



# Red Hat Integration 2020-Q4

## Getting Started with Camel Kafka Connector

TECHNOLOGY PREVIEW - Using Camel components as Kafka connectors



# Red Hat Integration 2020-Q4 Getting Started with Camel Kafka Connector

---

TECHNOLOGY PREVIEW - Using Camel components as Kafka connectors

## Legal Notice

Copyright © 2021 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This guide introduces Camel Kafka Connector, explains how to install into AMQ Streams and Kafka Connect on OpenShift, and how to get started with example Camel Kafka connectors. This guide also describes the Camel Kafka connectors that you can configure in this release.

## Table of Contents

<b>CHAPTER 1. INTRODUCTION TO CAMEL KAFKA CONNECTOR</b> .....	<b>3</b>
1.1. CAMEL KAFKA CONNECTOR OVERVIEW	3
1.2. CAMEL KAFKA CONNECTOR FEATURES	4
1.2.1. Platforms and components	4
1.2.2. Technology Preview features	4
1.2.3. Camel Kafka connectors	4
1.2.4. Camel data formats	5
1.3. CAMEL KAFKA CONNECTOR ARCHITECTURE	5
Kafka Connect concepts	6
Camel Kafka Connector configuration	6
1.4. CAMEL KAFKA CONNECTOR DISTRIBUTIONS	6
<b>CHAPTER 2. DEPLOYING CAMEL KAFKA CONNECTOR WITH AMQ STREAMS ON OPENSIFT</b> .....	<b>8</b>
2.1. CONFIGURING AUTHENTICATION WITH REGISTRY.REDHAT.IO	8
2.2. INSTALLING AMQ STREAMS AND KAFKA CONNECT S2I ON OPENSIFT	9
2.3. DEPLOYING CAMEL KAFKA CONNECTOR USING KAFKA CONNECT S2I ON OPENSIFT	11
<b>CHAPTER 3. DEPLOYING CAMEL KAFKA CONNECTOR DEVELOPER EXAMPLES</b> .....	<b>15</b>
3.1. DEPLOYING CAMEL KAFKA CONNECTOR EXAMPLES ON OPENSIFT	15
3.2. DEPLOYING CAMEL KAFKA CONNECTOR EXAMPLES ON RHEL	15
<b>CHAPTER 4. EXTENDING CAMEL KAFKA CONNECTOR</b> .....	<b>17</b>
4.1. CONFIGURING A CAMEL KAFKA CONNECTOR AGGREGATOR	17
4.2. WRITING A CUSTOM CAMEL KAFKA CONNECTOR AGGREGATOR	18
4.3. CONFIGURING CAMEL DATA FORMATS IN CAMEL KAFKA CONNECTOR	20
4.4. EXTENDING CAMEL KAFKA CONNECTORS USING MAVEN ARCHETYPES	21
<b>CHAPTER 5. CAMEL KAFKA CONNECTOR CONFIGURATION REFERENCE</b> .....	<b>24</b>
5.1. CAMEL-AWS2-KINESIS-KAFKA-CONNECTOR SINK CONFIGURATION	25
5.2. CAMEL-AWS2-KINESIS-KAFKA-CONNECTOR SOURCE CONFIGURATION	28
5.3. CAMEL-AWS2-S3-KAFKA-CONNECTOR SINK CONFIGURATION	33
5.4. CAMEL-AWS2-S3-KAFKA-CONNECTOR SOURCE CONFIGURATION	39
5.5. CAMEL-AWS2-SNS-KAFKA-CONNECTOR SINK CONFIGURATION	49
5.6. CAMEL-AWS2-SQS-KAFKA-CONNECTOR SINK CONFIGURATION	53
5.7. CAMEL-AWS2-SQS-KAFKA-CONNECTOR SOURCE CONFIGURATION	59
5.8. CAMEL-CQL-KAFKA-CONNECTOR SINK CONFIGURATION	68
5.9. CAMEL-CQL-KAFKA-CONNECTOR SOURCE CONFIGURATION	70
5.10. CAMEL-ELASTICSEARCH-REST-KAFKA-CONNECTOR SINK CONFIGURATION	74
5.11. CAMEL-FILE-KAFKA-CONNECTOR SINK CONFIGURATION	78
5.12. CAMEL-HDFS-KAFKA-CONNECTOR SINK CONFIGURATION	85
5.13. CAMEL-HTTP-KAFKA-CONNECTOR SINK CONFIGURATION	89
5.14. CAMEL-JDBC-KAFKA-CONNECTOR SINK CONFIGURATION	98
5.15. CAMEL-JMS-KAFKA-CONNECTOR SINK CONFIGURATION	101
5.16. CAMEL-JMS-KAFKA-CONNECTOR SOURCE CONFIGURATION	124
5.17. CAMEL-MONGODB-KAFKA-CONNECTOR SINK CONFIGURATION	149
5.18. CAMEL-MONGODB-KAFKA-CONNECTOR SOURCE CONFIGURATION	153
5.19. CAMEL-NETTY-KAFKA-CONNECTOR SOURCE CONFIGURATION	157
5.20. CAMEL-SALESFORCE-KAFKA-CONNECTOR SOURCE CONFIGURATION	170
5.21. CAMEL-SLACK-KAFKA-CONNECTOR SOURCE CONFIGURATION	181
5.22. CAMEL-SYSLOG-KAFKA-CONNECTOR SOURCE CONFIGURATION	184
5.23. CAMEL-TIMER-KAFKA-CONNECTOR SOURCE CONFIGURATION	185



# CHAPTER 1. INTRODUCTION TO CAMEL KAFKA CONNECTOR

This chapter introduces the features, concepts, and distributions provided by Camel Kafka Connector:

- [Section 1.1, "Camel Kafka Connector overview"](#)
- [Section 1.2, "Camel Kafka Connector features"](#)
- [Section 1.3, "Camel Kafka Connector architecture"](#)
- [Section 1.4, "Camel Kafka Connector distributions"](#)



## IMPORTANT

Camel Kafka Connector is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production.

These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information about the support scope of Red Hat Technology Preview features, see <https://access.redhat.com/support/offerings/techpreview>.

## 1.1. CAMEL KAFKA CONNECTOR OVERVIEW

Apache Camel is a highly flexible open source integration framework for connecting a wide range of different systems, which is based on standard Enterprise Integration Patterns (EIPs). Apache Kafka Connect is the Kafka-native approach for connecting to external systems, which is specifically designed for event-driven architectures.

Camel Kafka Connector enables you to use standard Camel components as Kafka Connect connectors. This widens the scope of possible integrations beyond the external systems supported by Kafka Connect connectors alone. Camel Kafka Connector works as an adapter that makes the popular Camel component ecosystem available in Kafka-based AMQ Streams on OpenShift.

Camel Kafka Connector provides a user-friendly way to configure Camel components directly in the Kafka Connect framework. Using Camel Kafka Connector, you can leverage Camel components for integration with different systems by connecting to or from Camel Kafka sink or source connectors. You do not need to write any code, and can include the appropriate connector JARs in your Kafka Connect image and configure connector options using custom resources.

Camel Kafka Connector is built on Apache Camel Kafka Connector, which is a subproject of the Apache Camel open source community. Camel Kafka Connector is fully integrated with AMQ Streams and Kafka Connect, and is available on both OpenShift Container Platform and Red Hat Enterprise Linux.

Camel Kafka Connector is available with the Red Hat Integration - Camel K distribution for cloud-native integration on OpenShift. Camel K is a lightweight integration framework built from Apache Camel K that runs natively in the cloud on OpenShift. Camel K is specifically designed for serverless and microservice architectures.

### Additional resources

- [Apache Camel Kafka Connector GitHub project](#)

- [Apache Camel Kafka Connector website](#)
- [Deploying Red Hat Integration - Camel K on OpenShift](#)

## 1.2. CAMEL KAFKA CONNECTOR FEATURES

The Camel Kafka Connector Technology Preview includes the following main features:

### 1.2.1. Platforms and components

- OpenShift Container Platform 4.5 or 4.6
- Red Hat Enterprise Linux 7.x or 8.x
- AMQ Streams 1.5
- Kafka Connect 2.5
- Camel 3.5
- OpenJDK 8 or 11

### 1.2.2. Technology Preview features

- Selected Camel Kafka connectors
- Marshaling/unmarshalling of Camel data formats for sink and source connectors
- Aggregation for sink connectors
- Maven archetypes for extending connectors

### 1.2.3. Camel Kafka connectors

Table 1.1. Camel Kafka connectors in Technology Preview

Connector	Sink/source
Amazon Web Services (AWS2) Kinesis	Sink and source
Amazon Web Services (AWS2) S3	Sink and source
Amazon Web Services (AWS2) SNS	Sink only
Amazon Web Services (AWS2) SQS	Sink and source
Cassandra Query Language (CQL)	Sink and source
Elasticsearch	Sink only
File	Sink only

Connector	Sink/source
Hadoop Distributed File System (HDFS)	Sink only
Hypertext Transfer Protocol (HTTP)	Sink only
Java Database Connectivity (JDBC)	Sink only
Java Message Service (JMS)	Sink and source
MongoDB	Sink and source
Salesforce	Source only
Slack	Source only
Syslog	Source only
Timer	Source only

### 1.2.4. Camel data formats

The Camel Kafka Connector Technology Preview includes marshaling and unmarshaling of Camel data formats. For example, these formats include Apache Avro, Base64, Google Protobuf, JSON, SOAP, Zip file, and many more. You can configure marshaling and unmarshaling of Camel data formats using properties in your Camel Kafka Connector, configuration.

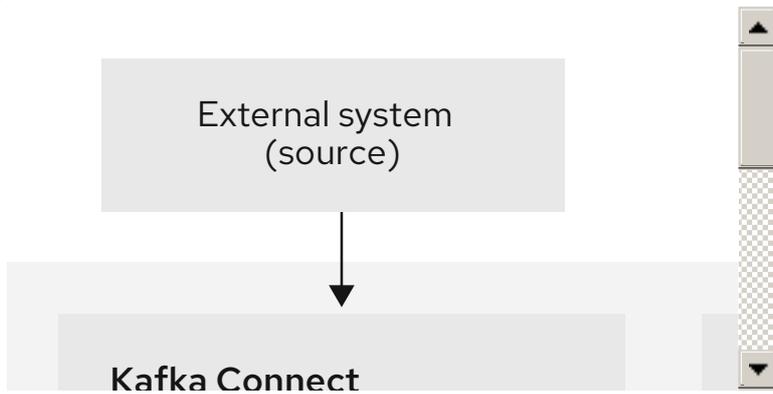
## 1.3. CAMEL KAFKA CONNECTOR ARCHITECTURE

AMQ Streams is a distributed and scalable streaming platform based on Apache Kafka that includes a publish/subscribe messaging broker. Kafka Connect provides a framework to integrate Kafka-based systems with external systems. Using Kafka Connect, you can configure *source* and *sink* connectors to stream data from external systems into and out of a Kafka broker.

Camel Kafka Connector reuses the flexibility of Camel components and makes them available in Kafka Connect as source and sink connectors that you can use to stream data into and out of AMQ Streams. For example, you can ingest data from Amazon Web Services for processing using an AWS S3 source connector, or consolidate events stored in Kafka into an Elasticsearch instance for analytics using an Elasticsearch sink connector.

The following diagram shows a simplified view of the Camel Kafka Connector cloud-native integration architecture based on AMQ Streams:

Figure 1.1. Camel Kafka Connector architecture



## Kafka Connect concepts

### Source connector

Source connectors work like consumers and pull data from external systems into Kafka topics to make the data available for stream processing. For example, these external source systems include Amazon Web Services or Java Message Service.

### Sink connector

Sink connectors work like producers and push data from Kafka topics into external systems for offline analysis. For example, these external sink systems include Cassandra, Syslog, or Elasticsearch.

### Sink/source task

Tasks are typically created by a sink or source connector and are responsible for handling the data.

### Key/value converter

Key/value converters can serialize/deserialize the key or value of a Kafka message in various formats.

### Transformer

Transformers can manipulate Kafka message content, for example, renaming fields or routing to topics based on values.

### Aggregator

Sink connectors can use an aggregator to batch up records before sending them to an external system.

## Camel Kafka Connector configuration

You can use Camel Kafka Connector configuration to specify the following:

- Kafka Connect configuration options
- Camel route definitions
- Camel configuration options

### Additional resources

- [Apache Kafka Connect user documentation](#)

## 1.4. CAMEL KAFKA CONNECTOR DISTRIBUTIONS

The Camel Kafka Connector distributions are bundled with Red Hat Integration - Camel K:

Table 1.2. Camel Kafka Connector available distributions

Distribution	Description	Location
Maven repository	Maven artifacts for Camel Kafka Connector	<a href="#">Software Downloads &gt; Red Hat Integration</a>
Source code	Source code for Camel Kafka Connector	<a href="#">Software Downloads &gt; Red Hat Integration</a>
Demonstration examples	Camel Kafka Connector examples and Debezium community example	<ul style="list-style-type: none"><li>● <a href="#">Camel Kafka Connector examples</a></li><li>● <a href="#">Debezium PostgreSQL connector (Debezium community)</a></li></ul>

**NOTE**

You must have a subscription for Red Hat Integration and be logged into the Red Hat Customer Portal to access the Camel Kafka Connector distributions available with Red Hat Integration - Camel K.

## CHAPTER 2. DEPLOYING CAMEL KAFKA CONNECTOR WITH AMQ STREAMS ON OPENSIFT

This chapter explains how to install Camel Kafka Connector into AMQ Streams on OpenShift and how to get started with example connectors.

- [Section 2.1, “Configuring authentication with registry.redhat.io”](#)
- [Section 2.2, “Installing AMQ Streams and Kafka Connect S2I on OpenShift”](#)
- [Section 2.3, “Deploying Camel Kafka Connector using Kafka Connect S2I on OpenShift”](#)

### 2.1. CONFIGURING AUTHENTICATION WITH REGISTRY.REDHAT.IO

You must configure authentication with the **registry.redhat.io** container registry before you can use AMQ Streams and Kafka Connect Source-2-Image (S2I) to deploy Camel Kafka Connector on OpenShift.

#### Prerequisites

- You must have cluster administrator access to an OpenShift Container Platform cluster.
- You must have the OpenShift **oc** client tool installed. For more details, see the [OpenShift CLI documentation](#).

#### Procedure

1. Log into your OpenShift cluster as administrator, for example:

```
$ oc login --user system:admin --token=my-token --server=https://my-cluster.example.com:6443
```

2. Open the project in which you want to deploy Camel Kafka Connector, for example:

```
$ oc project myproject
```

3. Create a **docker-registry** secret using your Red Hat Customer Portal account, and replace **PULL\_SECRET\_NAME** with the name of the secret that you want to create:

```
$ oc create secret docker-registry PULL_SECRET_NAME \
  --docker-server=registry.redhat.io \
  --docker-username=CUSTOMER_PORTAL_USERNAME \
  --docker-password=CUSTOMER_PORTAL_PASSWORD \
  --docker-email=EMAIL_ADDRESS
```

You should see the following output:

```
secret/PULL_SECRET_NAME created
```



### IMPORTANT

You must create this pull secret in every OpenShift project namespace that will include the image streams and use **registry.redhat.io**.

4. Link the secret to your service account to use the secret for pulling images. The following example uses the **default** service account:

```
$ oc secrets link default PULL_SECRET_NAME --for=pull
```

The service account name must match the name that the service account Pod uses.

5. Link the secret to the **builder** service account in the namespace in which you plan to use Kafka Connect S2I:

```
$ oc secrets link builder PULL_SECRET_NAME
```



### NOTE

If you do not wish to use your Red Hat account username and password to create the pull secret, you should create an authentication token by using a registry service account.

#### Additional resources

- [Red Hat Registry authentication](#)
- [Red Hat Registry service accounts](#)

## 2.2. INSTALLING AMQ STREAMS AND KAFKA CONNECT S2I ON OPENSIFT

AMQ Streams and Kafka Connect with Source-2-Image (S2I) are required to install Camel Kafka Connector. If you do not already have AMQ Streams installed, you can install the AMQ Streams Operator on your OpenShift cluster from the OperatorHub. The OperatorHub is available from the OpenShift Container Platform web console and provides an interface for cluster administrators to discover and install Operators. For more details, see the [OpenShift documentation](#).

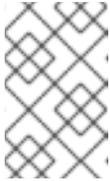
#### Prerequisites

- You must have cluster administrator access to an OpenShift Container Platform cluster.
- You must have authenticated with **registry.redhat.io** using the steps in [Section 2.1, “Configuring authentication with registry.redhat.io”](#).
- See [Using AMQ Streams on OpenShift](#) for detailed information on installing AMQ Streams and Kafka Connect S2I. This section shows a simple default example of installing using the OpenShift OperatorHub.

#### Procedure

1. In the OpenShift Container Platform web console, log in using an account with cluster administrator privileges.

2. Select your project from the **Project** drop-down in the toolbar, for example, **myproject**. This must be the project in which you have authenticated with **registry.redhat.io**.
3. In the left navigation menu, click **Operators > OperatorHub**.
4. In the **Filter by keyword** text box, enter **AMQ** to find the **Red Hat Integration - AMQ Streams Operator**.
5. Read the information about the Operator, and click **Install** to display the Operator subscription page.
6. Select your subscription settings, for example:
  - **Update Channel > stable**
  - **Installation Mode > A specific namespace on the cluster > myproject**
  - **Approval Strategy > Automatic**



#### NOTE

These settings depend on the specific requirements of your environment. For more details, see [OpenShift documentation on Adding Operators to a cluster](#).

7. Click **Install**, and wait a few moments until the Operator is ready for use.
8. Create a new Kafka broker cluster:
  - a. Under **Red Hat Integration - AMQ Streams > Provided APIs > Kafka**, click **Create Instance** to create a new Kafka broker cluster.
  - b. Edit the custom resource definition as appropriate, and click **Create**.



#### IMPORTANT

The default example creates a Kafka cluster with 3 Zookeeper nodes and 3 Kafka nodes with **ephemeral** storage. This temporary storage is suitable for development and testing only, and not for a production environment. For more details, see [Using AMQ Streams on OpenShift](#).

9. Create a new Kafka Connect S2I cluster:
  - a. Under **Red Hat Integration - AMQ Streams > Provided APIs > Kafka Connect S2I**, click **Create Instance** to create a new Kafka Connect cluster with OpenShift Source-2-Image support.
  - b. Edit the custom resource definition as appropriate, and click **Create**. For more details on using Kafka Connect with S2I, see [Using AMQ Streams on OpenShift](#).
10. Select **Workloads > Pods** to verify that the deployed resources are running on OpenShift.

#### Additional resources

- [OpenShift documentation on Adding Operators to a cluster](#)

## 2.3. DEPLOYING CAMEL KAFKA CONNECTOR USING KAFKA CONNECT S2I ON OPENSIFT

This section explains how to use Kafka Connect Source-2-Image (S2I) with AMQ Streams to add your Camel Kafka connectors to your existing Docker-based Kafka Connect image and to build a new image. This section also shows how to create an instance of a Camel Kafka connector plug-in using an example AWS2 S3 Camel Kafka connector.

### Prerequisites

- You must have cluster administrator access to an OpenShift Container Platform cluster.
- You must have installed AMQ Streams and Kafka Connect with S2I support on your OpenShift cluster. For more details, see [Section 2.2, "Installing AMQ Streams and Kafka Connect S2I on OpenShift"](#).
- You must have downloaded Camel Kafka Connector from [Software Downloads > Red Hat Integration](#).
- You must have access to an Amazon S3 bucket.

### Procedure

1. Log into your OpenShift cluster as administrator, for example:

```
$ oc login --user system:admin --token=my-token --server=https://my-cluster.example.com:6443
```

2. Change to the project in which Kafka Connect S2I is installed:

```
$ oc project myproject
```

3. Add your downloaded Camel Kafka connectors to the existing Kafka Connect Docker image build, and wait for the new image build to be configured with the new connectors. For example:

```
$ oc start-build my-connect-cluster-connect --from-dir=./camel-kafka-connector/connectors/ --follow
Uploading directory "camel-kafka-connector/connectors" as binary input for the build ...
...
Uploading finished
build.build.openshift.io/my-connect-cluster-connect-2 started
Receiving source from STDIN as archive ...
Caching blobs under "/var/cache/blobs".
Getting image source signatures
...
Writing manifest to image destination
Storing signatures
Generating dockerfile with builder image image-registry.openshift-image-registry.svc:5000/myproject/my-connect-cluster-connect-source@sha256:12d5ed92510941f1569faa449665e9fc6ea544e67b7ae189ec6b8df434e121f4
STEP 1: FROM image-registry.openshift-image-registry.svc:5000/myproject/my-connect-cluster-connect-source@sha256:12d5ed92510941f1569faa449665e9fc6ea544e67b7ae189ec6b8df434e121f
```

```

4
STEP 2: LABEL "io.openshift.build.image"="image-registry.openshift-image-
registry.svc:5000/myproject/my-connect-cluster-connect-
source@sha256:12d5ed92510941f1569faa449665e9fc6ea544e67b7ae189ec6b8df434e121f4"
"io.openshift.build.source-location"="/tmp/build/inputs"
STEP 3: ENV OPENSIFT_BUILD_NAME="my-connect-cluster-connect-2"
OPENSIFT_BUILD_NAMESPACE="myproject"
STEP 4: USER root
STEP 5: COPY upload/src /tmp/src
STEP 6: RUN chown -R 1001:0 /tmp/src
STEP 7: USER 1001
STEP 8: RUN /opt/kafka/s2i/assemble
Assembling plugins into custom plugin directory /tmp/kafka-plugins
Moving plugins to /tmp/kafka-plugins
STEP 9: CMD /opt/kafka/s2i/run
STEP 10: COMMIT temp.builder.openshift.io/myproject/my-connect-cluster-connect-
2:d0873588
Getting image source signatures
...
Writing manifest to image destination
Storing signatures
...
Pushing image image-registry.openshift-image-registry.svc:5000/myproject/my-connect-
cluster-connect:latest ...
Getting image source signatures
...
Writing manifest to image destination
Storing signatures
Successfully pushed image-registry.openshift-image-registry.svc:5000/myproject/my-
connect-cluster-
connect@sha256:9db57d33df6d0494ea6ee6e4696fcf79eb81aabee0bbc180dec5324d33e7ed
a
Push successful

```

4. Check that Camel Kafka Connector is available in your Kafka Connect cluster as follows:

```

$ oc exec -i `oc get pods --field-selector status.phase=Running -l strimzi.io/name=my-
connect-cluster-connect -o=jsonpath='{.items[0].metadata.name}'` -- curl -s http://my-
connect-cluster-connect-api:8083/connector-plugins

```

You should see something like the following output:

```

[{"class":"org.apache.kafka.connect.file.FileStreamSinkConnector","type":"sink","version":"2.5
.0.redhat-00003"},
{"class":"org.apache.kafka.connect.file.FileStreamSourceConnector","type":"source","version"
:"2.5.0.redhat-00003"},
{"class":"org.apache.kafka.connect.mirror.MirrorCheckpointConnector","type":"source","versi
on":"1"},
{"class":"org.apache.kafka.connect.mirror.MirrorHeartbeatConnector","type":"source","version
":"1"},
{"class":"org.apache.kafka.connect.mirror.MirrorSourceConnector","type":"source","version":"
1"}]]

```

5. Use the following annotation to enable instantiating Camel Kafka connectors using a specific custom resource:

```
$ oc annotate kafkaconnects2i my-connect-cluster strimzi.io/use-connector-resources=true
kafkaconnects2i.kafka.strimzi.io/my-connect-cluster annotated
```



## IMPORTANT

When the **use-connector-resources** option is enabled, do not use the Kafka Connect API server. The Kafka Connect Operator will revert any changes that you make.

6. Create the connector instance by creating a specific custom resource that includes your connector configuration. The following example shows the configuration for an AWS2 S3 connector plug-in:

```
$ oc apply -f - << EOF
apiVersion: kafka.strimzi.io/v1alpha1
kind: KafkaConnector
metadata:
  name: s3-source-connector
  namespace: myproject
  labels:
    strimzi.io/cluster: my-connect-cluster
spec:
  class: org.apache.camel.kafkaconnector.aws2s3.CamelAws2s3SourceConnector
  tasksMax: 1
  config:
    key.converter: org.apache.kafka.connect.storage.StringConverter
    value.converter: org.apache.kafka.connect.storage.StringConverter
    topics: s3-topic
    camel.source.path.bucketNameOrArn: camel-kafka-connector
    camel.source.maxPollDuration: 10000
    camel.component.aws2-s3.accessKey: xxxx
    camel.component.aws2-s3.secretKey: yyyy
    camel.component.aws2-s3.region: region
EOF
kafkaconnector.kafka.strimzi.io/s3-source-connector created
```

7. Check the status of your connector using the following example command:

```
$ oc exec -i `oc get pods --field-selector status.phase=Running -l strimzi.io/name=my-
connect-cluster-connect -o=jsonpath='{.items[0].metadata.name}'` -- curl -s http://my-
connect-cluster-connect-api:8083/connectors/s3-source-connector/status
```

8. Connect to your AWS Console, and upload a file to the **camel-kafka-connector** AWS S3 bucket to activate the Camel Kafka route.
9. You can run the Kafka console consumer to see the messages received from the topic as follows:

```
oc exec -i -c kafka my-cluster-kafka-0 -- bin/kafka-console-consumer.sh --bootstrap-server
localhost:9092 --topic s3-topic --from-beginning
CONTENTS_OF_FILE
CONTENTS_OF_FILE
...
```

### Additional resources

- [Apache Camel Kafka Connector installation instructions on OpenShift](#)
- [Using AMQ Streams on OpenShift](#)

## CHAPTER 3. DEPLOYING CAMEL KAFKA CONNECTOR DEVELOPER EXAMPLES

Camel Kafka Connector provides demonstration examples for selected connectors, which are available from <https://github.com/jboss-fuse/camel-kafka-connector-examples>. This chapter provides details on how to deploy these examples based on your Camel Kafka Connector installation platform:

- [Section 3.1, “Deploying Camel Kafka Connector examples on OpenShift”](#)
- [Section 3.2, “Deploying Camel Kafka Connector examples on RHEL”](#)

### 3.1. DEPLOYING CAMEL KAFKA CONNECTOR EXAMPLES ON OPENSHIFT

This section describes how to deploy Camel Kafka Connector demonstration examples for selected connectors on OpenShift.

#### Prerequisites

- Scroll down to see the *OpenShift - What is needed* section in each of the readmes shown in the Procedure that follows.

#### Procedure

1. Go to the GitHub readme for one of the following examples:
  - [AWS2 S3 connectors \(excluding MinIO source\)](#)
  - [AWS2 SNS sink connector](#)
  - [AWS2 SQS connectors](#)
2. Scroll down to the *OpenShift* section of the readme for your chosen example.
3. Perform the steps described in the readme to run the example.

#### Additional resources

- [Using AMQ Streams on OpenShift](#)

### 3.2. DEPLOYING CAMEL KAFKA CONNECTOR EXAMPLES ON RHEL

This section describes how to deploy Camel Kafka Connector demonstration examples for selected connectors on Red Hat Enterprise Linux.

#### Prerequisites

- See the *What is needed* section in each of the readmes shown in the Procedure that follows.

#### Procedure

1. Go to the GitHub readme for one of the following examples:

- [AWS2 S3 connectors](#)
- [AWS2 SNS sink connector](#)
- [AWS2 SQS connectors](#)
- [CQL connectors](#)

2. Perform the steps described in the readme to run the example.

#### **Additional resources**

- [Using AMQ Streams on RHEL](#)
- [Debezium PostgreSQL connector and Apache Camel Kafka Connector \(community example\)](#)

## CHAPTER 4. EXTENDING CAMEL KAFKA CONNECTOR

This chapter explains how to extend and customize Camel Kafka connectors and components. Camel Kafka Connector provides an easy way to configure Camel components directly in the Kafka Connect framework, without needing to write code. However, in some scenarios, you might want to extend and customize Camel Kafka Connector for specific use cases.

- [Section 4.1, “Configuring a Camel Kafka connector aggregator”](#)
- [Section 4.2, “Writing a custom Camel Kafka connector aggregator”](#)
- [Section 4.3, “Configuring Camel data formats in Camel Kafka Connector”](#)
- [Section 4.4, “Extending Camel Kafka connectors using Maven archetypes”](#)

### 4.1. CONFIGURING A CAMEL KAFKA CONNECTOR AGGREGATOR

In some scenarios using a Camel Kafka sink connector, you might want to add an aggregator to batch up your Kafka records before sending them to the external sink system. Typically, this involves defining a specific batch size and timeout for aggregation of records. When complete, the aggregate record is sent to the external system.

You can configure aggregation settings in your Camel Kafka Connector properties using one of the aggregators provided by Apache Camel, or you can implement a custom aggregator in Java. This section describes how to configure the Camel aggregator settings in your Camel Kafka Connector properties.

#### Prerequisites

- You must have installed Camel Kafka Connector, for example, see [Section 2.2, “Installing AMQ Streams and Kafka Connect S2I on OpenShift”](#).
- You must have deployed your sink connector, for example, see [Section 2.3, “Deploying Camel Kafka Connector using Kafka Connect S2I on OpenShift”](#). This section shows an example using the AWS S3 sink connector.

#### Procedure

- Configure your sink connector and aggregator settings in Camel Kafka Connector properties, depending on your installation platform:

##### OpenShift

The following example shows the AWS S3 sink connector and aggregator configuration in a custom resource:

```
oc apply -f - << EOF
apiVersion: kafka.strimzi.io/v1alpha1
kind: KafkaConnector
metadata:
  name: s3-sink-connector
  namespace: myproject
  labels:
    strimzi.io/cluster: my-connect-cluster
spec:
  class: org.apache.camel.kafkaconnector.aws2s3.CamelAws2s3SinkConnector
```

```

tasksMax: 1
config:
  key.converter: org.apache.kafka.connect.storage.StringConverter
  value.converter: org.apache.kafka.connect.storage.StringConverter
  topics: s3-topic
  camel.sink.path.bucketNameOrArn: camel-kafka-connector
  camel.sink.endpoint.keyName: ${date:now:yyyyMMdd-HH:mm:ssSSS}-${exchangeId}
  # Camel aggregator settings
  camel.beans.aggregate:
class:org.apache.camel.kafkaconnector.aggregator.StringAggregator
  camel.beans.aggregation.size: 10
  camel.beans.aggregation.timeout: 5000
  camel.component.aws2-s3.accessKey: xxxx
  camel.component.aws2-s3.secretKey: yyyy
  camel.component.aws2-s3.region: region
EOF

```

## Red Hat Enterprise Linux

The following example shows the AWS S3 sink connector and aggregator configuration in the **CamelAwws3SinkConnector.properties** file:

```

name=CamelAWS2S3SinkConnector
connector.class=org.apache.camel.kafkaconnector.aws2s3.CamelAws2s3SinkConnector
key.converter=org.apache.kafka.connect.storage.StringConverter
value.converter=org.apache.kafka.connect.storage.StringConverter

topics=mytopic

camel.sink.path.bucketNameOrArn=camel-kafka-connector

camel.component.aws2-s3.access-key=xxxx
camel.component.aws2-s3.secret-key=yyyy
camel.component.aws2-s3.region=eu-west-1

camel.sink.endpoint.keyName=${date:now:yyyyMMdd-HH:mm:ssSSS}-${exchangeId}

# Camel aggregator settings
camel.beans.aggregate=class:org.apache.camel.kafkaconnector.aggregator.StringAggregat
or
camel.beans.aggregation.size=10
camel.beans.aggregation.timeout=5000

```

### Additional resources

- [Demonstration example of AWS2 S3 sink connector with aggregator](#)
- [Demonstration example of AWS2 S3 sink connector with zip aggregator](#)
- [Apache Camel Kafka Connector aggregation](#)

## 4.2. WRITING A CUSTOM CAMEL KAFKA CONNECTOR AGGREGATOR

In some scenarios using a Camel Kafka sink connector, you might want to add an aggregator to batch up

your Kafka records before sending them to the external sink system. Typically, this involves defining a specific batch size and timeout for aggregation of records. When complete, the aggregate record is sent to the external system.

You can implement your own aggregator or configure one of the aggregators provided by Apache Camel. This section describes how to implement a custom aggregator in Java using the Camel **AggregationStrategy** class.

### Prerequisites

- You must have Red Hat Fuse installed.

### Procedure

1. Write your own custom aggregator by implementing the Camel **AggregationStrategy** class, for example:

```
package org.apache.camel.kafkaconnector.aggregator;

import org.apache.camel.AggregationStrategy;
import org.apache.camel.Exchange;
import org.apache.camel.Message;

public class StringAggregator implements AggregationStrategy {

    @Override
    public Exchange aggregate(Exchange oldExchange, Exchange newExchange) { 1
        // lets append the old body to the new body
        if (oldExchange == null) {
            return newExchange;
        }

        String body = oldExchange.getIn().getBody(String.class); 2
        if (body != null) {
            Message newIn = newExchange.getIn();
            String newBody = newIn.getBody(String.class);
            if (newBody != null) {
                body += System.lineSeparator() + newBody;
            }

            newIn.setBody(body);
        }
        return newExchange; 3
    }
}
```

- 1** The **oldExchange** and **newExchange** objects correspond to the Kafka records arriving at the aggregator.
- 2** In this case, each **newExchange** body will be concatenated with the **oldExchange** body and separated using the **System** line separator.
- 3** This process continues until the batch size is completed or the timeout is reached.

2. Add your custom aggregator code to your existing Camel Kafka connector. See [Section 4.4, “Extending Camel Kafka connectors using Maven archetypes”](#).

### Additional resources

- [Apache Camel Kafka Connector aggregators](#)

## 4.3. CONFIGURING CAMEL DATA FORMATS IN CAMEL KAFKA CONNECTOR

Camel Kafka Connector provides marshaling/unmarshaling of Camel data formats for sink and source connectors. For example, these formats include Apache Avro, Base64, Google Protobuf, JSON, SOAP, Zip file, and many more.

Typically, you would use a Camel **DataFormat** in your Camel DSL to marshal and unmarshal messages to and from different Camel data formats. For example, if you are receiving messages from a Camel File or JMS component and want to unmarshal the payload for further processing, you can use a **DataFormat** to implement this in the Camel DSL.

Using Camel Kafka Connector, you can simply configure marshaling and unmarshaling of Camel data formats using properties in your connector configuration. This section shows how to configure marshaling for the Camel Zip file data format using the **camel.sink.marshall: zipfile** property.

### Prerequisites

- You must have Camel Kafka Connector installed on OpenShift or Red Hat Enterprise Linux.
- You must have already built your connector starting from an archetype and edited your **pom.xml** to add the required dependencies. See [Section 4.4, “Extending Camel Kafka connectors using Maven archetypes”](#).

### Procedure

- Configure the connector settings for marshalling/unmarshalling the data format in your Camel Kafka Connector configuration, depending on your installation platform:

#### OpenShift

The following example shows the AWS S3 sink connector and Camel Zip data format configuration in a custom resource:

```
oc apply -f - << EOF
apiVersion: kafka.strimzi.io/v1alpha1
kind: KafkaConnector
metadata:
  name: s3-sink-connector
  namespace: myproject
  labels:
    strimzi.io/cluster: my-connect-cluster
spec:
  class: org.apache.camel.kafkaconnector.aws2s3.CamelAws2s3SinkConnector
  tasksMax: 1
  config:
    key.converter: org.apache.kafka.connect.storage.StringConverter
    value.converter: org.apache.kafka.connect.storage.StringConverter
```

```

topics: s3-topic
camel.sink.path.bucketNameOrArn: camel-kafka-connector
camel.sink.endpoint.keyName: ${date:now:yyyyMMdd-HHmssSSS}-
${exchangeld}.zip
# Camel data format setting
camel.sink.marshall: zipfile
camel.component.aws2-s3.accessKey: xxxx
camel.component.aws2-s3.secretKey: yyyy
camel.component.aws2-s3.region: region
EOF

```

### Red Hat Enterprise Linux

The following example shows the AWS S3 sink connector and Camel Zip data configuration in the **CamelAwss3SinkConnector.properties** file:

```

name=CamelAWS2S3SinkConnector
connector.class=org.apache.camel.kafkaconnector.aws2s3.CamelAws2s3SinkConnector
key.converter=org.apache.kafka.connect.storage.StringConverter
value.converter=org.apache.kafka.connect.storage.StringConverter

topics=mytopic

# Camel data format setting
camel.sink.marshall=zipfile

camel.sink.path.bucketNameOrArn=camel-kafka-connector

camel.component.aws2-s3.access-key=xxxx
camel.component.aws2-s3.secret-key=yyyy
camel.component.aws2-s3.region=eu-west-1

camel.sink.endpoint.keyName=${date:now:yyyyMMdd-HHmssSSS}-${exchangeld}.zip

```

### Additional resources

- [Demonstration example of AWS2 S3 sink connector with marshaling of Camel zip data format](#)
- [Apache Camel data format types](#)
- [Apache Camel DataFormat](#)

## 4.4. EXTENDING CAMEL KAFKA CONNECTORS USING MAVEN ARCHETYPES

In some scenarios, you might need to extend your Camel Kafka Connector system. For example, when using a sink connector, you might want to add a custom aggregator to batch up your Kafka records before sending them to the external sink system. Alternatively, you might want to configure a connector for marshaling or unmarshaling of Camel data formats, such as Apache Avro, Google Protobuf, JSON, or Zip file.

You can extend an existing Camel Kafka connector using the Maven **camel-kafka-connector-extensible-archetype**. An archetype is a Maven project template, which provides a consistent way of generating a project. This section describes how to use the archetype to create a Maven project to be

extended and how to add your project dependencies.



## NOTE

Using Maven archetypes to write additional Kafka Connect converters or transformers is not included in the Technology Preview and has community support only.

## Prerequisites

- You must have Apache Maven installed.

## Procedure

- Enter the **mvn archetype:generate** command to create a Maven project to extend Camel Kafka Connector. For example:

```
$ mvn archetype:generate -
DarchetypeGroupId=org.apache.camel.kafkaconnector.archetypes -
DarchetypeArtifactId=camel-kafka-connector-extensible-archetype -
DarchetypeVersion=CONNECTOR_VERSION
[INFO] Scanning for projects...
[INFO]
[INFO] -----< org.apache.maven:standalone-pom >-----
[INFO] Building Maven Stub Project (No POM) 1
[INFO] -----[ pom ]-----
[INFO]
[INFO] >>> maven-archetype-plugin:3.1.2:generate (default-cli) > generate-sources @
standalone-pom >>>
[INFO]
[INFO] <<< maven-archetype-plugin:3.1.2:generate (default-cli) < generate-sources @
standalone-pom <<<
[INFO]
[INFO]
[INFO] --- maven-archetype-plugin:3.1.2:generate (default-cli) @ standalone-pom ---
[INFO] Generating project in Interactive mode
[INFO] Archetype repository not defined. Using the one from
[org.apache.camel.kafkaconnector.archetypes:camel-kafka-connector-extensible-
archetype:0.4.0] found in catalog remote
```

- Enter values for each of the properties when prompted. The following example extends a **camel-aws2-s3-kafka-connector**:

```
Define value for property 'groupId': org.apache.camel.kafkaconnector.extended
Define value for property 'artifactId': myconnector-extended
Define value for property 'version' 1.0-SNAPSHOT: :
Define value for property 'package' org.apache.camel.kafkaconnector.extended: :
Define value for property 'camel-kafka-connector-name': camel-aws2-s3-kafka-connector
[INFO] Using property: camel-kafka-connector-version = CONNECTOR_VERSION
Confirm properties configuration:
groupId: org.apache.camel.kafkaconnector.extended
artifactId: myconnector-extended
version: 1.0-SNAPSHOT
package: org.apache.camel.kafkaconnector.extended
camel-kafka-connector-name: camel-aws2-s3-kafka-connector
camel-kafka-connector-version: CONNECTOR_VERSION
```

3. Enter **Y** to confirm your properties:

```

Y: : Y
[INFO] -----
[INFO] Using following parameters for creating project from Archetype: camel-kafka-
connector-extensible-archetype:CONNECTOR_VERSION
[INFO] -----
[INFO] Parameter: groupId, Value: org.apache.camel.kafkaconnector.extended
[INFO] Parameter: artifactId, Value: myconnector-extended
[INFO] Parameter: version, Value: 1.0-SNAPSHOT
[INFO] Parameter: package, Value: org.apache.camel.kafkaconnector.extended
[INFO] Parameter: packageInPathFormat, Value: org/apache/camel/kafkaconnector/extended
[INFO] Parameter: package, Value: org.apache.camel.kafkaconnector.extended
[INFO] Parameter: version, Value: 1.0-SNAPSHOT
[INFO] Parameter: groupId, Value: org.apache.camel.kafkaconnector.extended
[INFO] Parameter: camel-kafka-connector-name, Value: camel-aws2-s3-kafka-connector
[INFO] Parameter: camel-kafka-connector-version, Value: CONNECTOR_VERSION
[INFO] Parameter: artifactId, Value: myconnector-extended
[INFO] Project created from Archetype in dir: /home/workspace/myconnector-extended
[INFO] -----
[INFO] BUILD SUCCESS
[INFO] -----
[INFO] Total time: 05:44 min
[INFO] Finished at: 2020-09-04T08:55:00+02:00
[INFO] -----

```

4. Enter the dependencies that you need in the **pom.xml** for the created Maven project.
5. Build the Maven project to create a **.zip** or **tar.gz** file for your extended Camel Kafka connector:

```
mvn clean package
```

#### Additional resources

- [Apache Camel Kafka Connector archetypes](#)
- [Apache Maven](#)

## CHAPTER 5. CAMEL KAFKA CONNECTOR CONFIGURATION REFERENCE

This chapter provides reference information on the Camel Kafka connectors that you can configure using Camel Kafka Connector.



### IMPORTANT

This Technology Preview release includes a targeted subset of the available Apache Camel Kafka connectors. Additional connectors will be added to Camel Kafka Connector in future releases.

**Table 5.1. Camel Kafka Connector configuration**

Connector	Sink	Source
Amazon Web Services Kinesis	<a href="#">Camel AWS2 Kinesis sink connector</a>	<a href="#">Camel AWS2 Kinesis source connector</a>
Amazon Web Services S3	<a href="#">Camel AWS2 S3 sink connector</a>	<a href="#">Camel AWS2 s3 source connector</a>
Amazon Web Services SNS	<a href="#">Camel AWS2 SNS sink connector</a>	-
Amazon Web Services SQS	<a href="#">Camel AWS2 SQS sink connector</a>	<a href="#">Camel AWS2 SQS source connector</a>
Cassandra Query Language	<a href="#">Camel CQL sink connector</a>	<a href="#">Camel CQL source connector</a>
Elasticsearch	<a href="#">Camel Elasticsearch sink connector</a>	-
File	<a href="#">Camel file sink connector</a>	-
Hadoop Distributed File System	<a href="#">Camel HDFS sink connector</a>	-
Hypertext Transfer Protocol	<a href="#">Camel HTTP sink connector</a>	-
Java Database Connectivity	<a href="#">Camel JDBC sink connector</a>	-
Java Message Service	<a href="#">Camel JMS sink connector</a>	<a href="#">Camel JMS source connector</a>
MongoDB	<a href="#">Camel MongoDB sink connector</a>	<a href="#">Camel MongoDB source connector</a>
Salesforce	-	<a href="#">Camel Salesforce source connector</a>
Slack	-	<a href="#">Camel Slack source connector</a>

Connector	Sink	Source
Syslog	-	<a href="#">Camel syslog source connector</a>
Timer	-	<a href="#">Camel timer source connector</a>

## 5.1. CAMEL-AWS2-KINESIS-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-aws2-kinesis-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-aws2-kinesis-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.aws2kinesis.CamelAws2kinesisSinkConnector
```

The camel-aws2-kinesis sink connector supports 25 options, which are listed below.

Name	Description	Default	Priority
<b>camel.sink.path.stream Name</b>	Name of the stream	null	HIGH
<b>camel.sink.endpoint.amazonKinesisClient</b>	Amazon Kinesis client to use for all requests for this endpoint	null	MEDIUM
<b>camel.sink.endpoint.autoDiscoverClient</b>	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking	true	MEDIUM
<b>camel.sink.endpoint.proxyHost</b>	To define a proxy host when instantiating the Kinesis client	null	MEDIUM
<b>camel.sink.endpoint.proxyPort</b>	To define a proxy port when instantiating the Kinesis client	null	MEDIUM
<b>camel.sink.endpoint.proxyProtocol</b>	To define a proxy protocol when instantiating the Kinesis client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.region</b>	The region in which Kinesis Firehose client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id()	null	MEDIUM
<b>camel.sink.endpoint.trustAllCertificates</b>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.sink.endpoint.accessKey</b>	Amazon AWS Access Key	null	MEDIUM
<b>camel.sink.endpoint.secretKey</b>	Amazon AWS Secret Key	null	MEDIUM
<b>camel.component.aws2-kinesis.amazonKinesisClient</b>	Amazon Kinesis client to use for all requests for this endpoint	null	MEDIUM
<b>camel.component.aws2-kinesis.autoDiscoverClient</b>	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking	true	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-kinesis.configuration</code>	Component configuration	null	MEDIUM
<code>camel.component.aws2-kinesis.proxyHost</code>	To define a proxy host when instantiating the Kinesis client	null	MEDIUM
<code>camel.component.aws2-kinesis.proxyPort</code>	To define a proxy port when instantiating the Kinesis client	null	MEDIUM
<code>camel.component.aws2-kinesis.proxyProtocol</code>	To define a proxy protocol when instantiating the Kinesis client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<code>camel.component.aws2-kinesis.region</code>	The region in which Kinesis Firehose client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code>	null	MEDIUM
<code>camel.component.aws2-kinesis.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.component.aws2-kinesis.lazyStartProducer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<code>camel.component.aws2-kinesis.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.component.aws2-kinesis.accessKey</code>	Amazon AWS Access Key	null	MEDIUM
<code>camel.component.aws2-kinesis.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM

The camel-aws2-kinesis sink connector has no converters out of the box.

The camel-aws2-kinesis sink connector has no transforms out of the box.

The camel-aws2-kinesis sink connector has no aggregation strategies out of the box.

## 5.2. CAMEL-AWS2-KINESIS-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-aws2-kinesis-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-aws2-kinesis-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.aws2kinesis.CamelAws2kinesisSourceConnector
```

The camel-aws2-kinesis source connector supports 53 options, which are listed below.

Name	Description	Default	Priority
camel.source.path.streamName	Name of the stream	null	HIGH
camel.source.endpoint.amazonKinesisClient	Amazon Kinesis client to use for all requests for this endpoint	null	MEDIUM
camel.source.endpoint.autoDiscoverClient	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking	true	MEDIUM
camel.source.endpoint.proxyHost	To define a proxy host when instantiating the Kinesis client	null	MEDIUM
camel.source.endpoint.proxyPort	To define a proxy port when instantiating the Kinesis client	null	MEDIUM
camel.source.endpoint.proxyProtocol	To define a proxy protocol when instantiating the Kinesis client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
camel.source.endpoint.region	The region in which Kinesis Firehose client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id()	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.trustAllCertificates</b>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<b>camel.source.endpoint.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.source.endpoint.iteratorType</b>	Defines where in the Kinesis stream to start getting records One of: [AT_SEQUENCE_NUMBER] [AFTER_SEQUENCE_NUMBER] [TRIM_HORIZON] [LATEST] [AT_TIMESTAMP] [null]	"TRIM_HORIZON"	MEDIUM
<b>camel.source.endpoint.maxResultsPerRequest</b>	Maximum number of records that will be fetched in each poll	1	MEDIUM
<b>camel.source.endpoint.sendEmptyMessageWhenIdle</b>	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	MEDIUM
<b>camel.source.endpoint.sequenceNumber</b>	The sequence number to start polling from. Required if iteratorType is set to AFTER_SEQUENCE_NUMBER or AT_SEQUENCE_NUMBER	null	MEDIUM
<b>camel.source.endpoint.shardClosed</b>	Define what will be the behavior in case of shard closed. Possible value are ignore, silent and fail. In case of ignore a message will be logged and the consumer will restart from the beginning, in case of silent there will be no logging and the consumer will start from the beginning, in case of fail a ReachedClosedStateException will be raised One of: [ignore] [fail] [silent]	"ignore"	MEDIUM
<b>camel.source.endpoint.shardId</b>	Defines which shardId in the Kinesis stream to get records from	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: [InOnly] [InOut] [InOptionalOut]	null	MEDIUM
<b>camel.source.endpoint.pollStrategy</b>	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.	null	MEDIUM
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.source.endpoint.backoffErrorThreshold</b>	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.	null	MEDIUM
<b>camel.source.endpoint.backoffIdleThreshold</b>	The number of subsequent idle polls that should happen before the <code>backoffMultiplier</code> should kick-in.	null	MEDIUM
<b>camel.source.endpoint.backoffMultiplier</b>	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.	null	MEDIUM
<b>camel.source.endpoint.delay</b>	Milliseconds before the next poll.	500L	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.greedy</b>	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	MEDIUM
<b>camel.source.endpoint.initialDelay</b>	Milliseconds before the first poll starts.	1000L	MEDIUM
<b>camel.source.endpoint.repeatCount</b>	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0L	MEDIUM
<b>camel.source.endpoint.runLoggingLevel</b>	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"TRACE"	MEDIUM
<b>camel.source.endpoint.scheduledExecutorService</b>	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.	null	MEDIUM
<b>camel.source.endpoint.scheduler</b>	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler	"none"	MEDIUM
<b>camel.source.endpoint.schedulerProperties</b>	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.	null	MEDIUM
<b>camel.source.endpoint.startScheduler</b>	Whether the scheduler should be auto started.	true	MEDIUM
<b>camel.source.endpoint.timeUnit</b>	Time unit for initialDelay and delay options. One of: [NANOSECONDS] [MICROSECONDS] [MILLISECONDS] [SECONDS] [MINUTES] [HOURS] [DAYS]	"MILLISECONDS"	MEDIUM
<b>camel.source.endpoint.useFixedDelay</b>	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	MEDIUM
<b>camel.source.endpoint.accessKey</b>	Amazon AWS Access Key	null	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM
<code>camel.component.aws2-kinesis.amazonKinesisClient</code>	Amazon Kinesis client to use for all requests for this endpoint	null	MEDIUM
<code>camel.component.aws2-kinesis.autoDiscoverClient</code>	Setting the <code>autoDiscoverClient</code> mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking	true	MEDIUM
<code>camel.component.aws2-kinesis.configuration</code>	Component configuration	null	MEDIUM
<code>camel.component.aws2-kinesis.proxyHost</code>	To define a proxy host when instantiating the Kinesis client	null	MEDIUM
<code>camel.component.aws2-kinesis.proxyPort</code>	To define a proxy port when instantiating the Kinesis client	null	MEDIUM
<code>camel.component.aws2-kinesis.proxyProtocol</code>	To define a proxy protocol when instantiating the Kinesis client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<code>camel.component.aws2-kinesis.region</code>	The region in which Kinesis Firehose client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example <code>ap-east-1</code> ) You'll need to use the name <code>Region.EU_WEST_1.id()</code>	null	MEDIUM
<code>camel.component.aws2-kinesis.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.component.aws2-kinesis.bridgeErrorHandler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.aws2-kinesis.iteratorType</b>	Defines where in the Kinesis stream to start getting records One of: [AT_SEQUENCE_NUMBER] [AFTER_SEQUENCE_NUMBER] [TRIM_HORIZON] [LATEST] [AT_TIMESTAMP] [null]	"TRIM_HORIZON"	MEDIUM
<b>camel.component.aws2-kinesis.maxResultsPerRequest</b>	Maximum number of records that will be fetched in each poll	1	MEDIUM
<b>camel.component.aws2-kinesis.sequenceNumber</b>	The sequence number to start polling from. Required if iteratorType is set to AFTER_SEQUENCE_NUMBER or AT_SEQUENCE_NUMBER	null	MEDIUM
<b>camel.component.aws2-kinesis.shardClosed</b>	Define what will be the behavior in case of shard closed. Possible value are ignore, silent and fail. In case of ignore a message will be logged and the consumer will restart from the beginning, in case of silent there will be no logging and the consumer will start from the beginning, in case of fail a ReachedClosedStateException will be raised One of: [ignore] [fail] [silent]	"ignore"	MEDIUM
<b>camel.component.aws2-kinesis.shardId</b>	Defines which shardId in the Kinesis stream to get records from	null	MEDIUM
<b>camel.component.aws2-kinesis.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.component.aws2-kinesis.accessKey</b>	Amazon AWS Access Key	null	MEDIUM
<b>camel.component.aws2-kinesis.secretKey</b>	Amazon AWS Secret Key	null	MEDIUM

The camel-aws2-kinesis sink connector has no converters out of the box.

The camel-aws2-kinesis sink connector has no transforms out of the box.

The camel-aws2-kinesis sink connector has no aggregation strategies out of the box.

### 5.3. CAMEL-AWS2-S3-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-aws2-s3-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-aws2-s3-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.aws2s3.CamelAws2s3SinkConnector
```

The camel-aws2-s3 sink connector supports 61 options, which are listed below.

Name	Description	Default	Priority
camel.sink.path.bucketNameOrArn	Bucket name or ARN	null	HIGH
camel.sink.endpoint.amazonS3Client	Reference to a com.amazonaws.services.s3.AmazonS3 in the registry.	null	MEDIUM
camel.sink.endpoint.autoCreateBucket	Setting the autocreation of the S3 bucket bucketName. This will apply also in case of moveAfterRead option enabled and it will create the destinationBucket if it doesn't exist already.	true	MEDIUM
camel.sink.endpoint.autoDiscoverClient	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
camel.sink.endpoint.overrideEndpoint	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option	false	MEDIUM
camel.sink.endpoint.pojoRequest	If we want to use a POJO request as body or not	false	MEDIUM
camel.sink.endpoint.policy	The policy for this queue to set in the com.amazonaws.services.s3.AmazonS3#setBucketPolicy() method.	null	MEDIUM
camel.sink.endpoint.proxyHost	To define a proxy host when instantiating the SQS client	null	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.proxyPort</code>	Specify a proxy port to be used inside the client definition.	null	MEDIUM
<code>camel.sink.endpoint.proxyProtocol</code>	To define a proxy protocol when instantiating the S3 client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<code>camel.sink.endpoint.region</code>	The region in which S3 client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code>	null	MEDIUM
<code>camel.sink.endpoint.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.sink.endpoint.uriEndpointOverride</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option	null	MEDIUM
<code>camel.sink.endpoint.useIAMCredentials</code>	Set whether the S3 client should expect to load credentials on an EC2 instance or to expect static credentials to be passed in.	false	MEDIUM
<code>camel.sink.endpoint.customerAlgorithm</code>	Define the customer algorithm to use in case <code>CustomerKey</code> is enabled	null	MEDIUM
<code>camel.sink.endpoint.customerKeyId</code>	Define the id of Customer key to use in case <code>CustomerKey</code> is enabled	null	MEDIUM
<code>camel.sink.endpoint.customerKeyMD5</code>	Define the MD5 of Customer key to use in case <code>CustomerKey</code> is enabled	null	MEDIUM
<code>camel.sink.endpoint.deleteAfterWrite</code>	Delete file object after the S3 file has been uploaded	false	MEDIUM
<code>camel.sink.endpoint.keyName</code>	Setting the key name for an element in the bucket through endpoint parameter	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.multipartUpload</b>	If it is true, camel will upload the file with multi part format, the part size is decided by the option of partSize	false	MEDIUM
<b>camel.sink.endpoint.operation</b>	The operation to do in case the user don't want to do only an upload One of: [copyObject] [listObjects] [deleteObject] [deleteBucket] [listBuckets] [getObject] [getObjectRange]	null	MEDIUM
<b>camel.sink.endpoint.partSize</b>	Setup the partSize which is used in multi part upload, the default size is 25M.	26214400L	MEDIUM
<b>camel.sink.endpoint.storageClass</b>	The storage class to set in the com.amazonaws.services.s3.model.PutObjectRequest request.	null	MEDIUM
<b>camel.sink.endpoint.awsKMSKeyId</b>	Define the id of KMS key to use in case KMS is enabled	null	MEDIUM
<b>camel.sink.endpoint.useAwsKMS</b>	Define if KMS must be used or not	false	MEDIUM
<b>camel.sink.endpoint.useCustomerKey</b>	Define if Customer Key must be used or not	false	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.accessKey</code>	Amazon AWS Access Key	null	MEDIUM
<code>camel.sink.endpoint.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM
<code>camel.component.aws2-s3.amazonS3Client</code>	Reference to a <code>com.amazonaws.services.s3.AmazonS3</code> in the registry.	null	MEDIUM
<code>camel.component.aws2-s3.autoCreateBucket</code>	Setting the autocreation of the S3 bucket <code>bucketName</code> . This will apply also in case of <code>moveAfterRead</code> option enabled and it will create the <code>destinationBucket</code> if it doesn't exist already.	true	MEDIUM
<code>camel.component.aws2-s3.autoDiscoverClient</code>	Setting the <code>autoDiscoverClient</code> mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
<code>camel.component.aws2-s3.configuration</code>	The component configuration	null	MEDIUM
<code>camel.component.aws2-s3.overrideEndpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option	false	MEDIUM
<code>camel.component.aws2-s3.pojoRequest</code>	If we want to use a POJO request as body or not	false	MEDIUM
<code>camel.component.aws2-s3.policy</code>	The policy for this queue to set in the <code>com.amazonaws.services.s3.AmazonS3#setBucketPolicy()</code> method.	null	MEDIUM
<code>camel.component.aws2-s3.proxyHost</code>	To define a proxy host when instantiating the SQS client	null	MEDIUM
<code>camel.component.aws2-s3.proxyPort</code>	Specify a proxy port to be used inside the client definition.	null	MEDIUM
<code>camel.component.aws2-s3.proxyProtocol</code>	To define a proxy protocol when instantiating the S3 client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-s3.region</code>	The region in which S3 client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code>	null	MEDIUM
<code>camel.component.aws2-s3.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.component.aws2-s3.uriEndpointOverride</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option	null	MEDIUM
<code>camel.component.aws2-s3.useIAMCredentials</code>	Set whether the S3 client should expect to load credentials on an EC2 instance or to expect static credentials to be passed in.	false	MEDIUM
<code>camel.component.aws2-s3.customerAlgorithm</code>	Define the customer algorithm to use in case <code>CustomerKey</code> is enabled	null	MEDIUM
<code>camel.component.aws2-s3.customerKeyId</code>	Define the id of Customer key to use in case <code>CustomerKey</code> is enabled	null	MEDIUM
<code>camel.component.aws2-s3.customerKeyMD5</code>	Define the MD5 of Customer key to use in case <code>CustomerKey</code> is enabled	null	MEDIUM
<code>camel.component.aws2-s3.deleteAfterWrite</code>	Delete file object after the S3 file has been uploaded	false	MEDIUM
<code>camel.component.aws2-s3.keyName</code>	Setting the key name for an element in the bucket through endpoint parameter	null	MEDIUM
<code>camel.component.aws2-s3.lazyStartProducer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<code>camel.component.aws2-s3.multiPartUpload</code>	If it is true, camel will upload the file with multi part format, the part size is decided by the option of <code>partSize</code>	false	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-s3.operation</code>	The operation to do in case the user don't want to do only an upload One of: [copyObject] [listObjects] [deleteObject] [deleteBucket] [listBuckets] [getObject] [getObjectRange]	null	MEDIUM
<code>camel.component.aws2-s3.partSize</code>	Setup the partSize which is used in multi part upload, the default size is 25M.	26214400L	MEDIUM
<code>camel.component.aws2-s3.storageClass</code>	The storage class to set in the com.amazonaws.services.s3.model.PutObjectRequest request.	null	MEDIUM
<code>camel.component.aws2-s3.awsKMSKeyId</code>	Define the id of KMS key to use in case KMS is enabled	null	MEDIUM
<code>camel.component.aws2-s3.useAwsKMS</code>	Define if KMS must be used or not	false	MEDIUM
<code>camel.component.aws2-s3.useCustomerKey</code>	Define if Customer Key must be used or not	false	MEDIUM
<code>camel.component.aws2-s3.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.component.aws2-s3.accessKey</code>	Amazon AWS Access Key	null	MEDIUM
<code>camel.component.aws2-s3.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM

The camel-aws2-s3 sink connector supports 1 converters out of the box, which are listed below.

```
org.apache.camel.kafkaconnector.aws2s3.converters.S3ObjectConverter
```

The camel-aws2-s3 sink connector has no transforms out of the box.

The camel-aws2-s3 sink connector has no aggregation strategies out of the box.

## 5.4. CAMEL-AWS2-S3-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-aws2-s3-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
```

```
<artifactId>camel-aws2-s3-kafka-connector</artifactId>
<version>x.x.x</version>
<!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.aws2s3.CamelAws2s3SourceConnector
```

The camel-aws2-s3 source connector supports 85 options, which are listed below.

Name	Description	Default	Priority
<b>camel.source.path.bucketNameOrArn</b>	Bucket name or ARN	null	HIGH
<b>camel.source.endpoint.amazonS3Client</b>	Reference to a com.amazonaws.services.s3.AmazonS3 in the registry.	null	MEDIUM
<b>camel.source.endpoint.autoCreateBucket</b>	Setting the autocreation of the S3 bucket bucketName. This will apply also in case of moveAfterRead option enabled and it will create the destinationBucket if it doesn't exist already.	true	MEDIUM
<b>camel.source.endpoint.autoDiscoverClient</b>	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
<b>camel.source.endpoint.overrideEndpoint</b>	Set the need for overriding the endpoint. This option needs to be used in combination with uriEndpointOverride option	false	MEDIUM
<b>camel.source.endpoint.pojoRequest</b>	If we want to use a POJO request as body or not	false	MEDIUM
<b>camel.source.endpoint.policy</b>	The policy for this queue to set in the com.amazonaws.services.s3.AmazonS3#setBucketPolicy() method.	null	MEDIUM
<b>camel.source.endpoint.proxyHost</b>	To define a proxy host when instantiating the SQS client	null	MEDIUM
<b>camel.source.endpoint.proxyPort</b>	Specify a proxy port to be used inside the client definition.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.proxyProtocol</b>	To define a proxy protocol when instantiating the S3 client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<b>camel.source.endpoint.region</b>	The region in which S3 client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id()	null	MEDIUM
<b>camel.source.endpoint.trustAllCertificates</b>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<b>camel.source.endpoint.uriEndpointOverride</b>	Set the overriding uri endpoint. This option needs to be used in combination with overrideEndpoint option	null	MEDIUM
<b>camel.source.endpoint.useIAMCredentials</b>	Set whether the S3 client should expect to load credentials on an EC2 instance or to expect static credentials to be passed in.	false	MEDIUM
<b>camel.source.endpoint.customerAlgorithm</b>	Define the customer algorithm to use in case CustomerKey is enabled	null	MEDIUM
<b>camel.source.endpoint.customerKeyId</b>	Define the id of Customer key to use in case CustomerKey is enabled	null	MEDIUM
<b>camel.source.endpoint.customerKeyMD5</b>	Define the MD5 of Customer key to use in case CustomerKey is enabled	null	MEDIUM
<b>camel.source.endpoint.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.deleteAfterRead</b>	Delete objects from S3 after they have been retrieved. The delete is only performed if the Exchange is committed. If a rollback occurs, the object is not deleted. If this option is false, then the same objects will be retrieve over and over again on the polls. Therefore you need to use the Idempotent Consumer EIP in the route to filter out duplicates. You can filter using the <code>AWS2S3Constants#BUCKET_NAME</code> and <code>AWS2S3Constants#KEY</code> headers, or only the <code>AWS2S3Constants#KEY</code> header.	true	MEDIUM
<b>camel.source.endpoint.delimiter</b>	The delimiter which is used in the <code>com.amazonaws.services.s3.model.ListObjectsRequest</code> to only consume objects we are interested in.	null	MEDIUM
<b>camel.source.endpoint.destinationBucket</b>	Define the destination bucket where an object must be moved when <code>moveAfterRead</code> is set to true.	null	MEDIUM
<b>camel.source.endpoint.destinationBucketPrefix</b>	Define the destination bucket prefix to use when an object must be moved and <code>moveAfterRead</code> is set to true.	null	MEDIUM
<b>camel.source.endpoint.destinationBucketSuffix</b>	Define the destination bucket suffix to use when an object must be moved and <code>moveAfterRead</code> is set to true.	null	MEDIUM
<b>camel.source.endpoint.fileName</b>	To get the object from the bucket with the given file name	null	MEDIUM
<b>camel.source.endpoint.includeBody</b>	If it is true, the exchange body will be set to a stream to the contents of the file. If false, the headers will be set with the S3 object metadata, but the body will be null. This option is strongly related to <code>autocloseBody</code> option. In case of setting <code>includeBody</code> to true and <code>autocloseBody</code> to false, it will be up to the caller to close the <code>S3Object</code> stream. Setting <code>autocloseBody</code> to true, will close the <code>S3Object</code> stream automatically.	true	MEDIUM
<b>camel.source.endpoint.includeFolders</b>	If it is true, the folders/directories will be consumed. If it is false, they will be ignored, and Exchanges will not be created for those	true	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.maxConnections</b>	Set the maxConnections parameter in the S3 client configuration	60	MEDIUM
<b>camel.source.endpoint.maxMessagesPerPoll</b>	Gets the maximum number of messages as a limit to poll at each polling. Gets the maximum number of messages as a limit to poll at each polling. The default value is 10. Use 0 or a negative number to set it as unlimited.	10	MEDIUM
<b>camel.source.endpoint.moveAfterRead</b>	Move objects from S3 bucket to a different bucket after they have been retrieved. To accomplish the operation the destinationBucket option must be set. The copy bucket operation is only performed if the Exchange is committed. If a rollback occurs, the object is not moved.	false	MEDIUM
<b>camel.source.endpoint.prefix</b>	The prefix which is used in the com.amazonaws.services.s3.model.ListObjects Request to only consume objects we are interested in.	null	MEDIUM
<b>camel.source.endpoint.sendEmptyMessageWhenIdle</b>	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	MEDIUM
<b>camel.source.endpoint.autocloseBody</b>	If this option is true and includeBody is true, then the S3Object.close() method will be called on exchange completion. This option is strongly related to includeBody option. In case of setting includeBody to true and autocloseBody to false, it will be up to the caller to close the S3Object stream. Setting autocloseBody to true, will close the S3Object stream automatically.	true	MEDIUM
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: [InOnly] [InOut] [InOptionalOut]	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.pollStrategy</b>	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.	null	MEDIUM
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.source.endpoint.backoffErrorThreshold</b>	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.	null	MEDIUM
<b>camel.source.endpoint.backoffIdleThreshold</b>	The number of subsequent idle polls that should happen before the <code>backoffMultiplier</code> should kick-in.	null	MEDIUM
<b>camel.source.endpoint.backoffMultiplier</b>	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.	null	MEDIUM
<b>camel.source.endpoint.delay</b>	Milliseconds before the next poll.	500L	MEDIUM
<b>camel.source.endpoint.greedy</b>	If <code>greedy</code> is enabled, then the <code>ScheduledPollConsumer</code> will run immediately again, if the previous run polled 1 or more messages.	false	MEDIUM
<b>camel.source.endpoint.initialDelay</b>	Milliseconds before the first poll starts.	1000L	MEDIUM
<b>camel.source.endpoint.repeatCount</b>	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0L	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.runLoggingLevel</b>	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"TRACE"	MEDIUM
<b>camel.source.endpoint.scheduledExecutorService</b>	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.	null	MEDIUM
<b>camel.source.endpoint.scheduler</b>	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler	"none"	MEDIUM
<b>camel.source.endpoint.schedulerProperties</b>	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.	null	MEDIUM
<b>camel.source.endpoint.startScheduler</b>	Whether the scheduler should be auto started.	true	MEDIUM
<b>camel.source.endpoint.timeUnit</b>	Time unit for initialDelay and delay options. One of: [NANOSECONDS] [MICROSECONDS] [MILLISECONDS] [SECONDS] [MINUTES] [HOURS] [DAYS]	"MILLISECONDS"	MEDIUM
<b>camel.source.endpoint.useFixedDelay</b>	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	MEDIUM
<b>camel.source.endpoint.accessKey</b>	Amazon AWS Access Key	null	MEDIUM
<b>camel.source.endpoint.secretKey</b>	Amazon AWS Secret Key	null	MEDIUM
<b>camel.component.aws2-s3.amazonS3Client</b>	Reference to a com.amazonaws.services.s3.AmazonS3 in the registry.	null	MEDIUM
<b>camel.component.aws2-s3.autoCreateBucket</b>	Setting the autocreation of the S3 bucket bucketName. This will apply also in case of moveAfterRead option enabled and it will create the destinationBucket if it doesn't exist already.	true	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-s3.autoDiscoverClient</code>	Setting the <code>autoDiscoverClient</code> mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
<code>camel.component.aws2-s3.configuration</code>	The component configuration	null	MEDIUM
<code>camel.component.aws2-s3.overrideEndpoint</code>	Set the need for overriding the endpoint. This option needs to be used in combination with <code>uriEndpointOverride</code> option	false	MEDIUM
<code>camel.component.aws2-s3.pojoRequest</code>	If we want to use a POJO request as body or not	false	MEDIUM
<code>camel.component.aws2-s3.policy</code>	The policy for this queue to set in the <code>com.amazonaws.services.s3.AmazonS3#setBucketPolicy()</code> method.	null	MEDIUM
<code>camel.component.aws2-s3.proxyHost</code>	To define a proxy host when instantiating the SQS client	null	MEDIUM
<code>camel.component.aws2-s3.proxyPort</code>	Specify a proxy port to be used inside the client definition.	null	MEDIUM
<code>camel.component.aws2-s3.proxyProtocol</code>	To define a proxy protocol when instantiating the S3 client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<code>camel.component.aws2-s3.region</code>	The region in which S3 client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example <code>ap-east-1</code> ) You'll need to use the name <code>Region.EU_WEST_1.id()</code>	null	MEDIUM
<code>camel.component.aws2-s3.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.component.aws2-s3.uriEndpointOverride</code>	Set the overriding uri endpoint. This option needs to be used in combination with <code>overrideEndpoint</code> option	null	MEDIUM
<code>camel.component.aws2-s3.useIAMCredentials</code>	Set whether the S3 client should expect to load credentials on an EC2 instance or to expect static credentials to be passed in.	false	MEDIUM
<code>camel.component.aws2-s3.customerAlgorithm</code>	Define the customer algorithm to use in case <code>CustomerKey</code> is enabled	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-s3.customerKeyId</code>	Define the id of Customer key to use in case CustomerKey is enabled	null	MEDIUM
<code>camel.component.aws2-s3.customerKeyMD5</code>	Define the MD5 of Customer key to use in case CustomerKey is enabled	null	MEDIUM
<code>camel.component.aws2-s3.bridgeErrorHandler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<code>camel.component.aws2-s3.deleteAfterRead</code>	Delete objects from S3 after they have been retrieved. The delete is only performed if the Exchange is committed. If a rollback occurs, the object is not deleted. If this option is false, then the same objects will be retrieve over and over again on the polls. Therefore you need to use the Idempotent Consumer EIP in the route to filter out duplicates. You can filter using the <code>AWS2S3Constants#BUCKET_NAME</code> and <code>AWS2S3Constants#KEY</code> headers, or only the <code>AWS2S3Constants#KEY</code> header.	true	MEDIUM
<code>camel.component.aws2-s3.delimiter</code>	The delimiter which is used in the <code>com.amazonaws.services.s3.model.ListObjectsRequest</code> to only consume objects we are interested in.	null	MEDIUM
<code>camel.component.aws2-s3.destinationBucket</code>	Define the destination bucket where an object must be moved when <code>moveAfterRead</code> is set to true.	null	MEDIUM
<code>camel.component.aws2-s3.destinationBucketPrefix</code>	Define the destination bucket prefix to use when an object must be moved and <code>moveAfterRead</code> is set to true.	null	MEDIUM
<code>camel.component.aws2-s3.destinationBucketSuffix</code>	Define the destination bucket suffix to use when an object must be moved and <code>moveAfterRead</code> is set to true.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-s3.fileName</code>	To get the object from the bucket with the given file name	null	MEDIUM
<code>camel.component.aws2-s3.includeBody</code>	If it is true, the exchange body will be set to a stream to the contents of the file. If false, the headers will be set with the S3 object metadata, but the body will be null. This option is strongly related to <code>autocloseBody</code> option. In case of setting <code>includeBody</code> to true and <code>autocloseBody</code> to false, it will be up to the caller to close the <code>S3Object</code> stream. Setting <code>autocloseBody</code> to true, will close the <code>S3Object</code> stream automatically.	true	MEDIUM
<code>camel.component.aws2-s3.includeFolders</code>	If it is true, the folders/directories will be consumed. If it is false, they will be ignored, and Exchanges will not be created for those	true	MEDIUM
<code>camel.component.aws2-s3.moveAfterRead</code>	Move objects from S3 bucket to a different bucket after they have been retrieved. To accomplish the operation the <code>destinationBucket</code> option must be set. The copy bucket operation is only performed if the Exchange is committed. If a rollback occurs, the object is not moved.	false	MEDIUM
<code>camel.component.aws2-s3.prefix</code>	The prefix which is used in the <code>com.amazonaws.services.s3.model.ListObjectsRequest</code> to only consume objects we are interested in.	null	MEDIUM
<code>camel.component.aws2-s3.autocloseBody</code>	If this option is true and <code>includeBody</code> is true, then the <code>S3Object.close()</code> method will be called on exchange completion. This option is strongly related to <code>includeBody</code> option. In case of setting <code>includeBody</code> to true and <code>autocloseBody</code> to false, it will be up to the caller to close the <code>S3Object</code> stream. Setting <code>autocloseBody</code> to true, will close the <code>S3Object</code> stream automatically.	true	MEDIUM
<code>camel.component.aws2-s3.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.component.aws2-s3.accessKey</code>	Amazon AWS Access Key	null	MEDIUM

Name	Description	Default	Priority
camel.component.aws2-s3.secretKey	Amazon AWS Secret Key	null	MEDIUM

The camel-aws2-s3 sink connector supports 1 converters out of the box, which are listed below.

```
org.apache.camel.kafkaconnector.aws2s3.converters.S3ObjectConverter
```

The camel-aws2-s3 sink connector has no transforms out of the box.

The camel-aws2-s3 sink connector has no aggregation strategies out of the box.

## 5.5. CAMEL-AWS2-SNS-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-aws2-sns-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-aws2-sns-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.aws2sns.CamelAws2snsSinkConnector
```

The camel-aws2-sns sink connector supports 42 options, which are listed below.

Name	Description	Default	Priority
camel.sink.path.topicNameOrArn	Topic name or ARN	null	HIGH
camel.sink.endpoint.amazonSNSClient	To use the AmazonSNS as the client	null	MEDIUM
camel.sink.endpoint.autoCreateTopic	Setting the autocreation of the topic	true	MEDIUM
camel.sink.endpoint.autoDiscoverClient	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
camel.sink.endpoint.headerFilterStrategy	To use a custom HeaderFilterStrategy to map headers to/from Camel.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.km sMasterKeyId</b>	The ID of an AWS-managed customer master key (CMK) for Amazon SNS or a custom CMK.	null	MEDIUM
<b>camel.sink.endpoint.laz yStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.me ssageStructure</b>	The message structure to use such as json	null	MEDIUM
<b>camel.sink.endpoint.pol icy</b>	The policy for this queue	null	MEDIUM
<b>camel.sink.endpoint.pr oxyHost</b>	To define a proxy host when instantiating the SNS client	null	MEDIUM
<b>camel.sink.endpoint.pr oxyPort</b>	To define a proxy port when instantiating the SNS client	null	MEDIUM
<b>camel.sink.endpoint.pr oxyProtocol</b>	To define a proxy protocol when instantiating the SNS client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<b>camel.sink.endpoint.qu eueUrl</b>	The queueUrl to subscribe to	null	MEDIUM
<b>camel.sink.endpoint.reg ion</b>	The region in which SNS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id()	null	MEDIUM
<b>camel.sink.endpoint.ser verSideEncryptionEnab led</b>	Define if Server Side Encryption is enabled or not on the topic	false	MEDIUM
<b>camel.sink.endpoint.su bject</b>	The subject which is used if the message header 'CamelAwsSnsSubject' is not present.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.subscribeSNSstoSQS</code>	Define if the subscription between SNS Topic and SQS must be done or not	false	MEDIUM
<code>camel.sink.endpoint.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.sink.endpoint.basicPropertyBinding</code>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.sink.endpoint.synchronous</code>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<code>camel.sink.endpoint.accessKey</code>	Amazon AWS Access Key	null	MEDIUM
<code>camel.sink.endpoint.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM
<code>camel.component.aws2-sns.amazonSNSClient</code>	To use the AmazonSNS as the client	null	MEDIUM
<code>camel.component.aws2-sns.autoCreateTopic</code>	Setting the autocreation of the topic	true	MEDIUM
<code>camel.component.aws2-sns.autoDiscoverClient</code>	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
<code>camel.component.aws2-sns.configuration</code>	Component configuration	null	MEDIUM
<code>camel.component.aws2-sns.kmsMasterKeyId</code>	The ID of an AWS-managed customer master key (CMK) for Amazon SNS or a custom CMK.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.aws2-sns.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.component.aws2-sns.messageStructure</b>	The message structure to use such as json	null	MEDIUM
<b>camel.component.aws2-sns.policy</b>	The policy for this queue	null	MEDIUM
<b>camel.component.aws2-sns.proxyHost</b>	To define a proxy host when instantiating the SNS client	null	MEDIUM
<b>camel.component.aws2-sns.proxyPort</b>	To define a proxy port when instantiating the SNS client	null	MEDIUM
<b>camel.component.aws2-sns.proxyProtocol</b>	To define a proxy protocol when instantiating the SNS client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<b>camel.component.aws2-sns.queueUrl</b>	The queueUrl to subscribe to	null	MEDIUM
<b>camel.component.aws2-sns.region</b>	The region in which SNS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id()	null	MEDIUM
<b>camel.component.aws2-sns.serverSideEncryptionEnabled</b>	Define if Server Side Encryption is enabled or not on the topic	false	MEDIUM
<b>camel.component.aws2-sns.subject</b>	The subject which is used if the message header 'CamelAwsSnsSubject' is not present.	null	MEDIUM

Name	Description	Default	Priority
camel.component.aws2-sns.subscribeSNSToSQS	Define if the subscription between SNS Topic and SQS must be done or not	false	MEDIUM
camel.component.aws2-sns.trustAllCertificates	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
camel.component.aws2-sns.basicPropertyBinding	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
camel.component.aws2-sns.accessKey	Amazon AWS Access Key	null	MEDIUM
camel.component.aws2-sns.secretKey	Amazon AWS Secret Key	null	MEDIUM

The camel-aws2-sns sink connector has no converters out of the box.

The camel-aws2-sns sink connector has no transforms out of the box.

The camel-aws2-sns sink connector has no aggregation strategies out of the box.

## 5.6. CAMEL-AWS2-SQS-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-aws2-sqs-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-aws2-sqs-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.aws2sqs.CamelAws2sqsSinkConnector
```

The camel-aws2-sqs sink connector supports 56 options, which are listed below.

Name	Description	Default	Priority
camel.sink.path.queueNameOrArn	Queue name or ARN	null	HIGH

Name	Description	Default	Priority
<code>camel.sink.endpoint.amazonAWSHost</code>	The hostname of the Amazon AWS cloud.	"amazonaws.com"	MEDIUM
<code>camel.sink.endpoint.amazonSQSClient</code>	To use the AmazonSQS as client	null	MEDIUM
<code>camel.sink.endpoint.autoCreateQueue</code>	Setting the autocreation of the queue	true	MEDIUM
<code>camel.sink.endpoint.autoDiscoverClient</code>	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
<code>camel.sink.endpoint.headerFilterStrategy</code>	To use a custom HeaderFilterStrategy to map headers to/from Camel.	null	MEDIUM
<code>camel.sink.endpoint.protocol</code>	The underlying protocol used to communicate with SQS	"https"	MEDIUM
<code>camel.sink.endpoint.proxyProtocol</code>	To define a proxy protocol when instantiating the SQS client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<code>camel.sink.endpoint.queueOwnerAWSAccountId</code>	Specify the queue owner aws account id when you need to connect the queue with different account owner.	null	MEDIUM
<code>camel.sink.endpoint.region</code>	The region in which SQS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id()	null	MEDIUM
<code>camel.sink.endpoint.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.sink.endpoint.delaySeconds</code>	Delay sending messages for a number of seconds.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.messageDeduplicationStrategy</b>	Only for FIFO queues. Strategy for setting the messageDeduplicationId on the message. Can be one of the following options: useExchangeId, useContentBasedDeduplication. For the useContentBasedDeduplication option, no messageDeduplicationId will be set on the message. One of: [useExchangeId] [useContentBasedDeduplication]	"useExchangeId"	MEDIUM
<b>camel.sink.endpoint.messageGroupIdStrategy</b>	Only for FIFO queues. Strategy for setting the messageGroupId on the message. Can be one of the following options: useConstant, useExchangeId, usePropertyValue. For the usePropertyValue option, the value of property CamelAwsMessageGroupId will be used. One of: [useConstant] [useExchangeId] [usePropertyValue]	null	MEDIUM
<b>camel.sink.endpoint.operation</b>	The operation to do in case the user don't want to send only a message One of: [sendBatchMessage] [deleteMessage] [listQueues] [purgeQueue]	null	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.delayQueue</b>	Define if you want to apply delaySeconds option to the queue or on single messages	false	MEDIUM
<b>camel.sink.endpoint.queueUrl</b>	To define the queueUrl explicitly. All other parameters, which would influence the queueUrl, are ignored. This parameter is intended to be used, to connect to a mock implementation of SQS, for testing purposes.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.synchronous</code>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<code>camel.sink.endpoint.proxyHost</code>	To define a proxy host when instantiating the SQS client	null	MEDIUM
<code>camel.sink.endpoint.proxyPort</code>	To define a proxy port when instantiating the SQS client	null	MEDIUM
<code>camel.sink.endpoint.maximumMessageSize</code>	The maximumMessageSize (in bytes) an SQS message can contain for this queue.	null	MEDIUM
<code>camel.sink.endpoint.messageRetentionPeriod</code>	The messageRetentionPeriod (in seconds) a message will be retained by SQS for this queue.	null	MEDIUM
<code>camel.sink.endpoint.policy</code>	The policy for this queue	null	MEDIUM
<code>camel.sink.endpoint.receiveMessageWaitTimeSeconds</code>	If you do not specify WaitTimeSeconds in the request, the queue attribute ReceiveMessageWaitTimeSeconds is used to determine how long to wait.	null	MEDIUM
<code>camel.sink.endpoint.redrivePolicy</code>	Specify the policy that send message to DeadLetter queue. See detail at Amazon docs.	null	MEDIUM
<code>camel.sink.endpoint.accessKey</code>	Amazon AWS Access Key	null	MEDIUM
<code>camel.sink.endpoint.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM
<code>camel.component.aws2-sqs.amazonAWSHost</code>	The hostname of the Amazon AWS cloud.	"amazonaws.com"	MEDIUM
<code>camel.component.aws2-sqs.amazonSQSClient</code>	To use the AmazonSQS as client	null	MEDIUM
<code>camel.component.aws2-sqs.autoCreateQueue</code>	Setting the autocreation of the queue	true	MEDIUM
<code>camel.component.aws2-sqs.autoDiscoverClient</code>	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-sqs.configuration</code>	The AWS SQS default configuration	null	MEDIUM
<code>camel.component.aws2-sqs.protocol</code>	The underlying protocol used to communicate with SQS	"https"	MEDIUM
<code>camel.component.aws2-sqs.proxyProtocol</code>	To define a proxy protocol when instantiating the SQS client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<code>camel.component.aws2-sqs.queueOwnerAWSAccountid</code>	Specify the queue owner aws account id when you need to connect the queue with different account owner.	null	MEDIUM
<code>camel.component.aws2-sqs.region</code>	The region in which SQS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code>	null	MEDIUM
<code>camel.component.aws2-sqs.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.component.aws2-sqs.delaySeconds</code>	Delay sending messages for a number of seconds.	null	MEDIUM
<code>camel.component.aws2-sqs.lazyStartProducer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<code>camel.component.aws2-sqs.messageDeduplicationIdStrategy</code>	Only for FIFO queues. Strategy for setting the <code>messageDeduplicationId</code> on the message. Can be one of the following options: <code>useExchangeId</code> , <code>useContentBasedDeduplication</code> . For the <code>useContentBasedDeduplication</code> option, no <code>messageDeduplicationId</code> will be set on the message. One of: [ <code>useExchangeId</code> ] [ <code>useContentBasedDeduplication</code> ]	"useExchangeId"	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-sqs.messageGroupIdStrategy</code>	Only for FIFO queues. Strategy for setting the <code>messageGroupId</code> on the message. Can be one of the following options: <code>useConstant</code> , <code>useExchangeId</code> , <code>usePropertyValue</code> . For the <code>usePropertyValue</code> option, the value of property <code>CamelAwsMessageGroupId</code> will be used. One of: <code>[useConstant]</code> <code>[useExchangeId]</code> <code>[usePropertyValue]</code>	null	MEDIUM
<code>camel.component.aws2-sqs.operation</code>	The operation to do in case the user don't want to send only a message One of: <code>[sendBatchMessage]</code> <code>[deleteMessage]</code> <code>[listQueues]</code> <code>[purgeQueue]</code>	null	MEDIUM
<code>camel.component.aws2-sqs.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.component.aws2-sqs.delayQueue</code>	Define if you want to apply <code>delaySeconds</code> option to the queue or on single messages	false	MEDIUM
<code>camel.component.aws2-sqs.queueUrl</code>	To define the <code>queueUrl</code> explicitly. All other parameters, which would influence the <code>queueUrl</code> , are ignored. This parameter is intended to be used, to connect to a mock implementation of SQS, for testing purposes.	null	MEDIUM
<code>camel.component.aws2-sqs.proxyHost</code>	To define a proxy host when instantiating the SQS client	null	MEDIUM
<code>camel.component.aws2-sqs.proxyPort</code>	To define a proxy port when instantiating the SQS client	null	MEDIUM
<code>camel.component.aws2-sqs.maximumMessageSize</code>	The <code>maximumMessageSize</code> (in bytes) an SQS message can contain for this queue.	null	MEDIUM
<code>camel.component.aws2-sqs.messageRetentionPeriod</code>	The <code>messageRetentionPeriod</code> (in seconds) a message will be retained by SQS for this queue.	null	MEDIUM
<code>camel.component.aws2-sqs.policy</code>	The policy for this queue	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-sqs.receiveMessageWaitTimeSeconds</code>	If you do not specify <code>WaitTimeSeconds</code> in the request, the queue attribute <code>ReceiveMessageWaitTimeSeconds</code> is used to determine how long to wait.	null	MEDIUM
<code>camel.component.aws2-sqs.redrivePolicy</code>	Specify the policy that send message to DeadLetter queue. See detail at Amazon docs.	null	MEDIUM
<code>camel.component.aws2-sqs.accessKey</code>	Amazon AWS Access Key	null	MEDIUM
<code>camel.component.aws2-sqs.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM

The camel-aws2-sqs sink connector has no converters out of the box.

The camel-aws2-sqs sink connector has no transforms out of the box.

The camel-aws2-sqs sink connector has no aggregation strategies out of the box.

## 5.7. CAMEL-AWS2-SQS-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-aws2-sqs-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-aws2-sqs-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.aws2sqs.CamelAws2sqsSourceConnector
```

The camel-aws2-sqs source connector supports 91 options, which are listed below.

Name	Description	Default	Priority
<code>camel.source.path.queueNameOrArn</code>	Queue name or ARN	null	HIGH
<code>camel.source.endpoint.amazonAWSHost</code>	The hostname of the Amazon AWS cloud.	"amazonaws.com"	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.amazonSQSClient</code>	To use the AmazonSQS as client	null	MEDIUM
<code>camel.source.endpoint.autoCreateQueue</code>	Setting the autocreation of the queue	true	MEDIUM
<code>camel.source.endpoint.autoDiscoverClient</code>	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
<code>camel.source.endpoint.headerFilterStrategy</code>	To use a custom HeaderFilterStrategy to map headers to/from Camel.	null	MEDIUM
<code>camel.source.endpoint.protocol</code>	The underlying protocol used to communicate with SQS	"https"	MEDIUM
<code>camel.source.endpoint.proxyProtocol</code>	To define a proxy protocol when instantiating the SQS client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<code>camel.source.endpoint.queueOwnerAWSAccountId</code>	Specify the queue owner aws account id when you need to connect the queue with different account owner.	null	MEDIUM
<code>camel.source.endpoint.region</code>	The region in which SQS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name Region.EU_WEST_1.id()	null	MEDIUM
<code>camel.source.endpoint.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.source.endpoint.attributeNames</code>	A list of attribute names to receive when consuming. Multiple names can be separated by comma.	null	MEDIUM
<code>camel.source.endpoint.bridgeErrorHandler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.concurrentConsumers</code>	Allows you to use multiple threads to poll the sqs queue to increase throughput	1	MEDIUM
<code>camel.source.endpoint.defaultVisibilityTimeout</code>	The default visibility timeout (in seconds)	null	MEDIUM
<code>camel.source.endpoint.deleteAfterRead</code>	Delete message from SQS after it has been read	true	MEDIUM
<code>camel.source.endpoint.deleteIfFiltered</code>	Whether or not to send the DeleteMessage to the SQS queue if an exchange fails to get through a filter. If 'false' and exchange does not make it through a Camel filter upstream in the route, then don't send DeleteMessage.	true	MEDIUM
<code>camel.source.endpoint.extendMessageVisibility</code>	If enabled then a scheduled background task will keep extending the message visibility on SQS. This is needed if it takes a long time to process the message. If set to true <code>defaultVisibilityTimeout</code> must be set. See details at Amazon docs.	false	MEDIUM
<code>camel.source.endpoint.kmsDataKeyReusePeriodSeconds</code>	The length of time, in seconds, for which Amazon SQS can reuse a data key to encrypt or decrypt messages before calling AWS KMS again. An integer representing seconds, between 60 seconds (1 minute) and 86,400 seconds (24 hours). Default: 300 (5 minutes).	null	MEDIUM
<code>camel.source.endpoint.kmsMasterKeyId</code>	The ID of an AWS-managed customer master key (CMK) for Amazon SQS or a custom CMK.	null	MEDIUM
<code>camel.source.endpoint.maxMessagesPerPoll</code>	Gets the maximum number of messages as a limit to poll at each polling. Is default unlimited, but use 0 or negative number to disable it as unlimited.	null	MEDIUM
<code>camel.source.endpoint.messageAttributeNames</code>	A list of message attribute names to receive when consuming. Multiple names can be separated by comma.	null	MEDIUM
<code>camel.source.endpoint.sendEmptyMessageWhenIdle</code>	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	MEDIUM
<code>camel.source.endpoint.serverSideEncryptionEnabled</code>	Define if Server Side Encryption is enabled or not on the queue	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.visibilityTimeout</b>	The duration (in seconds) that the received messages are hidden from subsequent retrieve requests after being retrieved by a <code>ReceiveMessage</code> request to set in the <code>com.amazonaws.services.sqs.model.SetQueueAttributesRequest</code> . This only make sense if its different from <code>defaultVisibilityTimeout</code> . It changes the queue visibility timeout attribute permanently.	null	MEDIUM
<b>camel.source.endpoint.waitTimeSeconds</b>	Duration in seconds (0 to 20) that the <code>ReceiveMessage</code> action call will wait until a message is in the queue to include in the response.	null	MEDIUM
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at <code>WARN</code> or <code>ERROR</code> level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: <code>[InOnly]</code> <code>[InOut]</code> <code>[InOptionalOut]</code>	null	MEDIUM
<b>camel.source.endpoint.pollStrategy</b>	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an <code>Exchange</code> have been created and being routed in Camel.	null	MEDIUM
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.source.endpoint.delayQueue</b>	Define if you want to apply <code>delaySeconds</code> option to the queue or on single messages	false	MEDIUM
<b>camel.source.endpoint.queueUrl</b>	To define the <code>queueUrl</code> explicitly. All other parameters, which would influence the <code>queueUrl</code> , are ignored. This parameter is intended to be used, to connect to a mock implementation of <code>SQS</code> , for testing purposes.	null	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.proxyHost</code>	To define a proxy host when instantiating the SQS client	null	MEDIUM
<code>camel.source.endpoint.proxyPort</code>	To define a proxy port when instantiating the SQS client	null	MEDIUM
<code>camel.source.endpoint.maximumMessageSize</code>	The <code>maximumMessageSize</code> (in bytes) an SQS message can contain for this queue.	null	MEDIUM
<code>camel.source.endpoint.messageRetentionPeriod</code>	The <code>messageRetentionPeriod</code> (in seconds) a message will be retained by SQS for this queue.	null	MEDIUM
<code>camel.source.endpoint.policy</code>	The policy for this queue	null	MEDIUM
<code>camel.source.endpoint.receiveMessageWaitTimeSeconds</code>	If you do not specify <code>WaitTimeSeconds</code> in the request, the queue attribute <code>ReceiveMessageWaitTimeSeconds</code> is used to determine how long to wait.	null	MEDIUM
<code>camel.source.endpoint.redrivePolicy</code>	Specify the policy that send message to DeadLetter queue. See detail at Amazon docs.	null	MEDIUM
<code>camel.source.endpoint.backoffErrorThreshold</code>	The number of subsequent error polls (failed due some error) that should happen before the <code>backoffMultiplier</code> should kick-in.	null	MEDIUM
<code>camel.source.endpoint.backoffIdleThreshold</code>	The number of subsequent idle polls that should happen before the <code>backoffMultiplier</code> should kick-in.	null	MEDIUM
<code>camel.source.endpoint.backoffMultiplier</code>	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then <code>backoffIdleThreshold</code> and/or <code>backoffErrorThreshold</code> must also be configured.	null	MEDIUM
<code>camel.source.endpoint.delay</code>	Milliseconds before the next poll.	500L	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.greedy</code>	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	MEDIUM
<code>camel.source.endpoint.initialDelay</code>	Milliseconds before the first poll starts.	1000L	MEDIUM
<code>camel.source.endpoint.repeatCount</code>	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0L	MEDIUM
<code>camel.source.endpoint.runLoggingLevel</code>	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"TRACE"	MEDIUM
<code>camel.source.endpoint.scheduledExecutorService</code>	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.	null	MEDIUM
<code>camel.source.endpoint.scheduler</code>	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler	"none"	MEDIUM
<code>camel.source.endpoint.schedulerProperties</code>	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.	null	MEDIUM
<code>camel.source.endpoint.startScheduler</code>	Whether the scheduler should be auto started.	true	MEDIUM
<code>camel.source.endpoint.timeUnit</code>	Time unit for initialDelay and delay options. One of: [NANOSECONDS] [MICROSECONDS] [MILLISECONDS] [SECONDS] [MINUTES] [HOURS] [DAYS]	"MILLISECONDS"	MEDIUM
<code>camel.source.endpoint.useFixedDelay</code>	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	MEDIUM
<code>camel.source.endpoint.accessKey</code>	Amazon AWS Access Key	null	MEDIUM
<code>camel.source.endpoint.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-sqs.amazonAWSHost</code>	The hostname of the Amazon AWS cloud.	"amazonaws.com"	MEDIUM
<code>camel.component.aws2-sqs.amazonSQSClient</code>	To use the AmazonSQS as client	null	MEDIUM
<code>camel.component.aws2-sqs.autoCreateQueue</code>	Setting the autocreation of the queue	true	MEDIUM
<code>camel.component.aws2-sqs.autoDiscoverClient</code>	Setting the autoDiscoverClient mechanism, if true, the component will look for a client instance in the registry automatically otherwise it will skip that checking.	true	MEDIUM
<code>camel.component.aws2-sqs.configuration</code>	The AWS SQS default configuration	null	MEDIUM
<code>camel.component.aws2-sqs.protocol</code>	The underlying protocol used to communicate with SQS	"https"	MEDIUM
<code>camel.component.aws2-sqs.proxyProtocol</code>	To define a proxy protocol when instantiating the SQS client One of: [HTTP] [HTTPS]	"HTTPS"	MEDIUM
<code>camel.component.aws2-sqs.queueOwnerAWSAccountid</code>	Specify the queue owner aws account id when you need to connect the queue with different account owner.	null	MEDIUM
<code>camel.component.aws2-sqs.region</code>	The region in which SQS client needs to work. When using this parameter, the configuration will expect the lowercase name of the region (for example ap-east-1) You'll need to use the name <code>Region.EU_WEST_1.id()</code>	null	MEDIUM
<code>camel.component.aws2-sqs.trustAllCertificates</code>	If we want to trust all certificates in case of overriding the endpoint	false	MEDIUM
<code>camel.component.aws2-sqs.attributeNames</code>	A list of attribute names to receive when consuming. Multiple names can be separated by comma.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-sqs.bridgeErrorHandler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<code>camel.component.aws2-sqs.concurrentConsumers</code>	Allows you to use multiple threads to poll the sqs queue to increase throughput	1	MEDIUM
<code>camel.component.aws2-sqs.defaultVisibilityTimeout</code>	The default visibility timeout (in seconds)	null	MEDIUM
<code>camel.component.aws2-sqs.deleteAfterRead</code>	Delete message from SQS after it has been read	true	MEDIUM
<code>camel.component.aws2-sqs.deleteIfFiltered</code>	Whether or not to send the DeleteMessage to the SQS queue if an exchange fails to get through a filter. If 'false' and exchange does not make it through a Camel filter upstream in the route, then don't send DeleteMessage.	true	MEDIUM
<code>camel.component.aws2-sqs.extendMessageVisibility</code>	If enabled then a scheduled background task will keep extending the message visibility on SQS. This is needed if it takes a long time to process the message. If set to true <code>defaultVisibilityTimeout</code> must be set. See details at Amazon docs.	false	MEDIUM
<code>camel.component.aws2-sqs.kmsDataKeyReusePeriodSeconds</code>	The length of time, in seconds, for which Amazon SQS can reuse a data key to encrypt or decrypt messages before calling AWS KMS again. An integer representing seconds, between 60 seconds (1 minute) and 86,400 seconds (24 hours). Default: 300 (5 minutes).	null	MEDIUM
<code>camel.component.aws2-sqs.kmsMasterKeyId</code>	The ID of an AWS-managed customer master key (CMK) for Amazon SQS or a custom CMK.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-sqs.messageAttributeNames</code>	A list of message attribute names to receive when consuming. Multiple names can be separated by comma.	null	MEDIUM
<code>camel.component.aws2-sqs.serverSideEncryptionEnabled</code>	Define if Server Side Encryption is enabled or not on the queue	false	MEDIUM
<code>camel.component.aws2-sqs.visibilityTimeout</code>	The duration (in seconds) that the received messages are hidden from subsequent retrieve requests after being retrieved by a <code>ReceiveMessage</code> request to set in the <code>com.amazonaws.services.sqs.model.SetQueueAttributesRequest</code> . This only make sense if its different from <code>defaultVisibilityTimeout</code> . It changes the queue visibility timeout attribute permanently.	null	MEDIUM
<code>camel.component.aws2-sqs.waitTimeSeconds</code>	Duration in seconds (0 to 20) that the <code>ReceiveMessage</code> action call will wait until a message is in the queue to include in the response.	null	MEDIUM
<code>camel.component.aws2-sqs.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.component.aws2-sqs.delayQueue</code>	Define if you want to apply <code>delaySeconds</code> option to the queue or on single messages	false	MEDIUM
<code>camel.component.aws2-sqs.queueUrl</code>	To define the <code>queueUrl</code> explicitly. All other parameters, which would influence the <code>queueUrl</code> , are ignored. This parameter is intended to be used, to connect to a mock implementation of SQS, for testing purposes.	null	MEDIUM
<code>camel.component.aws2-sqs.proxyHost</code>	To define a proxy host when instantiating the SQS client	null	MEDIUM
<code>camel.component.aws2-sqs.proxyPort</code>	To define a proxy port when instantiating the SQS client	null	MEDIUM
<code>camel.component.aws2-sqs.maximumMessageSize</code>	The <code>maximumMessageSize</code> (in bytes) an SQS message can contain for this queue.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.aws2-sqs.messageRetentionPeriod</code>	The messageRetentionPeriod (in seconds) a message will be retained by SQS for this queue.	null	MEDIUM
<code>camel.component.aws2-sqs.policy</code>	The policy for this queue	null	MEDIUM
<code>camel.component.aws2-sqs.receiveMessageWaitTimeSeconds</code>	If you do not specify WaitTimeSeconds in the request, the queue attribute ReceiveMessageWaitTimeSeconds is used to determine how long to wait.	null	MEDIUM
<code>camel.component.aws2-sqs.redrivePolicy</code>	Specify the policy that send message to DeadLetter queue. See detail at Amazon docs.	null	MEDIUM
<code>camel.component.aws2-sqs.accessKey</code>	Amazon AWS Access Key	null	MEDIUM
<code>camel.component.aws2-sqs.secretKey</code>	Amazon AWS Secret Key	null	MEDIUM

The camel-aws2-sqs sink connector has no converters out of the box.

The camel-aws2-sqs sink connector has no transforms out of the box.

The camel-aws2-sqs sink connector has no aggregation strategies out of the box.

## 5.8. CAMEL-CQL-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-cql-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-cql-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.cql.CamelCqlSinkConnector
```

The camel-cql sink connector supports 19 options, which are listed below.

Name	Description	Default	Priority
<code>camel.sink.path.beanRef</code>	beanRef is defined using bean:id	null	MEDIUM
<code>camel.sink.path.hosts</code>	Hostname(s) cassandra server(s). Multiple hosts can be separated by comma.	null	MEDIUM
<code>camel.sink.path.port</code>	Port number of cassandra server(s)	null	MEDIUM
<code>camel.sink.path.keyspace</code>	Keyspace to use	null	MEDIUM
<code>camel.sink.endpoint.clusterName</code>	Cluster name	null	MEDIUM
<code>camel.sink.endpoint.consistencyLevel</code>	Consistency level to use One of: [ANY] [ONE] [TWO] [THREE] [QUORUM] [ALL] [LOCAL_ONE] [LOCAL_QUORUM] [EACH_QUORUM] [SERIAL] [LOCAL_SERIAL]	null	MEDIUM
<code>camel.sink.endpoint.cql</code>	CQL query to perform. Can be overridden with the message header with key CamelCqlQuery.	null	MEDIUM
<code>camel.sink.endpoint.datacenter</code>	Datacenter to use	"datacenter1"	MEDIUM
<code>camel.sink.endpoint.loadBalancingPolicyClass</code>	To use a specific LoadBalancingPolicyClass	null	MEDIUM
<code>camel.sink.endpoint.password</code>	Password for session authentication	null	MEDIUM
<code>camel.sink.endpoint.prepareStatements</code>	Whether to use PreparedStatements or regular Statements	true	MEDIUM
<code>camel.sink.endpoint.resultSetConversionStrategy</code>	To use a custom class that implements logic for converting ResultSet into message body ALL, ONE, LIMIT_10, LIMIT_100...	null	MEDIUM
<code>camel.sink.endpoint.session</code>	To use the Session instance (you would normally not use this option)	null	MEDIUM
<code>camel.sink.endpoint.username</code>	Username for session authentication	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.component.cql.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.component.cql.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

The camel-cql sink connector has no converters out of the box.

The camel-cql sink connector has no transforms out of the box.

The camel-cql sink connector has no aggregation strategies out of the box.

## 5.9. CAMEL-CQL-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-cql-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-cql-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.cql.CamelCqlSourceConnector
```

The camel-cql source connector supports 37 options, which are listed below.

Name	Description	Default	Priority
<b>camel.source.path.beanRef</b>	beanRef is defined using bean:id	null	MEDIUM
<b>camel.source.path.hosts</b>	Hostname(s) cassandra server(s). Multiple hosts can be separated by comma.	null	MEDIUM
<b>camel.source.path.port</b>	Port number of cassandra server(s)	null	MEDIUM
<b>camel.source.path.keyspace</b>	Keyspace to use	null	MEDIUM
<b>camel.source.endpoint.clusterName</b>	Cluster name	null	MEDIUM
<b>camel.source.endpoint.consistencyLevel</b>	Consistency level to use One of: [ANY] [ONE] [TWO] [THREE] [QUORUM] [ALL] [LOCAL_ONE] [LOCAL_QUORUM] [EACH_QUORUM] [SERIAL] [LOCAL_SERIAL]	null	MEDIUM
<b>camel.source.endpoint.cql</b>	CQL query to perform. Can be overridden with the message header with key CamelCqlQuery.	null	MEDIUM
<b>camel.source.endpoint.datacenter</b>	Datacenter to use	"datacenter1"	MEDIUM
<b>camel.source.endpoint.loadBalancingPolicyClasses</b>	To use a specific LoadBalancingPolicyClass	null	MEDIUM
<b>camel.source.endpoint.password</b>	Password for session authentication	null	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.prepareStatements</code>	Whether to use PreparedStatements or regular Statements	true	MEDIUM
<code>camel.source.endpoint.resultSetConversionStrategy</code>	To use a custom class that implements logic for converting ResultSet into message body ALL, ONE, LIMIT_10, LIMIT_100...	null	MEDIUM
<code>camel.source.endpoint.session</code>	To use the Session instance (you would normally not use this option)	null	MEDIUM
<code>camel.source.endpoint.username</code>	Username for session authentication	null	MEDIUM
<code>camel.source.endpoint.bridgeErrorHandler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<code>camel.source.endpoint.sendEmptyMessageWhenIdle</code>	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	MEDIUM
<code>camel.source.endpoint.exceptionHandler</code>	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<code>camel.source.endpoint.exchangePattern</code>	Sets the exchange pattern when the consumer creates an exchange. One of: [InOnly] [InOut] [InOptionalOut]	null	MEDIUM
<code>camel.source.endpoint.pollStrategy</code>	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.source.endpoint.backoffErrorThreshold</b>	The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.	null	MEDIUM
<b>camel.source.endpoint.backoffIdleThreshold</b>	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.	null	MEDIUM
<b>camel.source.endpoint.backoffMultiplier</b>	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.	null	MEDIUM
<b>camel.source.endpoint.delay</b>	Milliseconds before the next poll.	500L	MEDIUM
<b>camel.source.endpoint.greedy</b>	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	MEDIUM
<b>camel.source.endpoint.initialDelay</b>	Milliseconds before the first poll starts.	1000L	MEDIUM
<b>camel.source.endpoint.repeatCount</b>	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0L	MEDIUM
<b>camel.source.endpoint.runLoggingLevel</b>	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"TRACE"	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.scheduledExecutorService</b>	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.	null	MEDIUM
<b>camel.source.endpoint.scheduler</b>	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler	"none"	MEDIUM
<b>camel.source.endpoint.schedulerProperties</b>	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.	null	MEDIUM
<b>camel.source.endpoint.startScheduler</b>	Whether the scheduler should be auto started.	true	MEDIUM
<b>camel.source.endpoint.timeUnit</b>	Time unit for initialDelay and delay options. One of: [NANOSECONDS] [MICROSECONDS] [MILLISECONDS] [SECONDS] [MINUTES] [HOURS] [DAYS]	"MILLISECONDS"	MEDIUM
<b>camel.source.endpoint.useFixedDelay</b>	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	MEDIUM
<b>camel.component.cql.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.component.cql.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

The camel-cql sink connector has no converters out of the box.

The camel-cql sink connector has no transforms out of the box.

The camel-cql sink connector has no aggregation strategies out of the box.

## 5.10. CAMEL-ELASTICSEARCH-REST-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-elasticsearch-rest-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-elasticsearch-rest-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.elasticsearchrest.CamelElasticsearchrestSinkConnector
```

The camel-elasticsearch-rest sink connector supports 33 options, which are listed below.

Name	Description	Default	Priority
camel.sink.path.clusterName	Name of the cluster	null	HIGH
camel.sink.endpoint.connectionTimeout	The time in ms to wait before connection will timeout.	30000	MEDIUM
camel.sink.endpoint.disconnect	Disconnect after it finish calling the producer	false	MEDIUM
camel.sink.endpoint.enableSniffer	Enable automatically discover nodes from a running Elasticsearch cluster	false	MEDIUM
camel.sink.endpoint.enableSSL	Enable SSL	false	MEDIUM
camel.sink.endpoint.from	Starting index of the response.	null	MEDIUM
camel.sink.endpoint.hostAddresses	Comma separated list with ip:port formatted remote transport addresses to use.	null	HIGH
camel.sink.endpoint.indexName	The name of the index to act against	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.maxRetryTimeout</b>	The time in ms before retry	30000	MEDIUM
<b>camel.sink.endpoint.operation</b>	What operation to perform One of: [Index] [Update] [Bulk] [BulkIndex] [GetById] [MultiGet] [MultiSearch] [Delete] [DeleteIndex] [Search] [Exists] [Ping]	null	MEDIUM
<b>camel.sink.endpoint.scrollKeepAliveMs</b>	Time in ms during which elasticsearch will keep search context alive	60000	MEDIUM
<b>camel.sink.endpoint.size</b>	Size of the response.	null	MEDIUM
<b>camel.sink.endpoint.sniffAfterFailureDelay</b>	The delay of a sniff execution scheduled after a failure (in milliseconds)	60000	MEDIUM
<b>camel.sink.endpoint.sniffInterval</b>	The interval between consecutive ordinary sniff executions in milliseconds. Will be honoured when sniffOnFailure is disabled or when there are no failures between consecutive sniff executions	300000	MEDIUM
<b>camel.sink.endpoint.socketTimeout</b>	The timeout in ms to wait before the socket will timeout.	30000	MEDIUM
<b>camel.sink.endpoint.useScroll</b>	Enable scroll usage	false	MEDIUM
<b>camel.sink.endpoint.waitForActiveShards</b>	Index creation waits for the write consistency number of shards to be available	1	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.synchronous</code>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<code>camel.component.elasticsearch-rest.lazyStartProducer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<code>camel.component.elasticsearch-rest.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.component.elasticsearch-rest.client</code>	To use an existing configured Elasticsearch client, instead of creating a client per endpoint. This allow to customize the client with specific settings.	null	MEDIUM
<code>camel.component.elasticsearch-rest.connectionTimeout</code>	The time in ms to wait before connection will timeout.	30000	MEDIUM
<code>camel.component.elasticsearch-rest.enableSniffer</code>	Enable automatically discover nodes from a running Elasticsearch cluster	"false"	MEDIUM
<code>camel.component.elasticsearch-rest.hostAddresses</code>	Comma separated list with ip:port formatted remote transport addresses to use. The ip and port options must be left blank for hostAddresses to be considered instead.	null	MEDIUM
<code>camel.component.elasticsearch-rest.maxRetryTimeout</code>	The time in ms before retry	30000	MEDIUM
<code>camel.component.elasticsearch-rest.sniffAfterFailureDelay</code>	The delay of a sniff execution scheduled after a failure (in milliseconds)	60000	MEDIUM

Name	Description	Default	Priority
<code>camel.component.elasticsearch-rest.snifferInterval</code>	The interval between consecutive ordinary sniff executions in milliseconds. Will be honoured when <code>sniffOnFailure</code> is disabled or when there are no failures between consecutive sniff executions	300000	MEDIUM
<code>camel.component.elasticsearch-rest.socketTimeout</code>	The timeout in ms to wait before the socket will timeout.	30000	MEDIUM
<code>camel.component.elasticsearch-rest.enableSSL</code>	Enable SSL	"false"	MEDIUM
<code>camel.component.elasticsearch-rest.password</code>	Password for authenticate	null	MEDIUM
<code>camel.component.elasticsearch-rest.user</code>	Basic authenticate user	null	MEDIUM

The camel-elasticsearch-rest sink connector has no converters out of the box.

The camel-elasticsearch-rest sink connector has no transforms out of the box.

The camel-elasticsearch-rest sink connector has no aggregation strategies out of the box.

## 5.11. CAMEL-FILE-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-file-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-file-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following `connector.class`

```
connector.class=org.apache.camel.kafkaconnector.file.CamelFileSinkConnector
```

The camel-file sink connector supports 27 options, which are listed below.

Name	Description	Default	Priority
<code>camel.sink.path.directoryName</code>	The starting directory	null	HIGH

Name	Description	Default	Priority
<b>camel.sink.endpoint.charset</b>	This option is used to specify the encoding of the file. You can use this on the consumer, to specify the encodings of the files, which allow Camel to know the charset it should load the file content in case the file content is being accessed. Likewise when writing a file, you can use this option to specify which charset to write the file as well. Do mind that when writing the file Camel may have to read the message content into memory to be able to convert the data into the configured charset, so do not use this if you have big messages.	null	MEDIUM
<b>camel.sink.endpoint.doneFileName</b>	Producer: If provided, then Camel will write a 2nd done file when the original file has been written. The done file will be empty. This option configures what file name to use. Either you can specify a fixed name. Or you can use dynamic placeholders. The done file will always be written in the same folder as the original file. Consumer: If provided, Camel will only consume files if a done file exists. This option configures what file name to use. Either you can specify a fixed name. Or you can use dynamic placeholders. The done file is always expected in the same folder as the original file. Only <code>\\${file.name}</code> and <code>\\${file.name.next}</code> is supported as dynamic placeholders.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.file Name</b>	Use Expression such as File Language to dynamically set the filename. For consumers, it's used as a filename filter. For producers, it's used to evaluate the filename to write. If an expression is set, it take precedence over the CamelFileName header. (Note: The header itself can also be an Expression). The expression options support both String and Expression types. If the expression is a String type, it is always evaluated using the File Language. If the expression is an Expression type, the specified Expression type is used - this allows you, for instance, to use OGNL expressions. For the consumer, you can use it to filter filenames, so you can for instance consume today's file using the File Language syntax: mydata- <code>\\${date:now:yyyyMMdd}.txt</code> . The producers support the CamelOverrideFileName header which takes precedence over any existing CamelFileName header; the CamelOverrideFileName is a header that is used only once, and makes it easier as this avoids to temporary store CamelFileName and have to restore it afterwards.	null	MEDIUM
<b>camel.sink.endpoint.appendChars</b>	Used to append characters (text) after writing files. This can for example be used to add new lines or other separators when writing and appending to existing files. To specify new-line (slash-n or slash-r) or tab (slash-t) characters then escape with an extra slash, eg slash-slash-n.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.fileExist</b>	What to do if a file already exists with the same name. Override, which is the default, replaces the existing file. - Append - adds content to the existing file. - Fail - throws a <code>GenericFileOperationException</code> , indicating that there is already an existing file. - Ignore - silently ignores the problem and does not override the existing file, but assumes everything is okay. - Move - option requires to use the <code>moveExisting</code> option to be configured as well. The option <code>eagerDeleteTargetFile</code> can be used to control what to do if an moving the file, and there exists already an existing file, otherwise causing the move operation to fail. The Move option will move any existing files, before writing the target file. - TryRename is only applicable if <code>tempFileName</code> option is in use. This allows to try renaming the file from the temporary name to the actual name, without doing any exists check. This check may be faster on some file systems and especially FTP servers. One of: [Override] [Append] [Fail] [Ignore] [Move] [TryRename]	"Override"	MEDIUM
<b>camel.sink.endpoint.flatten</b>	Flatten is used to flatten the file name path to strip any leading paths, so it's just the file name. This allows you to consume recursively into sub-directories, but when you eg write the files to another directory they will be written in a single directory. Setting this to true on the producer enforces that any file name in <code>CamelFileName</code> header will be stripped for any leading paths.	false	MEDIUM
<b>camel.sink.endpoint.jailStartingDirectory</b>	Used for jailing (restricting) writing files to the starting directory (and sub) only. This is enabled by default to not allow Camel to write files to outside directories (to be more secured out of the box). You can turn this off to allow writing files to directories outside the starting directory, such as parent or root folders.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.moveExisting</b>	Expression (such as File Language) used to compute file name to use when fileExist=Move is configured. To move files into a backup subdirectory just enter backup. This option only supports the following File Language tokens: file:name, file:name.ext, file:name.noext, file:onlyname, file:onlyname.noext, file:ext, and file:parent. Notice the file:parent is not supported by the FTP component, as the FTP component can only move any existing files to a relative directory based on current dir as base.	null	MEDIUM
<b>camel.sink.endpoint.tempFileName</b>	The same as tempPrefix option but offering a more fine grained control on the naming of the temporary filename as it uses the File Language. The location for tempFileName is relative to the final file location in the option 'fileName', not the target directory in the base uri. For example if option fileName includes a directory prefix: dir/finalFilename then tempFileName is relative to that subdirectory dir.	null	MEDIUM
<b>camel.sink.endpoint.tempPrefix</b>	This option is used to write the file using a temporary name and then, after the write is complete, rename it to the real name. Can be used to identify files being written and also avoid consumers (not using exclusive read locks) reading in progress files. Is often used by FTP when uploading big files.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.allowNullBody</b>	Used to specify if a null body is allowed during file writing. If set to true then an empty file will be created, when set to false, and attempting to send a null body to the file component, a <code>GenericFileWriteException</code> of 'Cannot write null body to file.' will be thrown. If the <code>fileExist</code> option is set to 'Override', then the file will be truncated, and if set to append the file will remain unchanged.	false	MEDIUM
<b>camel.sink.endpoint.chmod</b>	Specify the file permissions which is sent by the producer, the <code>chmod</code> value must be between 000 and 777; If there is a leading digit like in 0755 we will ignore it.	null	MEDIUM
<b>camel.sink.endpoint.chmodDirectory</b>	Specify the directory permissions used when the producer creates missing directories, the <code>chmod</code> value must be between 000 and 777; If there is a leading digit like in 0755 we will ignore it.	null	MEDIUM
<b>camel.sink.endpoint.eagerDeleteTargetFile</b>	Whether or not to eagerly delete any existing target file. This option only applies when you use <code>fileExists=Override</code> and the <code>tempFileName</code> option as well. You can use this to disable (set it to false) deleting the target file before the temp file is written. For example you may write big files and want the target file to exist during the temp file is being written. This ensure the target file is only deleted until the very last moment, just before the temp file is being renamed to the target filename. This option is also used to control whether to delete any existing files when <code>fileExist=Move</code> is enabled, and an existing file exists. If this option <code>copyAndDeleteOnRenameFails</code> false, then an exception will be thrown if an existing file existed, if its true, then the existing file is deleted before the move operation.	true	MEDIUM
<b>camel.sink.endpoint.forceWrites</b>	Whether to force syncing writes to the file system. You can turn this off if you do not want this level of guarantee, for example if writing to logs / audit logs etc; this would yield better performance.	true	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.keepLastModified</code>	Will keep the last modified timestamp from the source file (if any). Will use the <code>Exchange.FILE_LAST_MODIFIED</code> header to locate the timestamp. This header can contain either a <code>java.util.Date</code> or long with the timestamp. If the timestamp exists and the option is enabled it will set this timestamp on the written file. Note: This option only applies to the file producer. You cannot use this option with any of the <code>\ftp</code> producers.	false	MEDIUM
<code>camel.sink.endpoint.moveExistingFileStrategy</code>	Strategy (Custom Strategy) used to move file with special naming token to use when <code>fileExist=Move</code> is configured. By default, there is an implementation used if no custom strategy is provided	null	MEDIUM
<code>camel.sink.endpoint.autoCreate</code>	Automatically create missing directories in the file's pathname. For the file consumer, that means creating the starting directory. For the file producer, it means the directory the files should be written to.	true	MEDIUM
<code>camel.sink.endpoint.basicPropertyBinding</code>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.sink.endpoint.bufferSize</code>	Buffer size in bytes used for writing files (or in case of FTP for downloading and uploading files).	131072	MEDIUM
<code>camel.sink.endpoint.copyAndDeleteOnRenameFail</code>	Whether to fallback and do a copy and delete file, in case the file could not be renamed directly. This option is not available for the FTP component.	true	MEDIUM
<code>camel.sink.endpoint.renameUsingCopy</code>	Perform rename operations using a copy and delete strategy. This is primarily used in environments where the regular rename operation is unreliable (e.g. across different file systems or networks). This option takes precedence over the <code>copyAndDeleteOnRenameFail</code> parameter that will automatically fall back to the copy and delete strategy, but only after additional delays.	false	MEDIUM
<code>camel.sink.endpoint.synchronous</code>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.file.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.component.file.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

The camel-file sink connector has no converters out of the box.

The camel-file sink connector has no transforms out of the box.

The camel-file sink connector has no aggregation strategies out of the box.

## 5.12. CAMEL-HDFS-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-hdfs-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-hdfs-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.hdfs.CamelHdfsSinkConnector
```

The camel-hdfs sink connector supports 32 options, which are listed below.

Name	Description	Default	Priority
<b>camel.sink.path.hostName</b>	HDFS host to use	null	HIGH
<b>camel.sink.path.port</b>	HDFS port to use	8020	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.path.path</code>	The directory path to use	null	HIGH
<code>camel.sink.endpoint.connectOnStartup</code>	Whether to connect to the HDFS file system on starting the producer/consumer. If false then the connection is created on-demand. Notice that HDFS may take up till 15 minutes to establish a connection, as it has hardcoded 45 x 20 sec redelivery. By setting this option to false allows your application to startup, and not block for up till 15 minutes.	true	MEDIUM
<code>camel.sink.endpoint.fileSystemType</code>	Set to LOCAL to not use HDFS but local java.io.File instead. One of: [LOCAL] [HDFS]	"HDFS"	MEDIUM
<code>camel.sink.endpoint.fileType</code>	The file type to use. For more details see Hadoop HDFS documentation about the various files types. One of: [NORMAL_FILE] [SEQUENCE_FILE] [MAP_FILE] [BLOOMMAP_FILE] [ARRAY_FILE]	"NORMAL_FILE"	MEDIUM
<code>camel.sink.endpoint.keyType</code>	The type for the key in case of sequence or map files. One of: [NULL] [BOOLEAN] [BYTE] [INT] [FLOAT] [LONG] [DOUBLE] [TEXT] [BYTES]	"NULL"	MEDIUM
<code>camel.sink.endpoint.namedNodes</code>	A comma separated list of named nodes (e.g. srv11.example.com:8020,srv12.example.com:8020)	null	MEDIUM
<code>camel.sink.endpoint.owner</code>	The file owner must match this owner for the consumer to pickup the file. Otherwise the file is skipped.	null	MEDIUM
<code>camel.sink.endpoint.valueType</code>	The type for the key in case of sequence or map files One of: [NULL] [BOOLEAN] [BYTE] [INT] [FLOAT] [LONG] [DOUBLE] [TEXT] [BYTES]	"BYTES"	MEDIUM
<code>camel.sink.endpoint.append</code>	Append to existing file. Notice that not all HDFS file systems support the append option.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.override</b>	Whether to overwrite existing files with the same name	true	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.blockSize</b>	The size of the HDFS blocks	67108864L	MEDIUM
<b>camel.sink.endpoint.bufferSize</b>	The buffer size used by HDFS	4096	MEDIUM
<b>camel.sink.endpoint.checkIdleInterval</b>	How often (time in millis) in to run the idle checker background task. This option is only in use if the splitter strategy is IDLE.	500	MEDIUM
<b>camel.sink.endpoint.chunkSize</b>	When reading a normal file, this is split into chunks producing a message per chunk.	4096	MEDIUM
<b>camel.sink.endpoint.compressionCodec</b>	The compression codec to use One of: [DEFAULT] [GZIP] [BZIP2]	"DEFAULT"	MEDIUM
<b>camel.sink.endpoint.compressionType</b>	The compression type to use (is default not in use) One of: [NONE] [RECORD] [BLOCK]	"NONE"	MEDIUM
<b>camel.sink.endpoint.openedSuffix</b>	When a file is opened for reading/writing the file is renamed with this suffix to avoid to read it during the writing phase.	"opened"	MEDIUM
<b>camel.sink.endpoint.readSuffix</b>	Once the file has been read is renamed with this suffix to avoid to read it again.	"read"	MEDIUM
<b>camel.sink.endpoint.replication</b>	The HDFS replication factor	3	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.splitStrategy</b>	In the current version of Hadoop opening a file in append mode is disabled since it's not very reliable. So, for the moment, it's only possible to create new files. The Camel HDFS endpoint tries to solve this problem in this way: If the split strategy option has been defined, the hdfs path will be used as a directory and files will be created using the configured UuidGenerator. Every time a splitting condition is met, a new file is created. The splitStrategy option is defined as a string with the following syntax: splitStrategy=ST:value,ST:value,... where ST can be: BYTES a new file is created, and the old is closed when the number of written bytes is more than value MESSAGES a new file is created, and the old is closed when the number of written messages is more than value IDLE a new file is created, and the old is closed when no writing happened in the last value milliseconds	null	MEDIUM
<b>camel.sink.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.sink.endpoint.kerberosConfigFileLocation</b>	The location of the kerb5.conf file ( <a href="https://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html">https://web.mit.edu/kerberos/krb5-1.12/doc/admin/conf_files/krb5_conf.html</a> )	null	MEDIUM
<b>camel.sink.endpoint.kerberosKeytabLocation</b>	The location of the keytab file used to authenticate with the kerberos nodes (contains pairs of kerberos principals and encrypted keys (which are derived from the Kerberos password))	null	MEDIUM
<b>camel.sink.endpoint.kerberosUsername</b>	The username used to authenticate with the kerberos nodes	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.hdfs.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.component.hdfs.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.component.hdfs.jAASConfiguration</b>	To use the given configuration for security with JAAS.	null	MEDIUM
<b>camel.component.hdfs.kerberosConfigFile</b>	To use kerberos authentication, set the value of the 'java.security.krb5.conf' environment variable to an existing file. If the environment variable is already set, warn if different than the specified parameter	null	MEDIUM

The camel-hdfs sink connector has no converters out of the box.

The camel-hdfs sink connector has no transforms out of the box.

The camel-hdfs sink connector has no aggregation strategies out of the box.

## 5.13. CAMEL-HTTP-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-http-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-http-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.http.CamelHttpSinkConnector
```

The camel-http sink connector supports 80 options, which are listed below.

Name	Description	Default	Priority
<code>camel.sink.path.httpUri</code>	The url of the HTTP endpoint to call.	null	HIGH
<code>camel.sink.endpoint.disableStreamCache</code>	Determines whether or not the raw input stream from Servlet is cached or not (Camel will read the stream into a in memory/overflow to file, Stream caching) cache. By default Camel will cache the Servlet input stream to support reading it multiple times to ensure it Camel can retrieve all data from the stream. However you can set this option to true when you for example need to access the raw stream, such as streaming it directly to a file or other persistent store. DefaultHttpBinding will copy the request input stream into a stream cache and put it into message body if this option is false to support reading the stream multiple times. If you use Servlet to bridge/proxy an endpoint then consider enabling this option to improve performance, in case you do not need to read the message payload multiple times. The http producer will by default cache the response body stream. If setting this option to true, then the producers will not cache the response body stream but use the response stream as-is as the message body.	false	MEDIUM
<code>camel.sink.endpoint.headerFilterStrategy</code>	To use a custom HeaderFilterStrategy to filter header to and from Camel message.	null	MEDIUM
<code>camel.sink.endpoint.httpBinding</code>	To use a custom HttpBinding to control the mapping between Camel message and HttpClient.	null	MEDIUM
<code>camel.sink.endpoint.bridgeEndpoint</code>	If the option is true, HttpProducer will ignore the Exchange.HTTP_URI header, and use the endpoint's URI for request. You may also set the option <code>throwExceptionOnFailure</code> to be false to let the HttpProducer send all the fault response back.	false	MEDIUM
<code>camel.sink.endpoint.chunked</code>	If this option is false the Servlet will disable the HTTP streaming and set the content-length header on the response	true	MEDIUM
<code>camel.sink.endpoint.clearExpiredCookies</code>	Whether to clear expired cookies before sending the HTTP request. This ensures the cookies store does not keep growing by adding new cookies which is newer removed when they are expired.	true	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.connectionClose</code>	Specifies whether a Connection Close header must be added to HTTP Request. By default <code>connectionClose</code> is false.	false	MEDIUM
<code>camel.sink.endpoint.copyHeaders</code>	If this option is true then IN exchange headers will be copied to OUT exchange headers according to copy strategy. Setting this to false, allows to only include the headers from the HTTP response (not propagating IN headers).	true	MEDIUM
<code>camel.sink.endpoint.customHostHeader</code>	To use custom host header for producer. When not set in query will be ignored. When set will override host header derived from url.	null	MEDIUM
<code>camel.sink.endpoint.httpMethod</code>	Configure the HTTP method to use. The <code>HttpMethod</code> header cannot override this option if set. One of: [GET] [POST] [PUT] [DELETE] [HEAD] [OPTIONS] [TRACE] [PATCH]	null	MEDIUM
<code>camel.sink.endpoint.ignoreResponseBody</code>	If this option is true, The http producer won't read response body and cache the input stream	false	MEDIUM
<code>camel.sink.endpoint.lazyStartProducer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<code>camel.sink.endpoint.preserveHostHeader</code>	If the option is true, <code>HttpProducer</code> will set the Host header to the value contained in the current exchange Host header, useful in reverse proxy applications where you want the Host header received by the downstream server to reflect the URL called by the upstream client, this allows applications which use the Host header to generate accurate URL's for a proxied service	false	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.throwExceptionOnFailure</b>	Option to disable throwing the <code>HttpOperationFailedException</code> in case of failed responses from the remote server. This allows you to get all responses regardless of the HTTP status code.	true	MEDIUM
<b>camel.sink.endpoint.transferException</b>	If enabled and an Exchange failed processing on the consumer side, and if the caused Exception was send back serialized in the response as a <code>application/x-java-serialized-object</code> content type. On the producer side the exception will be deserialized and thrown as is, instead of the <code>HttpOperationFailedException</code> . The caused exception is required to be serialized. This is by default turned off. If you enable this then be aware that Java will deserialize the incoming data from the request to Java and that can be a potential security risk.	false	MEDIUM
<b>camel.sink.endpoint.cookieHandler</b>	Configure a cookie handler to maintain a HTTP session	null	MEDIUM
<b>camel.sink.endpoint.cookieStore</b>	To use a custom <code>CookieStore</code> . By default the <code>BasicCookieStore</code> is used which is an in-memory only cookie store. Notice if <code>bridgeEndpoint=true</code> then the cookie store is forced to be a noop cookie store as cookie shouldn't be stored as we are just bridging (eg acting as a proxy). If a <code>cookieHandler</code> is set then the cookie store is also forced to be a noop cookie store as cookie handling is then performed by the <code>cookieHandler</code> .	null	MEDIUM
<b>camel.sink.endpoint.deleteWithBody</b>	Whether the HTTP DELETE should include the message body or not. By default HTTP DELETE do not include any HTTP body. However in some rare cases users may need to be able to include the message body.	false	MEDIUM
<b>camel.sink.endpoint.getWithBody</b>	Whether the HTTP GET should include the message body or not. By default HTTP GET do not include any HTTP body. However in some rare cases users may need to be able to include the message body.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.okStatusCodeRange</b>	The status codes which are considered a success response. The values are inclusive. Multiple ranges can be defined, separated by comma, e.g. 200-204,209,301-304. Each range must be a single number or from-to with the dash included.	"200-299"	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.clientBuilder</b>	Provide access to the http client request parameters used on new RequestConfig instances used by producers or consumers of this endpoint.	null	MEDIUM
<b>camel.sink.endpoint.clientConnectionManager</b>	To use a custom HttpClientConnectionManager to manage connections	null	MEDIUM
<b>camel.sink.endpoint.connectionsPerRoute</b>	The maximum number of connections per route.	20	MEDIUM
<b>camel.sink.endpoint.httpClient</b>	Sets a custom HttpClient to be used by the producer	null	MEDIUM
<b>camel.sink.endpoint.httpClientConfigurer</b>	Register a custom configuration strategy for new HttpClient instances created by producers or consumers such as to configure authentication mechanisms etc.	null	MEDIUM
<b>camel.sink.endpoint.httpClientOptions</b>	To configure the HttpClient using the key/values from the Map.	null	MEDIUM
<b>camel.sink.endpoint.httpClientContext</b>	To use a custom HttpContext instance	null	MEDIUM
<b>camel.sink.endpoint.mapHttpMessageBody</b>	If this option is true then IN exchange Body of the exchange will be mapped to HTTP body. Setting this to false will avoid the HTTP mapping.	true	MEDIUM
<b>camel.sink.endpoint.mapHttpMessageFormUrlEncodedBody</b>	If this option is true then IN exchange Form Encoded body of the exchange will be mapped to HTTP. Setting this to false will avoid the HTTP Form Encoded body mapping.	true	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.mapHttpMessageHeaders</code>	If this option is true then IN exchange Headers of the exchange will be mapped to HTTP headers. Setting this to false will avoid the HTTP Headers mapping.	true	MEDIUM
<code>camel.sink.endpoint.maxTotalConnections</code>	The maximum number of connections.	200	MEDIUM
<code>camel.sink.endpoint.synchronous</code>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<code>camel.sink.endpoint.useSystemProperties</code>	To use System Properties as fallback for configuration	false	MEDIUM
<code>camel.sink.endpoint.proxyAuthDomain</code>	Proxy authentication domain to use with NTLM	null	MEDIUM
<code>camel.sink.endpoint.proxyAuthHost</code>	Proxy authentication host	null	MEDIUM
<code>camel.sink.endpoint.proxyAuthMethod</code>	Proxy authentication method to use One of: [Basic] [Digest] [NTLM]	null	MEDIUM
<code>camel.sink.endpoint.proxyAuthNtHost</code>	Proxy authentication domain (workstation name) to use with NTLM	null	MEDIUM
<code>camel.sink.endpoint.proxyAuthPassword</code>	Proxy authentication password	null	MEDIUM
<code>camel.sink.endpoint.proxyAuthPort</code>	Proxy authentication port	null	MEDIUM
<code>camel.sink.endpoint.proxyAuthScheme</code>	Proxy authentication scheme to use One of: [http] [https]	null	MEDIUM
<code>camel.sink.endpoint.proxyAuthUsername</code>	Proxy authentication username	null	MEDIUM
<code>camel.sink.endpoint.proxyHost</code>	Proxy hostname to use	null	MEDIUM
<code>camel.sink.endpoint.proxyPort</code>	Proxy port to use	null	MEDIUM
<code>camel.sink.endpoint.authDomain</code>	Authentication domain to use with NTLM	null	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.authenticationPreemptive</code>	If this option is true, camel-http sends preemptive basic authentication to the server.	false	MEDIUM
<code>camel.sink.endpoint.authenticationHost</code>	Authentication host to use with NTLM	null	MEDIUM
<code>camel.sink.endpoint.authenticationMethod</code>	Authentication methods allowed to use as a comma separated list of values Basic, Digest or NTLM.	null	MEDIUM
<code>camel.sink.endpoint.authenticationMethodPriority</code>	Which authentication method to prioritize to use, either as Basic, Digest or NTLM. One of: [Basic] [Digest] [NTLM]	null	MEDIUM
<code>camel.sink.endpoint.authenticationPassword</code>	Authentication password	null	MEDIUM
<code>camel.sink.endpoint.authenticationUsername</code>	Authentication username	null	MEDIUM
<code>camel.sink.endpoint.sslContextParameters</code>	To configure security using <code>SSLContextParameters</code> . Important: Only one instance of <code>org.apache.camel.util.jsse.SSLContextParameters</code> is supported per <code>HttpComponent</code> . If you need to use 2 or more different instances, you need to define a new <code>HttpComponent</code> per instance you need.	null	MEDIUM
<code>camel.sink.endpoint.x509HostnameVerifier</code>	To use a custom <code>X509HostnameVerifier</code> such as <code>DefaultHostnameVerifier</code> or <code>NoopHostnameVerifier</code>	null	MEDIUM
<code>camel.component.http.cookieStore</code>	To use a custom <code>org.apache.http.client.CookieStore</code> . By default the <code>org.apache.http.impl.client.BasicCookieStore</code> is used which is an in-memory only cookie store. Notice if <code>bridgeEndpoint=true</code> then the cookie store is forced to be a noop cookie store as cookie shouldn't be stored as we are just bridging (eg acting as a proxy).	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.http.azyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.component.http.allowJavaSerializedObject</b>	Whether to allow java serialization when a request uses context-type=application/x-java-serialized-object. This is by default turned off. If you enable this then be aware that Java will deserialize the incoming data from the request to Java and that can be a potential security risk.	false	MEDIUM
<b>camel.component.http.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.component.http.clientConnectionManager</b>	To use a custom and shared HttpClientConnectionManager to manage connections. If this has been configured then this is always used for all endpoints created by this component.	null	MEDIUM
<b>camel.component.http.connectionsPerRoute</b>	The maximum number of connections per route.	20	MEDIUM
<b>camel.component.http.connectionTimeToLive</b>	The time for connection to live, the time unit is millisecond, the default value is always keep alive.	null	MEDIUM
<b>camel.component.http.httpBinding</b>	To use a custom HttpBinding to control the mapping between Camel message and HttpClient.	null	MEDIUM
<b>camel.component.http.httpClientConfigurer</b>	To use the custom HttpClientConfigurer to perform configuration of the HttpClient that will be used.	null	MEDIUM
<b>camel.component.http.httpConfiguration</b>	To use the shared HttpConfiguration as base configuration.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.http.httpContext</code>	To use a custom <code>org.apache.http.protocol.HttpContext</code> when executing requests.	null	MEDIUM
<code>camel.component.http.maxTotalConnections</code>	The maximum number of connections.	200	MEDIUM
<code>camel.component.http.headerFilterStrategy</code>	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message.	null	MEDIUM
<code>camel.component.http.proxyAuthDomain</code>	Proxy authentication domain to use	null	MEDIUM
<code>camel.component.http.proxyAuthHost</code>	Proxy authentication host	null	MEDIUM
<code>camel.component.http.proxyAuthMethod</code>	Proxy authentication method to use One of: [Basic] [Digest] [NTLM]	null	MEDIUM
<code>camel.component.http.proxyAuthNtHost</code>	Proxy authentication domain (workstation name) to use with NTLM	null	MEDIUM
<code>camel.component.http.proxyAuthPassword</code>	Proxy authentication password	null	MEDIUM
<code>camel.component.http.proxyAuthPort</code>	Proxy authentication port	null	MEDIUM
<code>camel.component.http.proxyAuthUsername</code>	Proxy authentication username	null	MEDIUM
<code>camel.component.http.sslContextParameters</code>	To configure security using <code>SSLContextParameters</code> . Important: Only one instance of <code>org.apache.camel.support.jsse.SSLContextParameters</code> is supported per <code>HttpComponent</code> . If you need to use 2 or more different instances, you need to define a new <code>HttpComponent</code> per instance you need.	null	MEDIUM
<code>camel.component.http.useGlobalSslContextParameters</code>	Enable usage of global SSL context parameters.	false	MEDIUM
<code>camel.component.http.x509HostnameVerifier</code>	To use a custom <code>X509HostnameVerifier</code> such as <code>DefaultHostnameVerifier</code> or <code>NoopHostnameVerifier</code> .	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.http.connectionRequestTimeout</b>	The timeout in milliseconds used when requesting a connection from the connection manager. A timeout value of zero is interpreted as an infinite timeout. A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	MEDIUM
<b>camel.component.http.connectTimeout</b>	Determines the timeout in milliseconds until a connection is established. A timeout value of zero is interpreted as an infinite timeout. A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	MEDIUM
<b>camel.component.http.socketTimeout</b>	Defines the socket timeout in milliseconds, which is the timeout for waiting for data or, put differently, a maximum period inactivity between two consecutive data packets). A timeout value of zero is interpreted as an infinite timeout. A negative value is interpreted as undefined (system default).	-1	MEDIUM

The camel-http sink connector has no converters out of the box.

The camel-http sink connector has no transforms out of the box.

The camel-http sink connector has no aggregation strategies out of the box.

## 5.14. CAMEL-JDBC-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-jdbc-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-jdbc-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.jdbc.CamelJdbcSinkConnector
```

The camel-jdbc sink connector supports 19 options, which are listed below.

Name	Description	Default	Priority
<b>camel.sink.path.dataSourceName</b>	Name of DataSource to lookup in the Registry. If the name is dataSource or default, then Camel will attempt to lookup a default DataSource from the registry, meaning if there is a only one instance of DataSource found, then this DataSource will be used.	null	HIGH
<b>camel.sink.endpoint.allowNamedParameters</b>	Whether to allow using named parameters in the queries.	true	MEDIUM
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.outputClass</b>	Specify the full package and class name to use as conversion when outputType=SelectOne or SelectList.	null	MEDIUM
<b>camel.sink.endpoint.outputType</b>	Determines the output the producer should use. One of: [SelectOne] [SelectList] [StreamList]	"SelectList"	MEDIUM
<b>camel.sink.endpoint.parameters</b>	Optional parameters to the java.sql.Statement. For example to set maxRows, fetchSize etc.	null	MEDIUM
<b>camel.sink.endpoint.readSize</b>	The default maximum number of rows that can be read by a polling query. The default value is 0.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.resetAutoCommit</b>	Camel will set the autoCommit on the JDBC connection to be false, commit the change after executed the statement and reset the autoCommit flag of the connection at the end, if the resetAutoCommit is true. If the JDBC connection doesn't support to reset the autoCommit flag, you can set the resetAutoCommit flag to be false, and Camel will not try to reset the autoCommit flag. When used with XA transactions you most likely need to set it to false so that the transaction manager is in charge of committing this tx.	true	MEDIUM
<b>camel.sink.endpoint.transactionacted</b>	Whether transactions are in use.	false	MEDIUM
<b>camel.sink.endpoint.useGetBytesForBlob</b>	To read BLOB columns as bytes instead of string data. This may be needed for certain databases such as Oracle where you must read BLOB columns as bytes.	false	MEDIUM
<b>camel.sink.endpoint.useHeadersAsParameters</b>	Set this option to true to use the prepareStatementStrategy with named parameters. This allows to define queries with named placeholders, and use headers with the dynamic values for the query placeholders.	false	MEDIUM
<b>camel.sink.endpoint.useJDBC4ColumnNameAndLabelSemantics</b>	Sets whether to use JDBC 4 or JDBC 3.0 or older semantic when retrieving column name. JDBC 4.0 uses columnName to get the column name where as JDBC 3.0 uses both columnName or columnLabel. Unfortunately JDBC drivers behave differently so you can use this option to work out issues around your JDBC driver if you get problem using this component This option is default true.	true	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.beanRowMapper</b>	To use a custom org.apache.camel.component.jdbc.BeanRowMapper when using outputClass. The default implementation will lower case the row names and skip underscores, and dashes. For example CUST_ID is mapped as custId.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.prepareStatementStrategy</code>	Allows the plugin to use a custom <code>org.apache.camel.component.jdbc.JdbcPreparedStatementStrategy</code> to control preparation of the query and prepared statement.	null	MEDIUM
<code>camel.sink.endpoint.synchronous</code>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<code>camel.component.jdbc.dataSource</code>	To use the <code>DataSource</code> instance instead of looking up the data source by name from the registry.	null	MEDIUM
<code>camel.component.jdbc.lazyStartProducer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow <code>CamelContext</code> and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<code>camel.component.jdbc.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

The camel-jdbc sink connector has no converters out of the box.

The camel-jdbc sink connector has no transforms out of the box.

The camel-jdbc sink connector has no aggregation strategies out of the box.

## 5.15. CAMEL-JMS-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-jms-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-jms-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.jms.CamelJmsSinkConnector
```

The camel-jms sink connector supports 143 options, which are listed below.

Name	Description	Default	Priority
<b>camel.sink.path.destinationType</b>	The kind of destination to use One of: [queue] [topic] [temp-queue] [temp-topic]	"queue"	MEDIUM
<b>camel.sink.path.destinationName</b>	Name of the queue or topic to use as destination	null	HIGH
<b>camel.sink.endpoint.clientId</b>	Sets the JMS client ID to use. Note that this value, if specified, must be unique and can only be used by a single JMS connection instance. It is typically only required for durable topic subscriptions. If using Apache ActiveMQ you may prefer to use Virtual Topics instead.	null	MEDIUM
<b>camel.sink.endpoint.connectionFactory</b>	The connection factory to be use. A connection factory must be configured either on the component or endpoint.	null	MEDIUM
<b>camel.sink.endpoint.disableReplyTo</b>	Specifies whether Camel ignores the JMSReplyTo header in messages. If true, Camel does not send a reply back to the destination specified in the JMSReplyTo header. You can use this option if you want Camel to consume from a route and you do not want Camel to automatically send back a reply message because another component in your code handles the reply message. You can also use this option if you want to use Camel as a proxy between different message brokers and you want to route message from one system to another.	false	MEDIUM
<b>camel.sink.endpoint.durableSubscriptionName</b>	The durable subscriber name for specifying durable topic subscriptions. The clientId option must be configured as well.	null	MEDIUM
<b>camel.sink.endpoint.jmsMessageType</b>	Allows you to force the use of a specific javax.jms.Message implementation for sending JMS messages. Possible values are: Bytes, Map, Object, Stream, Text. By default, Camel would determine which JMS message type to use from the In body type. This option allows you to specify it. One of: [Bytes] [Map] [Object] [Stream] [Text]	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.testConnectionOnStartup</b>	Specifies whether to test the connection on startup. This ensures that when Camel starts that all the JMS consumers have a valid connection to the JMS broker. If a connection cannot be granted then Camel throws an exception on startup. This ensures that Camel is not started with failed connections. The JMS producers is tested as well.	false	MEDIUM
<b>camel.sink.endpoint.deliveryDelay</b>	Sets delivery delay to use for send calls for JMS. This option requires JMS 2.0 compliant broker.	-1L	MEDIUM
<b>camel.sink.endpoint.deliveryMode</b>	Specifies the delivery mode to be used. Possibles values are those defined by javax.jms.DeliveryMode. NON_PERSISTENT = 1 and PERSISTENT = 2. One of: [1] [2]	null	MEDIUM
<b>camel.sink.endpoint.deliveryPersistent</b>	Specifies whether persistent delivery is used by default.	true	MEDIUM
<b>camel.sink.endpoint.explicitQosEnabled</b>	Set if the deliveryMode, priority or timeToLive qualities of service should be used when sending messages. This option is based on Spring's JmsTemplate. The deliveryMode, priority and timeToLive options are applied to the current endpoint. This contrasts with the preserveMessageQos option, which operates at message granularity, reading QoS properties exclusively from the Camel In message headers.	"false"	MEDIUM
<b>camel.sink.endpoint.formatDateHeadersToIso8601</b>	Sets whether JMS date properties should be formatted according to the ISO 8601 standard.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.preserveMessageQos</b>	Set to true, if you want to send message using the QoS settings specified on the message, instead of the QoS settings on the JMS endpoint. The following three headers are considered JMSPriority, JMSDeliveryMode, and JMSExpiration. You can provide all or only some of them. If not provided, Camel will fall back to use the values from the endpoint instead. So, when using this option, the headers override the values from the endpoint. The explicitQosEnabled option, by contrast, will only use options set on the endpoint, and not values from the message header.	false	MEDIUM
<b>camel.sink.endpoint.priority</b>	Values greater than 1 specify the message priority when sending (where 0 is the lowest priority and 9 is the highest). The explicitQosEnabled option must also be enabled in order for this option to have any effect. One of: [1] [2] [3] [4] [5] [6] [7] [8] [9]	4	MEDIUM
<b>camel.sink.endpoint.replyToConcurrentConsumers</b>	Specifies the default number of concurrent consumers when doing request/reply over JMS. See also the maxMessagesPerTask option to control dynamic scaling up/down of threads.	1	MEDIUM
<b>camel.sink.endpoint.replyToMaxConcurrentConsumers</b>	Specifies the maximum number of concurrent consumers when using request/reply over JMS. See also the maxMessagesPerTask option to control dynamic scaling up/down of threads.	null	MEDIUM
<b>camel.sink.endpoint.replyToOnTimeoutMaxConcurrent Consumers</b>	Specifies the maximum number of concurrent consumers for continue routing when timeout occurred when using request/reply over JMS.	1	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.replyToOverride</code>	Provides an explicit ReplyTo destination in the JMS message, which overrides the setting of replyTo. It is useful if you want to forward the message to a remote Queue and receive the reply message from the ReplyTo destination.	null	MEDIUM
<code>camel.sink.endpoint.replyToType</code>	Allows for explicitly specifying which kind of strategy to use for replyTo queues when doing request/reply over JMS. Possible values are: Temporary, Shared, or Exclusive. By default Camel will use temporary queues. However if replyTo has been configured, then Shared is used by default. This option allows you to use exclusive queues instead of shared ones. See Camel JMS documentation for more details, and especially the notes about the implications if running in a clustered environment, and the fact that Shared reply queues has lower performance than its alternatives Temporary and Exclusive. One of: [Temporary] [Shared] [Exclusive]	null	MEDIUM
<code>camel.sink.endpoint.requestTimeout</code>	The timeout for waiting for a reply when using the InOut Exchange Pattern (in milliseconds). The default is 20 seconds. You can include the header CamelJmsRequestTimeout to override this endpoint configured timeout value, and thus have per message individual timeout values. See also the requestTimeoutCheckerInterval option.	20000L	MEDIUM
<code>camel.sink.endpoint.timeToLive</code>	When sending messages, specifies the time-to-live of the message (in milliseconds).	-1L	MEDIUM
<code>camel.sink.endpoint.allowAdditionalHeaders</code>	This option is used to allow additional headers which may have values that are invalid according to JMS specification. For example some message systems such as WMQ do this with header names using prefix JMS_IBM_MQMD_ containing values with byte array or other invalid types. You can specify multiple header names separated by comma, and use as suffix for wildcard matching.	null	MEDIUM
<code>camel.sink.endpoint.allowNullBody</code>	Whether to allow sending messages with no body. If this option is false and the message body is null, then a JMSEException is thrown.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.alwaysCopyMessage</b>	If true, Camel will always make a JMS message copy of the message when it is passed to the producer for sending. Copying the message is needed in some situations, such as when a <code>replyToDestinationSelectorName</code> is set (incidentally, Camel will set the <code>alwaysCopyMessage</code> option to true, if a <code>replyToDestinationSelectorName</code> is set)	false	MEDIUM
<b>camel.sink.endpoint.correlationProperty</b>	When using InOut exchange pattern use this JMS property instead of <code>JMSCorrelationID</code> JMS property to correlate messages. If set messages will be correlated solely on the value of this property <code>JMSCorrelationID</code> property will be ignored and not set by Camel.	null	MEDIUM
<b>camel.sink.endpoint.disableTimeToLive</b>	Use this option to force disabling time to live. For example when you do request/reply over JMS, then Camel will by default use the <code>requestTimeout</code> value as time to live on the message being sent. The problem is that the sender and receiver systems have to have their clocks synchronized, so they are in sync. This is not always so easy to archive. So you can use <code>disableTimeToLive=true</code> to not set a time to live value on the sent message. Then the message will not expire on the receiver system. See below in section About time to live for more details.	false	MEDIUM
<b>camel.sink.endpoint.forceSendOriginalMessage</b>	When using <code>mapJmsMessage=false</code> Camel will create a new JMS message to send to a new JMS destination if you touch the headers (get or set) during the route. Set this option to true to force Camel to send the original JMS message that was received.	false	MEDIUM
<b>camel.sink.endpoint.includeSentJMSMessageID</b>	Only applicable when sending to JMS destination using InOnly (eg fire and forget). Enabling this option will enrich the Camel Exchange with the actual <code>JMSMessageID</code> that was used by the JMS client when the message was sent to the JMS destination.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.replyToCacheLevelName</b>	Sets the cache level by name for the reply consumer when doing request/reply over JMS. This option only applies when using fixed reply queues (not temporary). Camel will by default use: CACHE_CONSUMER for exclusive or shared w/ replyToSelectorName. And CACHE_SESSION for shared without replyToSelectorName. Some JMS brokers such as IBM WebSphere may require to set the replyToCacheLevelName=CACHE_NONE to work. Note: If using temporary queues then CACHE_NONE is not allowed, and you must use a higher value such as CACHE_CONSUMER or CACHE_SESSION. One of: [CACHE_AUTO] [CACHE_CONNECTION] [CACHE_CONSUMER] [CACHE_NONE] [CACHE_SESSION]	null	MEDIUM
<b>camel.sink.endpoint.replyToDestinationSelectorName</b>	Sets the JMS Selector using the fixed name to be used so you can filter out your own replies from the others when using a shared queue (that is, if you are not using a temporary reply queue).	null	MEDIUM
<b>camel.sink.endpoint.streamMessageTypeEnabled</b>	Sets whether StreamMessage type is enabled or not. Message payloads of streaming kind such as files, InputStream, etc will either be sent as BytesMessage or StreamMessage. This option controls which kind will be used. By default BytesMessage is used which enforces the entire message payload to be read into memory. By enabling this option the message payload is read into memory in chunks and each chunk is then written to the StreamMessage until no more data.	false	MEDIUM
<b>camel.sink.endpoint.allowSerializedHeaders</b>	Controls whether or not to include serialized headers. Applies only when transferExchange is true. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	MEDIUM
<b>camel.sink.endpoint.artemisStreamingEnabled</b>	Whether optimizing for Apache Artemis streaming mode.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.asyncStartListener</b>	Whether to startup the JmsConsumer message listener asynchronously, when starting a route. For example if a JmsConsumer cannot get a connection to a remote JMS broker, then it may block while retrying and/or failover. This will cause Camel to block while starting routes. By setting this option to true, you will let routes startup, while the JmsConsumer connects to the JMS broker using a dedicated thread in asynchronous mode. If this option is used, then beware that if the connection could not be established, then an exception is logged at WARN level, and the consumer will not be able to receive messages; You can then restart the route to retry.	false	MEDIUM
<b>camel.sink.endpoint.asyncStopListener</b>	Whether to stop the JmsConsumer message listener asynchronously, when stopping a route.	false	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.destinationResolver</b>	A pluggable <code>org.springframework.jms.support.destination.DestinationResolver</code> that allows you to use your own resolver (for example, to lookup the real destination in a JNDI registry).	null	MEDIUM
<b>camel.sink.endpoint.errorHandler</b>	Specifies a <code>org.springframework.util.ErrorHandler</code> to be invoked in case of any uncaught exceptions thrown while processing a Message. By default these exceptions will be logged at the WARN level, if no errorHandler has been configured. You can configure logging level and whether stack traces should be logged using <code>errorHandlerLoggingLevel</code> and <code>errorHandlerLogStackTrace</code> options. This makes it much easier to configure, than having to code a custom errorHandler.	null	MEDIUM
<b>camel.sink.endpoint.exceptionListener</b>	Specifies the JMS Exception Listener that is to be notified of any underlying JMS exceptions.	null	MEDIUM
<b>camel.sink.endpoint.headerFilterStrategy</b>	To use a custom <code>HeaderFilterStrategy</code> to filter header to and from Camel message.	null	MEDIUM
<b>camel.sink.endpoint.idleConsumerLimit</b>	Specify the limit for the number of consumers that are allowed to be idle at any given time.	1	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.idleTaskExecutionLimit</b>	Specifies the limit for idle executions of a receive task, not having received any message within its execution. If this limit is reached, the task will shut down and leave receiving to other executing tasks (in the case of dynamic scheduling; see the <code>maxConcurrentConsumers</code> setting). There is additional doc available from Spring.	1	MEDIUM
<b>camel.sink.endpoint.includeAllJMSXProperties</b>	Whether to include all JMSXxxx properties when mapping from JMS to Camel Message. Setting this to true will include properties such as <code>JMSXAppID</code> , and <code>JMSXUserID</code> etc. Note: If you are using a custom <code>headerFilterStrategy</code> then this option does not apply.	false	MEDIUM
<b>camel.sink.endpoint.jmsKeyFormatStrategy</b>	Pluggable strategy for encoding and decoding JMS keys so they can be compliant with the JMS specification. Camel provides two implementations out of the box: default and passthrough. The default strategy will safely marshal dots and hyphens (. and -). The passthrough strategy leaves the key as is. Can be used for JMS brokers which do not care whether JMS header keys contain illegal characters. You can provide your own implementation of the <code>org.apache.camel.component.jms.JmsKeyFormatStrategy</code> and refer to it using the # notation. One of: [default] [passthrough]	null	MEDIUM
<b>camel.sink.endpoint.mapJmsMessage</b>	Specifies whether Camel should auto map the received JMS message to a suited payload type, such as <code>javax.jms.TextMessage</code> to a String etc.	true	MEDIUM
<b>camel.sink.endpoint.maxMessagesPerTask</b>	The number of messages per task. -1 is unlimited. If you use a range for concurrent consumers (eg min max), then this option can be used to set a value to eg 100 to control how fast the consumers will shrink when less work is required.	-1	MEDIUM
<b>camel.sink.endpoint.messageConverter</b>	To use a custom Spring <code>org.springframework.jms.support.converter.MessageConverter</code> so you can be in control how to map to/from a <code>javax.jms.Message</code> .	null	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.messageCreatedStrategy</code>	To use the given MessageCreatedStrategy which are invoked when Camel creates new instances of javax.jms.Message objects when Camel is sending a JMS message.	null	MEDIUM
<code>camel.sink.endpoint.messageIdEnabled</code>	When sending, specifies whether message IDs should be added. This is just a hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the message ID set to null; if the provider ignores the hint, the message ID must be set to its normal unique value.	true	MEDIUM
<code>camel.sink.endpoint.messageListenerContainerFactory</code>	Registry ID of the MessageListenerContainerFactory used to determine what org.springframework.jms.listener.AbstractMessageListenerContainer to use to consume messages. Setting this will automatically set consumerType to Custom.	null	MEDIUM
<code>camel.sink.endpoint.messageTimestampEnabled</code>	Specifies whether timestamps should be enabled by default on sending messages. This is just a hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the timestamp set to zero; if the provider ignores the hint the timestamp must be set to its normal value.	true	MEDIUM
<code>camel.sink.endpoint.publishSubNoLocal</code>	Specifies whether to inhibit the delivery of messages published by its own connection.	false	MEDIUM
<code>camel.sink.endpoint.receiveTimeout</code>	The timeout for receiving messages (in milliseconds).	1000L	MEDIUM
<code>camel.sink.endpoint.recoveryInterval</code>	Specifies the interval between recovery attempts, i.e. when a connection is being refreshed, in milliseconds. The default is 5000 ms, that is, 5 seconds.	5000L	MEDIUM
<code>camel.sink.endpoint.requestTimeoutCheckerInterval</code>	Configures how often Camel should check for timed out Exchanges when doing request/reply over JMS. By default Camel checks once per second. But if you must react faster when a timeout occurs, then you can lower this interval, to check more frequently. The timeout is determined by the option requestTimeout.	1000L	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.sink.endpoint.transferException</b>	If enabled and you are using Request Reply messaging (InOut) and an Exchange failed on the consumer side, then the caused Exception will be send back in response as a <code>javax.jms.ObjectMessage</code> . If the client is Camel, the returned Exception is rethrown. This allows you to use Camel JMS as a bridge in your routing - for example, using persistent queues to enable robust routing. Notice that if you also have <code>transferExchange</code> enabled, this option takes precedence. The caught exception is required to be serializable. The original Exception on the consumer side can be wrapped in an outer exception such as <code>org.apache.camel.RuntimeCamelException</code> when returned to the producer. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer!	false	MEDIUM
<b>camel.sink.endpoint.transferExchange</b>	You can transfer the exchange over the wire instead of just the body and headers. The following fields are transferred: In body, Out body, Fault body, In headers, Out headers, Fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level. You must enable this option on both the producer and consumer side, so Camel knows the payloads is an Exchange and not a regular payload. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer having to use compatible Camel versions!	false	MEDIUM
<b>camel.sink.endpoint.useMessageIDAsCorrelationID</b>	Specifies whether <code>JMSMessageID</code> should always be used as <code>JMSCorrelationID</code> for InOut messages.	false	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.waitForProvisionCorrelationToBeUpdatedCounter</code>	Number of times to wait for provisional correlation id to be updated to the actual correlation id when doing request/reply over JMS and when the option <code>useMessageIDAsCorrelationID</code> is enabled.	50	MEDIUM
<code>camel.sink.endpoint.waitForProvisionCorrelationToBeUpdatedThreadSleepingTime</code>	Interval in millis to sleep each time while waiting for provisional correlation id to be updated.	100L	MEDIUM
<code>camel.sink.endpoint.password</code>	Password to use with the <code>ConnectionFactory</code> . You can also configure username/password directly on the <code>ConnectionFactory</code> .	null	MEDIUM
<code>camel.sink.endpoint.username</code>	Username to use with the <code>ConnectionFactory</code> . You can also configure username/password directly on the <code>ConnectionFactory</code> .	null	MEDIUM
<code>camel.sink.endpoint.transacted</code>	Specifies whether to use transacted mode	false	MEDIUM
<code>camel.sink.endpoint.transactedInOut</code>	Specifies whether InOut operations (request reply) default to using transacted mode. If this flag is set to true, then Spring <code>JmsTemplate</code> will have <code>sessionTransacted</code> set to true, and the <code>acknowledgeMode</code> as <code>transacted</code> on the <code>JmsTemplate</code> used for InOut operations. Note from Spring JMS: that within a JTA transaction, the parameters passed to <code>createQueue</code> , <code>createTopic</code> methods are not taken into account. Depending on the Java EE transaction context, the container makes its own decisions on these values. Analogously, these parameters are not taken into account within a locally managed transaction either, since Spring JMS operates on an existing JMS Session in this case. Setting this flag to true will use a short local JMS transaction when running outside of a managed transaction, and a synchronized local JMS transaction in case of a managed transaction (other than an XA transaction) being present. This has the effect of a local JMS transaction being managed alongside the main transaction (which might be a native JDBC transaction), with the JMS transaction committing right after the main transaction.	false	MEDIUM

Name	Description	Default	Priority
<code>camel.sink.endpoint.lazyCreateTransactionManager</code>	If true, Camel will create a <code>JmsTransactionManager</code> , if there is no <code>transactionManager</code> injected when option <code>transacted=true</code> .	true	MEDIUM
<code>camel.sink.endpoint.transactionManager</code>	The Spring transaction manager to use.	null	MEDIUM
<code>camel.sink.endpoint.transactionName</code>	The name of the transaction to use.	null	MEDIUM
<code>camel.sink.endpoint.transactionTimeout</code>	The timeout value of the transaction (in seconds), if using <code>transacted</code> mode.	-1	MEDIUM
<code>camel.component.jms.clientId</code>	Sets the JMS client ID to use. Note that this value, if specified, must be unique and can only be used by a single JMS connection instance. It is typically only required for durable topic subscriptions. If using Apache ActiveMQ you may prefer to use Virtual Topics instead.	null	MEDIUM
<code>camel.component.jms.connectionFactory</code>	The connection factory to be use. A connection factory must be configured either on the component or endpoint.	null	MEDIUM
<code>camel.component.jms.disableReplyTo</code>	Specifies whether Camel ignores the <code>JMSReplyTo</code> header in messages. If true, Camel does not send a reply back to the destination specified in the <code>JMSReplyTo</code> header. You can use this option if you want Camel to consume from a route and you do not want Camel to automatically send back a reply message because another component in your code handles the reply message. You can also use this option if you want to use Camel as a proxy between different message brokers and you want to route message from one system to another.	false	MEDIUM
<code>camel.component.jms.durableSubscriptionName</code>	The durable subscriber name for specifying durable topic subscriptions. The <code>clientId</code> option must be configured as well.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.jmsMessageType</b>	Allows you to force the use of a specific <code>javax.jms.Message</code> implementation for sending JMS messages. Possible values are: Bytes, Map, Object, Stream, Text. By default, Camel would determine which JMS message type to use from the In body type. This option allows you to specify it. One of: [Bytes] [Map] [Object] [Stream] [Text]	null	MEDIUM
<b>camel.component.jms.testConnectionOnStartup</b>	Specifies whether to test the connection on startup. This ensures that when Camel starts that all the JMS consumers have a valid connection to the JMS broker. If a connection cannot be granted then Camel throws an exception on startup. This ensures that Camel is not started with failed connections. The JMS producers is tested as well.	false	MEDIUM
<b>camel.component.jms.deliveryDelay</b>	Sets delivery delay to use for send calls for JMS. This option requires JMS 2.0 compliant broker.	-1L	MEDIUM
<b>camel.component.jms.deliveryMode</b>	Specifies the delivery mode to be used. Possibles values are those defined by <code>javax.jms.DeliveryMode</code> . <code>NON_PERSISTENT = 1</code> and <code>PERSISTENT = 2</code> . One of: [1] [2]	null	MEDIUM
<b>camel.component.jms.deliveryPersistent</b>	Specifies whether persistent delivery is used by default.	true	MEDIUM
<b>camel.component.jms.explicitQosEnabled</b>	Set if the <code>deliveryMode</code> , <code>priority</code> or <code>timeToLive</code> qualities of service should be used when sending messages. This option is based on Spring's <code>JmsTemplate</code> . The <code>deliveryMode</code> , <code>priority</code> and <code>timeToLive</code> options are applied to the current endpoint. This contrasts with the <code>preserveMessageQos</code> option, which operates at message granularity, reading QoS properties exclusively from the Camel In message headers.	"false"	MEDIUM
<b>camel.component.jms.formatDateHeadersToIso8601</b>	Sets whether JMS date properties should be formatted according to the ISO 8601 standard.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.component.jms.reserveMessageQos</b>	Set to true, if you want to send message using the QoS settings specified on the message, instead of the QoS settings on the JMS endpoint. The following three headers are considered JMSPriority, JMSDeliveryMode, and JMSExpiration. You can provide all or only some of them. If not provided, Camel will fall back to use the values from the endpoint instead. So, when using this option, the headers override the values from the endpoint. The explicitQosEnabled option, by contrast, will only use options set on the endpoint, and not values from the message header.	false	MEDIUM
<b>camel.component.jms.priority</b>	Values greater than 1 specify the message priority when sending (where 0 is the lowest priority and 9 is the highest). The explicitQosEnabled option must also be enabled in order for this option to have any effect. One of: [1] [2] [3] [4] [5] [6] [7] [8] [9]	4	MEDIUM
<b>camel.component.jms.replyToConcurrentConsumers</b>	Specifies the default number of concurrent consumers when doing request/reply over JMS. See also the maxMessagesPerTask option to control dynamic scaling up/down of threads.	1	MEDIUM
<b>camel.component.jms.replyToMaxConcurrentConsumers</b>	Specifies the maximum number of concurrent consumers when using request/reply over JMS. See also the maxMessagesPerTask option to control dynamic scaling up/down of threads.	null	MEDIUM
<b>camel.component.jms.replyToOnTimeoutMaxConcurrentConsumers</b>	Specifies the maximum number of concurrent consumers for continue routing when timeout occurred when using request/reply over JMS.	1	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.replyToOverride</b>	Provides an explicit ReplyTo destination in the JMS message, which overrides the setting of replyTo. It is useful if you want to forward the message to a remote Queue and receive the reply message from the ReplyTo destination.	null	MEDIUM
<b>camel.component.jms.replyToType</b>	Allows for explicitly specifying which kind of strategy to use for replyTo queues when doing request/reply over JMS. Possible values are: Temporary, Shared, or Exclusive. By default Camel will use temporary queues. However if replyTo has been configured, then Shared is used by default. This option allows you to use exclusive queues instead of shared ones. See Camel JMS documentation for more details, and especially the notes about the implications if running in a clustered environment, and the fact that Shared reply queues has lower performance than its alternatives Temporary and Exclusive. One of: [Temporary] [Shared] [Exclusive]	null	MEDIUM
<b>camel.component.jms.requestTimeout</b>	The timeout for waiting for a reply when using the InOut Exchange Pattern (in milliseconds). The default is 20 seconds. You can include the header CamelJmsRequestTimeout to override this endpoint configured timeout value, and thus have per message individual timeout values. See also the requestTimeoutCheckerInterval option.	20000L	MEDIUM
<b>camel.component.jms.timeToLive</b>	When sending messages, specifies the time-to-live of the message (in milliseconds).	-1L	MEDIUM
<b>camel.component.jms.allowAdditionalHeaders</b>	This option is used to allow additional headers which may have values that are invalid according to JMS specification. For example some message systems such as WMQ do this with header names using prefix JMS_IBM_MQMD_ containing values with byte array or other invalid types. You can specify multiple header names separated by comma, and use as suffix for wildcard matching.	null	MEDIUM
<b>camel.component.jms.allowNullBody</b>	Whether to allow sending messages with no body. If this option is false and the message body is null, then an JMSEException is thrown.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.alwaysCopyMessage</b>	If true, Camel will always make a JMS message copy of the message when it is passed to the producer for sending. Copying the message is needed in some situations, such as when a <code>replyToDestinationSelectorName</code> is set (incidentally, Camel will set the <code>alwaysCopyMessage</code> option to true, if a <code>replyToDestinationSelectorName</code> is set)	false	MEDIUM
<b>camel.component.jms.correlationProperty</b>	When using InOut exchange pattern use this JMS property instead of <code>JMSCorrelationID</code> JMS property to correlate messages. If set messages will be correlated solely on the value of this property <code>JMSCorrelationID</code> property will be ignored and not set by Camel.	null	MEDIUM
<b>camel.component.jms.disableTimeToLive</b>	Use this option to force disabling time to live. For example when you do request/reply over JMS, then Camel will by default use the <code>requestTimeout</code> value as time to live on the message being sent. The problem is that the sender and receiver systems have to have their clocks synchronized, so they are in sync. This is not always so easy to archive. So you can use <code>disableTimeToLive=true</code> to not set a time to live value on the sent message. Then the message will not expire on the receiver system. See below in section About time to live for more details.	false	MEDIUM
<b>camel.component.jms.forceSendOriginalMessage</b>	When using <code>mapJmsMessage=false</code> Camel will create a new JMS message to send to a new JMS destination if you touch the headers (get or set) during the route. Set this option to true to force Camel to send the original JMS message that was received.	false	MEDIUM
<b>camel.component.jms.includeSentJMSMessageID</b>	Only applicable when sending to JMS destination using InOnly (eg fire and forget). Enabling this option will enrich the Camel Exchange with the actual <code>JMSMessageID</code> that was used by the JMS client when the message was sent to the JMS destination.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.replyToCacheLevelName</b>	Sets the cache level by name for the reply consumer when doing request/reply over JMS. This option only applies when using fixed reply queues (not temporary). Camel will by default use: CACHE_CONSUMER for exclusive or shared w/ replyToSelectorName. And CACHE_SESSION for shared without replyToSelectorName. Some JMS brokers such as IBM WebSphere may require to set the replyToCacheLevelName=CACHE_NONE to work. Note: If using temporary queues then CACHE_NONE is not allowed, and you must use a higher value such as CACHE_CONSUMER or CACHE_SESSION. One of: [CACHE_AUTO] [CACHE_CONNECTION] [CACHE_CONSUMER] [CACHE_NONE] [CACHE_SESSION]	null	MEDIUM
<b>camel.component.jms.replyToDestinationSelectorName</b>	Sets the JMS Selector using the fixed name to be used so you can filter out your own replies from the others when using a shared queue (that is, if you are not using a temporary reply queue).	null	MEDIUM
<b>camel.component.jms.streamMessageTypeEnabled</b>	Sets whether StreamMessage type is enabled or not. Message payloads of streaming kind such as files, InputStream, etc will either be sent as BytesMessage or StreamMessage. This option controls which kind will be used. By default BytesMessage is used which enforces the entire message payload to be read into memory. By enabling this option the message payload is read into memory in chunks and each chunk is then written to the StreamMessage until no more data.	false	MEDIUM
<b>camel.component.jms.allowAutoWiredConnectionFactory</b>	Whether to auto-discover ConnectionFactory from the registry, if no connection factory has been configured. If only one instance of ConnectionFactory is found then it will be used. This is enabled by default.	true	MEDIUM
<b>camel.component.jms.allowAutoWiredDestinationResolver</b>	Whether to auto-discover DestinationResolver from the registry, if no destination resolver has been configured. If only one instance of DestinationResolver is found then it will be used. This is enabled by default.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.allowSerializedHeaders</b>	Controls whether or not to include serialized headers. Applies only when transferExchange is true. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	MEDIUM
<b>camel.component.jms.artemisStreamingEnabled</b>	Whether optimizing for Apache Artemis streaming mode.	true	MEDIUM
<b>camel.component.jms.asyncStartListener</b>	Whether to startup the JmsConsumer message listener asynchronously, when starting a route. For example if a JmsConsumer cannot get a connection to a remote JMS broker, then it may block while retrying and/or failover. This will cause Camel to block while starting routes. By setting this option to true, you will let routes startup, while the JmsConsumer connects to the JMS broker using a dedicated thread in asynchronous mode. If this option is used, then beware that if the connection could not be established, then an exception is logged at WARN level, and the consumer will not be able to receive messages; You can then restart the route to retry.	false	MEDIUM
<b>camel.component.jms.asyncStopListener</b>	Whether to stop the JmsConsumer message listener asynchronously, when stopping a route.	false	MEDIUM
<b>camel.component.jms.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.component.jms.configuration</b>	To use a shared JMS configuration	null	MEDIUM
<b>camel.component.jms.destinationResolver</b>	A pluggable <code>org.springframework.jms.support.destination.DestinationResolver</code> that allows you to use your own resolver (for example, to lookup the real destination in a JNDI registry).	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.errorHandler</b>	Specifies a <code>org.springframework.util.ErrorHandler</code> to be invoked in case of any uncaught exceptions thrown while processing a Message. By default these exceptions will be logged at the WARN level, if no errorHandler has been configured. You can configure logging level and whether stack traces should be logged using <code>errorHandlerLoggingLevel</code> and <code>errorHandlerLogStackTrace</code> options. This makes it much easier to configure, than having to code a custom errorHandler.	null	MEDIUM
<b>camel.component.jms.exceptionListener</b>	Specifies the JMS Exception Listener that is to be notified of any underlying JMS exceptions.	null	MEDIUM
<b>camel.component.jms.idleConsumerLimit</b>	Specify the limit for the number of consumers that are allowed to be idle at any given time.	1	MEDIUM
<b>camel.component.jms.idleTaskExecutionLimit</b>	Specifies the limit for idle executions of a receive task, not having received any message within its execution. If this limit is reached, the task will shut down and leave receiving to other executing tasks (in the case of dynamic scheduling; see the <code>maxConcurrentConsumers</code> setting). There is additional doc available from Spring.	1	MEDIUM
<b>camel.component.jms.includeAllJMSXProperties</b>	Whether to include all JMSxxx properties when mapping from JMS to Camel Message. Setting this to true will include properties such as <code>JMSXAppID</code> , and <code>JMSXUserID</code> etc. Note: If you are using a custom <code>headerFilterStrategy</code> then this option does not apply.	false	MEDIUM
<b>camel.component.jms.jmsKeyFormatStrategy</b>	Pluggable strategy for encoding and decoding JMS keys so they can be compliant with the JMS specification. Camel provides two implementations out of the box: default and passthrough. The default strategy will safely marshal dots and hyphens (. and -). The passthrough strategy leaves the key as is. Can be used for JMS brokers which do not care whether JMS header keys contain illegal characters. You can provide your own implementation of the <code>org.apache.camel.component.jms.JmsKeyFormatStrategy</code> and refer to it using the # notation. One of: [default] [passthrough]	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.mapJmsMessage</b>	Specifies whether Camel should auto map the received JMS message to a suited payload type, such as <code>javax.jms.TextMessage</code> to a <code>String</code> etc.	true	MEDIUM
<b>camel.component.jms.maxMessagesPerTask</b>	The number of messages per task. -1 is unlimited. If you use a range for concurrent consumers (eg min max), then this option can be used to set a value to eg 100 to control how fast the consumers will shrink when less work is required.	-1	MEDIUM
<b>camel.component.jms.messageConverter</b>	To use a custom Spring <code>org.springframework.jms.support.converter.MessageConverter</code> so you can be in control how to map to/from a <code>javax.jms.Message</code> .	null	MEDIUM
<b>camel.component.jms.messageCreatedStrategy</b>	To use the given <code>MessageCreatedStrategy</code> which are invoked when Camel creates new instances of <code>javax.jms.Message</code> objects when Camel is sending a JMS message.	null	MEDIUM
<b>camel.component.jms.messageIdEnabled</b>	When sending, specifies whether message IDs should be added. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the message ID set to null; if the provider ignores the hint, the message ID must be set to its normal unique value.	true	MEDIUM
<b>camel.component.jms.messageListenerContainerFactory</b>	Registry ID of the <code>MessageListenerContainerFactory</code> used to determine what <code>org.springframework.jms.listener.AbstractMessageListenerContainer</code> to use to consume messages. Setting this will automatically set <code>consumerType</code> to <code>Custom</code> .	null	MEDIUM
<b>camel.component.jms.messageTimestampEnabled</b>	Specifies whether timestamps should be enabled by default on sending messages. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the timestamp set to zero; if the provider ignores the hint the timestamp must be set to its normal value.	true	MEDIUM
<b>camel.component.jms.pubSubNoLocal</b>	Specifies whether to inhibit the delivery of messages published by its own connection.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.queueBrowseStrategy</b>	To use a custom QueueBrowseStrategy when browsing queues	null	MEDIUM
<b>camel.component.jms.receiveTimeout</b>	The timeout for receiving messages (in milliseconds).	1000L	MEDIUM
<b>camel.component.jms.recoveryInterval</b>	Specifies the interval between recovery attempts, i.e. when a connection is being refreshed, in milliseconds. The default is 5000 ms, that is, 5 seconds.	5000L	MEDIUM
<b>camel.component.jms.requestTimeoutCheckerInterval</b>	Configures how often Camel should check for timed out Exchanges when doing request/reply over JMS. By default Camel checks once per second. But if you must react faster when a timeout occurs, then you can lower this interval, to check more frequently. The timeout is determined by the option requestTimeout.	1000L	MEDIUM
<b>camel.component.jms.transferException</b>	If enabled and you are using Request Reply messaging (InOut) and an Exchange failed on the consumer side, then the caused Exception will be send back in response as a <code>javax.jms.ObjectMessage</code> . If the client is Camel, the returned Exception is rethrown. This allows you to use Camel JMS as a bridge in your routing - for example, using persistent queues to enable robust routing. Notice that if you also have <code>transferExchange</code> enabled, this option takes precedence. The caught exception is required to be serializable. The original Exception on the consumer side can be wrapped in an outer exception such as <code>org.apache.camel.RuntimeCamelException</code> when returned to the producer. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer!	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.transferExchange</b>	You can transfer the exchange over the wire instead of just the body and headers. The following fields are transferred: In body, Out body, Fault body, In headers, Out headers, Fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level. You must enable this option on both the producer and consumer side, so Camel knows the payload is an Exchange and not a regular payload. Use this with caution as the data is using Java Object serialization and requires the receiver to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer having to use compatible Camel versions!	false	MEDIUM
<b>camel.component.jms.useMessageIDAsCorrelationID</b>	Specifies whether JMSMessageID should always be used as JMSCorrelationID for InOut messages.	false	MEDIUM
<b>camel.component.jms.waitForProvisionCorrelationToBeUpdatedCounter</b>	Number of times to wait for provisional correlation id to be updated to the actual correlation id when doing request/reply over JMS and when the option useMessageIDAsCorrelationID is enabled.	50	MEDIUM
<b>camel.component.jms.waitForProvisionCorrelationToBeUpdatedThreadSleepingTime</b>	Interval in millis to sleep each time while waiting for provisional correlation id to be updated.	100L	MEDIUM
<b>camel.component.jms.headerFilterStrategy</b>	To use a custom org.apache.camel.spi.HeaderFilterStrategy to filter header to and from Camel message.	null	MEDIUM
<b>camel.component.jms.password</b>	Password to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.	null	MEDIUM
<b>camel.component.jms.username</b>	Username to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.	null	MEDIUM
<b>camel.component.jms.transacted</b>	Specifies whether to use transacted mode	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.transactedInOut</b>	Specifies whether InOut operations (request reply) default to using transacted mode. If this flag is set to true, then Spring JmsTemplate will have sessionTransacted set to true, and the acknowledgeMode as transacted on the JmsTemplate used for InOut operations. Note from Spring JMS: that within a JTA transaction, the parameters passed to createQueue, createTopic methods are not taken into account. Depending on the Java EE transaction context, the container makes its own decisions on these values. Analogously, these parameters are not taken into account within a locally managed transaction either, since Spring JMS operates on an existing JMS Session in this case. Setting this flag to true will use a short local JMS transaction when running outside of a managed transaction, and a synchronized local JMS transaction in case of a managed transaction (other than an XA transaction) being present. This has the effect of a local JMS transaction being managed alongside the main transaction (which might be a native JDBC transaction), with the JMS transaction committing right after the main transaction.	false	MEDIUM
<b>camel.component.jms.lazyCreateTransactionManager</b>	If true, Camel will create a JmsTransactionManager, if there is no transactionManager injected when option transacted=true.	true	MEDIUM
<b>camel.component.jms.transactionManager</b>	The Spring transaction manager to use.	null	MEDIUM
<b>camel.component.jms.transactionName</b>	The name of the transaction to use.	null	MEDIUM
<b>camel.component.jms.transactionTimeout</b>	The timeout value of the transaction (in seconds), if using transacted mode.	-1	MEDIUM

The camel-jms sink connector has no converters out of the box.

The camel-jms sink connector has no transforms out of the box.

The camel-jms sink connector has no aggregation strategies out of the box.

## 5.16. CAMEL-JMS-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-jms-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-jms-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.jms.CamelJmsSourceConnector
```

The camel-jms source connector supports 143 options, which are listed below.

Name	Description	Default	Priority
<b>camel.source.path.destinationType</b>	The kind of destination to use One of: [queue] [topic] [temp-queue] [temp-topic]	"queue"	MEDIUM
<b>camel.source.path.destinationName</b>	Name of the queue or topic to use as destination	null	HIGH
<b>camel.source.endpoint.clientId</b>	Sets the JMS client ID to use. Note that this value, if specified, must be unique and can only be used by a single JMS connection instance. It is typically only required for durable topic subscriptions. If using Apache ActiveMQ you may prefer to use Virtual Topics instead.	null	MEDIUM
<b>camel.source.endpoint.connectionFactory</b>	The connection factory to be use. A connection factory must be configured either on the component or endpoint.	null	MEDIUM
<b>camel.source.endpoint.disableReplyTo</b>	Specifies whether Camel ignores the JMSReplyTo header in messages. If true, Camel does not send a reply back to the destination specified in the JMSReplyTo header. You can use this option if you want Camel to consume from a route and you do not want Camel to automatically send back a reply message because another component in your code handles the reply message. You can also use this option if you want to use Camel as a proxy between different message brokers and you want to route message from one system to another.	false	MEDIUM
<b>camel.source.endpoint.durableSubscriptionName</b>	The durable subscriber name for specifying durable topic subscriptions. The clientId option must be configured as well.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.jmsMessageType</b>	Allows you to force the use of a specific javax.jms.Message implementation for sending JMS messages. Possible values are: Bytes, Map, Object, Stream, Text. By default, Camel would determine which JMS message type to use from the In body type. This option allows you to specify it. One of: [Bytes] [Map] [Object] [Stream] [Text]	null	MEDIUM
<b>camel.source.endpoint.testConnectionOnStartup</b>	Specifies whether to test the connection on startup. This ensures that when Camel starts that all the JMS consumers have a valid connection to the JMS broker. If a connection cannot be granted then Camel throws an exception on startup. This ensures that Camel is not started with failed connections. The JMS producers is tested as well.	false	MEDIUM
<b>camel.source.endpoint.acknowledgementModeName</b>	The JMS acknowledgement name, which is one of: SESSION_TRANSACTED, CLIENT_ACKNOWLEDGE, AUTO_ACKNOWLEDGE, DUPS_OK_ACKNOWLEDGE One of: [SESSION_TRANSACTED] [CLIENT_ACKNOWLEDGE] [AUTO_ACKNOWLEDGE] [DUPS_OK_ACKNOWLEDGE]	"AUTO_ACKNOWLEDGE"	MEDIUM
<b>camel.source.endpoint.asyncConsumer</b>	Whether the JmsConsumer processes the Exchange asynchronously. If enabled then the JmsConsumer may pickup the next message from the JMS queue, while the previous message is being processed asynchronously (by the Asynchronous Routing Engine). This means that messages may be processed not 100% strictly in order. If disabled (as default) then the Exchange is fully processed before the JmsConsumer will pickup the next message from the JMS queue. Note if transacted has been enabled, then asyncConsumer=true does not run asynchronously, as transaction must be executed synchronously (Camel 3.0 may support async transactions).	false	MEDIUM
<b>camel.source.endpoint.autoStartup</b>	Specifies whether the consumer container should auto-startup.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.cacheLevel</b>	Sets the cache level by ID for the underlying JMS resources. See <code>cacheLevelName</code> option for more details.	null	MEDIUM
<b>camel.source.endpoint.cacheLevelName</b>	Sets the cache level by name for the underlying JMS resources. Possible values are: <code>CACHE_AUTO</code> , <code>CACHE_CONNECTION</code> , <code>CACHE_CONSUMER</code> , <code>CACHE_NONE</code> , and <code>CACHE_SESSION</code> . The default setting is <code>CACHE_AUTO</code> . See the Spring documentation and Transactions Cache Levels for more information. One of: <code>[CACHE_AUTO]</code> <code>[CACHE_CONNECTION]</code> <code>[CACHE_CONSUMER]</code> <code>[CACHE_NONE]</code> <code>[CACHE_SESSION]</code>	"CACHE_AUTO"	MEDIUM
<b>camel.source.endpoint.concurrentConsumers</b>	Specifies the default number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.	1	MEDIUM
<b>camel.source.endpoint.maxConcurrentConsumers</b>	Specifies the maximum number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToMaxConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.	null	MEDIUM
<b>camel.source.endpoint.replyTo</b>	Provides an explicit <code>ReplyTo</code> destination, which overrides any incoming value of <code>Message.getJMSReplyTo()</code> .	null	MEDIUM
<b>camel.source.endpoint.replyToDeliveryPersistent</b>	Specifies whether to use persistent delivery by default for replies.	true	MEDIUM
<b>camel.source.endpoint.selector</b>	Sets the JMS selector to use	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.subscriptionDurable</b>	Set whether to make the subscription durable. The durable subscription name to be used can be specified through the <code>subscriptionName</code> property. Default is false. Set this to true to register a durable subscription, typically in combination with a <code>subscriptionName</code> value (unless your message listener class name is good enough as subscription name). Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the <code>pubSubDomain</code> flag as well.	false	MEDIUM
<b>camel.source.endpoint.subscriptionName</b>	Set the name of a subscription to create. To be applied in case of a topic (pub-sub domain) with a shared or durable subscription. The subscription name needs to be unique within this client's JMS client id. Default is the class name of the specified message listener. Note: Only 1 concurrent consumer (which is the default of this message listener container) is allowed for each subscription, except for a shared subscription (which requires JMS 2.0).	null	MEDIUM
<b>camel.source.endpoint.subscriptionShared</b>	Set whether to make the subscription shared. The shared subscription name to be used can be specified through the <code>subscriptionName</code> property. Default is false. Set this to true to register a shared subscription, typically in combination with a <code>subscriptionName</code> value (unless your message listener class name is good enough as subscription name). Note that shared subscriptions may also be durable, so this flag can (and often will) be combined with <code>subscriptionDurable</code> as well. Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the <code>pubSubDomain</code> flag as well. Requires a JMS 2.0 compatible message broker.	false	MEDIUM
<b>camel.source.endpoint.acceptMessagesWhileStopping</b>	Specifies whether the consumer accept messages while it is stopping. You may consider enabling this option, if you start and stop JMS routes at runtime, while there are still messages enqueued on the queue. If this option is false, and you stop the JMS route, then messages may be rejected, and the JMS broker would have to attempt redeliveries, which yet again may be rejected, and eventually the message may be moved at a dead letter queue on the JMS broker. To avoid this its recommended to enable this option.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.allowReplyManagerQuickStop</b>	Whether the DefaultMessageListenerContainer used in the reply managers for request-reply messaging allow the DefaultMessageListenerContainer.runningAllowed flag to quick stop in case JmsConfiguration#isAcceptMessagesWhileStopping is enabled, and org.apache.camel.CamelContext is currently being stopped. This quick stop ability is enabled by default in the regular JMS consumers but to enable for reply managers you must enable this flag.	false	MEDIUM
<b>camel.source.endpoint.consumerType</b>	The consumer type to use, which can be one of: Simple, Default, or Custom. The consumer type determines which Spring JMS listener to use. Default will use org.springframework.jms.listener.DefaultMessageListenerContainer, Simple will use org.springframework.jms.listener.SimpleMessageListenerContainer. When Custom is specified, the MessageListenerContainerFactory defined by the messageListenerContainerFactory option will determine what org.springframework.jms.listener.AbstractMessageListenerContainer to use. One of: [Simple] [Default] [Custom]	"Default"	MEDIUM
<b>camel.source.endpoint.defaultTaskExecutorType</b>	Specifies what default TaskExecutor type to use in the DefaultMessageListenerContainer, for both consumer endpoints and the ReplyTo consumer of producer endpoints. Possible values: SimpleAsync (uses Spring's SimpleAsyncTaskExecutor) or ThreadPool (uses Spring's ThreadPoolTaskExecutor with optimal values - cached threadpool-like). If not set, it defaults to the previous behaviour, which uses a cached thread pool for consumer endpoints and SimpleAsync for reply consumers. The use of ThreadPool is recommended to reduce thread trash in elastic configurations with dynamically increasing and decreasing concurrent consumers. One of: [ThreadPool] [SimpleAsync]	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.eagerLoadingOfProperties</b>	Enables eager loading of JMS properties and payload as soon as a message is loaded which generally is inefficient as the JMS properties may not be required but sometimes can catch early any issues with the underlying JMS provider and the use of JMS properties. See also the option <code>eagerPoisonBody</code> .	false	MEDIUM
<b>camel.source.endpoint.eagerPoisonBody</b>	If <code>eagerLoadingOfProperties</code> is enabled and the JMS message payload (JMS body or JMS properties) is poison (cannot be read/mapped), then set this text as the message body instead so the message can be processed (the cause of the poison are already stored as exception on the Exchange). This can be turned off by setting <code>eagerPoisonBody=false</code> . See also the option <code>eagerLoadingOfProperties</code> .	"Poison JMS message due to <code>{exception.message}</code> "	MEDIUM
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: <code>[InOnly]</code> <code>[InOut]</code> <code>[InOptionalOut]</code>	null	MEDIUM
<b>camel.source.endpoint.exposeListenerSession</b>	Specifies whether the listener session should be exposed when consuming messages.	false	MEDIUM
<b>camel.source.endpoint.replyToSameDestination Allowed</b>	Whether a JMS consumer is allowed to send a reply message to the same destination that the consumer is using to consume from. This prevents an endless loop by consuming and sending back the same message to itself.	false	MEDIUM
<b>camel.source.endpoint.taskExecutor</b>	Allows you to specify a custom task executor for consuming messages.	null	MEDIUM
<b>camel.source.endpoint.allowSerializedHeaders</b>	Controls whether or not to include serialized headers. Applies only when <code>transferExchange</code> is true. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.artemisStreamingEnabled</code>	Whether optimizing for Apache Artemis streaming mode.	true	MEDIUM
<code>camel.source.endpoint.asyncStartListener</code>	Whether to startup the <code>JmsConsumer</code> message listener asynchronously, when starting a route. For example if a <code>JmsConsumer</code> cannot get a connection to a remote JMS broker, then it may block while retrying and/or failover. This will cause Camel to block while starting routes. By setting this option to true, you will let routes startup, while the <code>JmsConsumer</code> connects to the JMS broker using a dedicated thread in asynchronous mode. If this option is used, then beware that if the connection could not be established, then an exception is logged at WARN level, and the consumer will not be able to receive messages; You can then restart the route to retry.	false	MEDIUM
<code>camel.source.endpoint.asyncStopListener</code>	Whether to stop the <code>JmsConsumer</code> message listener asynchronously, when stopping a route.	false	MEDIUM
<code>camel.source.endpoint.basicPropertyBinding</code>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.source.endpoint.destinationResolver</code>	A pluggable <code>org.springframework.jms.support.destination.DestinationResolver</code> that allows you to use your own resolver (for example, to lookup the real destination in a JNDI registry).	null	MEDIUM
<code>camel.source.endpoint.errorHandler</code>	Specifies a <code>org.springframework.util.ErrorHandler</code> to be invoked in case of any uncaught exceptions thrown while processing a <code>Message</code> . By default these exceptions will be logged at the WARN level, if no <code>errorHandler</code> has been configured. You can configure logging level and whether stack traces should be logged using <code>errorHandlerLoggingLevel</code> and <code>errorHandlerLogStackTrace</code> options. This makes it much easier to configure, than having to code a custom <code>errorHandler</code> .	null	MEDIUM
<code>camel.source.endpoint.exceptionListener</code>	Specifies the JMS Exception Listener that is to be notified of any underlying JMS exceptions.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.headerFilterStrategy</b>	To use a custom HeaderFilterStrategy to filter header to and from Camel message.	null	MEDIUM
<b>camel.source.endpoint.idleConsumerLimit</b>	Specify the limit for the number of consumers that are allowed to be idle at any given time.	1	MEDIUM
<b>camel.source.endpoint.idleTaskExecutionLimit</b>	Specifies the limit for idle executions of a receive task, not having received any message within its execution. If this limit is reached, the task will shut down and leave receiving to other executing tasks (in the case of dynamic scheduling; see the maxConcurrentConsumers setting). There is additional doc available from Spring.	1	MEDIUM
<b>camel.source.endpoint.includeAllJMSXProperties</b>	Whether to include all JMSxxx properties when mapping from JMS to Camel Message. Setting this to true will include properties such as JMSXAppID, and JMSXUserID etc. Note: If you are using a custom headerFilterStrategy then this option does not apply.	false	MEDIUM
<b>camel.source.endpoint.jmsKeyFormatStrategy</b>	Pluggable strategy for encoding and decoding JMS keys so they can be compliant with the JMS specification. Camel provides two implementations out of the box: default and passthrough. The default strategy will safely marshal dots and hyphens (. and -). The passthrough strategy leaves the key as is. Can be used for JMS brokers which do not care whether JMS header keys contain illegal characters. You can provide your own implementation of the org.apache.camel.component.jms.JmsKeyFormatStrategy and refer to it using the # notation. One of: [default] [passthrough]	null	MEDIUM
<b>camel.source.endpoint.mapJmsMessage</b>	Specifies whether Camel should auto map the received JMS message to a suited payload type, such as javax.jms.TextMessage to a String etc.	true	MEDIUM
<b>camel.source.endpoint.maxMessagesPerTask</b>	The number of messages per task. -1 is unlimited. If you use a range for concurrent consumers (eg min max), then this option can be used to set a value to eg 100 to control how fast the consumers will shrink when less work is required.	-1	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.messageConverter</b>	To use a custom Spring <code>org.springframework.jms.support.converter.MessageConverter</code> so you can be in control how to map to/from a <code>javax.jms.Message</code> .	null	MEDIUM
<b>camel.source.endpoint.messageCreatedStrategy</b>	To use the given <code>MessageCreatedStrategy</code> which are invoked when Camel creates new instances of <code>javax.jms.Message</code> objects when Camel is sending a JMS message.	null	MEDIUM
<b>camel.source.endpoint.messageIdEnabled</b>	When sending, specifies whether message IDs should be added. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the message ID set to null; if the provider ignores the hint, the message ID must be set to its normal unique value.	true	MEDIUM
<b>camel.source.endpoint.messageListenerContainerFactory</b>	Registry ID of the <code>MessageListenerContainerFactory</code> used to determine what <code>org.springframework.jms.listener.AbstractMessageListenerContainer</code> to use to consume messages. Setting this will automatically set <code>consumerType</code> to Custom.	null	MEDIUM
<b>camel.source.endpoint.messageTimestampEnabled</b>	Specifies whether timestamps should be enabled by default on sending messages. This is just an hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the timestamp set to zero; if the provider ignores the hint the timestamp must be set to its normal value.	true	MEDIUM
<b>camel.source.endpoint.pubSubNoLocal</b>	Specifies whether to inhibit the delivery of messages published by its own connection.	false	MEDIUM
<b>camel.source.endpoint.receiveTimeout</b>	The timeout for receiving messages (in milliseconds).	1000L	MEDIUM
<b>camel.source.endpoint.recoveryInterval</b>	Specifies the interval between recovery attempts, i.e. when a connection is being refreshed, in milliseconds. The default is 5000 ms, that is, 5 seconds.	5000L	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.requestTimeoutCheckerInterval</b>	Configures how often Camel should check for timed out Exchanges when doing request/reply over JMS. By default Camel checks once per second. But if you must react faster when a timeout occurs, then you can lower this interval, to check more frequently. The timeout is determined by the option requestTimeout.	1000L	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.source.endpoint.transferException</b>	If enabled and you are using Request Reply messaging (InOut) and an Exchange failed on the consumer side, then the caused Exception will be send back in response as a javax.jms.ObjectMessage. If the client is Camel, the returned Exception is rethrown. This allows you to use Camel JMS as a bridge in your routing - for example, using persistent queues to enable robust routing. Notice that if you also have transferExchange enabled, this option takes precedence. The caught exception is required to be serializable. The original Exception on the consumer side can be wrapped in an outer exception such as org.apache.camel.RuntimeCamelException when returned to the producer. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer!	false	MEDIUM
<b>camel.source.endpoint.transferExchange</b>	You can transfer the exchange over the wire instead of just the body and headers. The following fields are transferred: In body, Out body, Fault body, In headers, Out headers, Fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level. You must enable this option on both the producer and consumer side, so Camel knows the payloads is an Exchange and not a regular payload. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer having to use compatible Camel versions!	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.useMessageIDAsCorrelationID</b>	Specifies whether JMSMessageID should always be used as JMSCorrelationID for InOut messages.	false	MEDIUM
<b>camel.source.endpoint.waitForProvisionCorrelationToBeUpdatedCounter</b>	Number of times to wait for provisional correlation id to be updated to the actual correlation id when doing request/reply over JMS and when the option useMessageIDAsCorrelationID is enabled.	50	MEDIUM
<b>camel.source.endpoint.waitForProvisionCorrelationToBeUpdatedThreadSleepingTime</b>	Interval in millis to sleep each time while waiting for provisional correlation id to be updated.	100L	MEDIUM
<b>camel.source.endpoint.errorHandlerLoggingLevel</b>	Allows to configure the default errorHandler logging level for logging uncaught exceptions. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"WARN"	MEDIUM
<b>camel.source.endpoint.errorHandlerLogStackTrace</b>	Allows to control whether stacktraces should be logged or not, by the default errorHandler.	true	MEDIUM
<b>camel.source.endpoint.password</b>	Password to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.	null	MEDIUM
<b>camel.source.endpoint.username</b>	Username to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.	null	MEDIUM
<b>camel.source.endpoint.transacted</b>	Specifies whether to use transacted mode	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.transactedInOut</b>	Specifies whether InOut operations (request reply) default to using transacted mode. If this flag is set to true, then Spring JmsTemplate will have sessionTransacted set to true, and the acknowledgeMode as transacted on the JmsTemplate used for InOut operations. Note from Spring JMS: that within a JTA transaction, the parameters passed to createQueue, createTopic methods are not taken into account. Depending on the Java EE transaction context, the container makes its own decisions on these values. Analogously, these parameters are not taken into account within a locally managed transaction either, since Spring JMS operates on an existing JMS Session in this case. Setting this flag to true will use a short local JMS transaction when running outside of a managed transaction, and a synchronized local JMS transaction in case of a managed transaction (other than an XA transaction) being present. This has the effect of a local JMS transaction being managed alongside the main transaction (which might be a native JDBC transaction), with the JMS transaction committing right after the main transaction.	false	MEDIUM
<b>camel.source.endpoint.lazyCreateTransactionManager</b>	If true, Camel will create a JmsTransactionManager, if there is no transactionManager injected when option transacted=true.	true	MEDIUM
<b>camel.source.endpoint.transactionManager</b>	The Spring transaction manager to use.	null	MEDIUM
<b>camel.source.endpoint.transactionName</b>	The name of the transaction to use.	null	MEDIUM
<b>camel.source.endpoint.transactionTimeout</b>	The timeout value of the transaction (in seconds), if using transacted mode.	-1	MEDIUM
<b>camel.component.jms.clientId</b>	Sets the JMS client ID to use. Note that this value, if specified, must be unique and can only be used by a single JMS connection instance. It is typically only required for durable topic subscriptions. If using Apache ActiveMQ you may prefer to use Virtual Topics instead.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.connectionFactory</b>	The connection factory to be use. A connection factory must be configured either on the component or endpoint.	null	MEDIUM
<b>camel.component.jms.disableReplyTo</b>	Specifies whether Camel ignores the JMSReplyTo header in messages. If true, Camel does not send a reply back to the destination specified in the JMSReplyTo header. You can use this option if you want Camel to consume from a route and you do not want Camel to automatically send back a reply message because another component in your code handles the reply message. You can also use this option if you want to use Camel as a proxy between different message brokers and you want to route message from one system to another.	false	MEDIUM
<b>camel.component.jms.durableSubscriptionName</b>	The durable subscriber name for specifying durable topic subscriptions. The clientId option must be configured as well.	null	MEDIUM
<b>camel.component.jms.jmsMessageType</b>	Allows you to force the use of a specific javax.jms.Message implementation for sending JMS messages. Possible values are: Bytes, Map, Object, Stream, Text. By default, Camel would determine which JMS message type to use from the In body type. This option allows you to specify it. One of: [Bytes] [Map] [Object] [Stream] [Text]	null	MEDIUM
<b>camel.component.jms.testConnectionOnStartup</b>	Specifies whether to test the connection on startup. This ensures that when Camel starts that all the JMS consumers have a valid connection to the JMS broker. If a connection cannot be granted then Camel throws an exception on startup. This ensures that Camel is not started with failed connections. The JMS producers is tested as well.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.acknowledgementModeName</b>	The JMS acknowledgement name, which is one of: SESSION_TRANSACTED, CLIENT_ACKNOWLEDGE, AUTO_ACKNOWLEDGE, DUPS_OK_ACKNOWLEDGE One of: [SESSION_TRANSACTED] [CLIENT_ACKNOWLEDGE] [AUTO_ACKNOWLEDGE] [DUPS_OK_ACKNOWLEDGE]	"AUTO_ACKNOWLEDGE"	MEDIUM
<b>camel.component.jms.asyncConsumer</b>	Whether the JmsConsumer processes the Exchange asynchronously. If enabled then the JmsConsumer may pickup the next message from the JMS queue, while the previous message is being processed asynchronously (by the Asynchronous Routing Engine). This means that messages may be processed not 100% strictly in order. If disabled (as default) then the Exchange is fully processed before the JmsConsumer will pickup the next message from the JMS queue. Note if transacted has been enabled, then asyncConsumer=true does not run asynchronously, as transaction must be executed synchronously (Camel 3.0 may support async transactions).	false	MEDIUM
<b>camel.component.jms.autoStartup</b>	Specifies whether the consumer container should auto-startup.	true	MEDIUM
<b>camel.component.jms.cacheLevel</b>	Sets the cache level by ID for the underlying JMS resources. See cacheLevelName option for more details.	null	MEDIUM
<b>camel.component.jms.cacheLevelName</b>	Sets the cache level by name for the underlying JMS resources. Possible values are: CACHE_AUTO, CACHE_CONNECTION, CACHE_CONSUMER, CACHE_NONE, and CACHE_SESSION. The default setting is CACHE_AUTO. See the Spring documentation and Transactions Cache Levels for more information. One of: [CACHE_AUTO] [CACHE_CONNECTION] [CACHE_CONSUMER] [CACHE_NONE] [CACHE_SESSION]	"CACHE_AUTO"	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.concurrentConsumers</b>	Specifies the default number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.	1	MEDIUM
<b>camel.component.jms.maxConcurrentConsumers</b>	Specifies the maximum number of concurrent consumers when consuming from JMS (not for request/reply over JMS). See also the <code>maxMessagesPerTask</code> option to control dynamic scaling up/down of threads. When doing request/reply over JMS then the option <code>replyToMaxConcurrentConsumers</code> is used to control number of concurrent consumers on the reply message listener.	null	MEDIUM
<b>camel.component.jms.replyTo</b>	Provides an explicit <code>ReplyTo</code> destination, which overrides any incoming value of <code>Message.getJMSReplyTo()</code> .	null	MEDIUM
<b>camel.component.jms.replyToDeliveryPersistent</b>	Specifies whether to use persistent delivery by default for replies.	true	MEDIUM
<b>camel.component.jms.selector</b>	Sets the JMS selector to use	null	MEDIUM
<b>camel.component.jms.subscriptionDurable</b>	Set whether to make the subscription durable. The durable subscription name to be used can be specified through the <code>subscriptionName</code> property. Default is false. Set this to true to register a durable subscription, typically in combination with a <code>subscriptionName</code> value (unless your message listener class name is good enough as subscription name). Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the <code>pubSubDomain</code> flag as well.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.subscriptionName</b>	Set the name of a subscription to create. To be applied in case of a topic (pub-sub domain) with a shared or durable subscription. The subscription name needs to be unique within this client's JMS client id. Default is the class name of the specified message listener. Note: Only 1 concurrent consumer (which is the default of this message listener container) is allowed for each subscription, except for a shared subscription (which requires JMS 2.0).	null	MEDIUM
<b>camel.component.jms.subscriptionShared</b>	Set whether to make the subscription shared. The shared subscription name to be used can be specified through the subscriptionName property. Default is false. Set this to true to register a shared subscription, typically in combination with a subscriptionName value (unless your message listener class name is good enough as subscription name). Note that shared subscriptions may also be durable, so this flag can (and often will) be combined with subscriptionDurable as well. Only makes sense when listening to a topic (pub-sub domain), therefore this method switches the pubSubDomain flag as well. Requires a JMS 2.0 compatible message broker.	false	MEDIUM
<b>camel.component.jms.acceptMessagesWhileStopping</b>	Specifies whether the consumer accept messages while it is stopping. You may consider enabling this option, if you start and stop JMS routes at runtime, while there are still messages enqueued on the queue. If this option is false, and you stop the JMS route, then messages may be rejected, and the JMS broker would have to attempt redeliveries, which yet again may be rejected, and eventually the message may be moved at a dead letter queue on the JMS broker. To avoid this its recommended to enable this option.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.allowReplyManagerQuickStop</b>	Whether the DefaultMessageListenerContainer used in the reply managers for request-reply messaging allow the DefaultMessageListenerContainer.runningAllowed flag to quick stop in case JmsConfiguration#isAcceptMessagesWhileStopping is enabled, and org.apache.camel.CamelContext is currently being stopped. This quick stop ability is enabled by default in the regular JMS consumers but to enable for reply managers you must enable this flag.	false	MEDIUM
<b>camel.component.jms.consumerType</b>	The consumer type to use, which can be one of: Simple, Default, or Custom. The consumer type determines which Spring JMS listener to use. Default will use org.springframework.jms.listener.DefaultMessageListenerContainer, Simple will use org.springframework.jms.listener.SimpleMessageListenerContainer. When Custom is specified, the MessageListenerContainerFactory defined by the messageListenerContainerFactory option will determine what org.springframework.jms.listener.AbstractMessageListenerContainer to use. One of: [Simple] [Default] [Custom]	"Default"	MEDIUM
<b>camel.component.jms.defaultTaskExecutorType</b>	Specifies what default TaskExecutor type to use in the DefaultMessageListenerContainer, for both consumer endpoints and the ReplyTo consumer of producer endpoints. Possible values: SimpleAsync (uses Spring's SimpleAsyncTaskExecutor) or ThreadPool (uses Spring's ThreadPoolTaskExecutor with optimal values - cached threadpool-like). If not set, it defaults to the previous behaviour, which uses a cached thread pool for consumer endpoints and SimpleAsync for reply consumers. The use of ThreadPool is recommended to reduce thread trash in elastic configurations with dynamically increasing and decreasing concurrent consumers. One of: [ThreadPool] [SimpleAsync]	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.eagerLoadingOfProperties</b>	Enables eager loading of JMS properties and payload as soon as a message is loaded which generally is inefficient as the JMS properties may not be required but sometimes can catch early any issues with the underlying JMS provider and the use of JMS properties. See also the option <code>eagerPoisonBody</code> .	false	MEDIUM
<b>camel.component.jms.eagerPoisonBody</b>	If <code>eagerLoadingOfProperties</code> is enabled and the JMS message payload (JMS body or JMS properties) is poison (cannot be read/mapped), then set this text as the message body instead so the message can be processed (the cause of the poison are already stored as exception on the Exchange). This can be turned off by setting <code>eagerPoisonBody=false</code> . See also the option <code>eagerLoadingOfProperties</code> .	"Poison JMS message due to <code>{exception.message}</code> "	MEDIUM
<b>camel.component.jms.exposeListenerSession</b>	Specifies whether the listener session should be exposed when consuming messages.	false	MEDIUM
<b>camel.component.jms.replyToSameDestinationAllowed</b>	Whether a JMS consumer is allowed to send a reply message to the same destination that the consumer is using to consume from. This prevents an endless loop by consuming and sending back the same message to itself.	false	MEDIUM
<b>camel.component.jms.taskExecutor</b>	Allows you to specify a custom task executor for consuming messages.	null	MEDIUM
<b>camel.component.jms.allowAutoWiredConnectionFactory</b>	Whether to auto-discover <code>ConnectionFactory</code> from the registry, if no connection factory has been configured. If only one instance of <code>ConnectionFactory</code> is found then it will be used. This is enabled by default.	true	MEDIUM
<b>camel.component.jms.allowAutoWiredDestinationResolver</b>	Whether to auto-discover <code>DestinationResolver</code> from the registry, if no destination resolver has been configured. If only one instance of <code>DestinationResolver</code> is found then it will be used. This is enabled by default.	true	MEDIUM
<b>camel.component.jms.allowSerializedHeaders</b>	Controls whether or not to include serialized headers. Applies only when <code>transferExchange</code> is true. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.artemisStreamingEnabled</b>	Whether optimizing for Apache Artemis streaming mode.	true	MEDIUM
<b>camel.component.jms.asyncStartListener</b>	Whether to startup the JmsConsumer message listener asynchronously, when starting a route. For example if a JmsConsumer cannot get a connection to a remote JMS broker, then it may block while retrying and/or failover. This will cause Camel to block while starting routes. By setting this option to true, you will let routes startup, while the JmsConsumer connects to the JMS broker using a dedicated thread in asynchronous mode. If this option is used, then beware that if the connection could not be established, then an exception is logged at WARN level, and the consumer will not be able to receive messages; You can then restart the route to retry.	false	MEDIUM
<b>camel.component.jms.asyncStopListener</b>	Whether to stop the JmsConsumer message listener asynchronously, when stopping a route.	false	MEDIUM
<b>camel.component.jms.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.component.jms.configuration</b>	To use a shared JMS configuration	null	MEDIUM
<b>camel.component.jms.destinationResolver</b>	A pluggable <code>org.springframework.jms.support.destination.DestinationResolver</code> that allows you to use your own resolver (for example, to lookup the real destination in a JNDI registry).	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.errorHandler</b>	Specifies a <code>org.springframework.util.ErrorHandler</code> to be invoked in case of any uncaught exceptions thrown while processing a Message. By default these exceptions will be logged at the WARN level, if no errorHandler has been configured. You can configure logging level and whether stack traces should be logged using <code>errorHandlerLoggingLevel</code> and <code>errorHandlerLogStackTrace</code> options. This makes it much easier to configure, than having to code a custom errorHandler.	null	MEDIUM
<b>camel.component.jms.exceptionListener</b>	Specifies the JMS Exception Listener that is to be notified of any underlying JMS exceptions.	null	MEDIUM
<b>camel.component.jms.idleConsumerLimit</b>	Specify the limit for the number of consumers that are allowed to be idle at any given time.	1	MEDIUM
<b>camel.component.jms.idleTaskExecutionLimit</b>	Specifies the limit for idle executions of a receive task, not having received any message within its execution. If this limit is reached, the task will shut down and leave receiving to other executing tasks (in the case of dynamic scheduling; see the <code>maxConcurrentConsumers</code> setting). There is additional doc available from Spring.	1	MEDIUM
<b>camel.component.jms.includeAllJMSXProperties</b>	Whether to include all JMSxxx properties when mapping from JMS to Camel Message. Setting this to true will include properties such as <code>JMSXAppID</code> , and <code>JMSXUserID</code> etc. Note: If you are using a custom <code>headerFilterStrategy</code> then this option does not apply.	false	MEDIUM
<b>camel.component.jms.jmsKeyFormatStrategy</b>	Pluggable strategy for encoding and decoding JMS keys so they can be compliant with the JMS specification. Camel provides two implementations out of the box: default and passthrough. The default strategy will safely marshal dots and hyphens (. and -). The passthrough strategy leaves the key as is. Can be used for JMS brokers which do not care whether JMS header keys contain illegal characters. You can provide your own implementation of the <code>org.apache.camel.component.jms.JmsKeyFormatStrategy</code> and refer to it using the # notation. One of: [default] [passthrough]	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.jms.mapJmsMessage</code>	Specifies whether Camel should auto map the received JMS message to a suited payload type, such as <code>javax.jms.TextMessage</code> to a <code>String</code> etc.	<code>true</code>	MEDIUM
<code>camel.component.jms.maxMessagesPerTask</code>	The number of messages per task. -1 is unlimited. If you use a range for concurrent consumers (eg min max), then this option can be used to set a value to eg 100 to control how fast the consumers will shrink when less work is required.	<code>-1</code>	MEDIUM
<code>camel.component.jms.messageConverter</code>	To use a custom Spring <code>org.springframework.jms.support.converter.MessageConverter</code> so you can be in control how to map to/from a <code>javax.jms.Message</code> .	<code>null</code>	MEDIUM
<code>camel.component.jms.messageCreatedStrategy</code>	To use the given <code>MessageCreatedStrategy</code> which are invoked when Camel creates new instances of <code>javax.jms.Message</code> objects when Camel is sending a JMS message.	<code>null</code>	MEDIUM
<code>camel.component.jms.messageIdEnabled</code>	When sending, specifies whether message IDs should be added. This is just a hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the message ID set to null; if the provider ignores the hint, the message ID must be set to its normal unique value.	<code>true</code>	MEDIUM
<code>camel.component.jms.messageListenerContainerFactory</code>	Registry ID of the <code>MessageListenerContainerFactory</code> used to determine what <code>org.springframework.jms.listener.AbstractMessageListenerContainer</code> to use to consume messages. Setting this will automatically set <code>consumerType</code> to <code>Custom</code> .	<code>null</code>	MEDIUM
<code>camel.component.jms.messageTimestampEnabled</code>	Specifies whether timestamps should be enabled by default on sending messages. This is just a hint to the JMS broker. If the JMS provider accepts this hint, these messages must have the timestamp set to zero; if the provider ignores the hint the timestamp must be set to its normal value.	<code>true</code>	MEDIUM
<code>camel.component.jms.pubSubNoLocal</code>	Specifies whether to inhibit the delivery of messages published by its own connection.	<code>false</code>	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.queueBrowseStrategy</b>	To use a custom QueueBrowseStrategy when browsing queues	null	MEDIUM
<b>camel.component.jms.receiveTimeout</b>	The timeout for receiving messages (in milliseconds).	1000L	MEDIUM
<b>camel.component.jms.recoveryInterval</b>	Specifies the interval between recovery attempts, i.e. when a connection is being refreshed, in milliseconds. The default is 5000 ms, that is, 5 seconds.	5000L	MEDIUM
<b>camel.component.jms.requestTimeoutCheckerInterval</b>	Configures how often Camel should check for timed out Exchanges when doing request/reply over JMS. By default Camel checks once per second. But if you must react faster when a timeout occurs, then you can lower this interval, to check more frequently. The timeout is determined by the option requestTimeout.	1000L	MEDIUM
<b>camel.component.jms.transferException</b>	If enabled and you are using Request Reply messaging (InOut) and an Exchange failed on the consumer side, then the caused Exception will be send back in response as a javax.jms.ObjectMessage. If the client is Camel, the returned Exception is rethrown. This allows you to use Camel JMS as a bridge in your routing - for example, using persistent queues to enable robust routing. Notice that if you also have transferExchange enabled, this option takes precedence. The caught exception is required to be serializable. The original Exception on the consumer side can be wrapped in an outer exception such as org.apache.camel.RuntimeCamelException when returned to the producer. Use this with caution as the data is using Java Object serialization and requires the received to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer!	false	MEDIUM

Name	Description	Default	Priority
<code>camel.component.jms.transferExchange</code>	You can transfer the exchange over the wire instead of just the body and headers. The following fields are transferred: In body, Out body, Fault body, In headers, Out headers, Fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level. You must enable this option on both the producer and consumer side, so Camel knows the payload is an Exchange and not a regular payload. Use this with caution as the data is using Java Object serialization and requires the receiver to be able to deserialize the data at Class level, which forces a strong coupling between the producers and consumer having to use compatible Camel versions!	false	MEDIUM
<code>camel.component.jms.useMessageIDAsCorrelationID</code>	Specifies whether JMSMessageID should always be used as JMSCorrelationID for InOut messages.	false	MEDIUM
<code>camel.component.jms.waitForProvisionCorrelationToBeUpdatedCounter</code>	Number of times to wait for provisional correlation id to be updated to the actual correlation id when doing request/reply over JMS and when the option <code>useMessageIDAsCorrelationID</code> is enabled.	50	MEDIUM
<code>camel.component.jms.waitForProvisionCorrelationToBeUpdatedThreadSleepingTime</code>	Interval in millis to sleep each time while waiting for provisional correlation id to be updated.	100L	MEDIUM
<code>camel.component.jms.headerFilterStrategy</code>	To use a custom <code>org.apache.camel.spi.HeaderFilterStrategy</code> to filter header to and from Camel message.	null	MEDIUM
<code>camel.component.jms.errorHandlerLoggingLevel</code>	Allows to configure the default errorHandler logging level for logging uncaught exceptions. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"WARN"	MEDIUM
<code>camel.component.jms.errorHandlerLogStackTrace</code>	Allows to control whether stacktraces should be logged or not, by the default errorHandler.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.component.jms.password</b>	Password to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.	null	MEDIUM
<b>camel.component.jms.username</b>	Username to use with the ConnectionFactory. You can also configure username/password directly on the ConnectionFactory.	null	MEDIUM
<b>camel.component.jms.transacted</b>	Specifies whether to use transacted mode	false	MEDIUM
<b>camel.component.jms.transactedInOut</b>	Specifies whether InOut operations (request reply) default to using transacted mode. If this flag is set to true, then Spring JmsTemplate will have sessionTransacted set to true, and the acknowledgeMode as transacted on the JmsTemplate used for InOut operations. Note from Spring JMS: that within a JTA transaction, the parameters passed to createQueue, createTopic methods are not taken into account. Depending on the Java EE transaction context, the container makes its own decisions on these values. Analogously, these parameters are not taken into account within a locally managed transaction either, since Spring JMS operates on an existing JMS Session in this case. Setting this flag to true will use a short local JMS transaction when running outside of a managed transaction, and a synchronized local JMS transaction in case of a managed transaction (other than an XA transaction) being present. This has the effect of a local JMS transaction being managed alongside the main transaction (which might be a native JDBC transaction), with the JMS transaction committing right after the main transaction.	false	MEDIUM
<b>camel.component.jms.lazyCreateTransactionManager</b>	If true, Camel will create a JmsTransactionManager, if there is no transactionManager injected when option transacted=true.	true	MEDIUM
<b>camel.component.jms.transactionManager</b>	The Spring transaction manager to use.	null	MEDIUM
<b>camel.component.jms.transactionName</b>	The name of the transaction to use.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.jms.transactionTimeout</code>	The timeout value of the transaction (in seconds), if using transacted mode.	-1	MEDIUM

The camel-jms sink connector has no converters out of the box.

The camel-jms sink connector has no transforms out of the box.

The camel-jms sink connector has no aggregation strategies out of the box.

## 5.17. CAMEL-MONGODB-KAFKA-CONNECTOR SINK CONFIGURATION

When using camel-mongodb-kafka-connector as sink make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-mongodb-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Sink connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.mongodb.CamelMongodbSinkConnector
```

The camel-mongodb sink connector supports 26 options, which are listed below.

Name	Description	Default	Priority
<code>camel.sink.path.connectionBean</code>	Sets the connection bean reference used to lookup a client for connecting to a database.	null	HIGH
<code>camel.sink.endpoint.collection</code>	Sets the name of the MongoDB collection to bind to this endpoint	null	MEDIUM
<code>camel.sink.endpoint.collectionIndex</code>	Sets the collection index (JSON FORMAT : { field1 : order1, field2 : order2})	null	MEDIUM
<code>camel.sink.endpoint.createCollection</code>	Create collection during initialisation if it doesn't exist. Default is true.	true	MEDIUM
<code>camel.sink.endpoint.database</code>	Sets the name of the MongoDB database to target	null	MEDIUM
<code>camel.sink.endpoint.mongoConnection</code>	Sets the connection bean used as a client for connecting to a database.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.operation</b>	Sets the operation this endpoint will execute against MongoDB. One of: [findById] [findOneByQuery] [findAll] [findDistinct] [insert] [save] [update] [remove] [bulkWrite] [aggregate] [getDbStats] [getColStats] [count] [command]	null	MEDIUM
<b>camel.sink.endpoint.outputType</b>	Convert the output of the producer to the selected type : DocumentList Document or Mongolterable. DocumentList or Mongolterable applies to findAll and aggregate. Document applies to all other operations. One of: [DocumentList] [Document] [Mongolterable]	null	MEDIUM
<b>camel.sink.endpoint.lazyStartProducer</b>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<b>camel.sink.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.sink.endpoint.cursorRegenerationDelay</b>	MongoDB tailable cursors will block until new data arrives. If no new data is inserted, after some time the cursor will be automatically freed and closed by the MongoDB server. The client is expected to regenerate the cursor if needed. This value specifies the time to wait before attempting to fetch a new cursor, and if the attempt fails, how long before the next attempt is made. Default value is 1000ms.	1000L	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.dynamycity</b>	Sets whether this endpoint will attempt to dynamically resolve the target database and collection from the incoming Exchange properties. Can be used to override at runtime the database and collection specified on the otherwise static endpoint URI. It is disabled by default to boost performance. Enabling it will take a minimal performance hit.	false	MEDIUM
<b>camel.sink.endpoint.readPreference</b>	Configure how MongoDB clients route read operations to the members of a replica set. Possible values are PRIMARY, PRIMARY_PREFERRED, SECONDARY, SECONDARY_PREFERRED or NEAREST One of: [PRIMARY] [PRIMARY_PREFERRED] [SECONDARY] [SECONDARY_PREFERRED] [NEAREST]	"PRIMARY"	MEDIUM
<b>camel.sink.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.sink.endpoint.writeConcern</b>	Configure the connection bean with the level of acknowledgment requested from MongoDB for write operations to a standalone mongod, replicaset or cluster. Possible values are ACKNOWLEDGED, W1, W2, W3, UNACKNOWLEDGED, JOURNALED or MAJORITY. One of: [ACKNOWLEDGED] [W1] [W2] [W3] [UNACKNOWLEDGED] [JOURNALED] [MAJORITY]	"ACKNOWLEDGED"	MEDIUM
<b>camel.sink.endpoint.writeResultAsHeader</b>	In write operations, it determines whether instead of returning WriteResult as the body of the OUT message, we transfer the IN message to the OUT and attach the WriteResult as a header.	false	MEDIUM
<b>camel.sink.endpoint.streamFilter</b>	Filter condition for change streams consumer.	null	MEDIUM
<b>camel.sink.endpoint.persistentId</b>	One tail tracking collection can host many trackers for several tailable consumers. To keep them separate, each tracker should have its own unique persistentId.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.sink.endpoint.persistentTailTracking</b>	Enable persistent tail tracking, which is a mechanism to keep track of the last consumed message across system restarts. The next time the system is up, the endpoint will recover the cursor from the point where it last stopped slurping records.	false	MEDIUM
<b>camel.sink.endpoint.tailTrackCollection</b>	Collection where tail tracking information will be persisted. If not specified, <code>MongoDbTailTrackingConfig#DEFAULT_COLLECTION</code> will be used by default.	null	MEDIUM
<b>camel.sink.endpoint.tailTrackDb</b>	Indicates what database the tail tracking mechanism will persist to. If not specified, the current database will be picked by default. Dynamicity will not be taken into account even if enabled, i.e. the tail tracking database will not vary past endpoint initialisation.	null	MEDIUM
<b>camel.sink.endpoint.tailTrackField</b>	Field where the last tracked value will be placed. If not specified, <code>MongoDbTailTrackingConfig#DEFAULT_FIELD</code> will be used by default.	null	MEDIUM
<b>camel.sink.endpoint.tailTrackIncreasingField</b>	Correlation field in the incoming record which is of increasing nature and will be used to position the tailing cursor every time it is generated. The cursor will be (re)created with a query of type: <code>tailTrackIncreasingField greater than lastValue</code> (possibly recovered from persistent tail tracking). Can be of type Integer, Date, String, etc. NOTE: No support for dot notation at the current time, so the field should be at the top level of the document.	null	MEDIUM
<b>camel.component.mongodb.mongoConnection</b>	Shared client used for connection. All endpoints generated from the component will share this connection client.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.mongodb.lazyStartProducer</code>	Whether the producer should be started lazy (on the first message). By starting lazy you can use this to allow CamelContext and routes to startup in situations where a producer may otherwise fail during starting and cause the route to fail being started. By deferring this startup to be lazy then the startup failure can be handled during routing messages via Camel's routing error handlers. Beware that when the first message is processed then creating and starting the producer may take a little time and prolong the total processing time of the processing.	false	MEDIUM
<code>camel.component.mongodb.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

The camel-mongodb sink connector has no converters out of the box.

The camel-mongodb sink connector has no transforms out of the box.

The camel-mongodb sink connector has no aggregation strategies out of the box.

## 5.18. CAMEL-MONGODB-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-mongodb-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-mongodb-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.mongodb.CamelMongoDbSourceConnector
```

The camel-mongodb source connector supports 29 options, which are listed below.

Name	Description	Default	Priority
<code>camel.source.path.connectionBean</code>	Sets the connection bean reference used to lookup a client for connecting to a database.	null	HIGH

Name	Description	Default	Priority
<b>camel.source.endpoint.collection</b>	Sets the name of the MongoDB collection to bind to this endpoint	null	MEDIUM
<b>camel.source.endpoint.collectionIndex</b>	Sets the collection index (JSON FORMAT : { field1 : order1, field2 : order2})	null	MEDIUM
<b>camel.source.endpoint.createCollection</b>	Create collection during initialisation if it doesn't exist. Default is true.	true	MEDIUM
<b>camel.source.endpoint.database</b>	Sets the name of the MongoDB database to target	null	MEDIUM
<b>camel.source.endpoint.mongoConnection</b>	Sets the connection bean used as a client for connecting to a database.	null	MEDIUM
<b>camel.source.endpoint.operation</b>	Sets the operation this endpoint will execute against MongoDB. One of: [findById] [findOneByQuery] [findAll] [findDistinct] [insert] [save] [update] [remove] [bulkWrite] [aggregate] [getDbStats] [getColStats] [count] [command]	null	MEDIUM
<b>camel.source.endpoint.outputType</b>	Convert the output of the producer to the selected type : DocumentList Document or Mongolterable. DocumentList or Mongolterable applies to findAll and aggregate. Document applies to all other operations. One of: [DocumentList] [Document] [Mongolterable]	null	MEDIUM
<b>camel.source.endpoint.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.source.endpoint.consumerType</b>	Consumer type.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: [InOnly] [InOut] [InOptionalOut]	null	MEDIUM
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.source.endpoint.cursorRegenerationDelay</b>	MongoDB tailable cursors will block until new data arrives. If no new data is inserted, after some time the cursor will be automatically freed and closed by the MongoDB server. The client is expected to regenerate the cursor if needed. This value specifies the time to wait before attempting to fetch a new cursor, and if the attempt fails, how long before the next attempt is made. Default value is 1000ms.	1000L	MEDIUM
<b>camel.source.endpoint.dynamicity</b>	Sets whether this endpoint will attempt to dynamically resolve the target database and collection from the incoming Exchange properties. Can be used to override at runtime the database and collection specified on the otherwise static endpoint URI. It is disabled by default to boost performance. Enabling it will take a minimal performance hit.	false	MEDIUM
<b>camel.source.endpoint.readPreference</b>	Configure how MongoDB clients route read operations to the members of a replica set. Possible values are PRIMARY, PRIMARY_PREFERRED, SECONDARY, SECONDARY_PREFERRED or NEAREST One of: [PRIMARY] [PRIMARY_PREFERRED] [SECONDARY] [SECONDARY_PREFERRED] [NEAREST]	"PRIMARY"	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.writeConcern</b>	Configure the connection bean with the level of acknowledgment requested from MongoDB for write operations to a standalone mongod, replicaset or cluster. Possible values are ACKNOWLEDGED, W1, W2, W3, UNACKNOWLEDGED, JOURNALED or MAJORITY. One of: [ACKNOWLEDGED] [W1] [W2] [W3] [UNACKNOWLEDGED] [JOURNALED] [MAJORITY]	"ACKNOWLEDGED"	MEDIUM
<b>camel.source.endpoint.writeResultAsHeader</b>	In write operations, it determines whether instead of returning WriteResult as the body of the OUT message, we transfer the IN message to the OUT and attach the WriteResult as a header.	false	MEDIUM
<b>camel.source.endpoint.streamFilter</b>	Filter condition for change streams consumer.	null	MEDIUM
<b>camel.source.endpoint.persistentId</b>	One tail tracking collection can host many trackers for several tailable consumers. To keep them separate, each tracker should have its own unique persistentId.	null	MEDIUM
<b>camel.source.endpoint.persistentTailTracking</b>	Enable persistent tail tracking, which is a mechanism to keep track of the last consumed message across system restarts. The next time the system is up, the endpoint will recover the cursor from the point where it last stopped slurping records.	false	MEDIUM
<b>camel.source.endpoint.tailTrackCollection</b>	Collection where tail tracking information will be persisted. If not specified, MongoDbTailTrackingConfig#DEFAULT_COLLECTION will be used by default.	null	MEDIUM
<b>camel.source.endpoint.tailTrackDb</b>	Indicates what database the tail tracking mechanism will persist to. If not specified, the current database will be picked by default. Dynamicity will not be taken into account even if enabled, i.e. the tail tracking database will not vary past endpoint initialisation.	null	MEDIUM
<b>camel.source.endpoint.tailTrackField</b>	Field where the last tracked value will be placed. If not specified, MongoDbTailTrackingConfig#DEFAULT_FIELD will be used by default.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.tailTrackIncreasingField</code>	Correlation field in the incoming record which is of increasing nature and will be used to position the tailing cursor every time it is generated. The cursor will be (re)created with a query of type: <code>tailTrackIncreasingField</code> greater than <code>lastValue</code> (possibly recovered from persistent tail tracking). Can be of type Integer, Date, String, etc. NOTE: No support for dot notation at the current time, so the field should be at the top level of the document.	null	MEDIUM
<code>camel.component.mongodb.mongoConnection</code>	Shared client used for connection. All endpoints generated from the component will share this connection client.	null	MEDIUM
<code>camel.component.mongodb.bridgeErrorHandler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<code>camel.component.mongodb.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

The camel-mongodb sink connector has no converters out of the box.

The camel-mongodb sink connector has no transforms out of the box.

The camel-mongodb sink connector has no aggregation strategies out of the box.

## 5.19. CAMEL-NETTY-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-netty-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-netty-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following `connector.class`

```
connector.class=org.apache.camel.kafkaconnector.netty.CamelNettySourceConnector
```

The camel-netty source connector supports 120 options, which are listed below.

Name	Description	Default	Priority
<b>camel.source.path.protocol</b>	The protocol to use which can be tcp or udp. One of: [tcp] [udp]	null	HIGH
<b>camel.source.path.host</b>	The hostname. For the consumer the hostname is localhost or 0.0.0.0. For the producer the hostname is the remote host to connect to	null	HIGH
<b>camel.source.path.port</b>	The host port number	null	HIGH
<b>camel.source.endpoint.disconnect</b>	Whether or not to disconnect(close) from Netty Channel right after use. Can be used for both consumer and producer.	false	MEDIUM
<b>camel.source.endpoint.keepAlive</b>	Setting to ensure socket is not closed due to inactivity	true	MEDIUM
<b>camel.source.endpoint.reuseAddress</b>	Setting to facilitate socket multiplexing	true	MEDIUM
<b>camel.source.endpoint.reuseChannel</b>	This option allows producers and consumers (in client mode) to reuse the same Netty Channel for the lifecycle of processing the Exchange. This is useful if you need to call a server multiple times in a Camel route and want to use the same network connection. When using this, the channel is not returned to the connection pool until the Exchange is done; or disconnected if the disconnect option is set to true. The reused Channel is stored on the Exchange as an exchange property with the key <code>NettyConstants#NETTY_CHANNEL</code> which allows you to obtain the channel during routing and use it as well.	false	MEDIUM
<b>camel.source.endpoint.sync</b>	Setting to set endpoint as one-way or request-response	true	MEDIUM
<b>camel.source.endpoint.tcpNoDelay</b>	Setting to improve TCP protocol performance	true	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.source.endpoint.broadcast</b>	Setting to choose Multicast over UDP	false	MEDIUM
<b>camel.source.endpoint.clientMode</b>	If the <code>clientMode</code> is true, netty consumer will connect the address as a TCP client.	false	MEDIUM
<b>camel.source.endpoint.reconnect</b>	Used only in <code>clientMode</code> in consumer, the consumer will attempt to reconnect on disconnection if this is enabled	true	MEDIUM
<b>camel.source.endpoint.reconnectInterval</b>	Used if <code>reconnect</code> and <code>clientMode</code> is enabled. The interval in milli seconds to attempt reconnection	10000	MEDIUM
<b>camel.source.endpoint.backlog</b>	Allows to configure a backlog for netty consumer (server). Note the backlog is just a best effort depending on the OS. Setting this option to a value such as 200, 500 or 1000, tells the TCP stack how long the accept queue can be. If this option is not configured, then the backlog depends on OS setting.	null	MEDIUM
<b>camel.source.endpoint.bossCount</b>	When netty works on nio mode, it uses default <code>bossCount</code> parameter from Netty, which is 1. User can use this option to override the default <code>bossCount</code> from Netty	1	MEDIUM
<b>camel.source.endpoint.bossGroup</b>	Set the <code>BossGroup</code> which could be used for handling the new connection of the server side across the <code>NettyEndpoint</code>	null	MEDIUM
<b>camel.source.endpoint.disconnectOnNoReply</b>	If <code>sync</code> is enabled then this option dictates <code>NettyConsumer</code> if it should disconnect where there is no reply to send back.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom ExceptionHandler. Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: [InOnly] [InOut] [InOptionalOut]	null	MEDIUM
<b>camel.source.endpoint.nettyServerBootstrapFactory</b>	To use a custom <code>NettyServerBootstrapFactory</code>	null	MEDIUM
<b>camel.source.endpoint.networkInterface</b>	When using UDP then this option can be used to specify a network interface by its name, such as <code>eth0</code> to join a multicast group.	null	MEDIUM
<b>camel.source.endpoint.noReplyLogLevel</b>	If <code>sync</code> is enabled this option dictates <code>NettyConsumer</code> which logging level to use when logging a there is no reply to send back. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"WARN"	MEDIUM
<b>camel.source.endpoint.serverClosedChannelExceptionCaughtLogLevel</b>	If the server ( <code>NettyConsumer</code> ) catches an <code>java.nio.channels.ClosedChannelException</code> then its logged using this logging level. This is used to avoid logging the closed channel exceptions, as clients can disconnect abruptly and then cause a flood of closed exceptions in the <code>Netty</code> server. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"DEBUG"	MEDIUM
<b>camel.source.endpoint.serverExceptionCaughtLog Level</b>	If the server ( <code>NettyConsumer</code> ) catches an exception then its logged using this logging level. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"WARN"	MEDIUM
<b>camel.source.endpoint.serverInitializerFactory</b>	To use a custom <code>ServerInitializerFactory</code>	null	MEDIUM
<b>camel.source.endpoint.usingExecutorService</b>	Whether to use ordered thread pool, to ensure events are processed orderly on the same channel.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.allowSerializedHeaders</b>	Only used for TCP when transferExchange is true. When set to true, serializable objects in headers and properties will be added to the exchange. Otherwise Camel will exclude any non-serializable objects and log it at WARN level.	false	MEDIUM
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.source.endpoint.channelGroup</b>	To use a explicit ChannelGroup.	null	MEDIUM
<b>camel.source.endpoint.nativeTransport</b>	Whether to use native transport instead of NIO. Native transport takes advantage of the host operating system and is only supported on some platforms. You need to add the netty JAR for the host operating system you are using. See more details at: <a href="http://netty.io/wiki/native-transport.html">http://netty.io/wiki/native-transport.html</a>	false	MEDIUM
<b>camel.source.endpoint.options</b>	Allows to configure additional netty options using option. as prefix. For example option.child.keepAlive=false to set the netty option child.keepAlive=false. See the Netty documentation for possible options that can be used.	null	MEDIUM
<b>camel.source.endpoint.receiveBufferSize</b>	The TCP/UDP buffer sizes to be used during inbound communication. Size is bytes.	65536	MEDIUM
<b>camel.source.endpoint.receiveBufferSizePredictor</b>	Configures the buffer size predictor. See details at Jetty documentation and this mail thread.	null	MEDIUM
<b>camel.source.endpoint.sendBufferSize</b>	The TCP/UDP buffer sizes to be used during outbound communication. Size is bytes.	65536	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.transferExchange</b>	Only used for TCP. You can transfer the exchange over the wire instead of just the body. The following fields are transferred: In body, Out body, fault body, In headers, Out headers, fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	MEDIUM
<b>camel.source.endpoint.udpByteArrayCodec</b>	For UDP only. If enabled the using byte array codec instead of Java serialization protocol.	false	MEDIUM
<b>camel.source.endpoint.workerCount</b>	When netty works on nio mode, it uses default workerCount parameter from Netty (which is <code>cpu_core_threads x 2</code> ). User can use this option to override the default workerCount from Netty.	null	MEDIUM
<b>camel.source.endpoint.workerGroup</b>	To use a explicit EventLoopGroup as the boss thread pool. For example to share a thread pool with multiple consumers or producers. By default each consumer or producer has their own worker pool with 2 x cpu count core threads.	null	MEDIUM
<b>camel.source.endpoint.allowDefaultCodec</b>	The netty component installs a default codec if both, encoder/decoder is null and textline is false. Setting <code>allowDefaultCodec</code> to false prevents the netty component from installing a default codec as the first element in the filter chain.	true	MEDIUM
<b>camel.source.endpoint.autoAppendDelimiter</b>	Whether or not to auto append missing end delimiter when sending using the textline codec.	true	MEDIUM
<b>camel.source.endpoint.decoderMaxLineLength</b>	The max line length to use for the textline codec.	1024	MEDIUM
<b>camel.source.endpoint.decoders</b>	A list of decoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.delimiter</b>	The delimiter to use for the textline codec. Possible values are LINE and NULL. One of: [LINE] [NULL]	"LINE"	MEDIUM
<b>camel.source.endpoint.encoders</b>	A list of encoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.	null	MEDIUM
<b>camel.source.endpoint.encoding</b>	The encoding (a charset name) to use for the textline codec. If not provided, Camel will use the JVM default Charset.	null	MEDIUM
<b>camel.source.endpoint.textline</b>	Only used for TCP. If no codec is specified, you can use this flag to indicate a text line based codec; if not specified or the value is false, then Object Serialization is assumed over TCP - however only Strings are allowed to be serialized by default.	false	MEDIUM
<b>camel.source.endpoint.enabledProtocols</b>	Which protocols to enable when using SSL	"TLSv1,TLSv1.1,TLSv1.2"	MEDIUM
<b>camel.source.endpoint.keyStoreFile</b>	Client side certificate keystore to be used for encryption	null	MEDIUM
<b>camel.source.endpoint.keyStoreFormat</b>	Keystore format to be used for payload encryption. Defaults to JKS if not set	null	MEDIUM
<b>camel.source.endpoint.keyStoreResource</b>	Client side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.	null	MEDIUM
<b>camel.source.endpoint.needClientAuth</b>	Configures whether the server needs client authentication when using SSL.	false	MEDIUM
<b>camel.source.endpoint.passphrase</b>	Password setting to use in order to encrypt/decrypt payloads sent using SSH	null	MEDIUM
<b>camel.source.endpoint.securityProvider</b>	Security provider to be used for payload encryption. Defaults to SunX509 if not set.	null	MEDIUM
<b>camel.source.endpoint.ssl</b>	Setting to specify whether SSL encryption is applied to this endpoint	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.sslClientCertHeaders</b>	When enabled and in SSL mode, then the Netty consumer will enrich the Camel Message with headers having information about the client certificate such as subject name, issuer name, serial number, and the valid date range.	false	MEDIUM
<b>camel.source.endpoint.sslContextParameters</b>	To configure security using SSLContextParameters	null	MEDIUM
<b>camel.source.endpoint.sslHandler</b>	Reference to a class that could be used to return an SSL Handler	null	MEDIUM
<b>camel.source.endpoint.trustStoreFile</b>	Server side certificate keystore to be used for encryption	null	MEDIUM
<b>camel.source.endpoint.trustStoreResource</b>	Server side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.	null	MEDIUM
<b>camel.component.netty.configuration</b>	To use the NettyConfiguration as configuration when creating endpoints.	null	MEDIUM
<b>camel.component.netty.disconnect</b>	Whether or not to disconnect(close) from Netty Channel right after use. Can be used for both consumer and producer.	false	MEDIUM
<b>camel.component.netty.keepAlive</b>	Setting to ensure socket is not closed due to inactivity	true	MEDIUM
<b>camel.component.netty.reuseAddress</b>	Setting to facilitate socket multiplexing	true	MEDIUM
<b>camel.component.netty.reuseChannel</b>	This option allows producers and consumers (in client mode) to reuse the same Netty Channel for the lifecycle of processing the Exchange. This is useful if you need to call a server multiple times in a Camel route and want to use the same network connection. When using this, the channel is not returned to the connection pool until the Exchange is done; or disconnected if the disconnect option is set to true. The reused Channel is stored on the Exchange as an exchange property with the key <code>NettyConstants#NETTY_CHANNEL</code> which allows you to obtain the channel during routing and use it as well.	false	MEDIUM

Name	Description	Default	Priority
<code>camel.component.netty.sync</code>	Setting to set endpoint as one-way or request-response	true	MEDIUM
<code>camel.component.netty.tcpNoDelay</code>	Setting to improve TCP protocol performance	true	MEDIUM
<code>camel.component.netty.bridgeErrorHandler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<code>camel.component.netty.broadcast</code>	Setting to choose Multicast over UDP	false	MEDIUM
<code>camel.component.netty.clientMode</code>	If the <code>clientMode</code> is true, netty consumer will connect the address as a TCP client.	false	MEDIUM
<code>camel.component.netty.reconnect</code>	Used only in <code>clientMode</code> in consumer, the consumer will attempt to reconnect on disconnection if this is enabled	true	MEDIUM
<code>camel.component.netty.reconnectInterval</code>	Used if <code>reconnect</code> and <code>clientMode</code> is enabled. The interval in milli seconds to attempt reconnection	10000	MEDIUM
<code>camel.component.netty.backlog</code>	Allows to configure a backlog for netty consumer (server). Note the backlog is just a best effort depending on the OS. Setting this option to a value such as 200, 500 or 1000, tells the TCP stack how long the accept queue can be. If this option is not configured, then the backlog depends on OS setting.	null	MEDIUM
<code>camel.component.netty.bossCount</code>	When netty works on nio mode, it uses default <code>bossCount</code> parameter from Netty, which is 1. User can use this option to override the default <code>bossCount</code> from Netty	1	MEDIUM
<code>camel.component.netty.bossGroup</code>	Set the <code>BossGroup</code> which could be used for handling the new connection of the server side across the <code>NettyEndpoint</code>	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.netty.disconnectOnNoReply</code>	If sync is enabled then this option dictates NettyConsumer if it should disconnect where there is no reply to send back.	true	MEDIUM
<code>camel.component.netty.executorService</code>	To use the given EventExecutorGroup.	null	MEDIUM
<code>camel.component.netty.maximumPoolSize</code>	Sets a maximum thread pool size for the netty consumer ordered thread pool. The default size is 2 x cpu_core plus 1. Setting this value to eg 10 will then use 10 threads unless 2 x cpu_core plus 1 is a higher value, which then will override and be used. For example if there are 8 cores, then the consumer thread pool will be 17. This thread pool is used to route messages received from Netty by Camel. We use a separate thread pool to ensure ordering of messages and also in case some messages will block, then netty worker threads (event loop) wont be affected.	null	MEDIUM
<code>camel.component.netty.nettyServerBootstrapFactory</code>	To use a custom NettyServerBootstrapFactory	null	MEDIUM
<code>camel.component.netty.networkInterface</code>	When using UDP then this option can be used to specify a network interface by its name, such as eth0 to join a multicast group.	null	MEDIUM
<code>camel.component.netty.noReplyLogLevel</code>	If sync is enabled this option dictates NettyConsumer which logging level to use when logging a there is no reply to send back. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"WARN"	MEDIUM
<code>camel.component.netty.serverClosedChannelExceptionCaughtLogLevel</code>	If the server (NettyConsumer) catches an java.nio.channels.ClosedChannelException then its logged using this logging level. This is used to avoid logging the closed channel exceptions, as clients can disconnect abruptly and then cause a flood of closed exceptions in the Netty server. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"DEBUG"	MEDIUM
<code>camel.component.netty.serverExceptionCaughtLog Level</code>	If the server (NettyConsumer) catches an exception then its logged using this logging level. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"WARN"	MEDIUM
<code>camel.component.netty.serverInitializerFactory</code>	To use a custom ServerInitializerFactory	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.netty.usingExecutorService</code>	Whether to use ordered thread pool, to ensure events are processed orderly on the same channel.	true	MEDIUM
<code>camel.component.netty.allowSerializedHeaders</code>	Only used for TCP when <code>transferExchange</code> is true. When set to true, serializable objects in headers and properties will be added to the exchange. Otherwise Camel will exclude any non-serializable objects and log it at WARN level.	false	MEDIUM
<code>camel.component.netty.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.component.netty.channelGroup</code>	To use a explicit <code>ChannelGroup</code> .	null	MEDIUM
<code>camel.component.netty.nativeTransport</code>	Whether to use native transport instead of NIO. Native transport takes advantage of the host operating system and is only supported on some platforms. You need to add the netty JAR for the host operating system you are using. See more details at: <a href="http://netty.io/wiki/native-transport.html">http://netty.io/wiki/native-transport.html</a>	false	MEDIUM
<code>camel.component.netty.options</code>	Allows to configure additional netty options using option. as prefix. For example <code>option.child.keepAlive=false</code> to set the netty option <code>child.keepAlive=false</code> . See the Netty documentation for possible options that can be used.	null	MEDIUM
<code>camel.component.netty.receiveBufferSize</code>	The TCP/UDP buffer sizes to be used during inbound communication. Size is bytes.	65536	MEDIUM
<code>camel.component.netty.receiveBufferSizePredictor</code>	Configures the buffer size predictor. See details at Jetty documentation and this mail thread.	null	MEDIUM
<code>camel.component.netty.sendBufferSize</code>	The TCP/UDP buffer sizes to be used during outbound communication. Size is bytes.	65536	MEDIUM

Name	Description	Default	Priority
<b>camel.component.netty.transferExchange</b>	Only used for TCP. You can transfer the exchange over the wire instead of just the body. The following fields are transferred: In body, Out body, fault body, In headers, Out headers, fault headers, exchange properties, exchange exception. This requires that the objects are serializable. Camel will exclude any non-serializable objects and log it at WARN level.	false	MEDIUM
<b>camel.component.netty.udpByteArrayCodec</b>	For UDP only. If enabled the using byte array codec instead of Java serialization protocol.	false	MEDIUM
<b>camel.component.netty.workerCount</b>	When netty works on nio mode, it uses default workerCount parameter from Netty (which is <code>cpu_core_threads x 2</code> ). User can use this option to override the default workerCount from Netty.	null	MEDIUM
<b>camel.component.netty.workerGroup</b>	To use a explicit EventLoopGroup as the boss thread pool. For example to share a thread pool with multiple consumers or producers. By default each consumer or producer has their own worker pool with <code>2 x cpu count core threads</code> .	null	MEDIUM
<b>camel.component.netty.allowDefaultCodec</b>	The netty component installs a default codec if both, encoder/decoder is null and textline is false. Setting <code>allowDefaultCodec</code> to false prevents the netty component from installing a default codec as the first element in the filter chain.	true	MEDIUM
<b>camel.component.netty.autoAppendDelimiter</b>	Whether or not to auto append missing end delimiter when sending using the textline codec.	true	MEDIUM
<b>camel.component.netty.decoderMaxLineLength</b>	The max line length to use for the textline codec.	1024	MEDIUM
<b>camel.component.netty.decoders</b>	A list of decoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.	null	MEDIUM
<b>camel.component.netty.delimiter</b>	The delimiter to use for the textline codec. Possible values are LINE and NULL. One of: [LINE] [NULL]	"LINE"	MEDIUM

Name	Description	Default	Priority
<code>camel.component.netty.encoders</code>	A list of encoders to be used. You can use a String which have values separated by comma, and have the values be looked up in the Registry. Just remember to prefix the value with # so Camel knows it should lookup.	null	MEDIUM
<code>camel.component.netty.encoding</code>	The encoding (a charset name) to use for the textline codec. If not provided, Camel will use the JVM default Charset.	null	MEDIUM
<code>camel.component.netty.textline</code>	Only used for TCP. If no codec is specified, you can use this flag to indicate a text line based codec; if not specified or the value is false, then Object Serialization is assumed over TCP - however only Strings are allowed to be serialized by default.	false	MEDIUM
<code>camel.component.netty.enabledProtocols</code>	Which protocols to enable when using SSL	"TLSv1,TLSv1.1,TLSv1.2"	MEDIUM
<code>camel.component.netty.keyStoreFile</code>	Client side certificate keystore to be used for encryption	null	MEDIUM
<code>camel.component.netty.keyStoreFormat</code>	Keystore format to be used for payload encryption. Defaults to JKS if not set	null	MEDIUM
<code>camel.component.netty.keyStoreResource</code>	Client side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.	null	MEDIUM
<code>camel.component.netty.needClientAuth</code>	Configures whether the server needs client authentication when using SSL.	false	MEDIUM
<code>camel.component.netty.passphrase</code>	Password setting to use in order to encrypt/decrypt payloads sent using SSH	null	MEDIUM
<code>camel.component.netty.securityProvider</code>	Security provider to be used for payload encryption. Defaults to SunX509 if not set.	null	MEDIUM
<code>camel.component.netty.ssl</code>	Setting to specify whether SSL encryption is applied to this endpoint	false	MEDIUM

Name	Description	Default	Priority
<code>camel.component.netty.sslClientCertHeaders</code>	When enabled and in SSL mode, then the Netty consumer will enrich the Camel Message with headers having information about the client certificate such as subject name, issuer name, serial number, and the valid date range.	false	MEDIUM
<code>camel.component.netty.sslContextParameters</code>	To configure security using SSLContextParameters	null	MEDIUM
<code>camel.component.netty.sslHandler</code>	Reference to a class that could be used to return an SSL Handler	null	MEDIUM
<code>camel.component.netty.trustStoreFile</code>	Server side certificate keystore to be used for encryption	null	MEDIUM
<code>camel.component.netty.trustStoreResource</code>	Server side certificate keystore to be used for encryption. Is loaded by default from classpath, but you can prefix with classpath:, file:, or http: to load the resource from different systems.	null	MEDIUM
<code>camel.component.netty.useGlobalSslContextParameters</code>	Enable usage of global SSL context parameters.	false	MEDIUM

The camel-netty sink connector has no converters out of the box.

The camel-netty sink connector has no transforms out of the box.

The camel-netty sink connector has no aggregation strategies out of the box.

## 5.20. CAMEL-SALESFORCE-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-salesforce-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-salesforce-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.salesforce.CamelSalesforceSourceConnector
```

The camel-salesforce source connector supports 116 options, which are listed below.

Name	Description	Default	Priority
<code>camel.source.path.topicName</code>	The name of the topic/channel to use	null	MEDIUM
<code>camel.source.endpoint.apexMethod</code>	APEX method name	null	MEDIUM
<code>camel.source.endpoint.apexQueryParams</code>	Query params for APEX method	null	MEDIUM
<code>camel.source.endpoint.apexUrl</code>	APEX method URL	null	MEDIUM
<code>camel.source.endpoint.apiVersion</code>	Salesforce API version.	"34.0"	MEDIUM
<code>camel.source.endpoint.backoffIncrement</code>	Backoff interval increment for Streaming connection restart attempts for failures beyond CometD auto-reconnect.	1000L	MEDIUM
<code>camel.source.endpoint.batchId</code>	Bulk API Batch ID	null	MEDIUM
<code>camel.source.endpoint.contentType</code>	Bulk API content type, one of XML, CSV, ZIP_XML, ZIP_CSV One of: [XML] [CSV] [JSON] [ZIP_XML] [ZIP_CSV] [ZIP_JSON]	null	MEDIUM
<code>camel.source.endpoint.defaultReplayId</code>	Default replayId setting if no value is found in initialReplayIdMap	null	MEDIUM
<code>camel.source.endpoint.format</code>	Payload format to use for Salesforce API calls, either JSON or XML, defaults to JSON One of: [JSON] [XML]	null	MEDIUM
<code>camel.source.endpoint.httpClient</code>	Custom Jetty Http Client to use to connect to Salesforce.	null	MEDIUM
<code>camel.source.endpoint.includeDetails</code>	Include details in Salesforce1 Analytics report, defaults to false.	null	MEDIUM
<code>camel.source.endpoint.initialReplayIdMap</code>	Replay IDs to start from per channel name.	null	MEDIUM
<code>camel.source.endpoint.instanceId</code>	Salesforce1 Analytics report execution instance ID	null	MEDIUM
<code>camel.source.endpoint.jobId</code>	Bulk API Job ID	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.limit</b>	Limit on number of returned records. Applicable to some of the API, check the Salesforce documentation.	null	MEDIUM
<b>camel.source.endpoint.maxBackoff</b>	Maximum backoff interval for Streaming connection restart attempts for failures beyond CometD auto-reconnect.	30000L	MEDIUM
<b>camel.source.endpoint.notFoundBehaviour</b>	Sets the behaviour of 404 not found status received from Salesforce API. Should the body be set to NULL NotFoundBehaviour#NULL or should a exception be signaled on the exchange NotFoundBehaviour#EXCEPTION - the default. One of: [EXCEPTION] [NULL]	"EXCEPTION"	MEDIUM
<b>camel.source.endpoint.notifyForFields</b>	Notify for fields, options are ALL, REFERENCED, SELECT, WHERE One of: [ALL] [REFERENCED] [SELECT] [WHERE]	null	MEDIUM
<b>camel.source.endpoint.notifyForOperationCreate</b>	Notify for create operation, defaults to false (API version = 29.0)	null	MEDIUM
<b>camel.source.endpoint.notifyForOperationDelete</b>	Notify for delete operation, defaults to false (API version = 29.0)	null	MEDIUM
<b>camel.source.endpoint.notifyForOperations</b>	Notify for operations, options are ALL, CREATE, EXTENDED, UPDATE (API version 29.0) One of: [ALL] [CREATE] [EXTENDED] [UPDATE]	null	MEDIUM
<b>camel.source.endpoint.notifyForOperationUndelete</b>	Notify for un-delete operation, defaults to false (API version = 29.0)	null	MEDIUM
<b>camel.source.endpoint.notifyForOperationUpdate</b>	Notify for update operation, defaults to false (API version = 29.0)	null	MEDIUM
<b>camel.source.endpoint.objectMapper</b>	Custom Jackson ObjectMapper to use when serializing/deserializing Salesforce objects.	null	MEDIUM
<b>camel.source.endpoint.rawPayload</b>	Use raw payload String for request and response (either JSON or XML depending on format), instead of DTOs, false by default	false	MEDIUM

Name	Description	Default	Priority
<code>camel.source.endpoint.reportId</code>	Salesforce Analytics report Id	null	MEDIUM
<code>camel.source.endpoint.reportMetadata</code>	Salesforce Analytics report metadata for filtering	null	MEDIUM
<code>camel.source.endpoint.resultId</code>	Bulk API Result ID	null	MEDIUM
<code>camel.source.endpoint.sObjectBlobFieldName</code>	SObject blob field name	null	MEDIUM
<code>camel.source.endpoint.sObjectClass</code>	Fully qualified SObject class name, usually generated using camel-salesforce-maven-plugin	null	MEDIUM
<code>camel.source.endpoint.sObjectFields</code>	SObject fields to retrieve	null	MEDIUM
<code>camel.source.endpoint.sObjectId</code>	SObject ID if required by API	null	MEDIUM
<code>camel.source.endpoint.sObjectIdName</code>	SObject external ID field name	null	MEDIUM
<code>camel.source.endpoint.sObjectIdValue</code>	SObject external ID field value	null	MEDIUM
<code>camel.source.endpoint.sObjectName</code>	SObject name if required or supported by API	null	MEDIUM
<code>camel.source.endpoint.sObjectQuery</code>	Salesforce SOQL query string	null	MEDIUM
<code>camel.source.endpoint.sObjectSearch</code>	Salesforce SOSL search string	null	MEDIUM
<code>camel.source.endpoint.updateTopic</code>	Whether to update an existing Push Topic when using the Streaming API, defaults to false	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.source.endpoint.replayId</b>	The replayId value to use when subscribing	null	MEDIUM
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom ExceptionHandler. Notice if the option bridgeErrorHandler is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: [InOnly] [InOut] [InOptionalOut]	null	MEDIUM
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.component.salesforce.apexMethod</b>	APEX method name	null	MEDIUM
<b>camel.component.salesforce.apexQueryParameters</b>	Query params for APEX method	null	MEDIUM
<b>camel.component.salesforce.apexUrl</b>	APEX method URL	null	MEDIUM
<b>camel.component.salesforce.apiVersion</b>	Salesforce API version.	"34.0"	MEDIUM
<b>camel.component.salesforce.backoffIncrement</b>	Backoff interval increment for Streaming connection restart attempts for failures beyond CometD auto-reconnect.	1000L	MEDIUM

Name	Description	Default	Priority
<code>camel.component.salesforce.batchId</code>	Bulk API Batch ID	null	MEDIUM
<code>camel.component.salesforce.contentType</code>	Bulk API content type, one of XML, CSV, ZIP_XML, ZIP_CSV One of: [XML] [CSV] [JSON] [ZIP_XML] [ZIP_CSV] [ZIP_JSON]	null	MEDIUM
<code>camel.component.salesforce.defaultReplayId</code>	Default replayId setting if no value is found in initialReplayIdMap	null	MEDIUM
<code>camel.component.salesforce.format</code>	Payload format to use for Salesforce API calls, either JSON or XML, defaults to JSON One of: [JSON] [XML]	null	MEDIUM
<code>camel.component.salesforce.httpClient</code>	Custom Jetty Http Client to use to connect to Salesforce.	null	MEDIUM
<code>camel.component.salesforce.httpClientConnectionTimeout</code>	Connection timeout used by the HttpClient when connecting to the Salesforce server.	60000L	MEDIUM
<code>camel.component.salesforce.httpClientIdleTimeout</code>	Timeout used by the HttpClient when waiting for response from the Salesforce server.	10000L	MEDIUM
<code>camel.component.salesforce.httpClientMaxContentLength</code>	Max content length of an HTTP response.	null	MEDIUM
<code>camel.component.salesforce.includeDetails</code>	Include details in Salesforce1 Analytics report, defaults to false.	null	MEDIUM
<code>camel.component.salesforce.initialReplayIdMap</code>	Replay IDs to start from per channel name.	null	MEDIUM
<code>camel.component.salesforce.instanceId</code>	Salesforce1 Analytics report execution instance ID	null	MEDIUM
<code>camel.component.salesforce.jobId</code>	Bulk API Job ID	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.salesforce.limit</code>	Limit on number of returned records. Applicable to some of the API, check the Salesforce documentation.	null	MEDIUM
<code>camel.component.salesforce.maxBackoff</code>	Maximum backoff interval for Streaming connection restart attempts for failures beyond CometD auto-reconnect.	30000L	MEDIUM
<code>camel.component.salesforce.notFoundBehaviour</code>	Sets the behaviour of 404 not found status received from Salesforce API. Should the body be set to NULL NotFoundBehaviour#NULL or should a exception be signaled on the exchange NotFoundBehaviour#EXCEPTION - the default. One of: [EXCEPTION] [NULL]	"EXCEPTION"	MEDIUM
<code>camel.component.salesforce.notifyForFields</code>	Notify for fields, options are ALL, REFERENCED, SELECT, WHERE One of: [ALL] [REFERENCED] [SELECT] [WHERE]	null	MEDIUM
<code>camel.component.salesforce.notifyForOperation Create</code>	Notify for create operation, defaults to false (API version = 29.0)	null	MEDIUM
<code>camel.component.salesforce.notifyForOperation Delete</code>	Notify for delete operation, defaults to false (API version = 29.0)	null	MEDIUM
<code>camel.component.salesforce.notifyForOperations</code>	Notify for operations, options are ALL, CREATE, EXTENDED, UPDATE (API version 29.0) One of: [ALL] [CREATE] [EXTENDED] [UPDATE]	null	MEDIUM
<code>camel.component.salesforce.notifyForOperation Undelete</code>	Notify for un-delete operation, defaults to false (API version = 29.0)	null	MEDIUM
<code>camel.component.salesforce.notifyForOperation Update</code>	Notify for update operation, defaults to false (API version = 29.0)	null	MEDIUM
<code>camel.component.salesforce.objectMapper</code>	Custom Jackson ObjectMapper to use when serializing/deserializing Salesforce objects.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.salesforce.packages</code>	In what packages are the generated DTO classes. Typically the classes would be generated using camel-salesforce-maven-plugin. Set it if using the generated DTOs to gain the benefit of using short SObject names in parameters/header values. Multiple packages can be separated by comma.	null	MEDIUM
<code>camel.component.salesforce.rawPayload</code>	Use raw payload String for request and response (either JSON or XML depending on format), instead of DTOs, false by default	false	MEDIUM
<code>camel.component.salesforce.reportId</code>	Salesforce Analytics report Id	null	MEDIUM
<code>camel.component.salesforce.reportMetadata</code>	Salesforce Analytics report metadata for filtering	null	MEDIUM
<code>camel.component.salesforce.resultId</code>	Bulk API Result ID	null	MEDIUM
<code>camel.component.salesforce.sObjectBlobFieldName</code>	SObject blob field name	null	MEDIUM
<code>camel.component.salesforce.sObjectClass</code>	Fully qualified SObject class name, usually generated using camel-salesforce-maven-plugin	null	MEDIUM
<code>camel.component.salesforce.sObjectFields</code>	SObject fields to retrieve	null	MEDIUM
<code>camel.component.salesforce.sObjectId</code>	SObject ID if required by API	null	MEDIUM
<code>camel.component.salesforce.sObjectIdName</code>	SObject external ID field name	null	MEDIUM
<code>camel.component.salesforce.sObjectIdValue</code>	SObject external ID field value	null	MEDIUM
<code>camel.component.salesforce.sObjectName</code>	SObject name if required or supported by API	null	MEDIUM
<code>camel.component.salesforce.sObjectQuery</code>	Salesforce SOQL query string	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.salesforce.sObjectSearch</code>	Salesforce SOSL search string	null	MEDIUM
<code>camel.component.salesforce.updateTopic</code>	Whether to update an existing Push Topic when using the Streaming API, defaults to false	false	MEDIUM
<code>camel.component.salesforce.config</code>	Global endpoint configuration - use to set values that are common to all endpoints	null	MEDIUM
<code>camel.component.salesforce.httpClientProperties</code>	Used to set any properties that can be configured on the underlying HTTP client. Have a look at properties of <code>SalesforceHttpClient</code> and the <code>Jetty HttpClient</code> for all available options.	null	MEDIUM
<code>camel.component.salesforce.longPollingTransportProperties</code>	Used to set any properties that can be configured on the <code>LongPollingTransport</code> used by the <code>BayeuxClient</code> (CometD) used by the streaming api	null	MEDIUM
<code>camel.component.salesforce.bridgeErrorHandler</code>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<code>camel.component.salesforce.basicPropertyBinding</code>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<code>camel.component.salesforce.httpProxyExcludedAddresses</code>	A list of addresses for which HTTP proxy server should not be used.	null	MEDIUM
<code>camel.component.salesforce.httpProxyHost</code>	Hostname of the HTTP proxy server to use.	null	MEDIUM
<code>camel.component.salesforce.httpProxyIncludedAddresses</code>	A list of addresses for which HTTP proxy server should be used.	null	MEDIUM
<code>camel.component.salesforce.httpProxyPort</code>	Port number of the HTTP proxy server to use.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.salesforce.httpProxySocks4</code>	If set to true the configures the HTTP proxy to use as a SOCKS4 proxy.	false	MEDIUM
<code>camel.component.salesforce.authenticationType</code>	Explicit authentication method to be used, one of USERNAME_PASSWORD, REFRESH_TOKEN or JWT. Salesforce component can auto-determine the authentication method to use from the properties set, set this property to eliminate any ambiguity. One of: [USERNAME_PASSWORD] [REFRESH_TOKEN] [JWT]	null	MEDIUM
<code>camel.component.salesforce.clientId</code>	OAuth Consumer Key of the connected app configured in the Salesforce instance setup. Typically a connected app needs to be configured but one can be provided by installing a package.	null	HIGH
<code>camel.component.salesforce.clientSecret</code>	OAuth Consumer Secret of the connected app configured in the Salesforce instance setup.	null	MEDIUM
<code>camel.component.salesforce.httpProxyAuthUri</code>	Used in authentication against the HTTP proxy server, needs to match the URI of the proxy server in order for the <code>httpProxyUsername</code> and <code>httpProxyPassword</code> to be used for authentication.	null	MEDIUM
<code>camel.component.salesforce.httpProxyPassword</code>	Password to use to authenticate against the HTTP proxy server.	null	MEDIUM
<code>camel.component.salesforce.httpProxyRealm</code>	Realm of the proxy server, used in preemptive Basic/Digest authentication methods against the HTTP proxy server.	null	MEDIUM
<code>camel.component.salesforce.httpProxySecure</code>	If set to false disables the use of TLS when accessing the HTTP proxy.	true	MEDIUM
<code>camel.component.salesforce.httpProxyUseDigestAuth</code>	If set to true Digest authentication will be used when authenticating to the HTTP proxy, otherwise Basic authorization method will be used	false	MEDIUM
<code>camel.component.salesforce.httpProxyUsername</code>	Username to use to authenticate against the HTTP proxy server.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.component.salesforce.instanceUrl</b>	URL of the Salesforce instance used after authentication, by default received from Salesforce on successful authentication	null	MEDIUM
<b>camel.component.salesforce.keystore</b>	KeyStore parameters to use in OAuth JWT flow. The KeyStore should contain only one entry with private key and certificate. Salesforce does not verify the certificate chain, so this can easily be a selfsigned certificate. Make sure that you upload the certificate to the corresponding connected app.	null	MEDIUM
<b>camel.component.salesforce.lazyLogin</b>	If set to true prevents the component from authenticating to Salesforce with the start of the component. You would generally set this to the (default) false and authenticate early and be immediately aware of any authentication issues.	false	MEDIUM
<b>camel.component.salesforce.loginConfig</b>	All authentication configuration in one nested bean, all properties set there can be set directly on the component as well	null	MEDIUM
<b>camel.component.salesforce.loginUrl</b>	URL of the Salesforce instance used for authentication, by default set to <a href="https://login.salesforce.com">https://login.salesforce.com</a>	"https://login.salesforce.com"	HIGH
<b>camel.component.salesforce.password</b>	Password used in OAuth flow to gain access to access token. It's easy to get started with password OAuth flow, but in general one should avoid it as it is deemed less secure than other flows. Make sure that you append security token to the end of the password if using one.	null	MEDIUM
<b>camel.component.salesforce.refreshToken</b>	Refresh token already obtained in the refresh token OAuth flow. One needs to setup a web application and configure a callback URL to receive the refresh token, or configure using the builtin callback at <a href="https://login.salesforce.com/services/oauth2/success">https://login.salesforce.com/services/oauth2/success</a> or <a href="https://test.salesforce.com/services/oauth2/success">https://test.salesforce.com/services/oauth2/success</a> and then retrieve the refresh_token from the URL at the end of the flow. Note that in development organizations Salesforce allows hosting the callback web application at localhost.	null	MEDIUM

Name	Description	Default	Priority
<code>camel.component.salesforce.sslContextParameters</code>	SSL parameters to use, see <code>SSLContextParameters</code> class for all available options.	null	MEDIUM
<code>camel.component.salesforce.useGlobalSslContextParameters</code>	Enable usage of global SSL context parameters	false	MEDIUM
<code>camel.component.salesforce.userName</code>	Username used in OAuth flow to gain access to access token. It's easy to get started with password OAuth flow, but in general one should avoid it as it is deemed less secure than other flows.	null	MEDIUM

The camel-salesforce sink connector has no converters out of the box.

The camel-salesforce sink connector has no transforms out of the box.

The camel-salesforce sink connector has no aggregation strategies out of the box.

## 5.21. CAMEL-SLACK-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-slack-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```
<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-slack-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>
```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.slack.CamelSlackSourceConnector
```

The camel-slack source connector supports 28 options, which are listed below.

Name	Description	Default	Priority
<code>camel.source.path.channel</code>	The channel name (syntax <code>#name</code> ) or slackuser (syntax <code>userName</code> ) to send a message directly to an user.	null	HIGH

Name	Description	Default	Priority
<b>camel.source.endpoint.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.source.endpoint.maxResults</b>	The Max Result for the poll	"10"	MEDIUM
<b>camel.source.endpoint.sendEmptyMessageWhenIdle</b>	If the polling consumer did not poll any files, you can enable this option to send an empty message (no body) instead.	false	MEDIUM
<b>camel.source.endpoint.serverUrl</b>	The Server URL of the Slack instance	"https://slack.com"	MEDIUM
<b>camel.source.endpoint.token</b>	The token to use	null	MEDIUM
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: <code>[InOnly]</code> <code>[InOut]</code> <code>[InOptionalOut]</code>	null	MEDIUM
<b>camel.source.endpoint.pollStrategy</b>	A pluggable <code>org.apache.camel.PollingConsumerPollingStrategy</code> allowing you to provide your custom implementation to control error handling usually occurred during the poll operation before an Exchange have been created and being routed in Camel.	null	MEDIUM
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.source.endpoint.backoffErrorThreshold</b>	The number of subsequent error polls (failed due some error) that should happen before the backoffMultiplier should kick-in.	null	MEDIUM
<b>camel.source.endpoint.backoffIdleThreshold</b>	The number of subsequent idle polls that should happen before the backoffMultiplier should kick-in.	null	MEDIUM
<b>camel.source.endpoint.backoffMultiplier</b>	To let the scheduled polling consumer backoff if there has been a number of subsequent idles/errors in a row. The multiplier is then the number of polls that will be skipped before the next actual attempt is happening again. When this option is in use then backoffIdleThreshold and/or backoffErrorThreshold must also be configured.	null	MEDIUM
<b>camel.source.endpoint.delay</b>	Milliseconds before the next poll.	500L	MEDIUM
<b>camel.source.endpoint.greedy</b>	If greedy is enabled, then the ScheduledPollConsumer will run immediately again, if the previous run polled 1 or more messages.	false	MEDIUM
<b>camel.source.endpoint.initialDelay</b>	Milliseconds before the first poll starts.	1000L	MEDIUM
<b>camel.source.endpoint.repeatCount</b>	Specifies a maximum limit of number of fires. So if you set it to 1, the scheduler will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0L	MEDIUM
<b>camel.source.endpoint.runLoggingLevel</b>	The consumer logs a start/complete log line when it polls. This option allows you to configure the logging level for that. One of: [TRACE] [DEBUG] [INFO] [WARN] [ERROR] [OFF]	"TRACE"	MEDIUM
<b>camel.source.endpoint.scheduledExecutorService</b>	Allows for configuring a custom/shared thread pool to use for the consumer. By default each consumer has its own single threaded thread pool.	null	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.scheduler</b>	To use a cron scheduler from either camel-spring or camel-quartz component. Use value spring or quartz for built in scheduler	"none"	MEDIUM
<b>camel.source.endpoint.schedulerProperties</b>	To configure additional properties when using a custom scheduler or any of the Quartz, Spring based scheduler.	null	MEDIUM
<b>camel.source.endpoint.startScheduler</b>	Whether the scheduler should be auto started.	true	MEDIUM
<b>camel.source.endpoint.timeUnit</b>	Time unit for initialDelay and delay options. One of: [NANOSECONDS] [MICROSECONDS] [MILLISECONDS] [SECONDS] [MINUTES] [HOURS] [DAYS]	"MILLISECONDS"	MEDIUM
<b>camel.source.endpoint.useFixedDelay</b>	Controls if fixed delay or fixed rate is used. See ScheduledExecutorService in JDK for details.	true	MEDIUM
<b>camel.component.slack.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.component.slack.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.component.slack.webhookUrl</b>	The incoming webhook URL	null	MEDIUM

The camel-slack sink connector has no converters out of the box.

The camel-slack sink connector has no transforms out of the box.

The camel-slack sink connector has no aggregation strategies out of the box.

## 5.22. CAMEL-SYSLOG-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-syslog-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```

<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-syslog-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>

```

The camel-syslog source connector supports is based on camel-netty source connector and supports all its options ; however has been already preconfigured and should be sufficient to provide the following properties:

Name	Description	Default	Priority
<b>camel.source.path.protocol</b>	The protocol to use which can be tcp or udp. One of: [tcp] [udp]	null	HIGH
<b>camel.source.path.host</b>	The hostname. For the consumer the hostname is localhost or 0.0.0.0. For the producer the hostname is the remote host to connect to	null	HIGH
<b>camel.source.path.port</b>	The host port number	null	HIGH

## 5.23. CAMEL-TIMER-KAFKA-CONNECTOR SOURCE CONFIGURATION

When using camel-timer-kafka-connector as source make sure to use the following Maven dependency to have support for the connector:

```

<dependency>
  <groupId>org.apache.camel.kafkaconnector</groupId>
  <artifactId>camel-timer-kafka-connector</artifactId>
  <version>x.x.x</version>
  <!-- use the same version as your Camel Kafka connector version -->
</dependency>

```

To use this Source connector in Kafka connect you'll need to set the following connector.class

```
connector.class=org.apache.camel.kafkaconnector.timer.CamelTimerSourceConnector
```

The camel-timer source connector supports 17 options, which are listed below.

Name	Description	Default	Priority
<b>camel.source.path.timerName</b>	The name of the timer	null	HIGH

Name	Description	Default	Priority
<b>camel.source.endpoint.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the <code>org.apache.camel.spi.ExceptionHandler</code> to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.source.endpoint.delay</b>	Delay before first event is triggered.	1000L	MEDIUM
<b>camel.source.endpoint.fixedRate</b>	Events take place at approximately regular intervals, separated by the specified period.	false	MEDIUM
<b>camel.source.endpoint.includeMetadata</b>	Whether to include metadata in the exchange such as fired time, timer name, timer count etc. This information is default included.	true	MEDIUM
<b>camel.source.endpoint.period</b>	If greater than 0, generate periodic events every period.	1000L	MEDIUM
<b>camel.source.endpoint.repeatCount</b>	Specifies a maximum limit of number of fires. So if you set it to 1, the timer will only fire once. If you set it to 5, it will only fire five times. A value of zero or negative means fire forever.	0L	MEDIUM
<b>camel.source.endpoint.exceptionHandler</b>	To let the consumer use a custom <code>ExceptionHandler</code> . Notice if the option <code>bridgeErrorHandler</code> is enabled then this option is not in use. By default the consumer will deal with exceptions, that will be logged at WARN or ERROR level and ignored.	null	MEDIUM
<b>camel.source.endpoint.exchangePattern</b>	Sets the exchange pattern when the consumer creates an exchange. One of: <code>[InOnly]</code> <code>[InOut]</code> <code>[InOptionalOut]</code>	null	MEDIUM
<b>camel.source.endpoint.basicPropertyBinding</b>	Whether the endpoint should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM
<b>camel.source.endpoint.daemon</b>	Specifies whether or not the thread associated with the timer endpoint runs as a daemon. The default value is true.	true	MEDIUM

Name	Description	Default	Priority
<b>camel.source.endpoint.pattern</b>	Allows you to specify a custom Date pattern to use for setting the time option using URI syntax.	null	MEDIUM
<b>camel.source.endpoint.synchronous</b>	Sets whether synchronous processing should be strictly used, or Camel is allowed to use asynchronous processing (if supported).	false	MEDIUM
<b>camel.source.endpoint.time</b>	A java.util.Date the first event should be generated. If using the URI, the pattern expected is: yyyy-MM-dd HH:mm:ss or yyyy-MM-dd'T'HH:mm:ss.	null	MEDIUM
<b>camel.source.endpoint.timer</b>	To use a custom Timer	null	MEDIUM
<b>camel.component.timer.bridgeErrorHandler</b>	Allows for bridging the consumer to the Camel routing Error Handler, which mean any exceptions occurred while the consumer is trying to pickup incoming messages, or the likes, will now be processed as a message and handled by the routing Error Handler. By default the consumer will use the org.apache.camel.spi.ExceptionHandler to deal with exceptions, that will be logged at WARN or ERROR level and ignored.	false	MEDIUM
<b>camel.component.timer.basicPropertyBinding</b>	Whether the component should use basic property binding (Camel 2.x) or the newer property binding with additional capabilities	false	MEDIUM

The camel-timer sink connector has no converters out of the box.

The camel-timer sink connector has no transforms out of the box.

The camel-timer sink connector has no aggregation strategies out of the box.