

Reactor Kafka Reference Guide

Rajini Sivaram, Mark Pollack

1.0.0.M3

Table of Contents

| | |
|--|----|
| Introduction | 1 |
| 1. Overview | 2 |
| 1.1. Apache Kafka | 2 |
| 1.2. Project Reactor | 2 |
| 1.3. Reactive API for Kafka | 2 |
| 2. Motivation | 3 |
| 2.1. Functional interface for Kafka | 3 |
| 2.2. Non-blocking Back-pressure | 3 |
| 2.3. End-to-end Reactive Pipeline | 3 |
| 2.4. Comparisons with other Kafka APIs | 3 |
| 2.4.1. Kafka Producer and Consumer APIs | 3 |
| 2.4.2. Kafka Connect API | 3 |
| 2.4.3. Kafka Streams API | 4 |
| 3. Getting Started | 5 |
| 3.1. Requirements | 5 |
| 3.2. Quick Start | 5 |
| 3.2.1. Start Kafka | 5 |
| 3.2.2. Run Reactor Kafka Samples | 6 |
| Sample Producer | 6 |
| Sample Consumer | 8 |
| 3.2.3. Building Reactor Kafka Applications | 9 |
| 4. Additional Resources | 10 |
| 4.1. Getting help | 10 |
| 4.2. Resources | 10 |
| 5. New & Noteworthy | 11 |
| 5.1. What's new in Reactor Kafka 1.0 | 11 |
| Reference Documentation | 11 |
| 6. Reactor Kafka API | 12 |
| 6.1. Overview | 12 |
| 6.2. Reactive Kafka Sender | 12 |
| 6.2.1. Error handling | 14 |
| 6.2.2. Send without result metadata | 14 |
| 6.2.3. Threading model | 16 |
| 6.2.4. Non-blocking back-pressure | 16 |
| 6.2.5. Closing the KafkaSender | 16 |
| 6.2.6. Access to the underlying <code>KafkaProducer</code> | 16 |
| 6.3. Reactive Kafka Receiver | 17 |
| 6.3.1. Subscribing to wildcard patterns | 18 |

| | |
|--|----|
| 6.3.2. Manual assignment of topic partitions | 18 |
| 6.3.3. Controlling commit frequency | 19 |
| 6.3.4. Auto-acknowledgement of batches of records | 19 |
| 6.3.5. Disabling automatic commits | 20 |
| 6.3.6. At-most-once delivery | 20 |
| 6.3.7. Partition assignment and revocation listeners | 21 |
| 6.3.8. Controlling start offsets for consuming records | 21 |
| 6.3.9. Consumer lifecycle | 22 |
| 7. Sample Scenarios | 23 |
| 7.1. Sending records to Kafka | 23 |
| 7.2. Replaying records from Kafka topics | 23 |
| 7.3. Reactive pipeline with Kafka sink | 24 |
| 7.4. Reactive pipeline with Kafka source | 24 |
| 7.5. Reactive pipeline with Kafka source and sink | 25 |
| 7.6. At-most-once delivery | 25 |
| 7.7. Fan-out with Multiple Streams | 26 |
| 7.8. Concurrent Processing with Partition-Based Ordering | 27 |

Introduction

Chapter 1. Overview

1.1. Apache Kafka

[Kafka](#) is a scalable, high-performance distributed messaging engine. Low latency, high throughput messaging capability combined with fault-tolerance have made Kafka a popular messaging service as well as a powerful streaming platform for processing real-time streams of events.

Apache Kafka provides three main APIs:

- Producer/Consumer API to publish messages to Kafka topics and consume messages from Kafka topics
- Connector API to pull data from existing data storage systems to Kafka or push data from Kafka topics to other data systems
- Streams API for transforming and analyzing real-time streams of events published to Kafka

1.2. Project Reactor

[Reactor](#) is a highly optimized reactive library for building efficient, non-blocking applications on the JVM based on the [Reactive Streams Specification](#). Reactor based applications can sustain very high throughput message rates and operate with a very low memory footprint, making it suitable for building efficient event-driven applications using the microservices architecture.

Reactor implements two publishers [Flux<T>](#) and [Mono<T>](#), both of which support non-blocking back-pressure. This enables exchange of data between threads with well-defined memory usage, avoiding unnecessary intermediate buffering or blocking.

1.3. Reactive API for Kafka

[Reactor Kafka](#) is a reactive API for Kafka based on Reactor and the Kafka Producer/Consumer API. Reactor Kafka API enables messages to be published to Kafka and consumed from Kafka using functional APIs with non-blocking back-pressure and very low overheads. This enables applications using Reactor to use Kafka as a message bus or streaming platform and integrate with other systems to provide an end-to-end reactive pipeline.

Chapter 2. Motivation

2.1. Functional interface for Kafka

Reactor Kafka is a functional Java API for Kafka. For applications that are written in functional style, this API enables Kafka interactions to be integrated easily without requiring non-functional asynchronous produce or consume APIs to be incorporated into the application logic.

2.2. Non-blocking Back-pressure

The Reactor Kafka API benefits from non-blocking back-pressure provided by Reactor. For example, in a pipeline, where messages received from an external source (e.g. an HTTP proxy) are published to Kafka, back-pressure can be applied easily to the whole pipeline, limiting the number of messages in-flight and controlling memory usage. Messages flow through the pipeline as they are available, with Reactor taking care of limiting the flow rate to avoid overflow, keeping application logic simple.

2.3. End-to-end Reactive Pipeline

The value proposition for Reactor Kafka is the efficient utilization of resources in applications with multiple external interactions where Kafka is one of the external systems. End-to-end reactive pipelines benefit from non-blocking back-pressure and efficient use of threads, enabling a large number of concurrent requests to be processed efficiently. The optimizations provided by Project Reactor enable development of reactive applications with very low overheads and predictable capacity planning to deliver low-latency, high-throughput pipelines.

2.4. Comparisons with other Kafka APIs

Reactor Kafka is not intended to replace any of the existing Kafka APIs. Instead, it is aimed at providing an alternative API for reactive event-driven applications.

2.4.1. Kafka Producer and Consumer APIs

For non-reactive applications, Kafka's Producer/Consumer API provides a low latency interface to publish messages to Kafka and consume messages from Kafka.

Applications using Kafka as a message bus using this API may consider switching to Reactor Kafka if the application is implemented in a functional style.

2.4.2. Kafka Connect API

[Kafka Connect](#) provides a simple interface to migrate messages from an external data system (e.g. a database) to one or more Kafka topics. Using existing connectors, this migration can be performed without writing any new code.

Applications using the connector API may consider using Reactor Kafka if a reactive API is available

for the external data system and transformations are required for the data. When transformations involve other I/O (e.g. to obtain additional information from another database), a reactive pipeline benefits from end-to-end non-blocking back-pressure provided by Reactor. Messages from/to different Kafka partitions can be processed in parallel, improving throughput by avoiding blocking for I/O. The pull model in Reactor controls the pace of messages flowing through the pipeline, enabling efficient use of threads and memory without the need for overflow handling in the application.

2.4.3. Kafka Streams API

[Kafka Streams](#) provides lightweight APIs to build stream processing applications that process data stored in Kafka using standard streaming concepts and transformation primitives. Using a simple threading model, the streams API avoids the need for back-pressure. This model works well in cases where transformations do not involve external interactions.

Reactor Kafka is useful for streams applications which process data from Kafka and use external interactions (e.g. get additional data for records from a database) for transformations. In this case, Reactor can provide end-to-end non-blocking back-pressure combined with better utilization of resources if all external interactions use the reactive model.

Chapter 3. Getting Started

3.1. Requirements

You need Java JRE installed (Java 8 or later).

You need [Apache Kafka](https://kafka.apache.org/downloads.html) installed (0.11.0.0 or later). Kafka can be downloaded from kafka.apache.org/downloads.html. Note that the Apache Kafka client library used with Reactor Kafka should be 0.11.0.0 or later and the broker version should be 0.10.1.0 or later.

3.2. Quick Start

This quick start tutorial sets up a single node Zookeeper and Kafka and runs the sample reactive producer and consumer. Instructions to set up multi-broker clusters are available [here](#).

3.2.1. Start Kafka

If you haven't yet downloaded Kafka, download Kafka version [0.11.0.0](#) or later.

Unzip the release and set KAFKA_DIR to the installation directory. For example,

```
> tar -zxf kafka_2.11-0.11.0.0.tgz -C /opt
> export KAFKA_DIR=/opt/kafka_2.11-0.11.0.0
```

Start a single-node Zookeeper instance using the Zookeeper installation included in the Kafka download:

```
> $KAFKA_DIR/bin/zookeeper-server-start.sh $KAFKA_DIR/config/zookeeper.properties >
/tmp/zookeeper.log &
```

Start a single-node Kafka instance:

```
> $KAFKA_DIR/bin/kafka-server-start.sh $KAFKA_DIR/config/server.properties >
/tmp/kafka.log &
```

Create a Kafka topic:

```
> $KAFKA_DIR/bin/kafka-topics.sh --zookeeper localhost:2181 --create --replication
-factor 1 --partitions 2 --topic demo-topic
Created topic "demo-topic".
```

Check that Kafka topic was created successfully:


```
> $KAFKA_DIR/bin/kafka-topics.sh --zookeeper localhost:2181 --describe
Topic:demo-topic    PartitionCount:2    ReplicationFactor:1 Configs:
    Topic: demo-topic    Partition: 0    Leader: 0    Replicas: 0 Isr: 0
    Topic: demo-topic    Partition: 1    Leader: 0    Replicas: 0 Isr: 0
```

3.2.2. Run Reactor Kafka Samples

Download and build Reactor Kafka from github.com/reactor/reactor-kafka/.

```
> git clone https://github.com/reactor/reactor-kafka
> cd reactor-kafka
> ./gradlew jar
```

Set **CLASSPATH** for running reactor-kafka samples. **CLASSPATH** can be obtained using the classpath task of the samples sub-project.

```
> export CLASSPATH=`./gradlew -q :reactor-kafka-samples:classpath`
```

Sample Producer

See github.com/reactor/reactor-kafka/blob/master/reactor-kafka-samples/src/main/java/reactor/kafka/samples/SampleProducer.java for sample producer code.

Run sample producer:

```
> $KAFKA_DIR/bin/kafka-run-class.sh reactor.kafka.samples.SampleProducer
Message 2 sent successfully, topic-partition=demo-topic-1 offset=0
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 3 sent successfully, topic-partition=demo-topic-1 offset=1
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 4 sent successfully, topic-partition=demo-topic-1 offset=2
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 6 sent successfully, topic-partition=demo-topic-1 offset=3
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 7 sent successfully, topic-partition=demo-topic-1 offset=4
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 10 sent successfully, topic-partition=demo-topic-1 offset=5
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 11 sent successfully, topic-partition=demo-topic-1 offset=6
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 12 sent successfully, topic-partition=demo-topic-1 offset=7
timestamp=13:33:16:717 GMT 30 Nov 2016
Message 13 sent successfully, topic-partition=demo-topic-1 offset=8
timestamp=13:33:16:717 GMT 30 Nov 2016
Message 14 sent successfully, topic-partition=demo-topic-1 offset=9
timestamp=13:33:16:717 GMT 30 Nov 2016
Message 16 sent successfully, topic-partition=demo-topic-1 offset=10
timestamp=13:33:16:717 GMT 30 Nov 2016
Message 17 sent successfully, topic-partition=demo-topic-1 offset=11
timestamp=13:33:16:717 GMT 30 Nov 2016
Message 20 sent successfully, topic-partition=demo-topic-1 offset=12
timestamp=13:33:16:717 GMT 30 Nov 2016
Message 1 sent successfully, topic-partition=demo-topic-0 offset=0
timestamp=13:33:16:712 GMT 30 Nov 2016
Message 5 sent successfully, topic-partition=demo-topic-0 offset=1
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 8 sent successfully, topic-partition=demo-topic-0 offset=2
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 9 sent successfully, topic-partition=demo-topic-0 offset=3
timestamp=13:33:16:716 GMT 30 Nov 2016
Message 15 sent successfully, topic-partition=demo-topic-0 offset=4
timestamp=13:33:16:717 GMT 30 Nov 2016
Message 18 sent successfully, topic-partition=demo-topic-0 offset=5
timestamp=13:33:16:717 GMT 30 Nov 2016
Message 19 sent successfully, topic-partition=demo-topic-0 offset=6
timestamp=13:33:16:717 GMT 30 Nov 2016
```

The sample producer sends 20 messages to Kafka topic `demo-topic` using the default partitioner. The partition and offset of each published message is output to console. As shown in the sample output above, the order of results may be different from the order of messages published. Results are delivered in order for each partition, but results from different partitions may be interleaved. In the sample, message index is included as correlation metadata to match each result to its corresponding message.

Sample Consumer

See github.com/reactor/reactor-kafka/blob/master/reactor-kafka-samples/src/main/java/reactor/kafka/samples/SampleConsumer.java for sample consumer code.

Run sample consumer:

```
> $KAFKA_DIR/bin/kafka-run-class.sh reactor.kafka.samples.SampleConsumer
Received message: topic-partition=demo-topic-1 offset=0 timestamp=13:33:16:716 GMT 30
Nov 2016 key=2 value=Message_2
Received message: topic-partition=demo-topic-1 offset=1 timestamp=13:33:16:716 GMT 30
Nov 2016 key=3 value=Message_3
Received message: topic-partition=demo-topic-1 offset=2 timestamp=13:33:16:716 GMT 30
Nov 2016 key=4 value=Message_4
Received message: topic-partition=demo-topic-1 offset=3 timestamp=13:33:16:716 GMT 30
Nov 2016 key=6 value=Message_6
Received message: topic-partition=demo-topic-1 offset=4 timestamp=13:33:16:716 GMT 30
Nov 2016 key=7 value=Message_7
Received message: topic-partition=demo-topic-1 offset=5 timestamp=13:33:16:716 GMT 30
Nov 2016 key=10 value=Message_10
Received message: topic-partition=demo-topic-1 offset=6 timestamp=13:33:16:716 GMT 30
Nov 2016 key=11 value=Message_11
Received message: topic-partition=demo-topic-1 offset=7 timestamp=13:33:16:717 GMT 30
Nov 2016 key=12 value=Message_12
Received message: topic-partition=demo-topic-1 offset=8 timestamp=13:33:16:717 GMT 30
Nov 2016 key=13 value=Message_13
Received message: topic-partition=demo-topic-1 offset=9 timestamp=13:33:16:717 GMT 30
Nov 2016 key=14 value=Message_14
Received message: topic-partition=demo-topic-1 offset=10 timestamp=13:33:16:717 GMT 30
Nov 2016 key=16 value=Message_16
Received message: topic-partition=demo-topic-1 offset=11 timestamp=13:33:16:717 GMT 30
Nov 2016 key=17 value=Message_17
Received message: topic-partition=demo-topic-1 offset=12 timestamp=13:33:16:717 GMT 30
Nov 2016 key=20 value=Message_20
Received message: topic-partition=demo-topic-0 offset=0 timestamp=13:33:16:712 GMT 30
Nov 2016 key=1 value=Message_1
Received message: topic-partition=demo-topic-0 offset=1 timestamp=13:33:16:716 GMT 30
Nov 2016 key=5 value=Message_5
Received message: topic-partition=demo-topic-0 offset=2 timestamp=13:33:16:716 GMT 30
Nov 2016 key=8 value=Message_8
Received message: topic-partition=demo-topic-0 offset=3 timestamp=13:33:16:716 GMT 30
Nov 2016 key=9 value=Message_9
Received message: topic-partition=demo-topic-0 offset=4 timestamp=13:33:16:717 GMT 30
Nov 2016 key=15 value=Message_15
Received message: topic-partition=demo-topic-0 offset=5 timestamp=13:33:16:717 GMT 30
Nov 2016 key=18 value=Message_18
Received message: topic-partition=demo-topic-0 offset=6 timestamp=13:33:16:717 GMT 30
Nov 2016 key=19 value=Message_19
```

The sample consumer consumes messages from topic **demo-topic** and outputs the messages to

console. The 20 messages published by the Producer sample should appear on the console. As shown in the output above, messages are consumed in order for each partition, but messages from different partitions may be interleaved.

3.2.3. Building Reactor Kafka Applications

To build your own application using the Reactor Kafka API, you need to include a dependency to Reactor Kafka.

For gradle:

```
dependencies {  
    compile "io.projectreactor.kafka:reactor-kafka:1.0.0.M3"  
}
```

For maven:

```
<dependency>  
    <groupId>io.projectreactor.kafka</groupId>  
    <artifactId>reactor-kafka</artifactId>  
    <version>1.0.0.M3</version>  
</dependency>
```

Chapter 4. Additional Resources

4.1. Getting help

If you are having trouble with Reactor Kafka, we'd like to help.

Report bugs in Reactor Kafka at github.com/reactor/reactor-kafka/issues.

Reactor Kafka is open source and the code and documentation are available at github.com/reactor/reactor-kafka.

4.2. Resources

- [Reactor Kafka on github](#)
- [Apache Kafka](#)
- [Project Reactor](#)
- [Reactor Core](#)
- [Reactive Streams Specification](#)
- [Understanding Reactive types](#)
- [Lite Rx API Hands-on](#)
- [Reactor by Example](#)

Chapter 5. New & Noteworthy

5.1. What's new in Reactor Kafka 1.0

- First release of Reactor Kafka

Reference Documentation

Chapter 6. Reactor Kafka API

6.1. Overview

This section describes the reactive API for producing and consuming messages using Apache Kafka. There are two main interfaces in Reactor Kafka:

1. `reactor.kafka.sender.KafkaSender` for publishing messages to Kafka
2. `reactor.kafka.receiver.KafkaReceiver` for consuming messages from Kafka

Full API for Reactor Kafka is available in the [javadocs](#).

The project uses [Reactor Core](#) to expose a "Reactive Streams" API.

6.2. Reactive Kafka Sender

Outbound messages are sent to Kafka using `reactor.kafka.sender.KafkaSender`. Senders are thread-safe and can be shared across multiple threads to improve throughput. A `KafkaSender` is associated with one `KafkaProducer` that is used to transport messages to Kafka.

A `KafkaSender` is created with an instance of sender configuration options `reactor.kafka.sender.SenderOptions`. Changes made to `SenderOptions` after the creation of `KafkaSender` will not be used by the `KafkaSender`. The properties of `SenderOptions` such as a list of bootstrap Kafka brokers and serializers are passed down to the underlying `KafkaProducer`. The properties may be configured on the `SenderOptions` instance at creation time or by using the setter `SenderOptions#producerProperty`. Other configuration options for the reactive `KafkaSender` like the maximum number of in-flight messages can also be configured before the `KafkaSender` instance is created.

The generic types of `SenderOptions<K, V>` and `KafkaSender<K, V>` are the key and value types of producer records published using the `KafkaSender` and corresponding serializers must be set on the `SenderOptions` instance before the `KafkaSender` is created.

```
Map<String, Object> producerProps = new HashMap<>();
producerProps.put(ProducerConfig.BOOTSTRAP_SERVERS_CONFIG, bootstrapServers);
producerProps.put(ProducerConfig.KEY_SERIALIZER_CLASS_CONFIG, IntegerSerializer.class);
producerProps.put(ProducerConfig.VALUE_SERIALIZER_CLASS_CONFIG, StringSerializer.class);

SenderOptions<Integer, String> senderOptions =
    SenderOptions.<Integer, String>create(producerProps)
        .maxInFlight(1024);
```

- ① Specify properties for underlying `KafkaProducer`
- ② Configure options for reactive `KafkaSender`

Once the required options have been configured on the options instance, a new `KafkaSender` instance can be created with the options already configured in `senderOptions`.

```
KafkaSender<Integer, String> sender = KafkaSender.create(senderOptions);
```

The `KafkaSender` is now ready to send messages to Kafka. The underlying `KafkaProducer` instance is created lazily when the first message is ready to be sent. At this point, a `KafkaSender` instance has been created, but no connections to Kafka have been made yet.

Let's now create a sequence of messages to send to Kafka. Each outbound message to be sent to Kafka is represented as a `SenderRecord`. A `SenderRecord` is a Kafka `ProducerRecord` with additional correlation metadata for matching send results to records. `ProducerRecord` consists of a key/value pair to send to Kafka and the name of the Kafka topic to send the message to. Producer records may also optionally specify a partition to send the message to or use the configured partitioner to choose a partition. Timestamp may also be optionally specified in the record and if not specified, the current timestamp will be assigned by the Producer. The additional correlation metadata included in `SenderRecord` is not sent to Kafka, but is included in the `SendResult` generated for the record when the send operation completes or fails. Since results of sends to different partitions may be interleaved, the correlation metadata enables results to be matched to their corresponding record.

A `Flux<SenderRecord>` of records is created for sending to Kafka. For beginners, [Lite Rx API Hands-on](#) provides a hands-on tutorial on using the Reactor classes `Flux` and `Mono`.

```
Flux<SenderRecord<Integer, String, Integer>> outboundFlux =  
    Flux.range(1, 10)  
        .map(i -> SenderRecord.create(topic, partition, timestamp, i, "Message_" + i,  
i));
```

The code segment above creates a sequence of messages to send to Kafka, using the message index as correlation metadata in each `SenderRecord`. The outbound Flux can now be sent to Kafka using the `KafkaSender` created earlier.

The code segment below sends the records to Kafka and prints out the response metadata received from Kafka and the correlation metadata for each record. The final `subscribe()` in the code block requests upstream to send the records to Kafka and the response metadata received from Kafka flow downstream. As each result is received, the record metadata from Kafka along with the correlation metadata identifying the record is printed out to console by the `onNext` handler. The response from Kafka includes the partition to which the record was sent as well as the offset at the which the record was appended, if available. When records are sent to multiple partitions, responses arrive in order for each partition, but responses from different partitions may be interleaved.


```

sender.send(outboundFlux) ①
    .doOnError(e-> log.error("Send failed", e)) ②
    .doOnNext(r -> System.out.printf("Message #%d send response: %s\n", r
    .correlationMetadata(), r.recordMetadata())) ③
    .subscribe(); ④

```

- ① Reactive send operation for the outbound Flux
- ② If Kafka send fails, log an error
- ③ Print metadata returned by Kafka and the message index in `correlationMetadata()`
- ④ Subscribe to trigger the actual flow of records from `outboundFlux` to Kafka.

See github.com/reactor/reactor-kafka/blob/master/reactor-kafka-samples/src/main/java/reactor/kafka/samples/SampleProducer.java for the full code listing of a sample producer.

6.2.1. Error handling

```

public SenderOptions<K, V> stopOnError(boolean stopOnError);

```

`SenderOptions#stopOnError()` specifies whether each send sequence should fail immediately if one record could not be delivered to Kafka after the configured number of retries or wait until all records have been processed. This can be used along with `ProducerConfig#ACKS_CONFIG` and `ProducerConfig#RETRIES_CONFIG` to configure the required quality of service.

```

<T> Flux<SenderResult<T>> send(Publisher<SenderRecord<K, V, T>> outboundRecords);

```

If `stopOnError` is false, a success or error response is returned for each outgoing record. For error responses, the exception from Kafka indicating the reason for send failure is set on `SenderResult` and can be retrieved using `SenderResult#exception()`. The Flux fails with an error after attempting to send all records published on `outboundRecords`. If `outboundRecords` is a non-terminating Flux, send continues to send records published on this Flux until the result Flux is explicitly cancelled by the user.

If `stopOnError` is true, a response is returned for the first failed send and the result Flux is terminated immediately with an error. Since multiple outbound messages may be in-flight at any time, it is possible that some messages are delivered successfully to Kafka after the first failure is detected. `SenderOptions#maxInFlight()` option may be configured to limit the number of messages in-flight at any time.

6.2.2. Send without result metadata

If individual results are not required for each send request, `ProducerRecord` can be sent to Kafka without providing correlation metadata using the `sendOutbound` method. `KafkaOutbound` is a fluent interface that enables sends to be chained together.

```
KafkaOutbound<K, V> sendOutbound(Publisher<? extends ProducerRecord<K, V>> records);
```

The send sequence is initiated by subscribing to the Mono obtained from `KafkaOutbound#then()`. The returned Mono completes successfully if all the outbound records are delivered successfully. The Mono terminates on the first send failure. If `outboundRecords` is a non-terminating Flux, records continue to be sent to Kafka unless a send fails or the returned Mono is cancelled.

```
sender.sendOutbound(Flux.range(1, 10)
    .map(i -> new ProducerRecord<Integer, String>(topic, i,
"Message_" + i))) ①
    .then() ②
    .doOnError(e -> e.printStackTrace()) ③
    .doOnSuccess(s -> System.out.println("Sends succeeded")) ④
    .subscribe(); ⑤
```

- ① Create `ProducerRecord` Flux. Records are not wrapped in `SenderRecord`
- ② Get the `Mono` to subscribe to for starting the message flow
- ③ Error indicates failure to send one or more records
- ④ Success indicates all records were published, individual partitions or offsets not returned
- ⑤ Subscribe to request the actual sends

Multiple sends can be chained together using a sequence of sends on the returned `KafkaOutbound`. When the Mono returned from `KafkaOutbound#then()` is subscribed to, the sends are invoked in sequence in the declaration order. The sequence is cancelled if any of the sends fail after the configured number of retries.

```
sender.sendOutbound(flux1) ①
    .send(flux2)
    .send(flux3)
    .then() ②
    .doOnError(e -> e.printStackTrace()) ③
    .doOnSuccess(s -> System.out.println("Sends succeeded")) ④
    .subscribe(); ⑤
```

- ① Sends `flux1`, `flux2` and `flux3` in order
- ② Get the `Mono` to subscribe to for starting the message flow sequence
- ③ Error indicates failure to send one or more records from any of the sends in the chain
- ④ Success indicates successful send of all records from the whole chain
- ⑤ Subscribe to initiate the sequence of sends in the chain

Note that in all cases the retries configured for the `KafkaProducer` are attempted and failures returned by the reactive `KafkaSender` indicate a failure to send after the configured number of retry attempts. Retries can result in messages being delivered out of order. The producer property `ProducerConfig#MAX_IN_FLIGHT_REQUESTS_PER_CONNECTION` may be set to one to avoid re-ordering.

6.2.3. Threading model

`KafkaProducer` uses a separate network thread for sending requests and processing responses. To ensure that the producer network thread is never blocked by applications while processing results, `KafkaSender` delivers responses to applications on a separate scheduler. By default, this is a single threaded pooled scheduler that is freed when no longer required. The scheduler can be overridden if required, for instance, to use a parallel scheduler when the Kafka sends are part of a larger pipeline. This is done on the `SenderOptions` instance before the `KafkaSender` instance is created using:

```
public SenderOptions<K, V> scheduler(Scheduler scheduler);
```

6.2.4. Non-blocking back-pressure

The number of in-flight sends can be controlled using the `maxInFlight` option. Requests for more elements from upstream are limited by the configured `maxInFlight` to ensure that the total number of requests at any time for which responses are pending are limited. Along with `buffer.memory` and `max.block.ms` options on `KafkaProducer`, `maxInFlight` enables control of memory and thread usage when `KafkaSender` is used in a reactive pipeline. This option can be configured on `SenderOptions` before the `KafkaSender` is created. Default value is 256. For small messages, a higher value will improve throughput.

```
public SenderOptions<K, V> maxInFlight(int maxInFlight);
```

6.2.5. Closing the KafkaSender

When the `KafkaSender` is no longer required, the `KafkaSender` instance can be closed. The underlying `KafkaProducer` is closed, closing all client connections and freeing all memory used by the producer.

```
sender.close();
```

6.2.6. Access to the underlying KafkaProducer

Reactive applications may sometimes require access to the underlying producer instance to perform actions that are not exposed by the `KafkaSender` interface. For example, an application might need to know the number of partitions in a topic in order to choose the partition to send a record to. Operations that are not provided directly by `KafkaSender` like `send` can be run on the underlying `KafkaProducer` using `KafkaSender#doOnProducer`.

```
sender.doOnProducer(producer -> producer.partitionsFor(topic))
    .doOnSuccess(partitions -> System.out.println("Partitions " + partitions))
    .subscribe();
```

User provided methods are executed asynchronously. A `Mono` is returned by `doOnProducer` which completes with the value returned by the user-provided function.

6.3. Reactive Kafka Receiver

Messages stored in Kafka topics are consumed using the reactive receiver `reactor.kafka.receiver.KafkaReceiver`. Each instance of `KafkaReceiver` is associated with a single instance of `KafkaConsumer`. `KafkaReceiver` is not thread-safe since the underlying `KafkaConsumer` cannot be accessed concurrently by multiple threads.

A receiver is created with an instance of receiver configuration options `reactor.kafka.receiver.ReceiverOptions`. Changes made to `ReceiverOptions` after the creation of the receiver instance will not be used by the `KafkaReceiver`. The properties of `ReceiverOptions` such as a list of bootstrap Kafka brokers and de-serializers are passed down to the underlying `KafkaConsumer`. These properties may be configured on the `ReceiverOptions` instance at creation time or by using the setter `ReceiverOptions#consumerProperty`. Other configuration options for the reactive `KafkaReceiver` including subscription topics must be added to options before the `KafkaReceiver` instance is created.

The generic types of `ReceiverOptions<K, V>` and `KafkaReceiver<K, V>` are the key and value types of consumer records consumed using the receiver and corresponding de-serializers must be set on the `ReceiverOptions` instance before the `KafkaReceiver` is created.

```
Map<String, Object> consumerProps = new HashMap<>();
consumerProps.put(ConsumerConfig.BOOTSTRAP_SERVERS_CONFIG, bootstrapServers);
consumerProps.put(ConsumerConfig.GROUP_ID_CONFIG, "sample-group");
consumerProps.put(ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG, IntegerDeserializer
.class);
consumerProps.put(ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG, StringDeserializer
.class);

ReceiverOptions<Integer, String> receiverOptions =
    ReceiverOptions.<Integer, String>create(consumerProps)           ①
        .subscription(Collections.singleton(topic));                ②
```

① Specify properties to be provided to `KafkaConsumer`

② Topics to subscribe to

Once the required configuration options have been configured on the options instance, a new `KafkaReceiver` instance can be created with these options to consume inbound messages. The code block below creates a receiver instance and creates an inbound Flux for the receiver. The underlying `KafkaConsumer` instance is created lazily later when the inbound Flux is subscribed to.

```
Flux<ReceiverRecord<Integer, String>> inboundFlux =
    KafkaReceiver.create(receiverOptions)
        .receive();
```

The inbound Kafka Flux is ready to be consumed. Each inbound message delivered by the Flux is represented as a `ReceiverRecord`. Each receiver record is a `ConsumerRecord` returned by `KafkaConsumer` along with a committable `ReceiverOffset` instance. The offset must be acknowledged after the message is processed since unacknowledged offsets will not be committed. If commit interval or commit batch size are configured, acknowledged offsets will be committed periodically. Offsets may also be committed manually using `ReceiverOffset#commit()` if finer grained control of commit operations is required.

```
inboundFlux.subscribe(r -> {  
    System.out.printf("Received message: %s\n", r);           ①  
    r.receiverOffset().acknowledge();                         ②  
});
```

① Prints each consumer record from Kafka

② Acknowledges that the record has been processed so that the offset may be committed

6.3.1. Subscribing to wildcard patterns

The example above subscribed to a single Kafka topic. The same API can be used to subscribe to more than one topic by specifying multiple topics in the collection provided to `ReceiverOptions#subscription()`. Subscription can also be made to a wildcard pattern by specifying a pattern to subscribe to. Group management in `KafkaConsumer` dynamically updates topic assignment when topics matching the pattern are created or deleted and assigns partitions of matching topics to available consumer instances.

```
receiverOptions = receiverOptions.subscription(Pattern.compile("demo.*")); ①
```

① Consume records from all topics starting with "demo"

Changes to `ReceiverOptions` must be made before the receiver instance is created. Altering the subscription deletes any existing subscriptions on the options instance.

6.3.2. Manual assignment of topic partitions

Partitions may be manually assigned to the receiver without using Kafka consumer group management.

```
receiverOptions = receiverOptions.assignment(Collections.singleton(new TopicPartition  
(topic, 0))); ①
```

① Consume from partition 0 of specified topic

Existing subscriptions and assignments on the options instance are deleted when a new assignment is specified. Every receiver created from this options instance with manual assignment consumes messages from all the specified partitions.

6.3.3. Controlling commit frequency

Commit frequency can be controlled using a combination of commit interval and commit batch size. Commits are performed when either the interval or batch size is reached. One or both of these options may be set on `ReceiverOptions` before the receiver instance is created. If commit interval is configured, at least one commit is scheduled within that interval if any records were consumed. If commit batch size is configured, a commit is scheduled when the configured number of records are consumed and acknowledged.

Manual acknowledgement of consumed records after processing along with automatic commits based on the configured commit frequency provides at-least-once delivery semantics. Messages are re-delivered if the consuming application crashes after message was dispatched but before it was processed and acknowledged. Only offsets explicitly acknowledged using `ReceiverOffset#acknowledge()` are committed. Note that acknowledging an offset acknowledges all previous offsets on the same partition. All acknowledged offsets are committed when partitions are revoked during rebalance and when the receive Flux is terminated.

Applications which require fine-grained control over the timing of commit operations can disable periodic commits and explicitly invoke `ReceiverOffset#commit()` when required to trigger a commit. This commit is asynchronous by default, but the application may invoke `Mono#block()` on the returned `Mono` to implement synchronous commits. Applications may batch commits by acknowledging messages as they are consumed and invoking `commit()` periodically to commit acknowledged offsets.

```
receiver.receive()
    .doOnNext(r -> {
        process(r);
        r.receiverOffset().commit().block();
    });
```

Note that committing an offset acknowledges and commits all previous offsets on that partition. All acknowledged offsets are committed when partitions are revoked during rebalance and when the receive Flux is terminated.

6.3.4. Auto-acknowledgement of batches of records

`KafkaReceiver#receiveAutoAck` returns a `Flux` of batches of records returned by each `KafkaConsumer#poll()`. The records in each batch are automatically acknowledged when the Flux corresponding to the batch terminates.

```
KafkaReceiver.create(receiverOptions)
    .receiveAutoAck()
    .concatMap(r -> r)
    .subscribe(r -> System.out.println("Received: " + r));
```

① Concatenate in order

② Print out each consumer record received, no explicit ack required

The maximum number of records in each batch can be controlled using the `KafkaConsumer` property `MAX_POLL_RECORDS`. This is used together with the fetch size and wait times configured on the `KafkaConsumer` to control the amount of data fetched from Kafka brokers in each poll. Each batch is returned as a Flux that is acknowledged after the Flux terminates. Acknowledged records are committed periodically based on the configured commit interval and batch size. This mode is simple to use since applications do not need to perform any acknowledge or commit actions. It is efficient as well and can be used for at-least-once delivery of messages.

6.3.5. Disabling automatic commits

Applications which don't require offset commits to Kafka may disable automatic commits by not acknowledging any records consumed using `KafkaReceiver#receive()`.

```
receiverOptions = ReceiverOptions.<Integer, String>create()
    .commitInterval(Duration.ZERO)           ①
    .commitBatchSize(0);                     ②
KafkaReceiver.create(receiverOptions)
    .receive()
    .subscribe(r -> process(r));           ③
```

- ① Disable periodic commits
- ② Disable commits based on batch size
- ③ Process records, but don't acknowledge

6.3.6. At-most-once delivery

Applications may disable automatic commits to avoid re-delivery of records. `ConsumerConfig#AUTO_OFFSET_RESET_CONFIG` can be configured to "latest" to consume only new records. But this could mean that an unpredictable number of records are not consumed if an application fails and restarts.

`KafkaReceiver#receiveAtmostOnce` can be used to consume records with at-most-once semantics with a configurable number of records-per-partition that may be lost if the application fails or crashes. Offsets are committed synchronously before the corresponding record is dispatched. Records are guaranteed not to be re-delivered even if the consuming application fails, but some records may not be processed if an application fails after the commit before the records could be processed.

This mode is expensive since each record is committed individually and records are not delivered until the commit operation succeeds. `ReceiverOptions#atmostOnceCommitCommitAheadSize` may be configured to reduce the cost of commits and avoid blocking before dispatch if the offset of the record has already been committed. By default, commit-ahead is disabled and at-most one record is lost per-partition if an application crashes. If commit-ahead is configured, the maximum number of records that may be lost per-partition is `ReceiverOptions#atmostOnceCommitCommitAheadSize + 1`.


```
KafkaReceiver.create(receiverOptions)
    .receiveAtMostOnce()
    .subscribe(r -> System.out.println("Received: " + r)); ①
```

① Process each consumer record, this record is not re-delivered if the processing fails

6.3.7. Partition assignment and revocation listeners

Applications can enable assignment and revocation listeners to perform any actions when partitions are assigned or revoked from a consumer.

When group management is used, assignment listeners are invoked whenever partitions are assigned to the consumer after a rebalance operation. When manual assignment is used, assignment listeners are invoked when the consumer is started. Assignment listeners can be used to seek to particular offsets in the assigned partitions so that messages are consumed from the specified offset.

When group management is used, revocation listeners are invoked whenever partitions are revoked from a consumer after a rebalance operation. When manual assignment is used, revocation listeners are invoked before the consumer is closed. Revocation listeners can be used to commit processed offsets when manual commits are used. Acknowledged offsets are automatically committed on revocation if automatic commits are enabled.

6.3.8. Controlling start offsets for consuming records

By default, receivers start consuming records from the last committed offset of each assigned partition. If a committed offset is not available, the offset reset strategy `ConsumerConfig#AUTO_OFFSET_RESET_CONFIG` configured for the `KafkaConsumer` is used to set the start offset to the earliest or latest offset on the partition. Applications can override offsets by seeking to new offsets in an assignment listener. Methods are provided on `ReceiverPartition` to seek to the earliest, latest or a specific offset in the partition.

```
void seekToBeginning();
void seekToEnd();
void seek(long offset);
```

For example, the following code block starts consuming messages from the latest offset.

```
receiverOptions = receiverOptions
    .addAssignListener(partitions -> partitions.forEach(p -> p.seekToEnd()))
    .subscription(Collections.singleton(topic));
KafkaReceiver.create(receiverOptions).receive().subscribe();
```

① Seek to the last offset in each assigned partition

6.3.9. Consumer lifecycle

Each `KafkaReceiver` instance is associated with a `KafkaConsumer` that is created when the inbound Flux returned by one of the receive methods in `KafkaReceiver` is subscribed to. The consumer is kept alive until the Flux completes. When the Flux completes, all acknowledged offsets are committed and the underlying consumer is closed.

Only one receive operation may be active in a `KafkaReceiver` at any one time. Any of the receive methods can be invoked after the receive Flux corresponding to the last receive is terminated.

Chapter 7. Sample Scenarios

This section shows sample code segments for typical scenarios where Reactor Kafka API may be used. Full code listing for these scenarios are included in the [samples sub-project](#).

7.1. Sending records to Kafka

See [KafkaSender API](#) for details on the KafkaSender API for sending outbound records to Kafka. The following code segment creates a simple pipeline that sends records to Kafka and processes the responses. The outbound flow is triggered when the returned Flux is subscribed to.

```
KafkaSender.create(SenderOptions.<Integer, String>create(producerProps).maxInFlight(512)) ①
    .send(outbound.map(r -> senderRecord(r)))
    .doOnNext(result -> processResponse(result)) ②
    .doOnError(e -> processError(e)); ③
```

- ① Create a sender with maximum 512 messages in-flight
- ② Send a sequence of sender records
- ③ Process send result when onNext is triggered

7.2. Replaying records from Kafka topics

See [KafkaReceiver API](#) for details on the KafkaReceiver API for consuming records from Kafka topics. The following code segment creates a Flux that replays all records on a topic and commits offsets after processing the messages. Manual acknowledgement provides at-least-once delivery semantics.

```
ReceiverOptions<Integer, String> options =
    ReceiverOptions.<Integer, String>create(consumerProps)
        .consumerProperty(ConsumerConfig.AUTO_OFFSET_RESET_CONFIG,
"earliest") ①
        .commitBatchSize(10)
        .subscription(Collections.singleton("demo-topic")); ②
KafkaReceiver.create(options)
    .receive()
    .doOnNext(r -> {
        processRecord(r); ④
        r.receiverOffset().acknowledge(); ⑤
    });
```

- ① Start consuming from first available offset on each partition if committed offsets are not

available

- ② Commit every 10 acknowledged messages
- ③ Topics to consume from
- ④ Process consumer record from Kafka
- ⑤ Acknowledge that record has been consumed

7.3. Reactive pipeline with Kafka sink

The code segment below consumes messages from an external source, performs some transformation and stores the output records in Kafka. Large number of retry attempts are configured on the Kafka producer so that transient failures don't impact the pipeline. Source commits are performed only after records are successfully written to Kafka.

```
senderOptions = senderOptions
    .producerProperty(ProducerConfig.ACKS_CONFIG, "all")           ①
    .producerProperty(ProducerConfig.RETRIES_CONFIG, Integer.MAX_VALUE) ②
    .maxInFlight(128);                                             ③
KafkaSender.create(senderOptions)
    .send(source.flux().map(r -> transform(r)))                    ④
    .doOnError(e-> log.error("Send failed, terminating.", e))      ⑤
    .doOnNext(r -> source.commit(r.correlationMetadata()));        ⑥
```

- ① Send is acknowledged by Kafka for acks=all after message is delivered to all in-sync replicas
- ② Large number of retries in the producer to cope with transient failures in brokers
- ③ Low in-flight count to avoid filling up producer buffer and blocking the pipeline, default stopOnError=true
- ④ Receive from external source, transform and send to Kafka
- ⑤ If a send fails, it indicates catastrophic error, fail the whole pipeline
- ⑥ Use correlation metadata in the sender record to commit source record

7.4. Reactive pipeline with Kafka source

The code segment below consumes records from Kafka topics, transforms the record and sends the output to an external sink. Kafka consumer offsets are committed after records are successfully output to sink.

```

receiverOptions = receiverOptions
    .commitInterval(Duration.ZERO)           ①
    .commitBatchSize(0)                     ②
    .subscription(Pattern.compile(topics));  ③
KafkaReceiver.create(receiverOptions)
    .receive()
    .publishOn(Schedulers.newSingle("sample", true))
    .concatMap(m -> sink.store(transform(m))
    .doOnSuccess(r -> m.receiverOffset().commit().block())
); ④ ⑤

```

- ① Disable periodic commits
- ② Disable commits by batch size
- ③ Wildcard subscription
- ④ Transform Kafka record and store in external sink
- ⑤ Synchronous commit after record is successfully delivered to sink

7.5. Reactive pipeline with Kafka source and sink

The code segment below consumes messages from Kafka topic, performs some transformation on the incoming messages and stores the result in some Kafka topics. Manual acknowledgement mode provides at-least-once semantics with messages acknowledged after the output records are delivered to Kafka. Acknowledged offsets are committed periodically based on the configured commit interval.

```

receiverOptions = receiverOptions
    .commitInterval(Duration.ofSeconds(10)) ①
    .subscription(Pattern.compile(topics));
sender.send(KafkaReceiver.create(receiverOptions)
    .receive()
    .map(m -> SenderRecord.create(transform(m.value()), m
    .receiverOffset())) ②
    .doOnNext(m -> m.correlationMetadata().acknowledge()); ③

```

- ① Configure interval for automatic commits
- ② Transform incoming record and create outbound record with transformed data in the payload and inbound offset as correlation metadata
- ③ Acknowledge the inbound offset using the offset instance in correlation metadata after outbound record is delivered to Kafka

7.6. At-most-once delivery

The code segment below demonstrates a flow with at-most once delivery. Producer does not wait for acks and does not perform any retries. Messages that cannot be delivered to Kafka on the first

attempt are dropped. `KafkaReceiver` commits offsets before delivery to the application to ensure that if the consumer restarts, messages are not redelivered. With replication factor 1 for topic partitions, this code can be used for at-most-once delivery.

```
senderOptions = senderOptions
    .producerProperty(ProducerConfig.ACKS_CONFIG, "0") ①
    .producerProperty(ProducerConfig.RETRIES_CONFIG, "0") ②
    .stopOnError(false); ③
receiverOptions = receiverOptions
    .subscription(Collections.singleton(sourceTopic));
KafkaSender.create(senderOptions)
    .send(KafkaReceiver.create(receiverOptions)
        .receiveAtmostOnce() ③
        .map(cr -> SenderRecord.create(transform(cr.value()),
            cr.offset())));
```

- ① Send with acks=0 completes when message is buffered locally, before it is delivered to Kafka broker
- ② No retries in producer
- ③ Ignore any error and continue to send remaining records
- ④ At-most-once receive

7.7. Fan-out with Multiple Streams

The code segment below demonstrates fan-out with the same records processed in multiple independent streams. Each stream is processed on a different thread and which transforms the input record and stores the output in a Kafka topic.

Reactor's `EmitterProcessor` is used to broadcast the input records from Kafka to multiple subscribers.

```
EmitterProcessor<Person> processor = EmitterProcessor.create(); ①
BlockingSink<Person> incoming = processor.connectSink(); ②
inputRecords = KafkaReceiver.create(receiverOptions)
    .receive()
    .doOnNext(m -> incoming.emit(m.value())); ③

outputRecords1 = processor.publishOn(scheduler1).map(p -> process1(p)); ④
outputRecords2 = processor.publishOn(scheduler2).map(p -> process2(p)); ⑤

Flux.merge(sender.send(outputRecords1), sender.send(outputRecords2))
    .doOnSubscribe(s -> inputRecords.subscribe())
    .subscribe(); ⑥
```

- ① Create publish/subscribe `EmitterProcessor` for fan-out of Kafka inbound records
- ② Create `BlockingSink` to which records are emitted

- ③ Receive from Kafka and emit to BlockingSink
- ④ Consume records on a scheduler, process and generate output records to send to Kafka
- ⑤ Add another processor for the same input data on a different scheduler
- ⑥ Merge the streams and subscribe to start the flow

7.8. Concurrent Processing with Partition-Based Ordering

The code segment below demonstrates a flow where messages are consumed from a Kafka topic, processed by multiple threads and the results stored in another Kafka topic. Messages are grouped by partition to guarantee ordering in message processing and commit operations. Messages from each partition are processed on a single thread.

```
Scheduler scheduler = Schedulers.newElastic("sample", 60, true);
KafkaReceiver.create(receiverOptions)
    .receive()
    .groupBy(m -> m.receiverOffset().topicPartition())
    .flatMap(partitionFlux ->
        partitionFlux.publishOn(scheduler)
            .map(r -> processRecord(partitionFlux.key(), r))
            .sample(Duration.ofMillis(5000))
            .concatMap(offset -> offset.commit()));
```

- ① Group by partition to guarantee ordering
- ② Commit periodically
- ③ Commit in sequence using concatMap