# Oracle® Cloud

# Using the Confluent Adapter with Oracle Integration 3

F87938-07
February 2025

Oracle Cloud Using the Confluent Adapter with Oracle Integration 3,

F87938-07



Primary Author: Oracle Corporation

# Contents

# Preface

This guide describes how to configure this adapter as a connection in an integration in Oracle Integration.

> **Note:**
>
> The use of this adapter may differ depending on the features you have, or whether your instance was provisioned using Standard or Enterprise edition. These differences are noted throughout this guide.

**Topics:**



## Audience

This guide is intended for developers who want to use this adapter in integrations in Oracle Integration.

## Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at `https://www.oracle.com/corporate/accessibility/`.

**Access to Oracle Support**

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit `https://support.oracle.com/portal/` or visit `Oracle Accessibility Learning and Support` if you are hearing impaired.

## Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and

the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

# Related Resources

See these Oracle resources:

- Oracle Cloud at `http://cloud.oracle.com`
- *Using Integrations in Oracle Integration 3*
- *Using the Oracle Mapper with Oracle Integration 3*
- Oracle Integration documentation on the Oracle Help Center.

# Conventions

The following text conventions are used in this document:

| Convention | Meaning |
|---|---|
| **boldface** | Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary. |
| *italic* | Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values. |
| `monospace` | Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter. |

1

# Understand the Confluent Adapter

Review the following topics to learn about the Confluent Adapter and how to use it as a connection in integrations in Oracle Integration. A typical workflow of adapter and integration tasks is also provided.

**Topics:**



## Confluent Adapter Capabilities

The Confluent Adapter enables you to create an integration in Oracle Integration that connects to a Confluent messaging system. The Confluent Adapter connects to the Confluent distributed publish-subscribe messaging system from Oracle Integration and allows for the publishing and consumption of messages from a Confluent Kafka topic.

The Confluent Adapter provides the following capabilities:

- Establishes a connection to the Confluent messaging system to enable messages to be published and consumed.
- Consumes messages from a Confluent topic and produces messages to a Confluent topic in the invoke (outbound) direction.
- Consumes messages from a topic based on a specified frequency in the trigger (inbound) direction.
- Enables you to browse the available metadata using the Adapter Endpoint Configuration Wizard (that is, the topics and partitions to which messages are published and consumed).
- Reads data from topics in the Kafka cluster using the Apache Kafka Streaming API.
- Supports a consumer group.
- Supports headers.
- Supports the following message structures:
  - Avro schema document
  - XML schema (XSD) and schema archive upload
  - Sample XML
    These schemas are applicable for the following scenarios:
    * Producing and consuming messages - Invoke connections (supported with the connectivity agent only)
    * Consuming messages - Trigger connections (supported with the connectivity agent only)

- Works only with the connectivity agent. Direct connectivity and connectivity to private endpoints is not supported.
- Supports the following security policies:
  - SASL Plain over SSL
  - SASL Plain
- Supports connectivity to a Confluent messaging system over SSL through the connectivity agent.
- Supports optionally configuring the Confluent Kafka producer to be transactional. This enables an application to send messages to multiple partitions atomically.

You can configure the Confluent Adapter as an invoke connection in an integration in Oracle Integration. The Confluent Adapter is one of many predefined adapters included with Oracle Integration. See the Adapters page in the Oracle Help Center.

# Confluent Adapter Restrictions

Note the following Confluent Adapter restrictions.

- Only GZIP compression (the default) is supported. Other forms of compression are not supported.
- The Confluent Adapter can only be used with the on-premises connectivity agent.

> **Note:**
>
> There are overall service limits for Oracle Integration. A service limit is the quota or allowance set on a resource. See Service Limits.

# What Application Version Is Supported?

For information about which application version is supported by this adapter, see the Connectivity Certification Matrix.

# Workflow to Create and Add a Confluent Adapter Connection to an Integration

You follow a very simple workflow to create a connection with an adapter and include the connection in an integration in Oracle Integration.

This table lists the workflow steps for both adapter tasks and overall integration tasks, and provides links to instructions for each step.

| Step | Description | More Information |
|---|---|---|
| 1 | Access Oracle Integration. | Go to your login URL. |
| 2 | Create the adapter connections for the applications you want to integrate. The connections can be reused in multiple integrations and are typically created by the administrator. | Create a Confluent Adapter Connection |

| Step | Description | More Information |
|---|---|---|
| 3 | Create the integration. When you do this, you add trigger (source) and invoke (target) connections to the integration. | Understand Integration Creation and Best Practices in *Using Integrations in Oracle Integration 3* and Add the Confluent Adapter Connection to an Integration |
| 4 | Map data between the trigger connection data structure and the invoke connection data structure. | Map Data in *Using Integrations in Oracle Integration 3* |
| 5 | (Optional) Create lookups that map the different values used by those applications to identify the same type of object (such as gender codes or country codes). | Manage Lookups in *Using Integrations in Oracle Integration 3* |
| 6 | Activate the integration. | Activate an Integration in *Using Integrations in Oracle Integration 3* |
| 7 | Monitor the integration on the dashboard. | Monitor Integrations During Runtime in *Using Integrations in Oracle Integration 3* |
| 8 | Track payload fields in messages during runtime. | Assign Business Identifiers for Tracking Fields in Messages and Track Integration Instances in *Using Integrations in Oracle Integration 3* |
| 9 | Manage errors at the integration level, connection level, or specific integration instance level. | Manage Errors in *Using Integrations in Oracle Integration 3* |

2

# Create a Confluent Adapter Connection

A connection is based on an adapter. You define connections to the specific cloud applications that you want to integrate.

**Topics:**



## Prerequisites for Creating a Connection

You must satisfy the following prerequisites to create a connection with the Confluent Adapter:



### Know the Host and Port of the Bootstrap Server

Know the host and port of the bootstrap server to use to connect to a list of Kafka brokers.

- Know the schema registry URL and the schema registry user name and password.
- Know the schema registry URL to connect to the schema registry to use the schemas.

### Obtain Security Policy Details

Obtain the following security policy details for the Confluent Adapter:

- If using the SASL Plain over SSL or SASL Plain security policy, know the SASL user name and password.
- To use the SASL Plain over SSL security policy, have the required certificates.

### Configure Confluent Kafka

To configure Confluent Kafka with the Confluent Adapter, you must obtain the following information to successfully configure the Confluent Adapter on the Connections page.

1. Generate the user name and password required for the Connections page at the following location:

   ```
   https://confluent.cloud/environments/env-dvgny/clusters/lkc-rjn91/api-keys
   ```

   a. Click **Add Key**.

   b. Enter the key in the **SASL Username** field and the secret key in the **SASL Password** field on the Connections page. See Configure Connection Security.

2. (Optional) Generate the truststore:

a. Generate the certificate.

```
 echo -n | openssl s_client -connect host:port | sed -ne  '/-BEGIN
CERTIFICATE-/,/-END CERTIFICATE-/p' > /tmp/server.cert
```

Where *host*:*port* is the combination of the bootstrap server and port.

b. Generate the truststore from the certificate created.

```
keytool -keystore conf_2.jks -alias ConfRoot
-import -file conf_server.cert -storetype JKS
```

> **Note:**
>
> Specify the `-storetype JKS` option when generating the truststore with a JDK version higher than 8.

For this example, `conf_2.jks` is the name of the truststore file to upload in the **TrustStore** field on the Connections page.

c. When prompted, enter a password. Remember the password because you must enter it in the **Truststore password** field on the Connections page.

**Copy JAR Files After Connectivity Agent Installation**

Confluent has version-specific and Kafka-compatible dependency libraries to support Avro serilizers and the schema registry. After installing the connectivity agent, copy the `kafka-schema-serializer`, `kafka-schema-registry-client`, and `kafka-avro-serializer` version 7.5.*x* JARs along with the dependent guava's `failureaccess` version 1.0.1 JAR to `agenthome/thirdparty/dependency` and restart the connectivity agent.

> **Note:**
>
> The Confluent Adapter supports Confluent Avro serializers/deserializers and Apache Kafka serializers/deserializers (String/ByteArray). It doesn't support other Confluent serializers/deserializers or any other serializers/deserializers.

# Create a Connection

Before you can build an integration, you must create the connections to the applications with which you want to share data.

> **Note:**
>
> You can also create a connection in the integration canvas. See why working with projects is preferred.

To create a connection in Oracle Integration:

1. Decide where to start:
   - Work in a project (see why working with projects is preferred).
     a. In the navigation pane, click **Projects**.
     b. Select the project name.
     c. Click **Integrations**.
     d. In the **Connections** section, click **Add** if no connections currently exist or **+** if connections already exist. The Create connection panel opens.
   - Work outside a project.
     a. In the navigation pane, click **Design**, then **Connections**.
     b. Click **Create**. The Create connection panel opens.
2. Select the adapter to use for this connection. To find the adapter, scroll through the list, or enter a partial or full name in the **Search** field.
3. Enter the information that describes this connection.

| Element | Description |
|---|---|
| **Name** | Enter a meaningful name to help others find your connection when they begin to create their own integrations. |
| **Identifier** | Automatically displays the name in capital letters that you entered in the **Name** field. If you modify the identifier name, don't include blank spaces (for example, `SALES OPPORTUNITY`). |
| **Role** | Select the role (direction) in which to use this connection.<br>**Note**: *Only* the roles supported by the adapter you selected are displayed for selection. Some adapters support all role combinations (trigger, invoke, or trigger and invoke). Other adapters support fewer role combinations.<br>When you select a role, only the connection properties and security policies appropriate to that role are displayed on the Connections page. If you select an adapter that supports both invoke and trigger, but select only one of those roles, you'll get an error when you try to drag the adapter into the section you didn't select.<br>For example, assume you configure a connection for the Oracle Service Cloud (RightNow) Adapter as only an **invoke**. Dragging the adapter to a **trigger** section in the integration produces an error. |
| **Keywords** | Enter optional keywords (tags). You can search on the connection keywords on the Connections page. |
| **Description** | Enter an optional description of the connection. |

| Element | Description |
|---|---|
| **Share with other projects** | **Note**: This field only appears if you are creating a connection in a project.<br>Select to make this connection publicly available in other projects. Connection sharing eliminates the need to create and maintain separate connections in different projects.<br>When you configure an adapter connection in a different project, the **Use a shared connection** field is displayed at the top of the Connections page. If the connection you are configuring matches the same type and role as the publicly available connection, you can select that connection to reference (inherit) its resources.<br>See Add and Share a Connection Across a Project. |

4. Click **Create**.

   Your connection is created. You're now ready to configure the connection properties, security policies, and (for some connections) access type.

5. Follow the steps to configure a connection.
   The connection property and connection security values are specific to each adapter. Your connection may also require configuration with an access type such as a private endpoint or an agent group.

6. Test the connection.

## Configure Connection Properties

Enter connection information so your application can process requests.

1. Go to the **Properties** section.
2. In the **Bootstrap Servers** field, specify the host and port to use to connect to a list of Kafka brokers. A Kafka cluster consists of one or more servers (Kafka brokers) running Kafka. Producers are processes that publish data (push messages) to Kafka topics within the broker. A consumer of topics pulls messages from a Kafka topic.
3. (Optional) Click **Optional properties**, and specify the schema registry URL.

> **Note:**
>
> The schema registry is required if you intend to use the Avro message structure.

## Configure Connection Security

Configure security for your Confluent Adapter connection.

1. Go to the **Security** section.
2. Select the security policy.

| Security Policy | Description |
|---|---|
| **SASL PLAIN over SSL** | Simple Authentication and Security Layer (SASL) is a framework for authentication and data security in Internet protocols. It separates authentication mechanisms from application protocols to enable any authentication mechanism supported by SASL to be used in any application protocol that uses SASL. Plain-text authentication assumes that the user name and password are submitted to the server in clear text. Therefore, this authentication method is only considered secure when using an encrypted connection. This security policy enables you to use SASL Plain with SSL encryption. |
| **SASL PLAIN** | Use SASL Plain without SSL encryption. |

3. Based on your security policy selection, enter the following details:

| If You Selected... | Specify These Details... |
|---|---|
| **SASL PLAIN over SSL** | • **SASL Username**: Enter the SASL user name.<br>• **SASL Password**: Enter the SASL password.<br>The remaining fields are optional:<br>• **Schema Registry Authentication**: Select an authentication type to use.<br>– **Basic Authentication**<br>– **None**<br>• **Schema Registry Username**: Specify the schema registry user name.<br>• **Schema Registry Password**: Specify the schema registry password.<br>• **TrustStore**: Select the check box, then click **Upload** to upload the truststore. See Configure Confluent Kafka.<br>• **TrustStore password**: Enter the TrustStore password. |
| **SASL PLAIN** | • **SASL Username**: Enter the SASL user name.<br>• **SASL Password**: Enter the SASL password.<br>The remaining fields are optional:<br>• **Schema Registry Authentication**: Select an authentication type to use.<br>– **Basic Authentication**<br>– **None**<br>• **Schema Registry Username**: Specify the schema registry user name.<br>• **Schema Registry Password**: Specify the schema registry password. |

# Configure the Endpoint Access Type

Configure access to your endpoint. Depending on the capabilities of the adapter you are configuring, options may appear to configure access to the public internet, to a private endpoint, or to an on-premises service hosted behind a fire wall.

### Select the Endpoint Access Type

1. Go to the **Access type** section.
2. Select the option for accessing your endpoint.

| Option | This Option Appears If Your Adapter Supports ... |
|---|---|
| **Connectivity agent** | Connections to on-premises endpoints through the connectivity agent.<br>a. Click **Associate agent group**.<br>The Associate agent group panel appears.<br>b. Select the agent group, and click **Use**.<br>To configure an agent group, you must download and install the on-premises connectivity agent. See Download and Run the Connectivity Agent Installer and About Creating Hybrid Integrations Using Oracle Integration in *Using Integrations in Oracle Integration 3*. |

# Test the Connection

Test your connection to ensure that it's configured successfully.

1. In the page title bar, click **Test**. What happens next depends on whether your adapter connection uses a Web Services Description Language (WSDL) file. Only some adapter connections use WSDLs.

| If Your Connection... | Then... |
|---|---|
| Doesn't use a WSDL | The test starts automatically and validates the inputs you provided for the connection. |
| Uses a WSDL | A dialog prompts you to select the type of connection testing to perform:<br>• **Validate and Test**: Performs a full validation of the WSDL, including processing of the imported schemas and WSDLs. Complete validation can take several minutes depending on the number of imported schemas and WSDLs. No requests are sent to the operations exposed in the WSDL.<br>• **Test**: Connects to the WSDL URL and performs a syntax check on the WSDL. No requests are sent to the operations exposed in the WSDL. |

2. Wait for a message about the results of the connection test.
   - If the test was successful, then the connection is configured properly.
   - If the test failed, then edit the configuration details you entered. Check for typos and verify URLs and credentials. Continue to test until the connection is successful.

3. When complete, click **Save**.

3

# Add the Confluent Adapter Connection to an Integration

When you drag the Confluent Adapter into the trigger or invoke area of an integration, the Adapter Endpoint Configuration Wizard is invoked. This wizard guides you through configuration of the Confluent Adapter endpoint properties.

The following sections describe the wizard pages that guide you through configuration of the Confluent Adapter as a trigger or invoke in an integration.

**Topics:**



## Basic Info Page

You can enter a name and description on the Basic Info page of each adapter in your integration.

| Element | Description |
|---|---|
| **What do you want to call your endpoint?** | Provide a meaningful name so that others can understand the responsibilities of this connection. You can include English alphabetic characters, numbers, underscores, and hyphens in the name. You can't include the following characters:<br>• No blank spaces (for example, `My Inbound Connection`)<br>• No special characters (for example, `#;83&` or `righ(t)now4`) except underscores and hyphens<br>• No multibyte characters |
| **What does this endpoint do?** | Enter an optional description of the connection's responsibilities. For example:<br>`This connection receives an inbound request to synchronize account information with the cloud application.` |
| **Messaging Type**<br>This field only appears when configuring the adapter in the trigger direction. | Select a messaging type:<br>• **Consumer**: Reads data from topics in the Kafka cluster using the Apache Kafka Consumer API.<br>• **Streaming**: Reads data from topics in the Kafka cluster using the Apache Kafka Streaming API. |

# Operations Page

Select the operation to perform.

| Element | Description |
|---|---|
| **What operation do you want to perform on a Kafka topic?** | • **Publish records to a Kafka topic**<br>• **Consume records from a Kafka topic**<br>• **Consume records from a Kafka topic by specifying offset**<br>• **Consume records from a Kafka topic by specifying multiple offsets and ranges** |

# Topic & Partition Page

Select the operation and topic on which to perform the operation and optionally specify the message structure.

- Configure the Confluent Adapter as an Invoke Connection
- Configure the Confluent Adapter as a Trigger Connection

**Configure the Confluent Adapter as an Invoke Connection**

| Element | Description |
|---|---|
| **Select a Topic** | Select the topic on which to perform the operation. You can also enter the beginning letters of the topic to filter the display of topics. A topic is a category in which applications can add, process, and reprocess messages. You subscribe to messages in topics. |
| **Specify the Partition**<br>(This field is only displayed if you select **Publish records to a Kafka topic** or **Consume records from a Kafka topic**.) | Specify the partition to which to push the selected topic. Kafka topics are divided into partitions that enable you to split data across multiple brokers. If you do not select a specific partition and use the **Default** selection, Kafka considers all available partitions and decides which one to use. |
| **Consumer Group**<br>(This field is only displayed if you select **Consume records from a Kafka topic**.) | Specify the consumer group to attach.<br>Consumers join a group by using the same group ID. Confluent assigns the partitions of a topic to the consumers in a group.<br>**Note**: When using a nondefault partition, specify the consumer group specific to the topic partition. |

| Element | Description |
|---|---|
| **Specify the option for consuming messages**<br>(This field is only displayed if you select **Consume records from a Kafka topic**.) | • **Read latest**: Reads the latest messages starting at the time at which the integration was activated.<br>• **Read from beginning**: Select to read messages from the beginning.<br>As an example, if you select to read from the beginning and have activated the integration, the first scheduled run picks up 20 records and the next scheduled run picks up the next 20 records. If the integration is then deactivated, edited, and reactivated, the next 20 records are picked up.<br>• **Read from specific offset**: Select to reset the consumer to a specific offset (partition-specific). Click **Add** to add the partition and the offset.<br>Once configured, if you edit it again, the following options are displayed:<br>• **Continue reading**: Reads the next message from where it left off before deactivation.<br>• **Reset to beginning**: Reads the message again from the beginning (may cause duplication).<br>• **Read from specific offset**: Select to reset the consumer to a specific offset (partition-specific). Click **Add** to add the partition and the offset (may cause duplication).<br>**Notes**:<br>• When you select **Read from beginning**, **Read latest**, or **Read from specific offset**, not all messages are guaranteed to be consumed in a single run. However, any remaining messages are consumed in subsequent runs.<br>• Selecting **Read from beginning**, **Reset to beginning**, or **Read from specific offset** resets the consumer group offset upon completing the wizard. |
| **Maximum Number of Records to be fetched**<br>(This field is only displayed if you select **Consume records from a Kafka topic** or **Consume records from a Kafka topic by specifying offset**.) | Specify the number of messages to read. The threshold for the complete message payload is 10 MB.<br>**Note**: This field specifies the upper boundary of records to fetch. It does not guarantee the specified amount of records to retrieve from the stream in a single run. Remaining messages are fetched in subsequent runs. |
| **Do you want to specify the message structure?** | Select **Yes** if you want to define the message structure to use on the Message Structure page of the wizard. Otherwise, select **No**. |
| **Do you want to specify the headers for the message?** | Select **Yes** if you want to define the message headers to use on the Headers page of the wizard. Otherwise, select **No**. |
| **Review and update advanced configurations** | Click **Edit** to open the Advanced Options section to enable or disable the transactional producer.<br>1. **Transaction Producer**: This field is only displayed if you select **Publish records to a Kafka topic**. This option provides the following capabilities:<br>a. If selected, the transactional producer enables an application to send messages to multiple partitions atomically.<br>b. If not selected, the Apache Kafka Adapter is configured as a nontransactional producer.<br>2. **Message Type**: This option defines the message type. Available options are **String** or **Bytes**. It defines the serializers to use for the message. This selection is applicable for the message key and value. |

### Configure the Confluent Adapter as a Trigger Connection

| Element | Description |
|---|---|
| **Select a Topic** | Select the topic on which to perform the operation. You can also enter the beginning letters of the topic to filter the display of topics. A topic is a category in which applications can add, process, and reprocess messages. You subscribe to messages in topics. |
| **Specify the Partition**<br>(This field is only displayed if you are reading data from topics using the Kafka Consumer API.) | Specify the partition to which to push the selected topic. Kafka topics are divided into partitions that enable you to split data across multiple brokers. If you do not select a specific partition and use the **Default** selection, Kafka considers all available partitions and decides which one to use. |
| **Consumer Group**<br>(This field is only displayed if you are reading data from topics using the Kafka Consumer API.) | Specify the consumer group to attach.<br>Consumers join a group by using the same group ID. Confluent assigns the partitions of a topic to the consumers in a group.<br>**Note**: When using a nondefault partition, specify the consumer group specific to the topic partition. |
| **Specify the option for consuming messages** | • **Read latest**: Reads the latest messages starting at the time at which the integration was activated.<br>• **Read from beginning**: Select to read messages from the beginning.<br>• **Read from specific offset**: Select to reset the consumer to a specific offset (partition-specific). Click **Add** to add the partition and the offset.<br>Once configured, if you edit it again, the following options are displayed:<br>• **Continue reading**: Reads the next message from where it left off before deactivation.<br>• **Reset to beginning**: Reads the message again from the beginning (may cause duplication).<br>• **Read from specific offset**: Select to reset the consumer to a specific offset (partition-specific). Click **Add** to add the partition and the offset (may cause duplication).<br>**Note**:<br>• When you select **Read from beginning**, **Read latest**, or **Read from specific offset**, not all messages are guaranteed to be consumed in a single run. However, any remaining messages are consumed in subsequent runs.<br>• Selecting **Read from beginning**, **Reset to beginning**, or **Read from specific offset** resets the consumer group offset upon completing the wizard. |
| **Application ID**<br>This field is only displayed if you are reading data from topics using the Kafka Streaming API.) | Specify the application ID to attach. The stream must have a unique ID. |
| **Polling Frequency (Sec)**<br>(This field is only displayed if you are reading data from topics using the Kafka Consumer API.) | Specify the frequency at which to fetch records. |
| **Maximum Number of Records to be fetched**<br>(This field is only displayed if you are reading data from topics using the Kafka Consumer API) | Specify the number of messages to read. The threshold for the complete message payload is 10 MB.<br>**Note**: This field specifies the upper boundary of records to fetch. It does not guarantee the specified amount of records to retrieve from the stream in a single run. Remaining messages are fetched in subsequent runs. |

| Element | Description |
|---|---|
| **Do you want to specify the message structure?** | Select **Yes** if you want to define the message structure to use on the Message Structure page of the wizard. Otherwise, select **No**. |
| **Do you want to specify the headers for the message?** | Select **Yes** if you want to define the message headers to use on the Headers page of the wizard. Otherwise, select **No**. |
| **Review and update advanced configurations** | Click **Edit** to open the Advanced Options section.<br>1. **Message Type**: This option defines the message type. Available options are **String** or **Bytes**. It defines the serializers to use for the message. This selection is applicable for the message key and value. |
| **Review and update rejected message configurations** | Define the topic for rejected messages. You can select the topic to which the rejected message is copied. |

# Message Structure Page

Select the message structure to use.

- Select the Message Structure
- If You Select Avro Schema (AVSC) document
- If You Select XML schema (XSD) document or Sample XML document

### Select the Message Structure

| Element | Description |
|---|---|
| **How would you like to specify the message structure?** | • **Avro Schema (AVSC) document**<br>• **XML Schema (XSD) document**<br>• **Sample XML document** |

### If You Select Avro Schema (AVSC) document

| Element | Description |
|---|---|
| **Schema Subject** | Specify the subject associated with the schema registry.<br>**Note**: This field is populated if the default schema is configured for Topic Value. |
| **Schema Version** | Specify the version associated with the schema registry. |

### If You Select XML schema (XSD) document or Sample XML document

| Element | Description |
|---|---|
| **Select File** | Select a file or drop one in the **Drag and Drop** field. Once selected, the file name is displayed in the **File Name** field. |
| **Element** | Select the element if you specified an XSD file. |

# Headers Page

Define the message headers structure to attach to the message. This page is displayed if you selected **Yes** for the **Do you want to specify the headers for the message?** field on the Topic & Partition page.

| Element | Description |
|---|---|
| **Specify Message Headers** | Specify headers and optional descriptions. |

# Summary Page

You can review the specified adapter configuration values on the Summary page.

| Element | Description |
|---|---|
| **Summary** | Displays a summary of the configuration values you defined on previous pages of the wizard.<br>The information that is displayed can vary by adapter. For some adapters, the selected business objects and operation name are displayed. For adapters for which a generated XSD file is provided, click the XSD link to view a read-only version of the file.<br>To return to a previous page to update any values, click the appropriate tab in the left panel or click **Go back**.<br>To cancel your configuration details, click **Cancel**. |

4

# Implement Common Patterns Using the Confluent Adapter

You can use the Confluent Adapter to implement the following common pattern.

**Topics:**



> **Note:**
>
> Oracle Integration offers a number of prebuilt integrations, known as *recipes*, that provide you with a head start in building your integrations. You can start with a recipe, and then customize it to fit your needs and requirements. Depending upon the solution provided, a variety of adapters are configured in the prebuilt integrations. See the Recipes and Accelerators page on the Oracle Help Center.

## Consume Messages from a Kafka Topic by Specifying Multiple Offsets and Ranges

You can configure an integration to use the Confluent Adapter to consume messages from a Kafka topic by specifying multiple offsets and ranges.

The following integration provides one example of how to implement this pattern:

- Create an application integration with a SOAP Adapter or REST Adapter as a trigger connection.
- Configure the Confluent Adapter to:
  - Consume records from a Kafka topic by specifying multiple offsets and ranges.
  - Specify the topic and message structure to use. For example, specify an XML schema (XSD) document and the headers to use for the message.
- Configure the mapper to perform appropriate source-to-target mappings between the SOAP Adapter or REST Adapter and the Confluent Adapter.
- Map or specify the partition and its offset range in the mapping (for example, the **Partition** can be `0` and the **Offset Range** format can be `6,10-11,16`).