Если вы видите что-то необычное, просто сообщите мне.

Streaming Data with Postgres + Kafka + Debezium: Part 2

By the end of this guide, we will have a functioning Streams app that will take input from a topic that Debezium feeds into, make a simple arithmetic operation on one of the columns, then output the result into a new topic.

Intro

This is a multi-part series introducing how we can use Kafka, Postgres, and Debezium to stream data stored in a relational database into Kafka. If you haven't read the first part of this series, you may do so here.

At the end of Part 1, we successfully demonstrated inserting a record into one of our tables and watching it stream into a Kafka topic. Now that we have the basics of data streaming set up, it's time to write a Kafka Streams app to do some custom data processing.

The source code for this example project is available here. What is a Streams App?

A Kafka Streams app allows you to transform data within the framework of Streams, Tables, and more. Think of it as an ultra-flexible data plumbing toolkit. You can receive data (as Kafka records) from topics, send data to topics, split data across topics, route data to external services or APIs, enrich data with data from outside sources (or even from data within Kafka, like a JOIN). The possibilities are endless.

This is a primarily hands-on guide, and it may be helpful to read up on the basics of the Streams framework here. Schema Design (CREATE TABLE)

We left off Part 1 with a simple database schema only meant to test our configuration. Let's generate some more SQL schema so we can begin building our app.

Ultimately, this app is going to provide an API that allows us to make a request and store financial trades. We provide the API with a ticker symbol, price, and quantity, and it will store our trade.

create table trade (id serial primary key, insert_time timestamp default NOW(), ticker varchar, price int, quantity int default 1);

One of the things I've learned working with and creating schemas involving money is that it's often best to avoid using floating points. Since this is just a pedagogical example, we're going to only support precision to the USD cent. Something that is valued at $10.99 will be stored as an 1099 (an integer). If we were to support alternative currencies, we could create a join table that stores metadata on what the price column actually represents. In this case, it would be a "cent".

With this basic schema we can model our financial trades. If we want to look through time and see when we traded what, this could be achieved with a pretty basic SELECT statement. However, if we want to provide materialized data that is calculated based on these rows it would be cumbersome to do so in SQL. This is a where Kafka's ability to generate data in real-time really shines. Building and Connecting a Streams App Setting up Java (Hello World!)

The canonical way to build a Kafka Streams app is through the official libraries, which are JVM and are easiest used with Java or Scala. In this tutorial we will use Java.

I'm not always a fan of using IDEs, especially big ones, but when working with Java I like to use IntelliJ. JetBrains provides a free community version that you can download.

Let's get to a basic Hello World Java app.

First, let's create a new Gradle-based project. Gradle is a package manager for JVM languages, and as an Android developer it happens to be the one with which I'm most familiar. Create a new Gradle project Create a new Gradle project

I created this project in the same directory as the Docker files we used in part one. Wait a short time after creating the project and you will see something like the following in the project sidebar. Create a new Java ClassFile Create a new Java Class/File

Our source code will go inside the java folder. Java loves namespacing, so let's first create a package. Right click on the java folder, and create a new package. I named mine com.arctype. Once the package is created, right click on it, and create a new Java class. I called mine MyStream. Inside of MyStream we will initiate all of the streams of our app. Open the class and add the following code to get to the Hello World stage.

package com.arctype;

public class MyStream { public static void main(String[] args) throws Exception { System.out.println("Hello World."); } }

The fast and easy-to-use SQL client for developers and teams

By default, this Java project has no runnable configurations. Let's add a runnable configuration that runs our MyStream class. Near the top right of the IDE you can press Add Configuration Create a run configuration Create a run configuration

With the configuration box open press the plus sign and select Application. Then link this runnable to our MyStream class by setting com.arctype.MyStream as the main class. You also need to set the classpath of module to arctype-kafka-streams.main. Select the main class for this configuration

Now that we have a run configuration, you can press the green Play button to build and execute our basic Java app. Press Play to run Press Play to run!

You should see an output similar to this:

12:40:15 PM: Executing task 'MyStream.main()'...

> “ Task :compileJava Task :processResources NO-SOURCE Task :classes

> “ Task :MyStream.main() Hello World.

BUILD SUCCESSFUL in 0s 2 actionable tasks: 2 executed 12:40:16 PM: Task execution finished 'MyStream.main()'.

Kafka Dependencies

Let's continue building the Streams app. We will need a few dependencies from the official Apache Kafka repositories. Open the build.gradle file and under the dependencies section add the following lines inside the dependencies section. The last line is a Java logging library.

```
compile 'org.apache.kafka:kafka-streams:2.2.0'
compile 'org.slf4j:slf4j-simple:1.7.21'
```

You may need to "sync" Gradle before you are able to reference these dependencies in your source code. Streams Configuration

Now, return to our MyStream class.

There are 3 major steps we will need to complete to have a running Streams app. We need to set various configuration parameters, build a topology by registering each of our streams, then create and run the Streams app.

1. Key/Value Configuration The Streams app needs to know where it can find Kafka. Like the console consumer tool we used in Part 1, you provide it with a reachable bootstrap server and from there the Kafka broker handles the rest.

To provide Streams with configuration, we use the Java Properties class which allows you to set key/value pairs. Here we will set the bootstrap server, the application id (multiple instances of the same application should share an ID so that Kafka knows how to coordinate them), and a few other defaults.

I created a configuration builder function within MyStream that returns a Java Properties object with the following configuration:

private final static String BOOTSTRAP_SERVERS = "localhost:29092"; private final static String APPLICATION_ID = "arctype-stream";

private static Properties makeProps() { final Properties props = new Properties();

```
    props.put(StreamsConfig.BOOTSTRAP_SERVERS_CONFIG, BOOTSTRAP_SERVERS);
    props.put(StreamsConfig.APPLICATION_ID_CONFIG, APPLICATION_ID);
```

```
    // You will learn more about SerDe's soon!
    props.put(StreamsConfig.DEFAULT_KEY_SERDE_CLASS_CONFIG, Serdes.String().getClass());
    props.put(StreamsConfig.DEFAULT_VALUE_SERDE_CLASS_CONFIG, Serdes.String().getClass());

    props.put(StreamsConfig.DEFAULT_DESERIALIZATION_EXCEPTION_HANDLER_CLASS_CONFIG,
LogAndContinueExceptionHandler.class);
    props.put(StreamsConfig.POLL_MS_CONFIG, 100);
    props.put(StreamsConfig.RETRIES_CONFIG, 100);
    return props;
```

}

2. Topology The other object we need to build is what is called a Topology. When you build a Topology, Kafka infers how data flows between topics and your consumers and producers. Each Stream you write will need to be added here.

As with the configuration builder, I created a function that builds a Topology object. Since we have yet to build our stream for processing our trades, I commented out what will soon be valid code. I also added a System.out.println that dumps out the description of the Topology that Kafka infers, simply because it is really interesting to see once you have a few streams set up.

```
private static Topology createTopology(Properties props) {
    StreamsBuilder builder = new StreamsBuilder();
    // Add your streams here.
    // TradeStream.build(builder);
    final Topology topology = builder.build();
    System.out.println(topology.describe());
    return topology;
}
```

3. Stream We're ready to build the actual KafkaStreams object. In our main() function, let's call the functions we just created to build a properties, topology, and streams object.

```
  final Properties props = makeProps();
  final Topology topology = createTopology(props);
  final KafkaStreams streams = new KafkaStreams(topology, props);
```

Now we need to start our stream. Below is the entire body of the main() function. Below the first 3 lines of object creation, the key line is streams.start();, which starts the streams app. The

countdown hook ensures the Streams app closes gracefully.

public static void main(String[] args) throws Exception { final Properties props = makeProps(); final Topology topology = createTopology(props); final KafkaStreams streams = new KafkaStreams(topology, props);

```
    final CountDownLatch latch = new CountDownLatch(1);
    Runtime.getRuntime().addShutdownHook(new Thread("streams-shutdown-hook") {
        @Override
        public void run(){
            streams.close();
            latch.countDown();
        }
    });

    try {
        streams.start();
        latch.await();
    } catch (Throwable e) {
        System.exit(1);
    }
    System.exit(0);
}
```

We technically have working Streams app now! It won't do anything, since we have not added any streams to its topology. However, now is a good time to test that it can connect to our Kafka cluster. After running docker-compose up on the project we defined in Part 1, go ahead and press the green Play button to compile and run our Streams app.

Kafka output is very verbose; keep an eye our for a line looking like this:

INFO org.apache.kafka.streams.KafkaStreams - stream-client [arctype-stream-9b8159f4-1983-49e4-9d6b-d4b1de20482b] State transition from REBALANCING to RUNNING

You'll see several log lines indicating that the state of the app has transitioned. If the app is able to connect and is working properly, the last state you'll see it transition to is RUNNING.

If you check the logs of the Docker Kafka cluster, you'll see various lines referencing our Streams app. I named mine arctype-stream, so that's what will show up as the cluster prepares to run the

streams app.

INFO [executor-Rebalance:Logging@66] - [GroupCoordinator 1]: Stabilized group arctype-stream generation 1 (__consumer_offsets-29)

Building our First Stream

Let's build our first Stream component. As a rudimentary example, let's produce a stream that duplicates all of our trades but doubles the price. SerDe (Serialization/Deserialization)

Serialization and deserialization are key components to building a Stream. Each message in a topic has a key and a value. When consuming or producing a Stream, we will need to provide SerDe's.

For example, we will first want to consume the topic that stores our data from the trade table. We will write a key deserializer that takes the JSON blob provided by Debezium and plucks the id (which is the primary key in our table) from the JSON and returns it as a key of type String. We will write a value deserializer that takes the entire JSON blog and maps it to a Java object (a Model).
The Collaborative SQL Editor

The serializers we write will work in reverse fashion. Our value serializer, for instance, will take our Java model, convert it into JSON, and return it as value of type String.

First, let's create a Java model to represent our trade SQL table.

// TradeModel.java

@JsonIgnoreProperties(ignoreUnknown = true) public class TradeModel { public Integer id; public String ticker; public Integer price; public Integer quantity; }

Key SerDe

Since our Key SerDe's are only focused on the id field, they are simpler and using a Factory pattern isn't necessary. It also provides us with a more direct look into what the serializers and deserializers are actually doing.

Deserializer We'll begin with a new class called IdDeserializer. To make a Deserializer we will implement, or conform to, Kafka's Deserializer interface. Once your class implements this interface, there are 3 methods you'll need to override: configure, deserialize, and close. We don't be doing

any operations that need configuration or cleanup, so we'll leave all but deserialize empty.

public class IdDeserializer implements Deserializer { private ObjectMapper objectMapper = new ObjectMapper();

```
@Override
public void configure(Map<String, ?> props, boolean isKey) { }

@Override
public void close() { }

@Override
public String deserialize(String topic, byte[] bytes) {
    if (bytes == null)
        return null;

    String id;
    try {
        Map payload = objectMapper.readValue(new String(bytes), Map.class);
        id = String.valueOf(payload.get("id"));
    } catch (Exception e) {
        throw new SerializationException(e);
    }
    return id;
}
```

}

bytes contains the data coming in from Debezium. We know this to be JSON since that's what we set Debezium to output. ObjectMapper is a utility from Jackson, which is a Java library for handling XML and JSON.

The code inside deserialize attempts to parse the JSON into an Map object, looks for the id key, and returns its value as a String. Now our Kafka record has a key!

Serializer For the purpose of this app, we know that all of our keys will be String's containing primary keys from Debezium. When the id is serialized (this will happen when you output records to a new topic), we can just pass the String through. This part gets trickier if you are, for example, outputting 2 records for every 1 input record. In some situations, like building a KTable, only the

latest record of the same key will be present.

Since the full source code for this project is available to download, I'll just show the important part of the serializer. Its implementation looks similar to the Deserializer.

@Override public byte[] serialize(String topic, String data) { return data.getBytes(); }

Value SerDe

SerDe's for our messages' values are more complex. They will convert JSON blobs to Java objects (models) and vice-versa. Since in a real app you will have many tables (and thus many models), it's best to write a factory that takes in a Java object, like this model, and returns a SerDe for the model. Factory

SerdeFactory is a generic factory that accept a Class and return SerDe's for it.

public class SerdeFactory { public static Serde createSerdeFor(Class clazz, boolean isKey) { Map<String, Object> serdeProps = new HashMap<>(); serdeProps.put("Class", clazz);

```
    Serializer<T> ser = new JsonSerializer<>();
    ser.configure(serdeProps, isKey);

    Deserializer<T> de = new JsonDeserializer<>();
    de.configure(serdeProps, isKey);

    return Serdes.serdeFrom(ser, de);
}
```

}

You may recall that the signature of Kafka's Deserializer initialization function looks like this:

public void configure(Map<String, ?> props, boolean isKey) { }

To pass in arbitrary configuration options, you build a key-value Map and pass it in. In this case, we'll pass in the Class of the model we receive when the factory is called. The factory then passes the Class to both the deserializer and serializer. Deserializer

Our generic JSON deserializer looks like this:

public class JsonDeserializer implements Deserializer { private ObjectMapper objectMapper = new ObjectMapper(); private Class clazz;

```
@SuppressWarnings("unchecked")
@Override
public void configure(Map<String, ?> props, boolean isKey) {
    clazz = (Class<T>) props.get("Class");
}

@Override
public void close() { }

@Override
public T deserialize(String topic, byte[] bytes) {
    if (bytes == null)
        return null;

    T data;
    Map payload;
    try {
        payload = objectMapper.readValue(new String(bytes), Map.class);
        // Debezium updates will contain a key "after" with the latest row contents.
        Map afterMap = (Map) payload.get("after");
        if (afterMap == null) {
             // Non-Debezium payloads
            data = objectMapper.readValue(objectMapper.writeValueAsBytes(payload), clazz);
        } else {
             // Incoming from Debezium
            data = objectMapper.readValue(objectMapper.writeValueAsBytes(afterMap), clazz);
        }

    } catch (Exception e) {
        throw new SerializationException(e);
    }
    return data;
}
```

}

On configure, we check the props bundle for the Class then store it in a private variable within the object called clazz. That way the JSON parser knows what Class to create and fill with data.

In the deserialize function, we first parse the JSON to a generic Map, and check if there is an after key. Incoming Debezium data will have an after key whose value is the most up-to-date information about the database row. If the JSON lacks an after key, we just attempt to deserialize the base JSON object into whatever Class this SerDe was instantiated with. It's an easy way to be able to use this SerDe with JSON data that isn't from Debezium. I would warn against using a SerDe this arbitrary in Production. KStream Operations

To capture the records from a Kafka topic, you can build a KStream. This will require the name of the topic and a SerDe as mentioned above. Once you have KStream you can begin transforming and routing data. Finally, we'll create our TradeStream class.

public class TradeStream { private final static String TRADE_TOPIC = "ARCTYPE.public.trade";

```
public static void build(StreamsBuilder builder) {
    final Serde<TradeModel> tradeModelSerde = SerdeFactory.createSerdeFor(TradeModel.class,
true);
    final Serde<String> idSerde = Serdes.serdeFrom(new IdSerializer(), new IdDeserializer());

    KStream<String, TradeModel> tradeModelKStream =
            builder.stream(TRADE_TOPIC, Consumed.with(idSerde, tradeModelSerde));

    tradeModelKStream.peek((key, value) -> {
        System.out.println(key.toString());
        System.out.println(value.toString());
    });
    tradeModelKStream.map((id, trade) -> {
        TradeModel tradeDoubled = new TradeModel();
        tradeDoubled.price = trade.price * 2;
        tradeDoubled.quantity = trade.quantity;
        tradeDoubled.ticker = trade.ticker;
        return new KeyValue<>(id, tradeDoubled);
    }).to("ARCTYPE.doubled-trades", Produced.with(idSerde, tradeModelSerde));
}
```

}

First, we build SerDe's using the objects we created in the previous section. Next, we build a KStream based on the topic ID that we know Debezium has created.

Lastly, we perform 2 operations. The .peek operation is great for debugging and getting some visibility into your Stream. It simply allows you to take a peek at the records that flow through. For testing, we'll print the key and value of all records to the console.

The .map operation allows you to elegantly map each record to some new object. Here we are going to map each Trade row to a new Trade row, but with the price multiplied by 2. After the object is created and returned, we will use .to to stream the data to a new topic (ARCTYPE.doubled-trades). Conclusion

Now, consuming topic ARCTYPE.doubled-trades will show the following output when you insert a record.

postgres=# insert into trade (ticker, price, quantity) values ('AAPL', 42069, 1);

Kafka Topic Output:

{"id":null,"ticker":"AAPL","price":84138,"quantity":1}

Stay tuned for the last part of this series where we will dive further into more Kafka Streams operations and the power of KStream, KGlobalTable, and KTable.

Revision #1
Created 2021-09-14 10:27:30 UTC by gasick
Updated 2023-04-16 19:26:41 UTC by gasick