

## EsperIO Reference

by [EsperTech Inc.](#)

Version 9.0.0

Copyright © 2024 EsperTech Inc.

---

### [Preface](#)

#### [1. Adapter Overview](#)

##### [1.1. The Adapter Interface](#)

#### [2. The File and CSV Input and Output Adapter](#)

##### [2.1. Data Flow Operators](#)

###### [2.1.1. Introduction](#)

###### [2.1.2. FileSink Operator](#)

###### [2.1.3. FileSource Operator](#)

##### [2.2. CSV Input Adapter API](#)

###### [2.2.1. Introduction](#)

###### [2.2.2. Playback of CSV-formatted Events](#)

###### [2.2.3. CSV Playback Options](#)

###### [2.2.4. Simulating Multiple Event Streams](#)

###### [2.2.5. Pausing and Resuming Operation](#)

#### [3. The Spring JMS Input and Output Adapter](#)

##### [3.1. Introduction](#)

##### [3.2. Runtime Configuration](#)

##### [3.3. Input Adapter](#)

###### [3.3.1. Spring Configuration](#)

###### [3.3.2. JMS Message Unmarshalling](#)

##### [3.4. Output Adapter](#)

###### [3.4.1. Spring Configuration](#)

###### [3.4.2. JMS Message Marshalling](#)

#### [4. The AMQP Input and Output Adapter](#)

##### [4.1. Introduction](#)

##### [4.2. AMQPSink Operator](#)

##### [4.3. AMQPSource Operator](#)

#### [5. The Kafka Adapter](#)

##### [5.1. Classpath Setup](#)

##### [5.2. Imports Setup](#)

##### [5.3. Input Adapter](#)

###### [5.3.1. Input Adapter Configuration and Start](#)

###### [5.3.2. Kafka Connectivity](#)

###### [5.3.3. Controlling Input Adapter Operation](#)

##### [5.4. Output Adapter](#)

###### [5.4.1. Output Adapter Configuration and Start](#)

###### [5.4.2. Kafka Connectivity](#)

###### [5.4.3. Controlling Output Adapter Operation](#)

#### [6. The HTTP Adapter](#)

##### [6.1. Adapter Overview](#)

##### [6.2. Getting Started](#)

###### [6.2.1. Plugin Loader Configuration](#)

###### [6.2.2. Configuration and Starting via API](#)

##### [6.3. HTTP Input Adapter](#)

###### [6.3.1. HTTP Service](#)

###### [6.3.2. Get Handlers](#)

###### [6.3.3. HTTP Input Limitations](#)

##### [6.4. HTTP Output Adapter](#)

###### [6.4.1. Triggered HTTP Get](#)

#### [7. The Socket Adapter](#)

##### [7.1. Getting Started](#)

###### [7.1.1. Plugin Loader Configuration](#)

###### [7.1.2. Configuration and Starting via API](#)

##### [7.2. Socket Service](#)

###### [7.2.1. Object Data Format](#)

###### [7.2.2. String CSV Data Format](#)

###### [7.2.3. String CSV Data Format With Property Order](#)

###### [7.2.4. JSON Formatted](#)

#### [8. The Relational Database Adapter](#)

##### [8.1. Adapter Overview](#)

##### [8.2. Getting Started](#)

###### [8.2.1. Plugin Loader Configuration](#)

###### [8.2.2. Configuration and Starting via API](#)

##### [8.3. JDBC Connections](#)

##### [8.4. Triggered DML Statement Execution](#)

##### [8.5. Triggered Update-Insert Execution](#)

## Preface

This document describes input and output adapters for the Esper Java event stream and complex event processor.

If you are new to Esper, the Esper reference manual should be your first stop.

If you are looking for information on a specific adapter, you are at the right spot.

## Chapter 1. Adapter Overview

### 1.1. The Adapter Interface

Input and output adapters to Esper provide the means of accepting events from various sources, and for making available events to destinations.



#### Note

Esper has a fully-evolved API and natively accepts various input event formats and natively produces various output event objects, as part of core Esper. It is therefore not necessary to use any of the adapters listed herein. Simply use the public Esper APIs directly in your code.

Most adapters present their own configuration as well as API. Some adapters also provide operators for use in data flows.

### 1.1. The Adapter Interface

The `Adapter` interface allows client applications to control the state of an input and output adapter. It provides state transition methods that each input and output adapter implements.

An input or output adapter is always in one of the following states:

- Opened - The begin state; The adapter is not generating or accepting events in this state
- Started - When the adapter is active, generating and accepting events
- Paused - When operation of the adapter is suspended
- Destroyed

The state transition table below outlines adapter states and, for each state, the valid state transitions:

**Table 1.1. Adapter State Transitions**

Start State	Method	Next State
Opened	<code>start()</code>	Started
Opened	<code>destroy()</code>	Destroyed
Started	<code>stop()</code>	Opened
Started	<code>pause()</code>	Paused
Started	<code>destroy()</code>	Destroyed
Paused	<code>resume()</code>	Started
Paused	<code>stop()</code>	Opened
Paused	<code>destroy()</code>	Destroyed

## Chapter 2. The File and CSV Input and Output Adapter

### 2.1. Data Flow Operators

- [2.1.1. Introduction](#)
- [2.1.2. FileSink Operator](#)
- [2.1.3. FileSource Operator](#)

### 2.2. CSV Input Adapter API

- [2.2.1. Introduction](#)
- [2.2.2. Playback of CSV-formatted Events](#)
- [2.2.3. CSV Playback Options](#)
- [2.2.4. Simulating Multiple Event Streams](#)
- [2.2.5. Pausing and Resuming Operation](#)

The file input and output adapter consists of:

1. File (including CSV) input and output utilizing data flow operators.
2. The CSV input adapter API.

This adapter only requires Esper runtime as a dependency and does not require Esper compiler.

### 2.1. Data Flow Operators

#### 2.1.1. Introduction

In order to use the File source and sink data flow operators, add `esperio-csv-version.jar` to your classpath and import the operator package.

For example:

```
Configuration configuration = new Configuration();
configuration.getCommon().addImport(FileSourceForge.class.getPackage().getName() + ".*");
```

The File input and output adapter provides the following data flow operators:

**Table 2.1. File Operators**

Operator	Description
FileSink	Write events to a file. See <a href="#">Section 2.1.2, "FileSink Operator"</a> .
FileSource	Read events from a file. See <a href="#">Section 2.1.3, "FileSource Operator"</a> .

#### 2.1.2. FileSink Operator

The FileSink operator receives input stream events, transforms events to comma-separated format and writes to a file.

The FileSink operator must have a single input stream.

The FileSink operator cannot declare any output streams.

Parameters for the FileSink operator are (required parameters are listed first):

**Table 2.2. FileSink Parameters**

Name	Description
file (required)	An expression returning the file name string. An absolute or relative file name.
classpathFile	An expression returning a boolean indicator whether the file name is found in a classpath directory, false by default.
append	An expression returning a boolean indicator whether to append or overwrite the file when it exists. false by default causes the existing file, if any, to be overwritten.

The following example declares a data flow that is triggered by MyMapEventType events from the event bus (type not declared here) and that writes to the file output.csv the CSV-formatted events:

```
create dataflow FileSinkSampleFlow
EventBusSource -> outstream<MyMapEventType> {}
FileSink(outstream) {
    file: 'output.csv',
    append: true
}
```

### 2.1.3. FileSource Operator

The FileSource operator reads from files, transforms file data and populates a data flow instance with events.

The FileSource operator cannot declare any input streams.

The FileSource operator must have at least one output stream. You can declare additional output streams to hold beginning-of-file and end-of-file indication.

Parameters for the FileSource operator are listed below, with the required parameters listed first:

**Table 2.3. FileSource Parameters**

Name	Description
file (required, or provide adapterInputSource)	An expression returning the file name string
adapterInputSource (required, or provide file)	An instance of AdapterInputSource if a file name cannot be provided.
classpathFile	An expression returning a boolean indicator whether the file is found in a classpath directory, false by default.
dateFormat	An expression returning the format to use when parsing dates; the default is SimpleDateFormat of yyyy-MM-dd'T'HH:mm:ss.SSS for Date and Calendar type properties.
format	An expression returning csv (the default) for comma-separate value or line for single-line.
hasTitleLine	An expression returning a boolean indicator whether a title line exists that the operator should read and parse to obtain event property names. For use with the csv format.
hasHeaderLine	An expression returning a boolean indicator whether a header line exists that the operator should skip. For use with the csv format.
numLoops	An expression returning a an integer value that instructs the runtime to restart reading the file upon encountering EOF, defaults to zero. For use with the csv format.
propertyNames	String array constant that provides a list of property names in the same order they appear in the file. For use with the csv format.
propertyNameLine	An expression returning the property name of the output event type that receives the line text of type string. For use with the line format.
propertyNameFile	An expression returning the property name of the output event type(s) that receive the file name of type string. For use with the line format.

The first output stream holds per-line output events. For use with the line format and if declaring two output streams, the second stream holds end-of-file indication. If declaring three output streams, the second stream holds beginning-of-file indication and the third stream holds end-of-file indication.

The line format requires that the output stream's event type is an object-array event type that features a single string-type property that the operator populates with each line of the file.

The file name (or adapterInputSource) may point to a zip file. If the file name ends with the literal zip the operator opens the zip file and uses the first packaged file. All other parameters including the format parameter for CSV or line-formatting then apply to the zipped file.

This example defines a data flow that consists of two operators that work together to read a file and send the resulting events into the runtime:

```
create dataflow SensorCSVFlow
FileSource -> sensorstream<TemperatureEventStream> {
    file: 'sensor_events.csv',
    propertyNames: ['sensor', 'temp', 'uptime'],
    numLoops: 3
}
EventBusSink(sensorstream) {}
```

The data flow above configures the FileSource operator to read the file sensor\_events.csv, populate the sensor, temp and uptime properties of the TemperatureEventStream event type (type definition not shown here) and make the output events available within the data flow under the name sensorstream.

The data flow above configures the EventBusSource operator to send the sensorstream events into the runtime for processing.

#### 2.1.3.1. FileSource Operator Detailed Example

This example shows the EPL and code to read and count lines in text files.

Below EPL defines an event type to each hold the file line text as well as to indicate the beginning and

end of a file (remove the semicolon if creating EPL individually and not as a module):

```
// for beginning-of-file events
create objectarray schema MyBOF (filename string);
// for end of file events
create objectarray schema MyEOF (filename string);
// for line text events
create objectarray schema MyLine (filename string, line string);
```

The next EPL statements count lines per file outputting the final line count only when the end-of-file is reached.

```
// Initiate a context partition for each file, terminate upon end-of-file
create context FileContext
  initiated by MyBOF as mybof
  terminated by MyEOF(filename=mybof.filename);

// For each file, count lines
context FileContext
  select context.mybof.filename as filename, count(*) as cnt
  from MyLine(filename=context.mybof.filename)
  output snapshot when terminated;
```

The below EPL defines a data flow that reads text files line-by-line and that send events into the runtime for processing.

```
create dataflow MyEOFEventFileReader
  FileSource -> mylines<MyLine>, mybof<MyBOF>, myeof<MyEOF> {
    format: 'line',
    propertyNameLine: 'line', // store the text in the event property 'line'
    propertyNameFile: 'filename' // store the file name in 'filename'
  }
  EventBusSink(mylines, mybof, myeof) {} // send events into runtime
```

The next sample code instantiates and runs data flows passing a file name:

```
EPDataFlowInstantiationOptions options = new EPDataFlowInstantiationOptions();
options.addParameterURI("FileSource/file", "myfile.txt");
EPDataFlowInstance instance = runtime.getDataFlowService()
    .instantiate("MyEOFEventFileReader", options);
instance.run();
```

## 2.2. CSV Input Adapter API

This chapter discusses the CSV input adapter API. CSV is an abbreviation for comma-separated values. CSV files are simple text files in which each line is a comma-separated list of values. CSV-formatted text can be read from many different input sources via `com.espertech.esperio.csv.AdapterInputSource`. Please consult the JavaDoc for additional information on `AdapterInputSource` and the CSV adapter.

### 2.2.1. Introduction

In summary the CSV input adapter API performs the following functions:

- Read events from an input source providing CSV-formatted text and send the events to an Esper runtime instance
  - Read from different types of input sources
  - Use a timestamp column to schedule events being sent into the runtime
  - Playback with options such as file looping, events per second and other options
  - Use the Esper runtime timer thread to read the CSV file
- Read multiple CSV files using a timestamp column to simulate events coming from different streams

The following formatting rules and restrictions apply to CSV-formatted text:

- Comment lines are prefixed with a single hash or pound # character
- Strings are placed in double quotes, e.g. "value"
- Escape rules follow common spreadsheet conventions, i.e. double quotes can be escaped via double quote
- A column header is required unless a property order is defined explicitly
- If a column header is used, properties are assumed to be of type String unless otherwise configured
- The value of the timestamp column, if one is given, must be in ascending order
- The adapter does not allocate new event types. Any event type names provided must exist and their bus-event-type modifier must be visible.

### 2.2.2. Playback of CSV-formatted Events

The adapter reads events from a CSV input source and sends events to an runtime using the class `com.espertech.esperio.csv.CSVInputAdapter`.

The below code snippet reads the CSV-formatted text file "simulation.csv" expecting the file in the classpath. The `AdapterInputSource` class can take other input sources.

```
AdapterInputSource source = new AdapterInputSource("simulation.csv");
(new CSVInputAdapter(runtime, source, "PriceEvent")).start();
```

To use the `CSVInputAdapter` without any options, the event type `PriceEvent` and its property names and value types must be known to the runtime. The next section elaborates on adapter options.

- Configure the runtime instance for a Map-based event type
- Place a header record in your CSV file that names each column as specified in the event type

The sample application code below shows all the steps to configure, via API, a Map-based event type and play the CSV file without setting any of the available options.

```

Map<String, Class> eventProperties = new HashMap<String, Class>();
eventProperties.put("symbol", String.class);
eventProperties.put("price", double.class);
eventProperties.put("volume", Integer.class);

Configuration configuration = new Configuration();
configuration.getCommon().addEventType("PriceEvent", eventProperties);

runtime = EPRuntimeProvider.getDefaultRuntime(configuration);

EPStatement stmt = compileDeploy(runtime, "select symbol, price, volume from PriceEvent.win:time(1 min)");

(new CSVInputAdapter(runtime, new AdapterInputSource(filename), "PriceEvent")).start();

```

The contents of a sample CSV file is shown next.

```

symbol,price,volume
IBM,55.5,1000

```

The next code snippet outlines using a `java.io.Reader` as an alternative input source :

```

String myCSV = "symbol, price, volume" + NEW_LINE + "IBM, 10.2, 10000";
StringReader reader = new StringReader(myCSV);
(new CSVInputAdapter(runtime, new AdapterInputSource(reader), "PriceEvent")).start();

```

In the previous code samples, the `PriceEvent` properties were defined programmatically with their correct types. It is possible to skip this step and use only a column header record. In such a case you must define property types in the header otherwise a type of `String` is assumed.

Consider the following:

```

symbol,double price, int volume
IBM,55.5,1000

symbol,price,volume
IBM,55.5,1000

```

The first CSV file defines explicit types in the column header while the second file does not. With the second file a statement like `select sum(volume) from PriceEvent.win:time(1 min)` will be rejected as in the second file `volume` is defaulted to type `String` - unless otherwise programmatically configured.

### 2.2.2.1. Using JavaBean POJO Events

The previous section used an event type based on `java.util.Map`. The adapter can also populate the CSV data into JavaBean events directly, as long as your event class provides setter-methods that follow JavaBean conventions. Note that `esperio` will ignore read-only properties i.e. if you have a read-only property `priceByVolume` it will not expect a corresponding column in the input file.

To use Java objects as events instead of Map-based event types, simply register the event type name for the Java class and provide the same name to the CSV adapter.

The below code snippet assumes that a `PriceEvent` class exists that exposes setter-methods for the three properties. The setter-methods are, for example, `setSymbol(String s)`, `setPrice(double p)` and `setVolume(long v)`.

```

Configuration configuration = new Configuration();
configuration.getCommon().addEventType("PriceEvent", PriceEvent.class);

runtime = EPRuntimeProvider.getDefaultRuntime(configuration);

EPStatement stmt = compileDeploy(runtime, "select symbol, price, volume from PriceEvent.win:time(1 min)");

(new CSVInputAdapter(runtime, new AdapterInputSource(filename), "PriceEvent")).start();

```

When using JavaBean POJO Events, the event properties types are known from the underlying event type configuration. The CSV file row header does not need to define column type explicitly.

### 2.2.2.2. Dealing with event with nested properties

Whether you use JavaBean POJO or Map-based event types, `EsperIO` provides support for nested event properties up to one level of nesting. The row header must then refer to the properties using a `propertyName.nestedPropertyName` syntax. There is no support for mapped or indexed properties.

For example consider the following:

```

public class Point {
    int x;
    int y;

    // with getters & setters
}
public class Figure {
    String name;
    Point point; // point.x and point.y are nested properties

    //with getters & setters
}

```

Or the equivalent representation with nested maps, assuming "Figure" is the declared event type name, the CSV file can contain the following row header:

```

name, point.x, point.y

```

### 2.2.3. CSV Playback Options

Use the `CSVInputAdapterSpec` class to set playback options. The following options are available:

- **Loop** - Reads the CSV input source in a loop; When the end is reached, the input adapter rewinds to the beginning
- **Events per second** - Controls the number of events per second that the adapter sends to the runtime
- **Property order** - Controls the order of event property values in the CSV input source, for use when the CSV input source does not have a header column
- **Property types** - Defines the property types, overriding property types of the event type.

- Runtime thread - Instructs the adapter to use the runtime timer thread to read the CSV input source and send events to the runtime
- External timer - Instructs the adapter to use the esper's external timer rather than the internal timer. See "Sending timer events" below
- Timestamp column name - Defines the name of the timestamp column in the CSV input source; The timestamp column must carry long-typed timestamp values relative to the current time; Use zero for the current time

The next code snippet shows the use of `CSVInputAdapterSpec` to set playback options.

```
CSVInputAdapterSpec spec = new CSVInputAdapterSpec(new AdapterInputSource(myURL), "PriceEvent");
spec.setEventsPerSec(1000);
spec.setLooping(true);

InputAdapter inputAdapter = new CSVInputAdapter(runtime, spec);
inputAdapter.start(); // method blocks unless runtime thread option is set
```

### 2.2.3.1. Sending timer events

The adapter can be instructed to use either esper's internal timer, or to drive timing itself by sending external timer events. If the internal timer is used, esperio will send all events in "real time". For example, if an input file contains the following data:

```
symbol,price,volume,timestamp
IBM,55.5,1000,2
GOOG,9.5,1000,3
MSFT,8.5,1000,3
JAVA,7.5,1000,1004
```

then esperio will sleep for 1001 milliseconds between sending the MSFT and JAVA events to the runtime.

If external timing is enabled then esperio will run through the input file at full speed without pausing. The algorithm used sends a time event after all events for a particular time have been received. For the above example file a time event for 2 will be sent after IBM, for 3 after MSFT and 1004 after JAVA. For many of use cases this gives a performance improvement.

### 2.2.4. Simulating Multiple Event Streams

The CSV input adapter can run simulations of events arriving in time-order from different input streams. Use the `AdapterCoordinator` as a specialized input adapter for coordinating multiple CSV input sources by timestamp.

The sample application code listed below simulates price and trade events arriving in timestamp order. Via the adapter the application reads two CSV-formatted files from a URL that each contain a timestamp column as well as price or trade events. The `AdapterCoordinator` uses the timestamp column to send events to the runtime in the exact ordering prescribed by the timestamp values.

```
AdapterInputSource sourceOne = new AdapterInputSource(new URL("FILE://prices.csv"));
CSVInputAdapterSpec inputOne = new CSVInputAdapterSpec(sourceOne, "PriceEvent");
inputOne.setTimestampColumn("timestamp");

AdapterInputSource sourceTwo = new AdapterInputSource(new URL("FILE://trades.csv"));
CSVInputAdapterSpec inputTwo = new CSVInputAdapterSpec(sourceTwo, "TradeEvent");
inputTwo.setTimestampColumn("timestamp");

AdapterCoordinator coordinator = new AdapterCoordinatorImpl(runtime, true);
coordinator.coordinate(new CSVInputAdapter(inputOne));
coordinator.coordinate(new CSVInputAdapter(inputTwo));
coordinator.start();
```

The `AdapterCoordinatorImpl` is provided with two parameters: the runtime instance, and a boolean value that instructs the adapter to use the runtime timer thread if set to true, and the adapter can use the application thread if the flag passed is false.

You may not set an event rate per second when using a timestamp column and time-order.

### 2.2.5. Pausing and Resuming Operation

The CSV adapter can employ the runtime timer thread of an Esper runtime instance to read and send events. This can be controlled via the `setUsingRuntimeThread` method on `CSVInputAdapterSpec`. We use that feature in the sample code below to pause and resume a running CSV input adapter.

```
CSVInputAdapterSpec spec = new CSVInputAdapterSpec(new AdapterInputSource(myURL), "PriceEvent");
spec.setEventsPerSec(100);
spec.setUsingRuntimeThread(true);

InputAdapter inputAdapter = new CSVInputAdapter(runtime, spec);
inputAdapter.start(); // method starts adapter and returns, non-blocking
Thread.sleep(5000); // sleep 5 seconds
inputAdapter.pause();
Thread.sleep(5000); // sleep 5 seconds
inputAdapter.resume();
Thread.sleep(5000); // sleep 5 seconds
inputAdapter.stop();
```

## Chapter 3. The Spring JMS Input and Output Adapter

### [3.1. Introduction](#)

### [3.2. Runtime Configuration](#)

### [3.3. Input Adapter](#)

#### [3.3.1. Spring Configuration](#)

#### [3.3.2. JMS Message Unmarshalling](#)

### [3.4. Output Adapter](#)

#### [3.4.1. Spring Configuration](#)

#### [3.4.2. JMS Message Marshalling](#)

This chapter discusses the input and output adapters for JMS based on the Spring `JmsTemplate` technology. For more information on Spring, and the latest version of Spring, please visit <http://www.springframework.org>.

The output adapter requires Esper compiler as well as Esper runtime as a dependency. The input adapter only requires Esper runtime as a dependency.

## 3.1. Introduction

Here are the steps to use the adapters:

1. Configure an Esper runtime instance to use a `SpringContextLoader` for loading input and output adapters, and point it to a Spring `JmsTemplate` configuration file.
2. Create a Spring `JmsTemplate` configuration file for your JMS provider and add all your input and output adapter entries in the same file.
3. For receiving events from a JMS destination into a runtime (input adapter):
  1. List the destination and un-marshalling class in the Spring configuration.
  2. Create EPL statements using the event type name matching the event objects or the Map-event type names received.
4. For sending events to a JMS destination (output adapter):
  1. Use the `insert-into` syntax naming the stream to insert-into using the same name as listed in the Spring configuration file
  2. Configure the Map event type of the stream in the runtime configuration

In summary the Spring JMS input adapter performs the following functions:

- Initialize from a given Spring configuration file in classpath or from a filename. The Spring configuration file sets all JMS parameters such as JMS connection factory, destination and listener pools.
- Attach to a JMS destination and listen to messages using the Spring class `org.springframework.jms.core.JmsTemplate`
- Unmarshal a JMS message and send into the configured runtime instance

The Spring JMS output adapter can:

- Initialize from a given Spring configuration file in classpath or from a filename, and attach to a JMS destination
- Act as a listener to one or more named streams populated via `insert-into` syntax by EPL statements
- Marshal events generated by a stream into a JMS message, and send to the given destination

## 3.2. Runtime Configuration

The Spring JMS input and output adapters are configured as part of the Esper runtime configuration. `EsperIO` supplies a `SpringContextLoader` class that loads a Spring configuration file which in turn configures the JMS input and output adapters. List the `SpringContextLoader` class as an adapter loader in the Esper configuration file as the below example shows. The configuration API can alternatively be used to configure one or more adapter loaders.

```
<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <!-- Sample configuration for an input/output adapter loader -->
    <plugin-loader name="MyLoader" class-name="com.espertech.esperio.jms.SpringContextLoader" />
    <!-- SpringApplicationContext translates into Spring ClassPathXmlApplicationContext
         or FileSystemXmlApplicationContext. Only one app-context of a sort can be used.
         When both attributes are used classpath and file, classpath prevails -->
    <init-arg name="classpath-app-context" value="spring\jms-spring.xml" />
    <init-arg name="file-app-context" value="spring\jms-spring.xml" />
  </plugin-loader>
</runtime>
</esper-configuration>
```

The loader loads the Spring configuration file from classpath via the `classpath-app-context` configuration, or from a file via `file-app-context`.

## 3.3. Input Adapter

### 3.3.1. Spring Configuration

The Spring configuration file must list input and output adapters to be initialized by `SpringContextLoader` upon runtime initialization. Please refer to your JMS provider documentation, and the Spring framework documentation on help to configure your specific JMS provider via Spring.

The next XML snippet shows a complete sample configuration for an input adapter. The sample includes the JMS configuration for an Apache ActiveMQ JMS provider.

```

<!-- Spring Application Context -->
<beans default-destroy-method="destroy">

    <!-- JMS ActiveMQ Connection Factory -->
    <bean id="jmsActiveMQFactory" class="org.apache.activemq.pool.PooledConnectionFactory">
        <property name="connectionFactory">
            <bean class="org.apache.activemq.ActiveMQConnectionFactory">
                <property name="brokerURL" value="tcp://localhost:61616"/>
            </bean>
        </property>
    </bean>

    <!-- ActiveMQ destination to use by default -->
    <bean id="defaultDestination"
        class="org.apache.activemq.command.ActiveMQQueue">
        <constructor-arg value="ESPER.QUEUE"/>
    </bean>

    <!-- Spring JMS Template for ActiveMQ -->
    <bean id="jmsActiveMQTemplate" class="org.springframework.jms.core.JmsTemplate">
        <property name="connectionFactory">
            <ref bean="jmsActiveMQFactory"/>
        </property>
        <property name="defaultDestination">
            <ref bean="defaultDestination"/>
        </property>
    </bean>

    <!-- Provides listener threads -->
    <bean id="listenerContainer"
        class="org.springframework.jms.listener.SimpleMessageListenerContainer">
        <property name="connectionFactory" ref="jmsActiveMQFactory"/>
        <property name="destination" ref="defaultDestination"/>
        <property name="messageListener" ref="jmsInputAdapter"/>
    </bean>

    <!-- Default unmarshaller -->
    <bean id="jmsMessageUnmarshaller"
        class="com.espertech.esperio.jms.JMSDefaultAnyMessageUnmarshaller"/>

    <!-- Input adapter -->
    <bean id="jmsInputAdapter" class="com.espertech.esperio.jms.SpringJMSTemplateInputAdapter">
        <property name="jmsTemplate">
            <ref bean="jmsActiveMQTemplate"/>
        </property>
        <property name="jmsMessageUnmarshaller">
            <ref bean="jmsMessageUnmarshaller"/>
        </property>
    </bean>

</beans>

```

This input adapter attaches to the JMS destination `ESPER.QUEUE` at an Apache MQ broker available at port `tcp://localhost:61616`. It configures an un-marshalling class as discussed next.

### 3.3.2. JMS Message Unmarshalling

EsperIO provides a class for unmarshaling JMS message instances into events for processing by an runtime in the class `JMSDefaultAnyMessageUnmarshaller`. The class unmarshals as follows:

- If the received `Message` is of type `javax.xml.MapMessage`, extract the event type name out of the message and send to the runtime via `sendEventMap(Map, name)`
- If the received `Message` is of type `javax.xml.ObjectMessage`, extract the `Serializable` out of the message and send to the runtime via `sendEventBean(Object, object.getClass().getSimpleName())`
- If the received `Message` is of type `javax.xml.TextMessage`, extract the text out of the message, parse it as JSON-formatted message according to the named JSON event type that was previously created by `create json schema` and send to the runtime via `sendEventJson(json, eventTypeName)` (see below for `eventTypeName`).
- Else the un-marshaller outputs a warning and ignores the message

For map messages, the unmarshaller must be made aware of the event type of events within `MapMessage` messages. This is achieved by the client application setting a well-defined property on the message: `InputAdapter.ESPERIO_MAP_EVENT_TYPE`. An example code snippet is:

```

MapMessage mapMessage = jmsSession.createMapMessage();
mapMessage.setObject(InputAdapter.ESPERIO_MAP_EVENT_TYPE, "MyInputEvent");

```

For text JSON messages, the unmarshaller must be made aware of the event type of events within `TextMessage` messages. This is achieved by the client application setting a well-defined property on the message: `InputAdapter.ESPERIO_JSON_EVENT_TYPE`. An example code snippet is:

```

MapMessage mapMessage = jmsSession.createMapMessage();
mapMessage.setString(InputAdapter.ESPERIO_JSON_EVENT_TYPE, "MyJsonEvent");

```

## 3.4. Output Adapter

### 3.4.1. Spring Configuration

The Spring configuration file lists all input and output adapters in one file. The `SpringContextLoader` upon runtime initialization starts all input and output adapters.

The next XML snippet shows a complete sample configuration of an output adapter. Please check with your JMS provider for the appropriate Spring class names and settings. Note that the input and output adapter Spring configurations can be in the same file.



```

<!-- Application Context -->
<beans default-destroy-method="destroy">

  <!-- JMS ActiveMQ Connection Factory -->
  <bean id="jmsActiveMQFactory" class="org.apache.activemq.pool.PooledConnectionFactory">
    <property name="connectionFactory">
      <bean class="org.apache.activemq.ActiveMQConnectionFactory">
        <property name="brokerURL" value="tcp://localhost:61616"/>
      </bean>
    </property>
  </bean>

  <!-- ActiveMQ destination to use by default -->
  <bean id="defaultDestination"
        class="org.apache.activemq.command.ActiveMQQueue">
    <constructor-arg value="ESPER.QUEUE"/>
  </bean>

  <!-- Spring JMS Template for ActiveMQ -->
  <bean id="jmsActiveMQTemplate" class="org.springframework.jms.core.JmsTemplate">
    <property name="connectionFactory">
      <ref bean="jmsActiveMQFactory"/>
    </property>
    <property name="defaultDestination">
      <ref bean="defaultDestination"/>
    </property>
    <property name="receiveTimeout">
      <value>30000</value>
    </property>
  </bean>

  <!-- Marshaller marshals events into map messages -->
  <bean id="jmsMessageMarshaller" class="com.esperitech.esperio.jms.JMSDefaultMapMessageMarshaller">
  <bean id="myCustomMarshaller" class="com.esperitech.esperio.jms.JMSDefaultMapMessageMarshaller">

  <!-- Output adapter puts it all together -->
  <bean id="jmsOutputAdapter" class="com.esperitech.esperio.jms.SpringJMSTemplateOutputAdapter">
    <property name="jmsTemplate">
      <ref bean="jmsActiveMQTemplate"/>
    </property>
    <property name="subscriptionMap">
      <map>
        <entry>
          <key><idref local="subscriptionOne"/></key>
          <ref bean="subscriptionOne"/>
        </entry>
        <entry>
          <key><idref local="subscriptionTwo"/></key>
          <ref bean="subscriptionTwo"/>
        </entry>
      </map>
    </property>
    <property name="jmsMessageMarshaller">
      <ref bean="jmsMessageMarshaller"/>
    </property>
  </bean>

  <bean id="subscriptionOne" class="com.esperitech.esperio.jms.JMSSubscription">
    <property name="eventName" value="MyOutputStream"/>
  </bean>

  <bean id="subscriptionTwo" class="com.esperitech.esperio.jms.JMSSubscription">
    <property name="eventName" value="MyOtherOutputStream"/>
    <property name="jmsMessageMarshaller">
      <ref bean="myCustomMarshaller"/>
    </property>
  </bean>

</beans>

```

### 3.4.2. JMS Message Marshalling

EsperIO provides a marshal implementation in the class `JMSDefaultMapMessageMarshaller`. This marshaller constructs a JMS `MapMessage` from any event received by copying event properties into the name-value pairs of the message. The configuration file makes it easy to configure a custom marshaller that adheres to the `com.esperitech.esperio.jms.JMSMessageMarshaller` interface.

Note that this marshaller uses `javax.jms.MapMessage` name-value pairs and not general `javax.jms.Message` properties. This means when you'll read the event properties back from the JMS `MapMessage`, you will have to use the `javax.jms.MapMessage.getObject(...)` method.

The `SpringJMSTemplateOutputAdapter` is configured with a list of subscription instances of type `JMSSubscription` as the sample configuration shows. Each subscription defines an event type name that must be configured and used in the `insert-into` syntax of a statement.

To connect the Spring JMS output adapter and the EPL statements producing events, use the `insert-into` syntax to direct events for output. Here is a sample statement that sends events into `MyOutputStream`:

```
insert into MyOutputStream select assetId, zone from RFIDEvent
```

The type `MyOutputStream` must be known to a runtime. The output adapter requires the name to be configured with the runtime instance, e.g.:

```

<esper-configuration xmlns="http://www.esperitech.com/schema/esper">
  <common>
    <event-type name="MyOutputStream">
      <java-util-map>
        <map-property name="assetId" class="String"/>
        <map-property name="zone" class="int"/>
      </java-util-map>
    </event-type>
  </common>
</esper-configuration>

```

## Chapter 4. The AMQP Input and Output Adapter

### [4.1. Introduction](#)

### [4.2. AMQPSink Operator](#)

### [4.3. AMQPSource Operator](#)

This chapter discusses the input and output adapters for AMQP. AMQP input and output utilizes data flow operators.

This adapter only requires Esper runtime as a dependency and does not require Esper compiler.

## 4.1. Introduction

In order to use the AMQP data flow operators, add `esperio-amqp-version.jar` to your classpath and import the operator package or class using the static or runtime configuration.

The following code snippet uses the runtime configuration API to import the AMQP adapter classes:

```
Configuration configuration = new Configuration();
configuration.getCommon().addImport(AMQPSource.class.getPackage().getName() + ".*");
```

The AMQP input and output adapter provides the following data flow operators:

**Table 4.1. AMQP Operators**

Operator	Description
AMQPSink	Send messages to an AMQP broker. See <a href="#">Section 4.2, "AMQPSink Operator"</a> .
AMQPSource	Receive messages from an AMQP broker. See <a href="#">Section 4.3, "AMQPSource Operator"</a> .

## 4.2. AMQPSink Operator

The AMQPSink operator receives input stream events, transforms events to AMQP messages and sends messages into an AMQP queue.

The AMQPSink operator must have a single input stream.

The AMQPSink operator cannot declare any output streams.

Parameters for the AMQPSink operator are:

**Table 4.2. AMQPSink Parameters**

Name	Description
host (required)	An expression returning the host name string
queueName	An expression returning the queue name string
collector (required)	Transformation class or instance for events to AMQP message
port	An expression returning the port number integer
username	An expression returning the user name string
password	An expression returning the password string (also see <code>systemProperties</code> or <code>data flow instantiation options</code> )
vhost	An expression returning the vhost string
exchange	An expression returning the exchange name string
routingKey	An expression returning the routing key string
logMessages	An expression returning a boolean indicator whether to log a text for each message (requires debug-level logging)
waitMsecNextMsg	An expression returning the number of milliseconds wait between messages, a long-typed value
declareDurable	An expression returning a boolean indicator whether durable, false by default
declareExclusive	An expression returning a boolean indicator whether exclusive, false by default
declareAutoDelete	An expression returning a boolean indicator whether auto-delete, true by default
declareAdditionalArgs	Map of additional arguments passed to AMQP of type <code>Map&lt;String, Object&gt;</code>

Either the `queueName` or the combination of `exchange` and `routingKey` are required parameters.

The collector is required and must be specified to transform events to AMQP messages. The collector instance must implement the interface `ObjectToAMQPCollector`. The adapter provides a default implementation `ObjectToAMQPCollectorSerializable` that employs default serialization.

For JSON formatted messages, the adapter provides the `ObjectToAMQPCollectorJson` JSON collector. The collector emits string-type messages. You must specify a JSON event type name e.g. `create json schema MyJsonEventType(...)` and `outstream<MyJsonEventType>`.

The following example declares a data flow that is triggered by `MyMapEventType` events from the event bus (type not declared here) that sends serialized messages to an AMQP queue:

```
create dataflow AMQPOutgoingDataFlow
EventBusSource -> outstream<MyMapEventType> {
  AMQPSink(outstream) {
    host: 'localhost',
    queueName: 'myqueue',
    collector: {class: 'ObjectToAMQPCollectorSerializable'}
  }
}
```

## 4.3. AMQPSource Operator

The AMQPSource operator receives AMQP messages from a queue, transforms messages and populates a data flow instance with events.

The AMQPSource operator cannot declare any input streams.

The AMQPSource operator must have a single output stream.

Parameters for the AMQPSource operator are listed below, with the required parameters listed first:

**Table 4.3. AMQPSource Parameters**

Name	Description
host (required)	An expression returning the host name string
queueName (required)	An expression returning the queue name string
collector (required)	Transformation class or instance for AMQP message to underlying event transformation
port	An expression returning the port number integer

Name	Description
username	An expression returning the user name string
password	An expression returning the password string (also see systemProperties or data flow instantiation options)
vhost	An expression returning the vhost string
exchange	An expression returning the exchange name string
routingKey	An expression returning the routing key string
logMessages	An expression returning a boolean indicator whether to log a text for each message (requires debug-level logging)
waitMsecNextMsg	An expression returning the number of milliseconds wait between messages, a long-typed value
declareDurable	An expression returning a boolean indicator whether durable, false by default
declareExclusive	An expression returning a boolean indicator whether exclusive, false by default
declareAutoDelete	An expression returning a boolean indicator whether auto-delete, true by default
declareAdditionalArgs	Map of additional arguments passed to AMQP of type Map<String, Object>
prefetchCount	An expression returning the prefetch count integer, defaults to 100
consumeAutoAck	An expression returning a boolean indicator whether to auto-ack, true by default

The collector is required and must be specified to transform AMQP messages to events. The collector instance must implement the interface `AMQPToObjectCollector`. The adapter provides a default implementation `AMQPToObjectCollectorSerializable` that employs default serialization.

For JSON formatted messages, the adapter provides the `AMQPToObjectCollectorJson` JSON collector. The collector deserializes the AMQP message bytes into a string using the JVM default encoding and parses the JSON message for processing as an event. You must specify a JSON event type name e.g. `create json schema MyJsonEventType (...) and outstream<MyJsonEventType>`.

The following example declares a data flow that receives AMQP messages from a queue, transforms each message and sends each message of type `MyMapEventType` into the event bus:

```
create dataflow AMQPIncomingDataFlow
  AMQPSource -> outstream<MyMapEventType> {
    host: 'localhost',
    queueName: 'myqueue',
    collector: {class: 'AMQPToObjectCollectorSerializable'},
    logMessages: true
  }
  EventBusSink(outstream) {}
```

## Chapter 5. The Kafka Adapter

### [5.1. Classpath Setup](#)

### [5.2. Imports Setup](#)

### [5.3. Input Adapter](#)

#### [5.3.1. Input Adapter Configuration and Start](#)

#### [5.3.2. Kafka Connectivity](#)

#### [5.3.3. Controlling Input Adapter Operation](#)

### [5.4. Output Adapter](#)

#### [5.4.1. Output Adapter Configuration and Start](#)

#### [5.4.2. Kafka Connectivity](#)

#### [5.4.3. Controlling Output Adapter Operation](#)

This chapter discusses the EsperIO Kafka input adapter.

This input adapter is for receiving events and event or runtime time from Kafka topics.

The scope of this input adapter is a local reader and is not meant for coordinated use by multiple servers, which is the scope of Esper Enterprise Edition. Please see Esper Enterprise Edition for information on the horizontal scale-out architecture based on Kafka (the scope of this input adapter is NOT horizontal scale-out).

## 5.1. Classpath Setup

Please add the `esperio-kafka-version.jar` jar file to your classpath.

Please also add `kafka-clients-version.jar` and the Kafka client dependencies to your classpath.

The EsperIO Kafka input adapter supports the new Kafka consumer only and requires Kafka client version 1.0.1 and higher.

The adapter only requires Esper runtime as a dependency.

## 5.2. Imports Setup

For use with the Kafka output adapter, and when using the `KafkaOutputDefault` annotation, please add the `KafkaOutputDefault` import. For example:

```
configuration.getCommon().addImport(KafkaOutputDefault.class);
```

## 5.3. Input Adapter

### 5.3.1. Input Adapter Configuration and Start

You may configure and start the EsperIO Kafka input adapter either as part of your Esper configuration file in the plugin loader section or via the adapter API.

The following example shows an Esper configuration file with all properties:

```

<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <plugin-loader name="KafkaInput" class-name="com.espertech.esperio.kafka.EsperIOKafkaInputAdapter" />
    <!--
      Kafka Consumer Properties: Passed-Through to Kafka Consumer.
    -->
    <init-arg name="bootstrap.servers" value="127.0.0.1:9092"/>
    <init-arg name="key.deserializer" value="org.apache.kafka.common.serialization.StringDeserializer"/>
    <init-arg name="value.deserializer" value="com.mycompany.MyCustomDeserializer"/>
    <init-arg name="group.id" value="my_group_id"/>

    <!--
      EsperIO Kafka Input Properties: Define subscription, topics, processor and timestamp extractor.
    -->
    <init-arg name="esperio.kafka.input.subscriber" value="com.espertech.esperio.kafka.EsperIOKafkaInputSubscriberByTopicList"/>
    <init-arg name="esperio.kafka.topics" value="my_topic"/>
    <init-arg name="esperio.kafka.input.processor" value="com.espertech.esperio.kafka.EsperIOKafkaInputProcessorDefaultClass"/>
    <init-arg name="esperio.kafka.input.timestampextractor" value="com.espertech.esperio.kafka.EsperIOKafkaInputTimestampExtractor"/>
  </plugin-loader>
</runtime>
</esper-configuration>

```

Alternatively the equivalent API calls to configure the adapter are:

```

Properties props = new Properties();

// Kafka Consumer Properties
props.put(ConsumerConfig.BOOTSTRAP_SERVERS_CONFIG, "127.0.0.1:9092");
props.put(ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG, org.apache.kafka.common.serialization.StringDeserializer.class);
props.put(ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG, com.mycompany.MyCustomDeserializer.class);
props.put(ConsumerConfig.GROUP_ID_CONFIG, "my_group_id");

// EsperIO Kafka Input Adapter Properties
props.put(EsperIOKafkaConfig.INPUT_SUBSCRIBER_CONFIG, EsperIOKafkaInputSubscriberByTopicList.class);
props.put(EsperIOKafkaConfig.TOPICS_CONFIG, "my_topic");
props.put(EsperIOKafkaConfig.INPUT_PROCESSOR_CONFIG, EsperIOKafkaInputProcessorDefaultClass.class);
props.put(EsperIOKafkaConfig.INPUT_TIMESTAMPEXTRACTOR_CONFIG, EsperIOKafkaInputTimestampExtractor.class);

Configuration config = new Configuration();
config.getRuntime().addPluginLoader("KafkaInput", EsperIOKafkaInputAdapterPlugin.class);

```

By adding the plug-in loader to the configuration as above the runtime automatically starts the adapter as part of runtime initialization.

Alternatively, the adapter can be started and stopped programatically as follows:

```

// start adapter
EsperIOKafkaInputAdapter adapter = new EsperIOKafkaInputAdapter(props, "default");
adapter.start();

// destroy the adapter when done
adapter.destroy();

```

### 5.3.2. Kafka Connectivity

All properties are passed to the Kafka consumer. This allows your application to add additional properties that are not listed here and according to Kafka consumer documentation.

Required properties are below. `ConsumerConfig` is part of the Kafka API in `org.apache.kafka.clients.consumer.ConsumerConfig`.

Table 5.1. Kafka Consumer Required Properties

Name	API Name	Description
bootstrap.servers	ConsumerConfig.BOOTSTRAP_SERVERS_CONFIG	Kafka bootstrap server list.
key.deserializer	ConsumerConfig.KEY_DESERIALIZER_CLASS_CONFIG	Fully-qualified class name of Kafka message key deserializer.
value.deserializer	ConsumerConfig.VALUE_DESERIALIZER_CLASS_CONFIG	Fully-qualified class name of Kafka message value deserializer.
group.id	ConsumerConfig.GROUP_ID_CONFIG	Application consumer group id.

### 5.3.3. Controlling Input Adapter Operation

The input adapter operation depends on the *subscriber* and *processor*.

The *subscriber* is responsible for calling Kafka consumer subscribe methods, i.e. calls Kafka API `consumer.subscribe(...)`.

The *processor* is responsible for processing Kafka API `ConsumerRecords` messages, i.e. implements `process(ConsumerRecords records)`.

Properties that define the subscriber and consumer are below. `EsperIOKafka` is part of the EsperIO Kafka API in `com.espertech.esperio.kafka.EsperIOKafkaConfig`.

Table 5.2. Kafka Input Adapter Properties

Name	API Name	Description
esperio.kafka.input.subscriber	EsperIOKafkaConfig.INPUT_SUBSCRIBER_CONFIG	Required property. Fully-qualified class name of subscriber that subscribes to Kafka topics. The class must implement the interface <code>EsperIOKafkaInputSubscriber</code> . You may use <code>com.espertech.esperio.kafka.EsperIOKafkaInputSubscriberByTopicList</code> to provide a topic list in <code>esperio.kafka.topics</code> .
esperio.kafka.topics	EsperIOKafkaConfig.TOPICS_CONFIG	Optional property and only required if using <code>EsperIOKafkaInputSubscriberByTopicList</code> . Specifies a comma-separated list of topic names to subscribe to.

Name	API Name	Description
esperio.kafka.input.processor	EsperIOKafkaConfig.INPUT_PROCESSOR_CONFIG	Required property. Fully-qualified class name of the Kafka consumer events into the runtime and may advance runtime time. The class must implement the interface <code>EsperIOKafkaInputProcessor</code> . You may use <code>com.espertech.esperio.kafka.EsperIOKafkaInputProcessorDefault</code> as the default event and time processing.
esperio.kafka.input.timestampextractor	EsperIOKafkaConfig.INPUT_TIMESTAMPEXTRACTOR_CONFIG	Optional property. Fully-qualified class name of the Kafka message time long-typed timestamp from a consumer record, for use with the processor. The class must implement the interface <code>EsperIOKafkaInputTimestampExtractor</code> . You may use <code>com.espertech.esperio.kafka.EsperIOKafkaInputTimestampExtractorDefault</code> which returns the time of each consumer record that it processes.
esperio.kafka.input.eventtypename	EsperIOKafkaConfig.INPUT_EVENTTYPE_NAME	Optional property. Event type name of the JSON event type; <code>EsperIOKafkaInputProcessorJson</code> only.

### 5.3.3.1. Subscriber

The subscriber is responsible for calling `consumer.subscribe(...)`.

The adapter provides a default implementation by name `EsperIOKafkaInputSubscriberByTopicList`. Your application may provide its own subscriber by implementing the simple `EsperIOKafkaInputSubscriber` interface.

This default implementation takes the value of `esperio.kafka.topics` and subscribes to each topic.

For reference, we provide the code of the default subscriber below (repository or source jar for full code):

```
public class EsperIOKafkaInputSubscriberByTopicList implements EsperIOKafkaInputSubscriber {
    public void subscribe(EsperIOKafkaInputSubscriberContext context) {
        String topicsCSV = EsperIOKafkaInputAdapter.getRequiredProperty(context.getProperties(), "topics");
        String[] topicNames = topicsCSV.split(",");
        List<String> topics = new ArrayList<>();
        for (String topicName : topicNames) {
            if (topicName.trim().length() > 0) {
                topics.add(topicName.trim());
            }
        }
        context.getConsumer().subscribe(topics);
    }
}
```

### 5.3.3.2. Processor

The processor is responsible for processing Kafka API `ConsumerRecords`.

The adapter provides a default implementation by name `EsperIOKafkaInputProcessorDefault`. Your application may provide its own processor by implementing the simple `EsperIOKafkaInputProcessor` interface.

This default processor can be configured with an optional timestamp extractor that obtains a timestamp for each consumer record. If no timestamp extractor is configured, the default processor does not advance time.

For reference, we provide the (slightly simplified) code of the default processor below (repository or source jar for full code):

```
public class EsperIOKafkaInputProcessorDefault implements EsperIOKafkaInputProcessor {
    private EPRuntime runtime;
    private EsperIOKafkaInputTimestampExtractor timestampExtractor;

    public void init(EsperIOKafkaInputProcessorContext context) {
        this.runtime = context.getRuntime();

        String timestampExtractorClassName = context.getProperties().getProperty("timestampExtractor");
        if (timestampExtractorClassName != null) {
            timestampExtractor = (EsperIOKafkaInputTimestampExtractor) JavaClassHelper.instantiate(timestampExtractorClassName);
        }
    }

    public void process(ConsumerRecords<Object, Object> records) {
        for (ConsumerRecord record : records) {
            if (timestampExtractor != null) {
                long timestamp = timestampExtractor.extract(record);
                // advances runtime time
                runtime.getEPRuntime().advanceTimeSpan(timestamp);
            }

            if (record.value() != null) {
                Object event = record.value();
                runtime.getEPRuntime().sendEventBean(event, event.getClass().getSimpleName());
            }
        }
    }

    public void close() {}
}
```

The default processor takes the message value and sends it as an event into the runtime. It assumes the event type name is the simple class name. The default processor takes the extracted time, if a timestamp extractor is provided, and sends a time span event to the runtime to advance runtime time.

For use with JSON, `EsperIO` provides the `EsperIOKafkaInputProcessorJson` processor. Please define the

esperio.kafka.input.eventtypename property that provides the JSON event type name.

You must provide your own processor if any additional event transformation is required or if using `epRuntime.sendEvent(Map/ObjectArray/Node)` or if the default behavior does not fit for other reasons.

## 5.4. Output Adapter

### 5.4.1. Output Adapter Configuration and Start

You may configure and start the EsperIO Kafka output adapter either as part of your Esper configuration file in the plugin loader section or via the adapter API.

The following example shows an Esper configuration file with all properties:

```
<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <plugin-loader name="KafkaOutput" class-name="com.espertech.esperio.kafka.EsperIOKafkaOutputAdapter">
      <!--
        Kafka Producer Properties: Passed-Through to Kafka Consumer.
      -->
      <init-arg name="bootstrap.servers" value="127.0.0.1:9092"/>
      <init-arg name="key.serializer" value="org.apache.kafka.common.serialization.StringSerializer"/>
      <init-arg name="value.serializer" value="com.mycompany.MyCustomSerializer"/>

      <!--
        EsperIO Kafka Output Properties: Define a flow controller.
      -->
      <init-arg name="esperio.kafka.output.flowcontroller" value="com.espertech.esperio.kafka.EsperIOKafkaOutputFlowController"/>
      <init-arg name="esperio.kafka.topics" value="my_topic"/>
    </plugin-loader>
  </runtime>
</esper-configuration>
```

Alternatively the equivalent API calls to configure the adapter are:

```
Properties props = new Properties();

// Kafka Producer Properties
props.put(ProducerConfig.BOOTSTRAP_SERVERS_CONFIG, "127.0.0.1:9092");
props.put(ProducerConfig.KEY_SERIALIZER_CLASS_CONFIG, org.apache.kafka.common.serialization.StringSerializer.class);
props.put(ProducerConfig.VALUE_SERIALIZER_CLASS_CONFIG, org.apache.kafka.common.serialization.StringSerializer.class);

// EsperIO Kafka Output Adapter Properties
props.put(EsperIOKafkaConfig.OUTPUT_FLOWCONTROLLER_CONFIG, EsperIOKafkaOutputFlowController.class);
props.put(EsperIOKafkaConfig.TOPICS_CONFIG, "my_topic");

Configuration config = new Configuration();
config.getRuntime().addPluginLoader("KafkaOutput", EsperIOKafkaOutputAdapterPlugin.class);
```

By adding the plug-in loader to the configuration as above the runtime automatically starts the adapter as part of runtime initialization.

Alternatively, the adapter can be started and stopped programmatically as follows:

```
// start adapter
EsperIOKafkaOutputAdapter adapter = new EsperIOKafkaOutputAdapter(props, "default");
adapter.start();

// destroy the adapter when done
adapter.destroy();
```

### 5.4.2. Kafka Connectivity

All properties are passed to the Kafka producer. This allows your application to add additional properties that are not listed here and according to Kafka producer documentation.

Required properties are below. `ProducerConfig` is part of the Kafka API in `org.apache.kafka.clients.producer.ProducerConfig`.

Table 5.3. Kafka Producer Required Properties

Name	API Name	Description
<code>bootstrap.servers</code>	<code>ProducerConfig.BOOTSTRAP_SERVERS_CONFIG</code>	Kafka bootstrap server list.
<code>key.serializer</code>	<code>ProducerConfig.KEY_SERIALIZER_CLASS_CONFIG</code>	Fully-qualified class name of Kafka message key serializer.
<code>value.serializer</code>	<code>ProducerConfig.VALUE_SERIALIZER_CLASS_CONFIG</code>	Fully-qualified class name of Kafka message value serializer.

### 5.4.3. Controlling Output Adapter Operation

The output adapter operation depends on the *flow controller*, which is responsible for attaching listeners to statements that send messages to Kafka topics.

Properties that define the flow controller are below. `EsperIOKafka` is part of the EsperIO Kafka API in `com.espertech.esperio.kafka.EsperIOKafkaConfig`.

Table 5.4. Kafka Output Adapter Properties

Name	API Name	Description
<code>esperio.kafka.output.flowcontroller</code>	<code>EsperIOKafkaConfig.OUTPUT_FLOWCONTROLLER_CONFIG</code>	Required property. Fully-qualified class name of flow controller that produces messages. The class must implement the interface <code>EsperIOKafkaOutputFlowController</code> . You may use <code>com.espertech.esperio.kafka.EsperIOKafkaOutputFlowController</code> and provide a topic list in <code>esperio.kafka.topics</code> .
<code>esperio.kafka.topics</code>	<code>EsperIOKafkaConfig.TOPICS_CONFIG</code>	Specifies a comma-separated list of topic names to produce messages to above flow controller and not required otherwise.

#### 5.4.3.1. Flow Controller

The flow controller is responsible for allocating a `KafkaProducer` and associating statement listeners to the producer, for listeners to send messages to Kafka topics.

The adapter provides a default implementation by name `EsperIOKafkaOutputFlowControllerByAnnotatedStmt`. Your application may provide its own subscriber by implementing the simple `EsperIOKafkaOutputFlowControllerContext` interface.

#### 5.4.3.1.1. Default Flow Controller `EsperIOKafkaOutputFlowControllerByAnnotatedStmt`

The flow controller takes the value of `esperio.kafka.topics` and produces a message to each topic for each statement listener output event.

The flow controller attaches a listener to all statements that have the `@KafkaOutputDefault` annotation. Please ensure that the annotation is part of your imports. The adapter considers all newly-created statements that have the annotation.

Thus please create the EPL as follows:

```
@KafkaOutputDefault select * from .....
```

The flow controller produces JSON output. It uses the runtime JSON renderer that can be obtained from `runtime.getRenderEventService().getJSONRenderer(statement.getEventType())`:

The statement listeners that the flow controller attaches do not provide a key or partition id to the producer. The listeners simply invoke `new ProducerRecord(topic, json)` for output event and each topic. The value serializer must be the string serializer.

For reference, please find the source code of the flow controller in the repository.

## Chapter 6. The HTTP Adapter

### [6.1. Adapter Overview](#)

### [6.2. Getting Started](#)

#### [6.2.1. Plugin Loader Configuration](#)

#### [6.2.2. Configuration and Starting via API](#)

### [6.3. HTTP Input Adapter](#)

#### [6.3.1. HTTP Service](#)

#### [6.3.2. Get Handlers](#)

#### [6.3.3. HTTP Input Limitations](#)

### [6.4. HTTP Output Adapter](#)

#### [6.4.1. Triggered HTTP Get](#)

This chapter discusses the EsperIO HTTP adapter.

The output adapter requires Esper compiler as well as Esper runtime as a dependency. The input adapter only requires Esper runtime as a dependency.

## 6.1. Adapter Overview

The EsperIO HTTP input and output adapter can be used to send events into an Esper runtime instance as well as perform HTTP requests triggered by output events generated by an Esper runtime instance.

To send events into an Esper runtime instance for processing you declare an HTTP service, which causes the adapter to expose an HTTP protocol server on the configured port to handle incoming requests. Your configuration then attaches Get handlers that receive Get requests that post events into the runtime with data from each request.

Output events generated by an Esper runtime instance can trigger an HTTP Get operation to a URI of your choice. For this purpose define a triggering event stream and the desired target URI and parameters.

## 6.2. Getting Started

You may configure the EsperIO HTTP adapter either as part of your Esper configuration file in the plugin loader section or via the adapter API. Add the `esperio-http-version.jar` file to your classpath.

For input adapter operation, add the `httpcore-version.jar` to your classpath. If using Java NIO add the `httpcore-nio-version.jar` to your classpath in addition.

For output adapter operation, add the `httpclient-version.jar` to your classpath.

A sample adapter configuration file is provided in `esperio-http-sample-config.xml` in the `etc` folder of the distribution. A configuration file must be valid according to schema `esperio-http-configuration-8-0.xsd`.

### 6.2.1. Plugin Loader Configuration

You may place the configuration XML directly into your Esper configuration file as the example below shows:

```
<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <plugin-loader name="EsperIOHTTPAdapter"
      class-name="com.espertech.esperio.http.EsperIOHTTPAdapterPlugin">
      <init-arg name="esperio.http.configuration.file"
        value="file:/path/esperio-http-sample-config.xml" />
    </plugin-loader>
  </runtime>
</esper-configuration>
```

Alternatively you can provide a URL in the Esper configuration file to point to your adapter configuration file:

```
<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <plugin-loader name="EsperIOHTTPAdapter"
      class-name="com.espertech.esperio.http.EsperIOHTTPAdapterPlugin">
      <init-arg name="esperio.http.configuration.file"
        value="file:/path/esperio-http-sample-config.xml" />
    </plugin-loader>
  </runtime>
</esper-configuration>
```

### 6.2.2. Configuration and Starting via API

If using Spring or if your application requires API access, the following code snippet configures and starts the adapter via API.

The class for configuring an EsperIO HTTP adapter is `com.espertech.esperio.http.config.ConfigurationHTTPAdapter`. The adapter class itself is `EsperIOHTTPAdapter`.

The code snippet below is a sample that configures using driver manager and starts the adapter via API:

```
ConfigurationHTTPAdapter adapterConfig = new ConfigurationHTTPAdapter();

// add additional configuration
Request request = new Request();
request.setStream("TriggerEvent");
request.setUri("http://localhost:8077/root");
adapterConfig.getRequests().add(request);

// start adapter
EsperIOHTTPAdapter httpAdapter = new EsperIOHTTPAdapter(adapterConfig, "runtimeURI");
httpAdapter.start();

// destroy the adapter when done
httpAdapter.destroy();
```

## 6.3. HTTP Input Adapter

### 6.3.1. HTTP Service

A service is required for the adapter to receive events via a HTTP client connection.

The synopsis is as follows:

```
<esperio-http-configuration>
  <service name="[name]" port="[port]" [nio="true|false"]/>
  <!-- add additional configuration here -->
</esperio-http-configuration>
```

The `name` attribute value is required and provides the name of the HTTP service for use in logging and for get-handlers as described below.

The `nio` attribute is optional and can be used to enable Java NIO (disabled by default).

If configuring via the adapter API or Spring, use the `com.espertech.esperio.http.config.Service` class.

An example XML to configure a service and single get-handler is:

```
<esperio-http-configuration>
  <service name="myservice" port="8079" nio="false"/>
  <get service="myservice" pattern="*/>
</esperio-http-configuration>
```

### 6.3.2. Get Handlers

One or more handlers for HTTP Get operations can be installed for a service and are used to receive events.

Define a `get` element in the adapter configuration file (or use the `GetRequest` class) for every handler to register for a service.

The synopsis is as follows:

```
<get service="[service]" pattern="[pattern]"/>
```

The `service` attribute value is required and provides the name of the HTTP service to register the Get operation handler for.

A value for the `pattern` attribute is required and may be either `*` for all URIs, `*[uri]` for all URIs ending with the given URI or `[uri]*` for all URI starting with the given URI.

A sample Get-handler configuration follows:

```
<get service="myservice" pattern="*/>
```

When posting events to the runtime, the Get request URI must contain a `stream` parameter that carries the name of the stream (event type) to insert into. Each event property to be populated in the input event must be part of the Get request parameter values.

For example, the URI `http://localhost:8079/sendevent?stream=MyFirewallEvent&name=Joe&changed=true` entered into a browser sends an input event of type `MyFirewallEvent` setting the `name` property of the event to "Joe" and the `changed` property of the event to true.

Note that if your target type is a Java object event, your event class must provide setter-methods according to JavaBean conventions. The event class should also provide a default constructor taking no parameters. If your event class does not have a default constructor, your application may configure a factory method via `ConfigurationEventTypeLegacy`.

### 6.3.3. HTTP Input Limitations

The XML DOM target type is not supported. There is no support for building a DOM object from the HTTP input and sending that as an event into the runtime.

## 6.4. HTTP Output Adapter

### 6.4.1. Triggered HTTP Get

This facility instructs the adapter to perform an HTTP Get request when a triggering event occurs, passing event properties as URI parameters.

Define a `request` element in the adapter configuration file (or use the `Request` class) for every HTTP Get to execute.

The synopsis is as follows:

```
<request stream="[stream]" uri="[uri_with_placeholders]"/>
```

A value for the `stream` attribute is required and provides the name of the stream that triggers the HTTP Get. The adapter expects a stream by this name to exist at adapter start time.

The `uri_with_placeholders` attribute value is required. You may place event property placeholders inside the URI to format the URI as needed. Placeholders are of the format `$(property_name)`.

A sample request configuration follows:



```
<request stream="TriggerFirewallStream" uri="http://myremotehost:80/root/event"/>
```

Assuming the `HttpTriggerStream` has event properties `name` and `ipaddress` then a sample `Get` request URI is as follows:

```
http://myremotehost:80/root/event?stream=TriggerFirewallStream&name=Joe&ipaddress=120.1.0.0
```

You may parameterize the URI via placeholders by placing `$(property_name)` and the special placeholder `$(stream)` into the URI string.

The next example configuration defines URI parameters via placeholder:

```
<request stream="TriggerFirewallStream" uri="http://myremotehost:80/root/$(stream)?violation&name=Joe&violationip=120.1.0.0"/>
```

The URI generated by the adapter:

```
http://myremotehost:80/root/TriggerFirewallStream?violation&name=Joe&violationip=120.1.0.0
```

## Chapter 7. The Socket Adapter

### [7.1. Getting Started](#)

#### [7.1.1. Plugin Loader Configuration](#)

#### [7.1.2. Configuration and Starting via API](#)

### [7.2. Socket Service](#)

#### [7.2.1. Object Data Format](#)

#### [7.2.2. String CSV Data Format](#)

#### [7.2.3. String CSV Data Format With Property Order](#)

#### [7.2.4. JSON Formatted](#)

This chapter discusses the EsperIO Socket adapter.

The EsperIO Socket input adapter can be used to send events into an Esper runtime via socket client, either as Java objects or as CSV name-value pair strings.

The adapter only requires Esper runtime as a dependency.

## 7.1. Getting Started

You may configure the EsperIO Socket adapter either as part of your Esper configuration file in the plugin loader section or via the adapter API. Add the `esperio-socket-version.jar` file to your classpath. There are no other dependent jar files required.

A sample adapter configuration file is provided in `esperio-socket-sample-config.xml` in the `etc` folder of the distribution. A configuration file must be valid according to schema `esperio-socket-configuration-8-0.xsd`.

### 7.1.1. Plugin Loader Configuration

You may place the configuration XML directly into your Esper configuration file as the example below shows:

```
<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <plugin-loader name="EsperIOSocketAdapter"
      class-name="com.espertech.esperio.socket.EsperIOSocketAdapterPlugin">
      <config-xml>
        <esperio-socket-configuration>
          <socket name="mysocketOne" port="7101" data="object"/>
          <socket name="mysocketTwo" port="7102" data="csv" hostname="myhost" backlog="20" />
        </esperio-socket-configuration>
      </config-xml>
    </plugin-loader>
  </runtime>
</esper-configuration>
```

Alternatively you can provide a URL in the Esper configuration file to point to your adapter configuration file:

```
<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <plugin-loader name="EsperIOSocketAdapter"
      class-name="com.espertech.esperio.socket.EsperIOSocketAdapterPlugin">
      <init-arg name="esperio.socket.configuration.file"
        value="file:/path/esperio-socket-sample-config.xml" />
    </plugin-loader>
  </runtime>
</esper-configuration>
```

### 7.1.2. Configuration and Starting via API

If using Spring or if your application requires API access, the following code snippet configures and starts the adapter via API.

The class for configuring an EsperIO Socket adapter is `com.espertech.esperio.socket.config.ConfigurationSocketAdapter`. The adapter class itself is `EsperIOSocketAdapter`.

The code snippet below is a sample that configures using driver manager and starts the adapter via API:

```

ConfigurationSocketAdapter adapterConfig = new ConfigurationSocketAdapter();

SocketConfig socket = new SocketConfig();
socket.setDataType(DataType.CSV);
socket.setPort(port);
adapterConfig.getSockets().put("SocketService", socket);

// start adapter
EsperIOSocketAdapter socketAdapter = new EsperIOSocketAdapter(adapterConfig, "runtimeURI");
socketAdapter.start();

// destroy the adapter when done
socketAdapter.destroy();

```

## 7.2. Socket Service

Add a socket configuration for each unique port that you want to expose a socket receive service for use by socket client connections.

The synopsis is as follows:

```

<esperio-socket-configuration>
  <socket name="[name]" port="[port]" data="[csv|object|property_ordered_csv]"
    [hostname="hostname"] [backlog="backlog"] [unescape="true|false"]/>
</esperio-socket-configuration>

```

The required *name* attribute provides the name of the socket service for use in logging.

The required *port* attribute provides the port that the socket service accepts client connections.

The required *data* attribute specifies whether the data arriving through the socket is formatted as a Java binary object stream or as CSV string values.

The optional *hostname* attribute can provide the host name passed to the server socket (*ServerSocket*).

The optional *backlog* attribute can provide the backlog number of connections passed to the server socket. This number defaults to 2 when a host name is passed but no backlog is provided.

The optional *unescape* attribute is false by default. When false the adapter does not unescape (Java escape rules) values. When true the adapter performs an unescape on all values.

If configuring via the adapter API or Spring, use the `com.esperitech.esperio.socket.config.SocketConfig` class.

### 7.2.1. Object Data Format

When sending events as Java objects, configure the *data* attribute value to `object` and use `ObjectOutputStream` to write events to the client socket. When sending a `java.util.Map` event, your `Map` must contain a `String` value for the key `stream` which must denote a configured `Map` event type.

This example XML configures a socket accepting client connections that provide Java objects:

```

<esperio-socket-configuration>
  <socket name="objectStreamSocket" port="8079" data="object"/>
</esperio-socket-configuration>

```

When object data type is configured, clients connections are expected to send `java.io.Serializable` or `java.io.Externalizable` objects using `ObjectOutputStream`.

Below is a block of sample code that for use in clients to the adapter. It assumes the `MyEvent` class implements either of the above interfaces:

```

// connect first
Socket requestSocket = new Socket("localhost", port);
ObjectOutputStream out = new ObjectOutputStream(requestSocket.getOutputStream());

// send a few events, here we send only one
out.writeObject(new MyEvent("Hello World"));
out.flush();

// Consider resetting the output stream from time-to-time, after sending a number of object
// This is because the stream may cache strings etc. . The reset is:
// out.reset();

// close when done
out.close();
requestSocket.close();

```

### 7.2.2. String CSV Data Format

When sending events as CSV strings, the format of the string should be:

```
stream=[type],[name]=[value] [,...] (newline)
```

The CSV string must end with a newline character: Each event is one line. Each CSV element must be in the `[name]=[value]` format. Your CSV must contain a value for `stream` which must denote a configured event type. The adapter parses each string value and populates an instance of the target type.

This next example XML configures a socket accepting client connections that provide events as CSV-formatted strings with name-value pairs :

```

<esperio-socket-configuration>
  <socket name="csvStreamSocket" port="8079" data="csv"/>
</esperio-socket-configuration>

```

A piece of client code that sends an event of type `MyEvent` may look as follows:

```

// connect first
String newline = System.getProperty("line.separator");
Socket requestSocket = new Socket("localhost", port);
BufferedWriter wr = new BufferedWriter(
    new OutputStreamWriter(socket.getOutputStream()));

// send a few events
wr.write("stream=MyEvent,price=20.0,upcCode=A0001");
wr.write(newline);
wr.flush();

// close when done
wr.close();
requestSocket.close();

```

Note that if your target type is a Java object event, your event class must provide setter-methods according to JavaBean conventions. The event class should also provide a default constructor taking no parameters. If your event class does not have a default constructor, your application may configure a factory method via `ConfigurationEventTypeLegacy`.

### 7.2.3. String CSV Data Format With Property Order

Similar to the string CSV data format as discussed earlier, this data format allows specifying a property order as well as the event type name.

The format of the string that represents an event is the CSV format (no `name=` or `stream=` texts are required):

```
value [, value [,...]] (newline)
```

This next example XML configures a socket accepting client connections that provide events as CSV-formatted strings with name-value pairs :

```

<esperio-socket-configuration>
  <socket name="csvStreamSocket" port="8079" data="property_ordered_csv" stream="MyEvent" />
</esperio-socket-configuration>

```

Set the `data` attribute to `property_ordered_csv`. The `stream` attribute is a required configuration and must contain the event type name. The `propertyOrder` attribute is also required and must contain the property names of properties of the event type, in the same order that the value for the property arrives on each line.

As part of the sample client code shown above, the following line sends an event with values 20.0 and A0001.

```

wr.write("20.0,A0001");
wr.write(newline);

```

### 7.2.4. JSON Formatted

The JSON-formatted message can be processed by following below format. Please make sure to use create json schema to declare the respective JSON event type.

The format of the string that represents an event in JSON format is:

```
stream=eventtypename,json=jsoncontext (newline)
```

The `eventtypename` is the JSON event type name. The `jsoncontext` is the JSON-formatted string without newline characters. The message must end with newline.

This XML snippet configures a socket accepting client connections that provide events as JSON-formatted strings:

```

<esperio-socket-configuration>
  <socket name="jsonStreamSocket" port="8079" data="json" stream="MyEvent"/>
</esperio-socket-configuration>

```

A sample message is:

```
stream=MyEvent,json={"price":20,"upcCode":"A0001"}
```

Each message must end with a newline character. The above example assumes a JSON event type such as `@public @buseventtype create json schema MyEvent(price double, upcCode string)`.

## Chapter 8. The Relational Database Adapter

### [8.1. Adapter Overview](#)

### [8.2. Getting Started](#)

#### [8.2.1. Plugin Loader Configuration](#)

#### [8.2.2. Configuration and Starting via API](#)

### [8.3. JDBC Connections](#)

### [8.4. Triggered DML Statement Execution](#)

### [8.5. Triggered Update-Insert Execution](#)

### [8.6. Executor Configuration](#)

### [8.7. Reading and Polling Database Tables](#)

#### [8.7.1. Polling and Startup SQL Queries](#)

This chapter discusses the EsperIO adapter for relational databases.

This adapter requires Esper compiler as well as Esper runtime as a dependency.

## 8.1. Adapter Overview

The EsperIO relational database adapter can write events to a database table.

If your application only reads from tables, the Esper jar file and Esper configuration suffices and there is no additional EsperIO DB adapter configuration or jar file required. Please see below tips for reading or polling tables.

The EsperIO DB adapter supports two means to write to a database table:

1. Execute a SQL DML (Data Manipulation, i.e. Update, Insert, Delete or stored procedure call) statement as a response to a triggering event.
2. Execute an Update-Insert: The adapter attempts an Update of a row by key and if unsuccessful

(update returns zero rows updated) the adapter performs an Insert.

The adapter also provides infrastructure for queuing database write requests for execution by a thread pool.

## 8.2. Getting Started

You may configure the EsperIO DB adapter either as part of your Esper configuration file in the plugin loader section or via the adapter API. Add the `esperio-db-version.jar` file to your classpath as well as the JDBC driver. There are no other dependent jar files required by the adapter.

A sample adapter configuration file is provided in `esperio-db-sample-config.xml` in the `etc` folder of the distribution. A configuration file must be valid according to schema `esperio-db-configuration-8-0.xsd`.

### 8.2.1. Plugin Loader Configuration

You may place the configuration XML directly into your Esper configuration file as the example below shows:

```
<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <plugin-loader name="EsperIODBAdapter"
      class-name="com.espertech.esperio.db.EsperIODBAdapterPlugin">
      <config-xml>
        <esperio-db-configuration>
          <!-- ....as outlined below or contents of esperio-db-sample-config.xml here... -->
        </esperio-db-configuration>
      </config-xml>
    </plugin-loader>
  </runtime>
</esper-configuration>
```

Alternatively you can provide a URL in the Esper configuration file to point to your adapter configuration file:

```
<esper-configuration xmlns="http://www.espertech.com/schema/esper">
  <runtime>
    <plugin-loader name="EsperIODBAdapter"
      class-name="com.espertech.esperio.db.EsperIODBAdapterPlugin">
      <init-arg name="esperio.db.configuration.file"
        value="file:/path/esperio-db-sample-config.xml" />
    </plugin-loader>
  </runtime>
</esper-configuration>
```

### 8.2.2. Configuration and Starting via API

If using Spring or if your application requires API access, the following code snippet configures and starts the adapter via API.

The class for configuring an EsperIO DB adapter is `com.espertech.esperio.db.config.ConfigurationDBAdapter`. The adapter class itself is `EsperIODBAdapter`.

The code snippet below is a sample that configures using driver manager and starts the adapter via API:

```
ConfigurationDBAdapter adapterConfig = new ConfigurationDBAdapter();
ConfigurationCommonDBRef configDB = new ConfigurationCommonDBRef();
configDB.setDriverManagerConnection("DRIVER", "URL", new Properties());
adapterConfig.getJdbcConnections().put("db1", configDB);

// add additional configuration such as DML and Upsert

// start adapter
EsperIODBAdapter dbAdapter = new EsperIODBAdapter(adapterConfig, "runtimeURI");
dbAdapter.start();
```

## 8.3. JDBC Connections

The configuration for the source of JDBC connections follows the Esper configuration. Please consult the Esper documentation or sample adapter configuration file for details.

Your configuration should set `auto-commit` to true thereby each change to database tables is automatically committed.

The adapter obtains a new connection for each operation and closes the connection after each operation. For optimum performance consider configuring a connection pool.

A sample JDBC connection configuration is shown in below XML. The API class is `ConfigurationCommonDBRef`. You may also configure a `DataSource` or `DataSource` factory as outlined in the Esper docs.

```
<esperio-db-configuration>
  <jdbc-connection name="db1">
    <drivermanager-connection class-name="com.mysql.jdbc.Driver" url="jdbc:mysql://localhost:3306/db1"
      user="root" password="password">
    <connection-settings auto-commit="true" catalog="TEST"/>
    </drivermanager-connection>
  </jdbc-connection>
  <!-- Add DML and Upsert configurations here, as below. -->
</esperio-db-configuration>
```

## 8.4. Triggered DML Statement Execution

This facility allows running a SQL DML (Data Manipulation) query, i.e. an Update, Insert, Delete query or a stored procedure when an event in a triggering stream occurs.

Define a `dml` element in the adapter configuration file (or use the `DMLQuery` class) for every query to execute.

The synopsis is as follows:

```

<dml connection="[connection]" stream="[stream]"
  [name="name"] [executor-name="executor"] [retry="count"] [retry-interval-sec="sec"]>
  <sql>[sql]</sql>
  <bindings>
    <parameter pos="[position]" property="[property_name]" />
    [...parameters]
  </bindings>
</dml>

```

The *connection* attribute value is required and provides the name of the configured JDBC connection.

A value for the *stream* attribute is required and provides the name of the stream that triggers the DML. The adapter expects a stream by this name to exist at adapter start time.

The *name* attribute is optional and is only used for logging errors.

The *executor-name* attribute is optional. If specified, it must be the name of an executor configuration. If specified, the adapter will use the executor service (queue and thread pool) for performing all DML work. If not specified, the adapter performs the DML work in the same thread.

The *retry* attribute is optional. If specified, the adapter will retry a given number of times in case an error is encountered. If *retry-interval-sec* is specified, the adapter waits the given number of seconds between retries.

The *sql* element is required and provides the SQL DML or stored procedure call to execute, with parameters as question mark (?).

The *bindings* element is required and provides the bindings for expression parameters.

The *parameter* element should occur as often as there are parameters in the SQL query. The *position* attribute starts at 1 and counts up for each parameter. The *property* parameter provide the name of the event property of the stream to use as the parameter value.

A sample DML configuration follows:

```

<dml connection="db1" stream="InsertToDBStream"
  name="MyInsertQuery" executor-name="queue1" retry="count">
  <sql>insert into MyEventStore(key1, value1, value2) values (?, ?, ?)</sql>
  <bindings>
    <parameter pos="1" property="eventProperty1" />
    <parameter pos="2" property="eventProperty2" />
    <parameter pos="3" property="eventProperty3" />
  </bindings>
</dml>

```

## 8.5. Triggered Update-Insert Execution

This facility allows running an SQL Update that is followed by an Insert if the Update did not update any rows.

Define an *upsert* element in the adapter configuration file (or use the *UpsertQuery* class) for every update-insert to execute.

The synopsis is as follows:

```

<upsert connection="[connection]" stream="[stream]" table-name="[table]"
  [name="name"] [executor-name="executor"] [retry="count"] [retry-interval-sec="sec"]>
  <keys>
    <column property="[property_name]" column="[column_name]" type="[sql_type]" />
    [...column]
  </keys>
  <values>
    <column property="[property_name]" column="[column_name]" type="[sql_type]" />
    [...column]
  </values>
</upsert>

```

The *connection* attribute value is required and provides the name of the configured JDBC connection.

A value for the *stream* attribute is required and provides the name of the stream that triggers the Update-Insert. The adapter expects a stream by this name to exist at adapter start time.

The *table* attribute value is required and provides the database table name.

The *name* attribute is optional and is only used for logging errors.

The *executor-name* attribute is optional. If specified, it must be the name of an executor configuration. If specified, the adapter will use the executor service (queue and thread pool) for performing all work. If not specified, the adapter performs the work in the same thread.

The *retry* attribute is optional. If specified, the adapter will retry a given number of times in case an error is encountered. If *retry-interval-sec* is specified, the adapter waits the given number of seconds between retries.

The *keys* element is required and provides the key columns of the table and the *values* element provides the list of value columns of the table.

The *column* element should occur as many as there are key and value columns in the table. The *property* attribute provides the name of the event property, the *column* attribute provides the database table column name and the *type* is any of the `java.sql.Types` names (case ignored).

A sample Update-Insert configuration follows:

```

<upsert connection="db1" stream="UpdateInsertDBTableTrigger" name="UpdateInsertSample"
  table-name="MyKeyedTable" executor-name="queue1" retry="3">
  <keys>
    <column property="eventProperty1" column="keyColumn1" type="varchar" />
    <column property="eventProperty2" column="keyColumn2" type="varchar" />
  </keys>
  <values>
    <column property="eventProperty3" column="valueColumn1" type="varchar" />
    <column property="eventProperty4" column="valueColumn2" type="integer" />
  </values>
</upsert>

```

## 8.6. Executor Configuration

Executors are named thread pools and queues that may be assigned to perform DML or update-insert work.

Define an *executor* element in the adapter configuration file (or use the *Executor* class) for every thread

pool and queue to declare.

Upon adapter start, for each executor the adapter starts the given number of threads and an unbound queue.

The synopsis is as follows:

```
<executors>
  <executor name="[name]" threads="[count]"/>
</executors>
```

The *name* attribute value is required and provides the name of the executor, while the *count* attribute specifies the number of threads for the thread pool.

An example executor configuration::

```
<executors>
  <executor name="threadPool1" threads="2"/>
</executors>
```

An application can obtain a handle to all thread pools and queues via the Esper runtime context:

```
ExecutorServices execs = (ExecutorServices)
runtime.getContext().lookup("EsperIOBAdapter/ExecutorServices");
```

## 8.7. Reading and Polling Database Tables

Herein we provide sample statements and documentation pointers to use Esper EPL for reading from database tables. If only reading and not writing to a database, no configuration or EsperIO jar is file required.

Please consult the Esper SQL access documentation for more information.

### 8.7.1. Polling and Startup SQL Queries

To execute an SQL query repeatedly, Esper provides the opportunity to join a pattern to an SQL statement. The pattern may provide a single interval or crontab schedule or may also contain multiple schedules or combinations thereof via the pattern `or` operator.

The sample query below simply executes every 10 seconds retrieving all rows from table `MyTable`:

```
select * from pattern[every timer:interval(10)], sql:dbl ['select * from MyTable']
```

To perform an incremental query, consider utilizing a variable to parameterize your SQL statement so that only new rows are returned.

The next EPL statements create a variable and pass the variable value to the query to poll for new rows only. It assumes the `timestamp` column in the `MyTable` table holds long-type millisecond values:

```
// Create a variable to hold the last poll timestamp
create variable long VarLastTimestamp = 0

// Poll every 15 seconds between 8am and 4pm based on variable value
insert into PollStream
select * from pattern[every timer:crontab(*, 8-16, *, *, *, */15)],
  sql:dbl ['select * from MyTable where timestamp > ${VarLastTimestamp}']

// Assign last value to variable
on PollStream set VarLastTimestamp = timestamp
```

A sample statement to read a table at startup time is below:

```
select * from pattern[timer:interval(0)], sql:dbl ['select * from MyTable']
```

## Chapter 9. XML and JSON Output

Esper supports output event rendering to JSON and XML directly in its output API, please see the Esper documentation set for more information.