.	
<code>spring.pulsar.administration.tls-allow-insecure-connection</code>	Whether the client accepts untrusted TLS certificates from the broker.	<code>false</code>

Name	Description	Default Value
<code>spring.pulsar.administration.tls-ciphers</code>	List of cipher suites. This is a named combination of authentication, encryption, MAC and key exchange algorithm used to negotiate the security settings for a network connection using TLS or SSL network protocol. By default, all the available cipher suites are supported.	
<code>spring.pulsar.administration.tls-hostname-verification-enable</code>	Whether the hostname is validated when the proxy creates a TLS connection with brokers.	<code>false</code>
<code>spring.pulsar.administration.tls-protocols</code>	List of SSL protocols used to generate the SSLContext. Allowed values in recent JVMs are TLS, TLSv1.3, TLSv1.2 and TLSv1.1.	
<code>spring.pulsar.administration.tls-trust-certs-file-path</code>	Path to the trusted TLS certificate file.	
<code>spring.pulsar.administration.tls-trust-store-password</code>	Store password for the key store file.	
<code>spring.pulsar.administration.tls-trust-store-path</code>	Location of the trust store file.	
<code>spring.pulsar.administration.tls-trust-store-type</code>	File format of the trust store file.	
<code>spring.pulsar.administration.use-key-store-tls</code>	Enable KeyStore instead of PEM type configuration if TLS is enabled.	<code>false</code>

Pulsar Reactive Sender Properties

Name	Description	Default Value
<code>spring.pulsar.reactive.sender.auto-update-partitions</code>	Whether partitioned producer automatically discover new partitions at runtime.	<code>true</code>
<code>spring.pulsar.reactive.sender.auto-update-partitions-interval</code>	Interval of partitions discovery updates.	<code>1m</code>

Name	Description	Default Value
<code>spring.pulsar.reactive.sender.batching-enabled</code>	Whether to automatically batch messages.	<code>true</code>
<code>spring.pulsar.reactive.sender.batching-max-bytes</code>	Maximum number of bytes permitted in a batch.	<code>128KB</code>
<code>spring.pulsar.reactive.sender.batching-max-messages</code>	Maximum number of messages to be batched.	<code>1000</code>
<code>spring.pulsar.reactive.sender.batching-max-publish-delay</code>	Time period within which the messages sent will be batched.	<code>1ms</code>
<code>spring.pulsar.reactive.sender.cache.expire-after-access</code>	Time period to expire unused entries in the cache.	<code>1m</code>
<code>spring.pulsar.reactive.sender.cache.initial-capacity</code>	Initial size of cache.	<code>50</code>
<code>spring.pulsar.reactive.sender.cache.maximum-size</code>	Maximum size of cache (entries).	<code>1000</code>
<code>spring.pulsar.reactive.sender.chunking-enabled</code>	Whether to split large-size messages into multiple chunks.	<code>false</code>
<code>spring.pulsar.reactive.sender.compression-type</code>	Message compression type.	
<code>spring.pulsar.reactive.sender.crypto-failure-action</code>	Action the producer will take in case of encryption failure.	
<code>spring.pulsar.reactive.sender.encryption-keys</code>	Names of the public encryption keys to use when encrypting data.	
<code>spring.pulsar.reactive.sender.hashing-scheme</code>	Message hashing scheme to choose the partition to which the message is published.	
<code>spring.pulsar.reactive.sender.initial-sequence-id</code>	Baseline for the sequence ids for messages published by the producer.	
<code>spring.pulsar.reactive.sender.lazy-start-partitioned-producers</code>	Whether producers in Shared mode register and connect immediately to the owner broker of each partition or start lazily on demand.	<code>false</code>

Name	Description	Default Value
<code>spring.pulsar.reactive.sender.max-pending-messages</code>	Maximum number of pending messages for the producer.	<code>1000</code>
<code>spring.pulsar.reactive.sender.max-pending-messages-across-partitions</code>	Maximum number of pending messages across all the partitions.	<code>50000</code>
<code>spring.pulsar.reactive.sender.message-routing-mode</code>	Message routing mode for a partitioned producer.	
<code>spring.pulsar.reactive.sender.multi-schema</code>	Whether the multiple schema mode is enabled.	<code>true</code>
<code>spring.pulsar.reactive.sender.producer-access-mode</code>	Type of access to the topic the producer requires.	
<code>spring.pulsar.reactive.sender.producer-name</code>	Name for the producer. If not assigned, a unique name is generated.	
<code>spring.pulsar.reactive.sender.properties.*</code>	Map of properties to add to the producer.	
<code>spring.pulsar.reactive.sender.round-robin-router-batching-partition-switch-frequency</code>		
<code>spring.pulsar.reactive.sender.send-timeout</code>	Time before a message has to be acknowledged by the broker.	<code>30s</code>
<code>spring.pulsar.reactive.sender.topic-name</code>	Topic the producer will publish to.	

Pulsar Reactive Consumer Properties

Name	Description	Default Value
<code>spring.pulsar.reactive.consumer.ack-timeout</code>	Timeout for unacked messages to be redelivered.	<code>0</code>
<code>spring.pulsar.reactive.consumer.ack-timeout-tick-time</code>	Precision for the ack timeout messages tracker.	<code>1s</code>

Name	Description	Default Value
<code>spring.pulsar.reactive.consumer.acknowledge-asynchronously</code>	When set to true, ignores the acknowledge operation completion and makes it asynchronous from the message consuming processing to improve performance by allowing the acknowledges and message processing to interleave. Defaults to true.	<code>true</code>
<code>spring.pulsar.reactive.consumer.acknowledge-scheduler-type</code>	Type of acknowledge scheduler.	
<code>spring.pulsar.reactive.consumer.acknowledgements-group-time</code>	Time to group acknowledgements before sending them to the broker.	<code>100ms</code>
<code>spring.pulsar.reactive.consumer.auto-ack-oldest-chunked-message-on-queue-full</code>	Whether to automatically drop outstanding un-acked messages if the queue is full.	<code>true</code>
<code>spring.pulsar.reactive.consumer.auto-update-partitions</code>	Whether the consumer auto-subscribes for partition increase. This is only for partitioned consumers.	<code>true</code>
<code>spring.pulsar.reactive.consumer.auto-update-partitions-interval</code>		<code>1m</code>
<code>spring.pulsar.reactive.consumer.batch-index-ack-enabled</code>	Whether batch index acknowledgement is enabled.	<code>false</code>
<code>spring.pulsar.reactive.consumer.consumer-name</code>	Consumer name to identify a particular consumer from the topic stats.	
<code>spring.pulsar.reactive.consumer.crypto-failure-action</code>	Action the consumer will take in case of decryption failure.	
<code>spring.pulsar.reactive.consumer.dead-letter-policy.dead-letter-topic</code>		
<code>spring.pulsar.reactive.consumer.dead-letter-policy.initial-subscription-name</code>		
<code>spring.pulsar.reactive.consumer.dead-letter-policy.max-redeliver-count</code>		
<code>spring.pulsar.reactive.consumer.dead-letter-policy.retry-letter-topic</code>		

Name	Description	Default Value
<code>spring.pulsar.reactive.consumer.expire-time-of-incomplete-chunked-message</code>	Time to expire incomplete chunks if the consumer won't be able to receive all chunks before in milliseconds.	<code>1m</code>
<code>spring.pulsar.reactive.consumer.max-pending-chunked-message</code>	Maximum number of chunked messages to be kept in memory.	<code>10</code>
<code>spring.pulsar.reactive.consumer.max-total-receiver-queue-size-across-partitions</code>	Maximum number of messages that a consumer can be pushed at once from a broker across all partitions.	<code>50000</code>
<code>spring.pulsar.reactive.consumer.negative-ack-redelivery-delay</code>	Delay before re-delivering messages that have failed to be processed.	<code>1m</code>
<code>spring.pulsar.reactive.consumer.priority-level</code>	Priority level for shared subscription consumers.	<code>0</code>
<code>spring.pulsar.reactive.consumer.properties</code>	Map of properties to add to the consumer.	
<code>spring.pulsar.reactive.consumer.read-compacted</code>	Whether to read messages from the compacted topic rather than the full message backlog.	<code>false</code>
<code>spring.pulsar.reactive.consumer.receiver-queue-size</code>	Number of messages that can be accumulated before the consumer calls "receive".	<code>1000</code>
<code>spring.pulsar.reactive.consumer.replicate-subscription-state</code>	Whether to replicate subscription state.	<code>false</code>
<code>spring.pulsar.reactive.consumer.retry-letter-topic-enable</code>	Whether the retry letter topic is enabled.	<code>false</code>
<code>spring.pulsar.reactive.consumer.subscription-initial-position</code>	Position where to initialize a newly created subscription.	
<code>spring.pulsar.reactive.consumer.subscription-mode</code>	Subscription mode to be used when subscribing to the topic.	
<code>spring.pulsar.reactive.consumer.subscription-name</code>	Subscription name for the consumer.	
<code>spring.pulsar.reactive.consumer.subscription-properties</code>	Map of properties to add to the subscription.	

Name	Description	Default Value
<code>spring.pulsar.reactive.consumer.subscription-type</code>	Subscription type to be used when subscribing to a topic.	
<code>spring.pulsar.reactive.consumer.topics</code>	Comma-separated list of topics the consumer subscribes to.	
<code>spring.pulsar.reactive.consumer.topics-pattern</code>	Pattern for topics the consumer subscribes to.	
<code>spring.pulsar.reactive.consumer.topics-pattern-auto-discovery-period</code>	Auto-discovery period for topics when topic pattern is used.	<code>1m</code>
<code>spring.pulsar.reactive.consumer.topics-pattern-subscription-mode</code>	Determines which topics the consumer should be subscribed to when using pattern subscriptions.	
<code>spring.pulsar.reactive.listener.handling-timeout</code>	Duration to wait before the message handling times out.	<code>2m</code>
<code>spring.pulsar.reactive.listener.schema-type</code>	SchemaType of the consumed messages.	
<code>spring.pulsar.reactive.listener.use-key-ordered-processing</code>	Whether per-key message ordering should be maintained when concurrent processing is used.	<code>false</code>

Pulsar Reactive Reader Properties

Name	Description	Default Value
<code>spring.pulsar.reactive.reader.crypto-failure-action</code>		
<code>spring.pulsar.reactive.reader.generated-subscription-name-prefix</code>		
<code>spring.pulsar.reactive.reader.key-hash-ranges</code>		
<code>spring.pulsar.reactive.reader.read-compacted</code>		
<code>spring.pulsar.reactive.reader.reader-name</code>		
<code>spring.pulsar.reactive.reader.receive-queue-size</code>		
<code>spring.pulsar.reactive.reader.subscription-name</code>		

Name	Description	Default Value
<code>spring.pulsar.reactive.reader.topic-names</code>		

Appendix B: Non-GA Versions

You can find snapshot or milestone versions of the dependencies in the following repositories:

Maven

```
<repositories>
  <repository>
    <id>spring-milestones</id>
    <name>Spring Milestones</name>
    <url>https://repo.spring.io/milestone</url>
    <snapshots>
      <enabled>>false</enabled>
    </snapshots>
  </repository>
  <repository>
    <id>spring-snapshots</id>
    <name>Spring Snapshots</name>
    <url>https://repo.spring.io/snapshot</url>
    <releases>
      <enabled>>false</enabled>
    </releases>
  </repository>
  <repository>
    <id>apache-snapshots</id>
    <name>Apache Snapshots</name>
    <url>https://repository.apache.org/content/repositories/snapshots</url>
    <releases>
      <enabled>>false</enabled>
    </releases>
  </repository>
</repositories>
```

```
repositories {  
    maven {  
        name = 'spring-milestones'  
        url = 'https://repo.spring.io/milestone'  
    }  
    maven {  
        name = 'spring-snapshots'  
        url = 'https://repo.spring.io/snapshot'  
    }  
    maven {  
        name = 'apache-snapshot'  
        url = 'https://repository.apache.org/content/repositories/snapshots'  
    }  
}
```

Appendix C: GraalVM Native Image Support

[GraalVM Native Images](#) are standalone executables that can be generated by processing compiled Java applications ahead-of-time. Native Images generally have a smaller memory footprint and start faster than their JVM counterparts.

Support

The required [AOT Runtime Hints](#) are built-in to Spring for Apache Pulsar so that it can seamlessly be used in native image based Spring applications.



The native image support in Spring for Apache Pulsar has been tested in basic scenarios, and we expect it to "just work". However, it is possible that more advanced use cases could surface the need to add additional runtime hints to your own application. If this occurs please file a [Github issue](#) with some details.

Next Steps

If you are interested in adding native image support to your own application then an excellent place to start is the [Spring Boot GraalVM Support](#) section of the Spring Boot reference docs.

Although there is no reference to Spring for Apache Pulsar in the aforementioned guide, you can find specific examples at the following coordinates:

- [Spring for Apache Pulsar \(imperative\)](#)
- [Spring for Apache Pulsar \(reactive\)](#)