
StreamReactor Documentation

Release 0.1

Andrew Stevenson

Sep 18, 2017

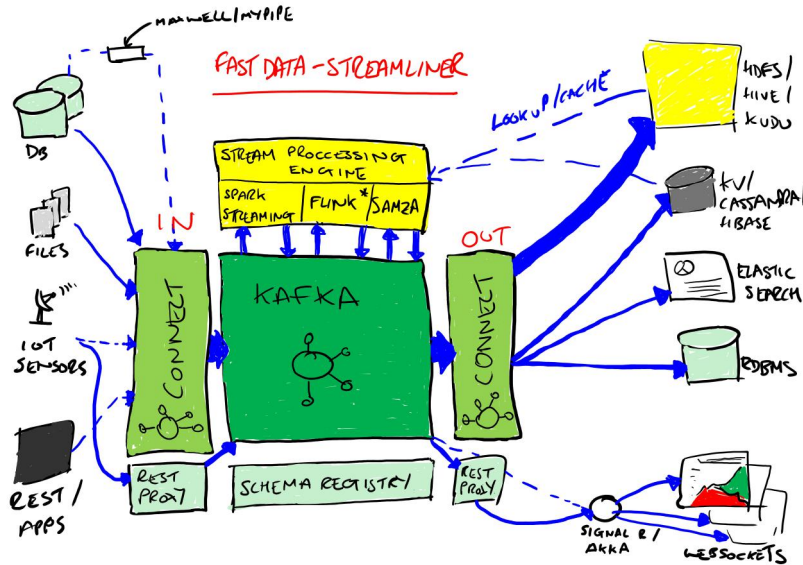
Contents

1	Components	3
	HTTP Routing Table	153

The Stream Reactor is a set of components to build a reference architecture for streaming data platforms. At its core is Kafka, with Kafka Connect providing a unified way to stream data in and out of the system.

The actual processing is left to streaming engines and libraries such as Spark Streaming, Apache Flink, Storm and Kafka Streams.

DataMountaineer provides a range of supporting components to the main technologies, mainly Kafka, Kafka Connect and the Confluent Platform.



Download [here](#).

CHAPTER 1

Components

Install

The Stream Reactor components are built around The Confluent Platform. They rely on the Kafka Brokers, Zookeepers and optionally the Schema Registry provided by this distribution.

The following releases are available:

- 0.2.3
- 0.2.2

Connector	Versions
BlockChain	Not applicable
Bloomberg	blpapi-3.8.8-2
Cassandra	Driver 3.0.0, Server 2.6.6
Druid	Tranquility 0.7.4
Elastic	Elastic 2.2.0, Elastic4s 2.3.0
HazelCast	HazelCast 3.6.0
HBase	HBase Server 1.2.0, HBase Client 1.2.0
InfluxDB	InfluxDB 2.3
JMS	javax.jms 1.1-rev-1, active-mq-code 1.26.0
Kudu	Kudu Client 0.9.0
MongoDB	MongoDB 3.3.0
Redis	Redis 2.8.1
ReThinkDB	ReThinkDB 2.3.3
VoltDB	VoltDB 6.4
Yahoo	yahoofinance-api 1.3.0

Install Confluent

Confluent can be downloaded for [here](#)

```
#make confluent home folder
mkdir confluent

#download confluent
wget http://packages.confluent.io/archive/3.0/confluent-3.0.1-2.11.tar.gz

#extract archive to confluent folder
tar -xvf confluent-3.0.1-2.11.tar.gz -C confluent

#setup variables
export CONFLUENT_HOME=~/.confluent/confluent-3.0.1
```

Start the Confluent platform.

```
#Start the confluent platform, we need kafka, zookeeper and the schema registry
bin/zookeeper-server-start etc/kafka/zookeeper.properties &
bin/kafka-server-start etc/kafka/server.properties &
bin/schema-registry-start etc/schema-registry/schema-registry.properties &
```

Stream Reactor Install

Download the latest release from [here](#).

Unpack the archive:

```
#Stream reactor release 0.2.3
tar xvf -C stream-reactor-0.2.3-cp-3.0.1.jar stream-reactor
```

Within the unpacked directory you will find the following structure:

```
stream-reactor-0.2-3.0.1
|-- LICENSE
|-- README.md
|-- bin
|   |-- cli.sh
|   `-- start-connect.sh
`-- libs
    |-- kafka-connect-cli-0.9-all.jar
    |-- kafka-connect-blockchain-0.2-3.0.1-all.jar
    |-- kafka-connect-bloomberg-0.2-3.0.1-all.jar
    |-- kafka-connect-cassandra-0.2-3.0.1-all.jar
    |-- kafka-connect-druid-0.2-3.0.1-all.jar
    |-- kafka-connect-elastic-0.2-3.0.1-all.jar
    |-- kafka-connect-hazelcast-0.2-3.0.1-all.jar
    |-- kafka-connect-hbase-0.2-3.0.1-all.jar
    |-- kafka-connect-influxdb-0.2-3.0.1-all.jar
    |-- kafka-connect-jms-0.2-3.0.1-all.jar
    |-- kafka-connect-kudu-0.2-3.0.1-all.jar
    |-- kafka-connect-redis-0.2-3.0.1-all.jar
    |-- kafka-connect-rethink-0.2-3.0.1-all.jar
    |-- kafka-connect-voltdb-0.2-3.0.1-all.jar
    |-- kafka-connect-yahoo-0.2-3.0.1-all.jar
    `-- kafka-socket-streamer-0.2-3.0.1-all.jar
```

The `libs` folder contains all the Stream Reactor Connector jars.

The `bin` folder contains the `start-connect.sh` script. This loads all the Stream Reactors jars onto the CLASS-PATH and starts Kafka Connect in distributed mode. The Confluent Platform, Zookeeper, Kafka and the Schema Registry must be started first.

Docker Install

All the Stream Reactor Connectors, Confluent and UI's for Connect, Schema Registry and topic browsing are available in Docker. The Docker images are available in [DockerHub](#) and maintained by our partner [Landoop](#)

Pull the latest images:

```
docker pull landoop/fast-data-dev
docker pull landoop/fast-data-dev-connect-cluster

#UI's
docker pull landoop/kafka-topics-ui
docker pull landoop/schema-registry-ui
```

Fast Data Dev

This is Docker image for development.

If you need

1. Kafka Broker
2. ZooKeeper
3. Schema Registry
4. Kafka REST Proxy
5. Kafka Connect Distributed
6. Certified DataMountaineer Connectors (ElasticSearch, Cassandra, Redis ..)
7. Landoop's Fast Data Web UIs : schema-registry , kafka-topics , kafka-connect and
8. Embedded integration tests with examples


Run with:

```
docker run --rm -it --net=host landoop/fast-data-dev
```

On Mac OSX run:

```
docker run --rm -it \
  -p 2181:2181 -p 3030:3030 -p 8081:8081 \
  -p 8082:8082 -p 8083:8083 -p 9092:9092 \
  -e ADV_HOST=127.0.0.1 \
  landoop/fast-data-dev
```

That's it. Your Broker is at localhost:9092, your Kafka REST Proxy at localhost:8082, your Schema Registry at localhost:8081, your Connect Distributed at localhost:8083, your ZooKeeper at localhost:2181 and at <http://localhost:3030> you will find Landoop's Web UIs for Kafka Topics and Schema Registry, as well as a Coyote test report.

**Kafka Development Environment**
docker container powered by Landoop

SCHEMAS

4

SCHEMA REGISTRY UI

Manage Avro schemas of your Kafka cluster

ENTER

TOPICS

11

KAFKA TOPICS UI

Browse topics and download data

ENTER

CONNECTORS

0

KAFKA CONNECT UI

Easy configure and manage connectors

COMING SOON


BROKERS

1

KAFKA BROKERS

Kafka management and monitoring

COMING SOON

 **RUNNING SERVICES**
fast-data-dev

✓

Confluent v3.0.1 - Kafka v0.10.0.1
1 broker, including kafka connect distributed, schema registry, kafka rest, etc

✓

Datamountaineer Stream Reactor v0.2.2
Source & Sink connectors supporting KCQL, ElasticSearch, Cassandra, Redis etc

✓

Landoop Fast Data platform v0.7
Manage schemas, view history, browse topics & configure your cluster from your browser.

SERVICES PORTS

9092 : Kafka Broker

8081 : Schema Registry

8082 : Kafka REST Proxy

8083 : Kafka Connect Distributed

2181 : ZooKeeper

3030 : Web Server

COYOTE HEALTH CHECKS

Your set up is being verified using the Coyote integration testing tool and generates working examples.

100%

Passed: 39 | Failed: 0

VIEW RESULTS

Find out more about Coyote awesome tool [here](#).


Report Issues & Stars!

Issue

2

Star

6

 powered by Landoop

Fast Data Dev Connect

This docker is targeted to more advanced users and is a special case since it doesn't set-up a Kafka cluster, instead it expects to find a Kafka Cluster with Schema Registry up and running.

The developer can then use this docker image to setup a connect-distributed cluster by just spawning a couple containers.

```
docker run -d --net=host \
  -e ID=01 \
  -e BS=broker1:9092,broker2:9092 \
  -e ZK=zk1:2181,zk2:2181 \
  -e SC=http://schema-registry:8081 \
  -e HOST=<IP OR FQDN>
landoop/fast-data-dev-connect-cluster
```

Things to look out for in configuration options:

1. It is important to give a full URL (including schema —<http://>) for schema registry.
2. ID should be unique to the Connect cluster you setup, for current and old instances. This is because Connect stores data in Brokers and Schema Registry. Thus even if you destroyed a Connect cluster, its data remain in your Kafka setup.
3. HOST should be set to an IP address or domain name that other connect instances and clients can use to reach the current instance. We chose not to try to autodetect this IP because such a feat would fail more often than not. Good choices are your local network ip (e.g 10.240.0.2) if you work inside a local network, your public ip (if you have one and want to use it) or a domain name that is resolvable by all the hosts you will use to talk to Connect.

If you don't want to run with `--net=host` you have to expose Connect's port which at default settings is 8083. There a `PORT` option, that allows you to set Connect's port explicitly if you can't use the default 8083. Please remember that it is important to expose Connect's port on the same port at the host. This is a choice we had to make for simplicity's sake.

```
docker run -d \
  -e ID=01 \
  -e BS=broker1:9092,broker2:9092 \
  -e ZK=zk1:2181,zk2:2181 \
  -e SC=http://schema-registry:8081 \
  -e HOST=<IP OR FQDN>
  -e PORT=8085
  -p 8085:8085
landoop/fast-data-dev-connect-cluster
```

Advanced

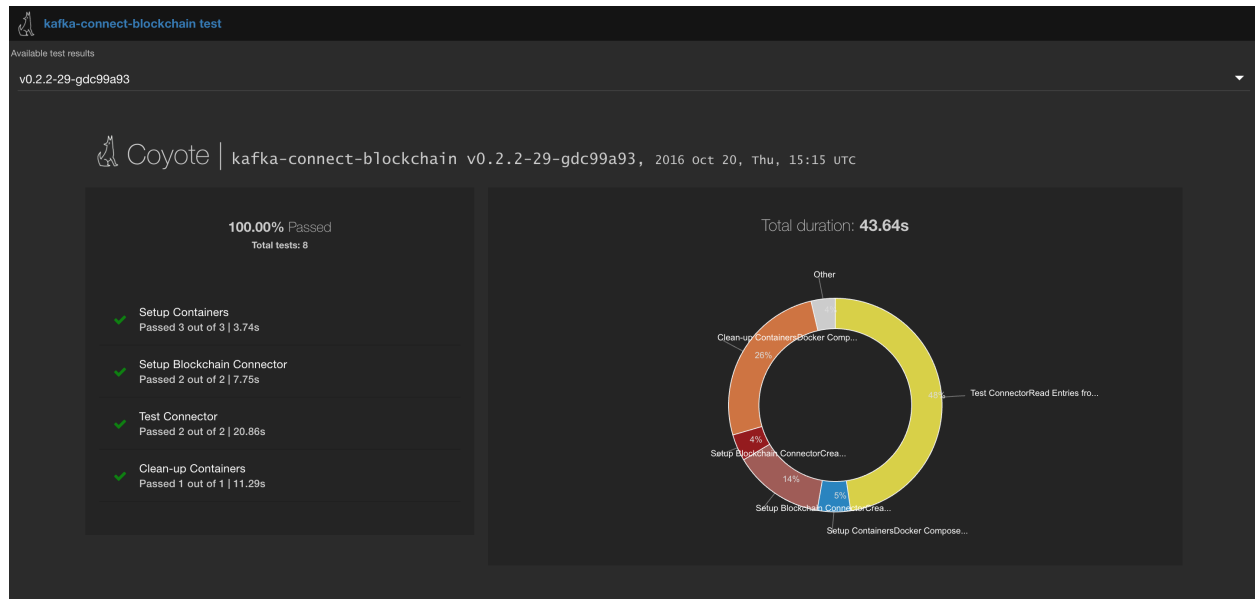
The container does not exit with CTRL+C. This is because we chose to pass control directly to Connect, so you check your logs via docker logs. You can stop it or kill it from another terminal.

Whilst the `PORT` variable sets the rest.port, the `HOST` variable sets the advertised host. This is the hostname that Connect will send to other Connect instances. By default Connect listens to all interfaces, so you don't have to worry as long as other instances can reach each instance via the advertised host.

Latest Test Results

To see the latest tests for the Connectors, in a docker, please visit Landoop's test github [here](#) Test results can be found [here](#).

An example for BlockChain is:



Kafka Connect

Kafka Connect is a tool to rapidly stream events in and out of Kafka. It has a narrow focus on data ingress in and egress out of the central nervous system of modern streaming frameworks. It is not an ETL and this separation of concerns allows developers to quickly build robust, durable and scalable pipelines in and out of Kafka.

Kafka connect forms an integral component in an ETL pipeline when combined with Kafka and a stream processing framework.

Modes

Kafka Connect can run either as a standalone process for testing and one-off jobs, or as a distributed, scalable, fault tolerant service supporting an entire organisations. This allows it to scale down to development, testing, and small production deployments with a low barrier to entry and low operational overhead, and to scale up to support a large organisations data pipeline.

Usually you'd run in distributed mode to get fault tolerance and better performance. In distributed mode you start Connect on multiple hosts and they join together to form a cluster. Connectors which are then submitted are distributed across the cluster.

For workers to join a Connect cluster, set the `group.id` in the `$CONFLUENT_HOME/etc/schema-registry/connect-avro-distributed.properties` file.

```
# The group ID is a unique identifier for the set of workers that form a single Kafka_
↪ Connect
# cluster
group.id=connect-cluster
```

Results

✓ Setup Containers				
Action	⌚ Time (sec)	>_ Command	Exit code	
✓ Docker Compose Pull	2.20	docker-compose pull	(ignore) 1	
✓ Docker Compose Up	1.09	docker-compose up -d	0	
✓ Check docker compose log	0.45	docker-compose logs	0	
Passed 3 out of 3 3.74 seconds				
✓ Setup Blockchain Connector				
Action	⌚ Time (sec)	>_ Command	Exit code	
✓ Create Topic	5.99	docker run --rm --network=kafkaconnectblockchain_default landoop/fast-data-dev kafka-topics --zookeeper fast-data-dev:2181 --topic blockchain --partition 1 --replication 1 --create	0	
✓ Create a Blockchain Distributed Connector	1.76	docker run --rm --network=kafkaconnectblockchain_default -i landoop/fast-data-dev curl -vs --stderr - -X POST -H "Content-Type: application/json" --data @- "http://fast-data-dev:8083/connectors"	0	
Passed 2 out of 2 7.75 seconds				
✓ Test Connector				
Action	⌚ Time (sec)	>_ Command	Exit code	
✓ Read Entries from Topic	20.81	timeout 20 docker run --rm --network=kafkaconnectblockchain_default landoop/fast-data-dev kafka-avro-console-consumer --zookeeper fast-data-dev:2181/confluent --bootstrap-server fast-data-dev:9092 --topic blockchain --from-beginning --new-consumer --property schema.registry.url=http://fast-data-dev:8081	(ignore) 124	
✓ Read First 3000 Lines of Connect Logs	0.05	docker exec kafkaconnectblockchain_fast-data-dev_1 head -n3000 /var/log/connect-distributed.log	0	
Passed 2 out of 2 20.86 seconds				

Schema Registry Support

DataMountaineer recommends all payloads in Kafka are Avro. Schema Registry provides a serving layer for your metadata. It provides a RESTful interface for storing and retrieving Avro schemas. It stores a versioned history of all schemas, provides multiple compatibility settings and allows evolution of schemas according to the configured compatibility setting. It provides serializers that plug into Kafka clients that handle schema storage and retrieval for Kafka messages that are sent in the Avro format.

All our Connectors support Avro and use the Confluent provided converters to translate the Avro into Kafka Connects internal `Struct` type to determine the schema and how to map onto the target sink store.

We have found some of the clients have already an infrastructure where they publish pure json on the topic and obviously the jump to follow the best practice and use schema registry is quite an ask. So we offer support for them as well for the following Sinks:

- ReThinkDB
- MongoDB
- InfluxDB
- DSE Cassandra Sink

We are upgrading the remaining Connectors. This allows plain text payloads with a json string.

Connectors

Kafka Connect Query Language

The Kafka Connect Query Language is implemented in antlr4 grammar files.

Why ?

While working on our sink/sources we ended up producing quite complex configuration in order to support the functionality required. Imagine a Sink where you Source from different topics and from each topic you want to cherry pick the payload fields or even rename them. Furthermore you might want the storage structure to be automatically created and/or even evolve or you might add new support for the likes of bucketing (Riak TS has one such scenario). Imagine the JDBC sink with a table which needs to be linked to two different topics and the fields in there need to be aligned with the table column names and the complex configuration involved ...or you can just write this

```
routes.query = "INSERT INTO transactions SELECT field1 as column1, field2 as column2, ↵
↵field3 FROM topic_A;
                INSERT INTO transactions SELECT fieldA1 as column1, fieldA2 as column2, ↵
↵fieldC FROM topic_B;"
```

Kafka Connect Query Language

There are two paths supported by this DSL. One is the INSERT and take the following form:

```
INSERT INTO $TARGET
SELECT *|columns
FROM $TOPIC_NAME
    [IGNORE columns]
    [AUTOCREATE]
    [PK columns]
    [AUTOEVOLVE]
    [BATCH = N]
    [CAPITALIZE]
    [PARTITIONBY cola[,colb]]
    [DISTRIBUTEBY cola[,colb]]
    [CLUSTERBY cola[,colb]]
    [WITHTIMESTAMP cola|sys_time()]
    [WITHFORMAT TEXT|JSON|AVRO|BINARY|OBJECT|MAP]
    [STOREAS $YOUR_TYPE ([key=value, .....])]
```

If you follow our connectors @Datamountaineer you will find depending on the Connect Sink only some of the the options are used. You will find all our documentation here

The second path is SELECT only. We have the Socket Streamer</> which allows you to peek into KAFKA via websocket and receive the payloads in real time!

```
SELECT *|columns
FROM $TOPIC_NAME
    [IGNORE columns]
    [WITHFORMAT TEXT|JSON|AVRO|BINARY]
    [WITHGROUP $YOUR_CONSUMER_GROUP]
    [WITHPARTITION (partition),[(partition, offset)]
    [SAMPLE $RECORDS_NUMBER EVERY $SLIDE_WINDOW]
```

Examples

```
SELECT field1 FROM mytopic // Project one avro field named field1
SELECT field1 AS newName // Project and renames a field
SELECT * FROM mytopic // Select everything - perfect for avro ↵
↵evolution
```

```

SELECT *, field1 AS newName FROM mytopic      // Select all & rename a field -
↪excellent for avro evolution
SELECT * FROM mytopic IGNORE badField        // Select all & ignore a field -
↪excellent for avro evolution
SELECT * FROM mytopic PK field1,field2        //Select all & with primary keys (for
↪the sources where primary keys are required)
SELECT * FROM mytopic AUTOCREATE              //Select all and create the target
↪Source (table for databases)
SELECT * FROM mytopic AUTOEVOLVE              //Select all & reflect the new fields
↪added to the avro payload into the target

```

Source Connectors

Source connectors load or stream data from external systems into Kafka.

Kafka Connect Blockchain

A Connector to hook into the live streaming providing a real time feed for new bitcoin blocks and transactions provided by www.blockchain.info. The connector subscribes to notification on blocks, transactions or an address and receives JSON objects describing a transaction or block when an event occurs. This json is then pushed via kafka connect to a kafka topic and therefore can be consumed either by a Sink or have a live stream processing using for example Kafka Streams.

Since it is a direct websocket connection the Source will only ever use one connector task at any point. There is no point spawning more and then having duplicate data.

One thing to remember is the subscription API from blockchain doesn't offer an option to start from a given timestamp. This means if the connect worker is down then you will miss some data.

The Sink connects to unconfirmed transactions!! Read more about the live data [here](#)

Prerequisites

- Confluent 3.0.1
- Java 1.8
- Scala 2.11

Confluent Setup

Follow the instructions [here](#).

Source Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools` cli to post in our distributed properties file for BlockChain.

```
bin/cli.sh create blockchain-source < conf/blockchain-source.properties

#Connector `blockchain-source`:
name=blockchain-source
connector.class=com.datamountaineer.streamreactor.connect.blockchain.source.
↳BlockchainSourceConnector
max.tasks=1
connect.blockchain.source.kafka.topic = blockchain-test
max.tasks=1
#task ids:
```

The `blockchain-source.properties` file defines:

1. The name of the source.
2. The Source class.
3. The max number of tasks the connector is allowed to created (1 task only).
4. The topics to write to.

If you switch back to the terminal you started the Connector in you should see the Blockchain Source being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
blockchain-source
```

```
[2016-08-21 20:31:36,398] INFO Finished starting connectors and tasks (org.apache.
↳ kafka.connect.runtime.distributed.DistributedHerder:769)
[2016-08-21 20:31:36,406] INFO
```

by Stefan Bocutiu

Test Records

Now we need to see records pushed on the topic. We can use the `kafka-avro-console-producer` to do this.

```
$ ./bin/kafka-avro-console-consumer --topic blockchain-test \
  --zookeeper localhost:2181 \
  --from-beginning
```

Now the console is reading blockchain transaction data which would print on the terminal.

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect Bloomberg

DOCS ARE WORKING IN PROGRESS

Kafka Connect Bloomberg is a Source connector to subscribe to Bloomberg feeds via the Bloomberg labs open API and write to Kafka.

Prerequisites

- Bloomberg subscription
- Confluent 3.0.1
- Java 1.8
- Scala 2.11

Setup

Confluent Setup

Follow the instructions [here](#).

Source Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools cli` to post in our distributed properties file for Redis. If you are using the *dockers* you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create bloomberg-source < conf/bloomberg-source.properties
#Connector name=name=`bloomberg-source`
connector.class=com.datamountaineer.streamreactor.connect.bloomberg.
↳BloombergSourceConnector
tasks.max=1
connect.bloomberg.server.host=localhost
connect.bloomberg.server.port=8194
connect.bloomberg.service.uri=/blp/mkdata
connect.bloomberg.subscriptions=AAPL US Equity:LAST_PRICE,BID,ASK;IBM US Equity:BID,
↳ASK,HIGH,LOW,OPEN
kafka.topic=bloomberg
connect.bloomberg.buffer.size=4096
#task ids: 0
```

The `bloomberg-source.properties` file defines:

1. The connector name.
2. The class containing the connector.
3. The number of tasks the connector is allowed to start.
4. The Bloomberg server host.
5. The Bloomberg server port.
6. The Bloomberg service uri.
7. The subscription keys to subscribe to.
8. The topic to write to.
9. The buffer size for the Bloomberg API to buffer events in.

If you switch back to the terminal you started the Connector in you should see the Bloomberg Source being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
bloomberg-source
```

Test Records

Now we need to see records pushed on the topic. We can use the `kafka-avro-console-producer` to do this.

```
$ ./bin/kafka-avro-console-consumer --topic blockchain-test \
  --zookeeper localhost:2181 \
  --from-beginning
```

Now the console is reading blockchain transaction data which would print on the terminal.

Features

The Source Connector allows subscriptions to BPIPE mkdata and refdata endpoints to feed data into Kafka.

Configurations

`connect.bloomberg.server.host`

The bloomberg endpoint to connect to.

- Data type : string
- Optional : no

`connect.bloomberg.server.port`

The Bloomberg endpoint to connect to.

- Data type : string
- Optional : no

`connect.bloomberg.service.uri`

Which Bloomberg service to connect to. Can be `//blp/mkdata` or `//blp/refdata`.

- Data type : string
- Optional : no

`connect.bloomberg.authentication.mode`

The mode to authentication against the Bloomberg server. Either `APPLICATION_ONLY` or `USER_AND_APPLICATION`.

- Data type : string
- Optional : no

`connect.bloomberg.subscriptions`

- Data type : string
- Optional : no

Specifies which ticker subscription to make. The format is `TICKER:FIELD,FIELD,..`; e.g. `AAPL US Equity:LAST_PRICE;IBM US Equity:BID`

`connect.bloomberg.buffer.size`

- Data type : int
- Optional : yes

- Default : 2048

The buffer accumulating the data updates received from Bloomberg. If not provided it will default to 2048. If the buffer is full and a new update will be received it won't be added to the buffer until it is first drained.

`connect.bloomberg.kafka.topic`

The topic to write to.

- Data type : string
- Optional : no

Example

```
name=bloomberg-source
connector.class=com.datamountaineer.streamreactor.connect.bloomberg.
↳BloombergSourceConnector
tasks.max=1
connect.bloomberg.server.host=localhost
connect.bloomberg.server.port=8194
connect.bloomberg.service.uri=/blp/mkdata
connect.bloomberg.subscriptions=AAPL US Equity:LAST_PRICE,BID,ASK;IBM US Equity:BID,
↳ASK,HIGH,LOW,OPEN
kafka.topic=bloomberg
connect.bloomberg.buffer.size=4096
```

Schema Evolution

TODO

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect Cassandra Source

Kafka Connect Cassandra is a Source Connector for reading data from Cassandra and writing to Kafka.

The Source supports:

1. *The KCQL routing querying* - Allows for table to topic routing.
2. Incremental mode
3. Bulk mode
4. Error policies for handling failures.

Prerequisites

- Cassandra 2.2.4
- Confluent 3.0.0
- Java 1.8
- Scala 2.11

Setup

Before we can do anything, including the QuickStart we need to install Cassandra and the Confluent platform.

Cassandra Setup

First download and install Cassandra if you don't have a compatible cluster available.

```
#make a folder for cassandra
mkdir cassandra

#Download Cassandra
wget http://apache.cs.uu.nl/cassandra/3.5/apache-cassandra-3.5-bin.tar.gz

#extract archive to cassandra folder
tar -xvf apache-cassandra-3.5-bin.tar.gz -C cassandra

#Set up environment variables
export CASSANDRA_HOME=~/.cassandra/apache-cassandra-3.5-bin
export PATH=$PATH:$CASSANDRA_HOME/bin

#Start Cassandra
sudo sh ~/.cassandra/bin/cassandra
```

Confluent Setup

Follow the instructions [here](#).

Source Connector

The Cassandra Source connector allows you to extract entries from Cassandra with the CQL driver and write them into a Kafka topic.

Each table specified in the configuration is polled periodically and each record from the result is converted to a Kafka Connect record. These records are then written to Kafka by the Kafka Connect framework.

The Source connector operates in two modes:

1. Bulk - Each table is selected in full each time it is polled.
2. Incremental - Each table is querying with lower and upper bounds to extract deltas.

In incremental mode the column used to identify new or delta rows has to be provided. This column must be of CQL Type Timestamp. Due to Cassandra's and CQL restrictions this should be a primary key or part of a composite primary keys. ALLOW_FILTERING can also be supplied as an configuration.

Note: TimeUUIDs are converted to strings. Use the [UUIDs](#) helpers to convert to Dates.

Source Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our [GitHub releases](#) page.

Test data

Once you have installed and started Cassandra create a table to extract records from. This snippet creates a table called orders and inserts 3 rows representing fictional orders or some options and futures on a trading platform.

Start the Cassandra cql shell

```
bin ./cqlsh
Connected to Test Cluster at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 3.0.2 | CQL spec 3.3.1 | Native protocol v4]
Use HELP for help.
cqlsh>
```

Execute the following:

```
CREATE KEYSPACE demo WITH REPLICATION = {'class' : 'SimpleStrategy', 'replication_
↪factor' : 3};
use demo;

create table orders (id int, created timeuuid, product text, qty int, price float,
↪PRIMARY KEY (id, created))
WITH CLUSTERING ORDER BY (created asc);


INSERT INTO orders (id, created, product, qty, price) VALUES (1, now(), 'OP-DAX-P-
↪20150201-95.7', 100, 94.2);
INSERT INTO orders (id, created, product, qty, price) VALUES (2, now(), 'OP-DAX-C-
↪20150201-100', 100, 99.5);
INSERT INTO orders (id, created, product, qty, price) VALUES (3, now(), 'FU-KOSPI-C-
↪20150201-100', 200, 150);

SELECT * FROM orders;
```

id	created	price	product	qty
1	17fa1050-137e-11e6-ab60-c9fbe0223a8f	94.2	OP-DAX-P-20150201-95.7	100
2	17fb6fe0-137e-11e6-ab60-c9fbe0223a8f	99.5	OP-DAX-C-20150201-100	100
3	17fbb00-137e-11e6-ab60-c9fbe0223a8f	150	FU-KOSPI-C-20150201-100	200

(3 rows)

(3 rows)



```
By Andrew Stevenson. (com.datamountaineer.streamreactor.connect.cassandra.source.
↳CassandraSourceTask:64)
[2016-05-06 13:34:41,193] INFO Attempting to connect to Cassandra cluster at_
↳localhost and create keyspace demo. (com.datamountaineer.streamreactor.connect.
↳cassandra.CassandraConnection$:49)
[2016-05-06 13:34:41,263] INFO Using username_password. (com.datamountaineer.
↳streamreactor.connect.cassandra.CassandraConnection$:83)
[2016-05-06 13:34:41,459] INFO Did not find Netty's native epoll transport in the_
↳classpath, defaulting to NIO. (com.datastax.driver.core.NettyUtil:83)
[2016-05-06 13:34:41,823] INFO Using data-center name 'datacenter1' for_
↳DCAwareRoundRobinPolicy (if this is incorrect, please provide the correct_
↳datacenter name with DCAwareRoundRobinPolicy constructor) (com.datastax.driver.core.
↳policies.DCAwareRoundRobinPolicy:95)
[2016-05-06 13:34:41,824] INFO New Cassandra host localhost/127.0.0.1:9042 added (com.
↳datastax.driver.core.Cluster:1475)
[2016-05-06 13:34:41,868] INFO Connection to Cassandra established. (com.
↳datamountaineer.streamreactor.connect.cassandra.source.CassandraSourceTask:87)
```

If you switch back to the terminal you started the Connector in you should see the Cassandra Source being accepted and the task starting and processing the 3 existing rows.

```
[2016-05-06 13:44:33,132] INFO Source task Thread[WorkerSourceTask-cassandra-source-
↳orders-0,5,main] finished initialization and start (org.apache.kafka.connect.
↳runtime.WorkerSourceTask:342)
[2016-05-06 13:44:33,137] INFO Query SELECT * FROM demo.orders WHERE created >_
↳maxTimeuuid(?) AND created <= minTimeuuid(?) ALLOW FILTERING executing with_
↳bindings (2016-05-06 09:23:28+0200, 2016-05-06 13:44:33+0200). (com.datamountaineer.
↳streamreactor.connect.cassandra.source.CassandraTableReader:156)
[2016-05-06 13:44:33,151] INFO Querying returning results for demo.orders. (com.
↳datamountaineer.streamreactor.connect.cassandra.source.CassandraTableReader:185)
[2016-05-06 13:44:33,160] INFO Processed 3 rows for table orders-topic.orders (com.
↳datamountaineer.streamreactor.connect.cassandra.source.CassandraTableReader:206)
[2016-05-06 13:44:33,160] INFO Found 3. Draining entries to batchSize 100. (com.
↳datamountaineer.streamreactor.connect.queues.QueueHelpers$:45)
[2016-05-06 13:44:33,197] WARN Error while fetching metadata with correlation id 0 :
↳{orders-topic=LEADER_NOT_AVAILABLE} (org.apache.kafka.clients.NetworkClient:582)
[2016-05-06 13:44:33,406] INFO Found 0. Draining entries to batchSize 100. (com.
↳datamountaineer.streamreactor.connect.queues.QueueHelpers$:45)
```

Check Kafka, 3 rows as before.

```
confluent-3.0.1/bin/kafka-avro-console-consumer \
--zookeeper localhost:2181 \
--topic orders-topic \
--from-beginning
{"id":{"int":1},"created":{"string":"Thu May 05 13:24:22 CEST 2016"},"price":{"float
↳":94.2},"product":{"string":"DAX-P-20150201-95.7"},"qty":{"int":100}}
{"id":{"int":2},"created":{"string":"Thu May 05 13:26:21 CEST 2016"},"price":{"float
↳":99.5},"product":{"string":"OP-DAX-C-20150201-100"},"qty":{"int":100}}
{"id":{"int":3},"created":{"string":"Thu May 05 13:26:44 CEST 2016"},"price":{"float
↳":150.0},"product":{"string":"FU-KOSPI-C-20150201-100"},"qty":{"int":200}}
```


The Source tasks will continue to poll but not pick up any new rows yet.

Inserting new data

Now lets insert a row into the Cassandra table. Start the CQL shell.

```
bin ./cqlsh
Connected to Test Cluster at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 3.0.2 | CQL spec 3.3.1 | Native protocol v4]
Use HELP for help.
```

Execute the following:

```
use demo;

INSERT INTO orders (id, created, product, qty, price) VALUES (4, now(), 'FU-
↳DATAMOUNTAINEER-C-20150201-100', 500, 10000);

SELECT * FROM orders;

id | created | price | product
↳ | qty
-----+-----+-----+-----
↳ +-----
1 | 17fa1050-137e-11e6-ab60-c9fbe0223a8f | 94.2 | OP-DAX-P-20150201-95.
↳7 | 100
2 | 17fb6fe0-137e-11e6-ab60-c9fbe0223a8f | 99.5 | OP-DAX-C-20150201-
↳100 | 100
4 | 02acf5d0-1380-11e6-ab60-c9fbe0223a8f | 10000 | FU-DATAMOUNTAINEER-C-20150201-
↳100 | 500
3 | 17fbbe00-137e-11e6-ab60-c9fbe0223a8f | 150 | FU-KOSPI-C-20150201-
↳100 | 200

(4 rows)
cqlsh:demo>
```

Check the logs.

```
[2016-05-06 13:45:33,134] INFO Query SELECT * FROM demo.orders WHERE created >↳
↳maxTimeuuid(?) AND created <= minTimeuuid(?) ALLOW FILTERING executing with↳
↳bindings (2016-05-06 13:31:37+0200, 2016-05-06 13:45:33+0200). (com.datamountaineer.
↳streamreactor.connect.cassandra.source.CassandraTableReader:156)
[2016-05-06 13:45:33,137] INFO Querying returning results for demo.orders. (com.
↳datamountaineer.streamreactor.connect.cassandra.source.CassandraTableReader:185)
[2016-05-06 13:45:33,138] INFO Processed 1 rows for table orders-topic.orders (com.
↳datamountaineer.streamreactor.connect.cassandra.source.CassandraTableReader:206)
[2016-05-06 13:45:33,138] INFO Found 0. Draining entries to batchSize 100. (com.
↳datamountaineer.streamreactor.connect.queues.QueueHelpers$:45)
```

Check Kafka.

```
confluent-3.0.1/bin/kafka-avro-console-consumer \
--zookeeper localhost:2181 \
--topic orders-topic \
--from-beginning

{"id":{"int":1},"created":{"string":"17fa1050-137e-11e6-ab60-c9fbe0223a8f"},"price":{"
↳"float":94.2},"product":{"string":"OP-DAX-P-20150201-95.7"},"qty":{"int":100}}
```

```
{ "id": { "int": 2 }, "created": { "string": "17fb6fe0-137e-11e6-ab60-c9fbe0223a8f" }, "price": {  
  ↪ "float": 99.5 }, "product": { "string": "OP-DAX-C-20150201-100" }, "qty": { "int": 100 } }  
{ "id": { "int": 3 }, "created": { "string": "17fbbe00-137e-11e6-ab60-c9fbe0223a8f" }, "price": {  
  ↪ "float": 150.0 }, "product": { "string": "FU-KOSPI-C-20150201-100" }, "qty": { "int": 200 } }  
{ "id": { "int": 4 }, "created": { "string": "02acf5d0-1380-11e6-ab60-c9fbe0223a8f" }, "price": {  
  ↪ "float": 10000.0 }, "product": { "string": "FU-DATAMOUNTAINEER-C-20150201-100" }, "qty": {  
  ↪ "int": 500 } }
```

Bingo, we have our extra row.

Features

Kafka Connect Query Language

Both connectors support **K**afka **C**onnect **Q**uery **L**anguage found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

Data Types

The Source connector supports copying tables in bulk and incrementally to Kafka.

The following CQL data types are supported:

CQL Type	Connect Data Type
TimeUUID	Optional String
UUID	Optional String
Inet	Optional String
Ascii	Optional String
Text	Optional String
Timestamp	Optional String
Date	Optional String
Tuple	Optional String
UDT	Optional String
Boolean	Optional Boolean
TinyInt	Optional Int8
SmallInt	Optional Int16
Int	Optional Int32
Decimal	Optional String
Float	Optional Float32
Counter	Optional Int64
BigInt	Optional Int64
VarInt	Optional Int64
Double	Optional Int64
Time	Optional Int64
Blob	Optional Bytes
Map	Optional String
List	Optional String
Set	Optional String

Note: For Map, List and Set the value is extracted from the Cassandra Row and inserted as a JSON string representation.

Modes

The Source connector runs in both bulk and incremental mode.

Each mode has a polling interval. This interval determines how often the readers execute queries against the Cassandra tables. It applies to both incremental and bulk modes. The `cassandra.import.mode` setting controls the import behaviour.

Incremental

In `incremental` mode the connector supports querying based on a column in the tables with CQL data type of `TimeUUID`.

Kafka Connect tracks the latest record it retrieved from each table, so it can start at the correct location on the next iteration (or in case of a crash). In this case the maximum value of the records returned by the result-set is tracked and stored in Kafka by the framework. If no offset is found for the table at startup a default timestamp of 1900-01-01 is used. This is then passed to a prepared statement containing a range query. For example:

```
SELECT * FROM demo.orders WHERE created > maxTimeuuid(?) AND created <= minTimeuuid(?)
```

Bulk

In `bulk` mode the connector extracts the full table, no where clause is attached to the query.

Warning: Watch out with the poll interval. After each interval the bulk query will be executed again.

Topic Routing

The Sink supports topic routing that allows mapping the messages from topics to a specific table. For example map a topic called “bloomberg_prices” to a table called “prices”. This mapping is set in the `connect.cassandra.source.kcql` option.

Error Policies

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don’t have adequate monitoring. Data is not lost as it’s still in Kafka subject to Kafka’s retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers..

Retry

Any error on write to the target database causes the `RetryIterable` exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.cassandra.source.max.retries` and the `connect.cassandra.source.retry.interval`.

Configurations

`connect.cassandra.contact.points`

Contact points (hosts) in Cassandra cluster.

- Data type: string
- Optional : no

`connect.cassandra.key.space`

Key space the tables to write belong to.

- Data type: string
- Optional : no

`connect.cassandra.port`

Port for the native Java driver.

- Data type: int
- Optional : yes
- Default : 9042

`connect.cassandra.username`

Username to connect to Cassandra with if `connect.cassandra.authentication.mode` is set to *username_password*.

- Data type: string
- Optional : yes

`connect.cassandra.password`

Password to connect to Cassandra with if `connect.cassandra.authentication.mode` is set to *username_password*.

- Data type: string
- Optional : yes

`connect.cassandra.ssl.enabled`

Enables SSL communication against SSL enable Cassandra cluster.

- Data type: boolean
- Optional : yes
- Default : false

```
connect.cassandra.trust.store.password
```

Password for truststore.

- Data type: string
- Optional : yes

```
connect.cassandra.key.store.path
```

Path to truststore.

- Data type: string
- Optional : yes

```
connect.cassandra.key.store.password
```

Password for key store.

- Data type: string
- Optional : yes

```
connect.cassandra.ssl.client.cert.auth
```

Path to keystore.

- Data type: string
- Optional : yes

```
connect.cassandra.import.poll.interval
```

The polling interval between queries against tables for bulk mode in milliseconds. Default is 1 minute.

- Data type: int
- Optional : yes
- Default : 10

Warning: WATCH OUT WITH BULK MODE AS MAY REPEATEDLY PULL IN THE SAME DATE.

```
connect.cassandra.import.mode
```

Either bulk or incremental.

- Data type : string
- Optional : no

```
connect.cassandra.source.kcql
```

Kafka connect query language expression. Allows for expressive table to topic routing, field selection and renaming. In incremental mode the timestampColumn can be specified by PK colName.

Examples:

```
INSERT INTO TOPIC1 SELECT * FROM TOPIC1 PK myTimeUUICol
```

- Data type : string
- Optional : no

Warning: The timestamp column must be of CQL Type TimeUUID.

`connect.cassandra.import.fetch.size`

The fetch size for the Cassandra driver to read.

- Data type : int
- Optional : yes
- Default : 1000

`connect.cassandra.source.task.buffer.size`

The size of the queue for buffering resultset records before write to Kafka.

- Data type : int
- Optional : yes
- Default : 10000

`connect.cassandra.source.task.batch.size`

The number of records the Source task should drain from the reader queue.

- Data type : int
- Optional : yes
- Default : 1000

`connect.cassandra.source.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.cassandra.source.max.retries` option. The `connect.cassandra.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: high
- Default: throw

`connect.cassandra.source.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.cassandra.source.error.policy` is set to `retry`.

- Type: string
- Importance: high
- Default: 10

`connect.cassandra.source.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.cassandra.source.error.policy` set to **RETRY**.

- Type: int
- Importance: medium

- Default : 60000 (1 minute)

Bulk Example

```
name=cassandra-source-orders-bulk
connector.class=com.datamountaineer.streamreactor.connect.cassandra.source.
↳CassandraSourceConnector
connect.cassandra.key.space=demo
connect.cassandra.source.kcql=INSERT INTO TABLE_X SELECT * FROM TOPIC_Y
connect.cassandra.import.mode=bulk
connect.cassandra.contact.points=localhost
connect.cassandra.username=cassandra
connect.cassandra.password=cassandra
```

Incremental Example

```
name=cassandra-source-orders-incremental
connector.class=com.datamountaineer.streamreactor.connect.cassandra.source.
↳CassandraSourceConnector
connect.cassandra.key.space=demo
connect.cassandra.source.kcql=INSERT INTO TABLE_X SELECT * FROM TOPIC_Y PK created
connect.cassandra.import.mode=incremental
connect.cassandra.contact.points=localhost
connect.cassandra.username=cassandra
connect.cassandra.password=cassandra
```

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

For the Source connector, at present no column selection is handled, every column from the table is queried to column additions and deletions are handled in accordance with the compatibility mode of the Schema Registry.

Future releases will support auto creation of tables and adding columns on changes to the topic schema.

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect CoAP Source

A Connector and Source to stream messages from a CoAP server and write them to a Kafka topic.

The Source supports:

1. DTLS secure clients.
2. Observable resources.

The Source Connector automatically converts the CoAP response into a Kafka Connect `Struct` to be store in Kafka as Avro or Json dependent on the Converters used in Connect. The schema can found [here](#).

The key of the `Struct` message sent to Kafka is made from the source defined in the message, the resource on the CoAP server and the message id.

Prerequisites

- Confluent 3.0.1
- Java 1.8
- Scala 2.11

Setup

Confluent Setup

Follow the instructions [here](#).

CoAP Setup

The connector uses [Californium](#) Java API under the hood. A resource is host at `coap://californium.eclipse.org:5683`. Copper, a [FireFox](#) browser addon is available so you can browse the server and resources.. This is an observable non confirmable resource. You can view messages via the browser addon by selecting the resource and clicking the `observe` button on the top menu.

Source Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our [GitHub releases](#) page.

Starting the Connector

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed more by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools cli` to post in our distributed properties file for MQTT. If you are using the [dockers](#) you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.


```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create coap-source < conf/coap-source.properties

#Connector name=`coap-source`
name = coap-source
tasks = 1
connector.class = com.datamountaineer.streamreactor.connect.coap.source.
↪CoapSourceConnector
connect.coap.uri = coap://californium.eclipse.org:5683
connect.coap.kcql = INSERT INTO coap-topic SELECT * FROM obs-pumping-non
#task ids: 0
```

The `coap-source.properties` file defines:

1. The name of the source.
2. The name number of tasks.
3. The class containing the connector.
4. The uri of the CoAP Server and port to connect to.
5. *The KCQL routing querying*.. This specifies the target topic and source resource on the CoAP server.

If you switch back to the terminal you started Kafka Connect in you should see the CoAP Source being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
coap-source
```

[illegible]

```
[2017-01-09 20:42:45,927] INFO Discovered resources /.well-known/core (com.
↳ datamountaineer.streamreactor.connect.coap.source.CoapReader:60)
[2017-01-09 20:42:45,927] INFO Discovered resources /large (com.datamountaineer.
↳ streamreactor.connect.coap.source.CoapReader:60)
[2017-01-09 20:42:45,928] INFO Discovered resources /large-create (com.
↳ datamountaineer.streamreactor.connect.coap.source.CoapReader:60)
[2017-01-09 20:42:45,928] INFO Discovered resources /large-post (com.datamountaineer.
↳ streamreactor.connect.coap.source.CoapReader:60)
[2017-01-09 20:42:45,928] INFO Discovered resources /large-separate (com.
↳ datamountaineer.streamreactor.connect.coap.source.CoapReader:60)
[2017-01-09 20:42:45,928] INFO Discovered resources /large-update (com.
↳ datamountaineer.streamreactor.connect.coap.source.CoapReader:60)
```

Check for records in Kafka

Check for records in Kafka with the console consumer.

```
confluent-3.0.1/bin/kafka-avro-console-consumer \
--zookeeper localhost:2181 \
--topic coap-topic \
--from-beginning
```

```
{ "message_id": { "int": 4803 }, "type": { "string": "ACK" }, "code": "4.04", "raw_code": { "int
↳ : 132 }, "rtt": { "long": 35 }, "is_last": { "boolean": true }, "is_notification": { "boolean
↳ : false }, "source": { "string": "idvm-infk-mattern04.inf.ethz.ch:5683" }, "destination": {
↳ "string": "" }, "timestamp": { "long": 0 }, "token": { "string": "b24774e37c2314a4" }, "is_
↳ duplicate": { "boolean": false }, "is_confirmable": { "boolean": false }, "is_rejected": {
↳ "boolean": false }, "is_acknowledged": { "boolean": false }, "is_canceled": { "boolean": false },
↳ "accept": { "int": -1 }, "block1": { "string": "" }, "block2": { "string": "" }, "content_format":
↳ { "int": -1 }, "etags": [], "location_path": { "string": "" }, "location_query": { "string": "" },
↳ "max_age": { "long": 60 }, "observe": null, "proxy_uri": null, "size_1": null, "size_2": null,
↳ "uri_host": null, "uri_port": null, "uri_path": { "string": "" }, "uri_query": { "string": "" },
↳ "payload": { "string": "" } }
{ "message_id": { "int": 4804 }, "type": { "string": "ACK" }, "code": "4.04", "raw_code": { "int
↳ : 132 }, "rtt": { "long": 34 }, "is_last": { "boolean": true }, "is_notification": { "boolean
↳ : false }, "source": { "string": "idvm-infk-mattern04.inf.ethz.ch:5683" }, "destination": {
↳ "string": "" }, "timestamp": { "long": 0 }, "token": { "string": "b24774e37c2314a4" }, "is_
↳ duplicate": { "boolean": false }, "is_confirmable": { "boolean": false }, "is_rejected": {
↳ "boolean": false }, "is_acknowledged": { "boolean": false }, "is_canceled": { "boolean": false },
↳ "accept": { "int": -1 }, "block1": { "string": "" }, "block2": { "string": "" }, "content_format":
↳ { "int": -1 }, "etags": [], "location_path": { "string": "" }, "location_query": { "string": "" },
↳ "max_age": { "long": 60 }, "observe": null, "proxy_uri": null, "size_1": null, "size_2": null,
↳ "uri_host": null, "uri_port": null, "uri_path": { "string": "" }, "uri_query": { "string": "" },
↳ "payload": { "string": "" } }
{ "message_id": { "int": 4805 }, "type": { "string": "ACK" }, "code": "4.04", "raw_code": { "int
↳ : 132 }, "rtt": { "long": 35 }, "is_last": { "boolean": true }, "is_notification": { "boolean
↳ : false }, "source": { "string": "idvm-infk-mattern04.inf.ethz.ch:5683" }, "destination": {
↳ "string": "" }, "timestamp": { "long": 0 }, "token": { "string": "b24774e37c2314a4" }, "is_
↳ duplicate": { "boolean": false }, "is_confirmable": { "boolean": false }, "is_rejected": {
↳ "boolean": false }, "is_acknowledged": { "boolean": false }, "is_canceled": { "boolean": false },
↳ "accept": { "int": -1 }, "block1": { "string": "" }, "block2": { "string": "" }, "content_format":
↳ { "int": -1 }, "etags": [], "location_path": { "string": "" }, "location_query": { "string": "" },
↳ "max_age": { "long": 60 }, "observe": null, "proxy_uri": null, "size_1": null, "size_2": null,
↳ "uri_host": null, "uri_port": null, "uri_path": { "string": "" }, "uri_query": { "string": "" },
↳ "payload": { "string": "" } }
{ "message_id": { "int": 4806 }, "type": { "string": "ACK" }, "code": "4.04", "raw_code": { "int
↳ : 132 }, "rtt": { "long": 35 }, "is_last": { "boolean": true }, "is_notification": { "boolean
↳ : false }, "source": { "string": "idvm-infk-mattern04.inf.ethz.ch:5683" }, "destination": {
↳ "string": "" }, "timestamp": { "long": 0 }, "token": { "string": "b24774e37c2314a4" }, "is_
↳ duplicate": { "boolean": false }, "is_confirmable": { "boolean": false }, "is_rejected": {
↳ "boolean": false }, "is_acknowledged": { "boolean": false }, "is_canceled": { "boolean": false },
↳ "accept": { "int": -1 }, "block1": { "string": "" }, "block2": { "string": "" }, "content_format":
↳ { "int": -1 }, "etags": [], "location_path": { "string": "" }, "location_query": { "string": "" },
↳ "max_age": { "long": 60 }, "observe": null, "proxy_uri": null, "size_1": null, "size_2": null,
↳ "uri_host": null, "uri_port": null, "uri_path": { "string": "" }, "uri_query": { "string": "" },
↳ "payload": { "string": "" } }
```

Features

1. Secure DTLS client connection.
2. Supports Observable resources to stream changes on a resource to Kafka.
3. Routing of data via KCQL to topics.
4. Automatic conversion of CoAP Response messages to Connect Structs.

Kafka Connect Query Language

Kafka **C**onnect **Q**uery **L**anguage found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The CoAP Source supports the following:

```
INSERT INTO <topic> SELECT * FROM <resource>
```

No selection of fields on the CoAP message is support. All the message attributes are mapped to predefined `Struct` representing the CoAP response message.

DTLS Client

The Connector use the [Californium](#) Java API and for secure connections use the Scandium security module provided by Californium. Scandium (Sc) is an implementation of Datagram Transport Layer Security 1.2, also known as [RFC 6347](#).

Please refer to the Californium [certification](#) repo page for more information.

DTLS Client connections can be enabled by setting the `connect.coap.keystore.pass` property. If set you must provide the following or the Connector will not start and throw a configuration exception:

- `connect.coap.keystore.pass`
- `connect.coap.keystore.path`
- `connect.coap.truststore.pass`
- `connect.coap.truststore.path`

Warning: The key and truststore must be available on the local disk of the worker task.

Loading specific certificates can be achieved by providing a comma separated list for the `connect.coap.certs` configuration option. The certificate chain can be set by the `connect.coap.cert.chain.key` configuration option.

Configurations

`connect.coap.uri`

Uri of the CoAP server.

- Data Type : string
- Importance: high
- Optional : no

`connect.coap.kcql`

The KCQL statement to select and route resources to topics.

- Data Type : string
- Importance: high
- Optional : no

`connect.coap.bind.port`

The port the DTLS connector will bind to on the Connector host.

- Data Type : int
- Importance: medium
- Optional : yes
- Default : 0

`connect.coap.bind.host`

The hostname the DTLS connector will bind to on the Connector host.

- Data Type : string
- Importance: medium
- Optional : yes
- Default : localhost

`connect.coap.keystore.pass`

The password of the key store

- Data Type : string
- Importance: medium
- Optional : yes
- Default : rootPass

`connect.coap.keystore.path`

The path to the keystore.

- Data Type : string
- Importance: medium
- Optional : yes
- Default :

`connect.coap.truststore.pass`

The password of the trust store

- Data Type : string
- Importance: medium

- Optional : yes
- Default : rootPass

`connect.coap.truststore.path`

The path to the truststore.

- Data Type : string
- Importance: medium
- Optional : yes
- Default :

`connect.coap.certs`

The certificates to load from the trust store.

- Data Type : list
- Importance: medium
- Optional : yes
- Default :

`connect.coap.cert.chain.key`

The key to use to get the certificate chain.

- Data Type : string
- Importance: medium
- Optional : yes
- Default : client

Schema Evolution

The schema is fixed.

The following schema is used for the key:

Name	Type
source	Optional string
source_resource	Optional String
message_id	Optional int32

The following schema is used for the payload:

Name	Type
message_id	Optional int32
type	Optional String
code	Optional String
raw_code	Optional int32
rtt	Optional int64
is_last	Optional boolean
is_notification	Optional boolean
source	Optional String
Continued on next page	

Table 1.1 – continued from previous page

destination	Optional String
timestamp	Optional int64
token	Optional String
is_duplicate	Optional boolean
is_confirmable	Optional boolean
is_rejected	Optional boolean
is_acknowledged	Optional boolean
is_canceled	Optional boolean
accept	Optional int32
block1	Optional String
block2	Optional String
content_format	Optional int32
etags	Array of Optional Strings
location_path	Optional String
location_query	Optional String
max_age	Optional int64
observe	Optional int32
proxy_uri	Optional String
size_1	Optional String
size_2	Optional String
uri_host	Optional String
uri_port	Optional int32
uri_path	Optional String
uri_query	Optional String
payload	Optional String

Kafka Connect ReThink

A Connector and Source to write events from ReThinkDB to Kafka. The connector subscribes to changefeeds on tables and streams the records to Kafka.

The Source supports:

1. *The KCQL routing querying* - Table to topic routing
2. Initialization (Read feed from start) via KCQL.
3. ReThinkDB type (add, delete, update).
4. ReThinkDB initial states.

Prerequisites

- Confluent 3.0.1
- RethinkDb 2.3.3
- Java 1.8
- Scala 2.11

Setup

Rethink Setup

Download and install RethinkDb. Follow the instruction [here](#) dependent on your operating system.

Confluent Setup

Follow the instructions [here](#).

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools cli` to post in our distributed properties file for ReThinkDB. If you are using the [dockers](#) you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create rethink-source < conf/rethink-source.properties
#Connector name=`rethink-source`
name=rethink-source
connect.rethink.source.host=localhost
connect.rethink.source.port=28015
connector.class=com.datamountaineer.streamreactor.connect.rethink.source.
↪ReThinkSourceConnector
tasks.max=1
connect.rethink.source.db=test
connect.rethink.sink.kcql=INSERT INTO rethink-topic SELECT * FROM source-test
#task ids: 0
```

The `rethink-source.properties` file defines:

1. The name of the source.
2. The name of the rethink host to connect to.
3. The rethink port to connect to.
4. The Source class.

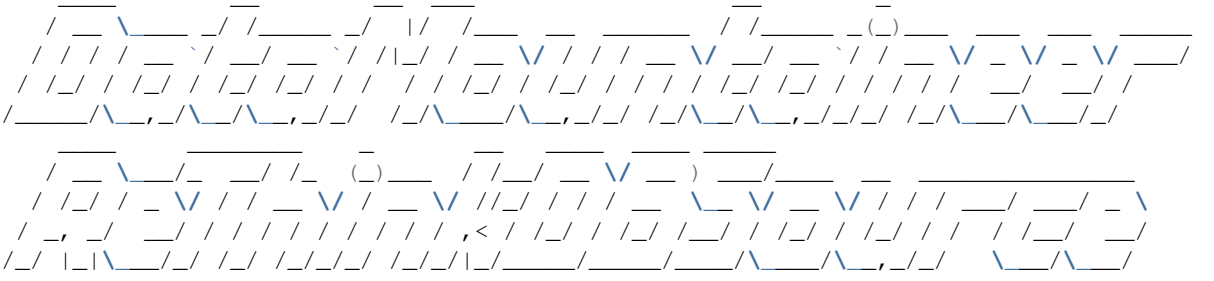
5. The max number of tasks the connector is allowed to created. The connector splits and groups the *connect.rethink.source.kcql* by the number of tasks to ensure a distribution based on allowed number of tasks and Source tables.
6. The ReThinkDB database to connect to.
7. *The KCQL routing querying.*

If you switch back to the terminal you started the Connector in you should see the ReThinkDB Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
rethink-source
```

```
[2016-10-05 12:09:35,414] INFO
```



By Andrew Stevenson (com.datamountaineer.streamreactor.connect.rethink.source.
↪ReThinkSourceTask:48)

```
[2016-10-05 12:09:35,420] INFO ReThinkSourceConfig values:  
connect.rethink.source.port = 28015  
connect.rethink.source.host = localhost  
connect.rethink.source.kcql = insert into rethink-topic select * from source-test  
connect.rethink.source.db = test
```

Test Records

Go to the ReThink Admin console <http://localhost:8080/#tables> and add a database called *test* and table called *source-test*. Then on the Data Explorer tab insert the following and hit run to insert the record into the table.

```
r.table('source_test').insert([
  { name: "datamountaineers-rule", tv_show: "Battlestar Galactica",
    posts: [
      {title: "Decommissioning speech3", content: "The Cylon War is long over..."},
      {title: "We are at war", content: "Moments ago, this ship received word..."},
      {title: "The new Earth", content: "The discoveries of the past few days..."}
    ]
  }
])
```

Check for records in Kafka

Check for records in Kafka with the console consumer..


```
confluent-3.0.1/bin/kafka-avro-console-consumer \
  --zookeeper localhost:2181 \
  --topic rethink-topic \
  --from-beginning

{"state":{"string":"initializing"},"old_val":null,"new_val":null,"type":{"string":
↪ "state"}}
{"state":{"string":"ready"},"old_val":null,"new_val":null,"type":{"string":"state"}
↪ }
{"state":null,"old_val":null,"new_val":{"string":{"tv_show=Battlestar Galactica,
↪ name=datamountaineers-rule, id=ec9d337e-ee07-4128-a830-22e4f055ce64, posts=[
↪ {title=Decommissioning speech3, content=The Cylon War is long over...}, {title=We
↪ are at war, content=Moments ago, this ship received word...}, {title=The new Earth,
↪ content=The discoveries of the past few days...}}]}"},"type":{"string":"add"}}
```

Features

The ReThinkDb Source writes change feed records from RethinkDb to Kafka.

The Source supports:

1. Table to topic routing
2. Initialization (Read feed from start)
3. ReThinkDB type (add, delete, update)
4. ReThinkDB initial states

Kafka Connect Query Language

Kafka **C**onnect **Q**uery **L**anguage found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The ReThink Source supports the following:

```
INSERT INTO <target table> SELECT <fields> FROM <source topic> <INITIALIZE>
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA
INSERT INTO tableA SELECT * FROM topicA

#Insert mode, select all fields from topicA and write to tableA, read from start
INSERT INTO tableA SELECT * FROM topicA INITIALIZE
```

Configurations

```
connect.rethink.source.kcql
```

Kafka connect query language expression. Allows for expressive topic to table routing, field selection and renaming. Fields to be used as the row key can be set by specifying the PK. The below example uses field1 as the primary key.

- Data type : string
- Importance: high

- Optional : no

Examples:

```
INSERT INTO TOPIC1 SELECT * FROM TABLE1; INSERT INTO TOPIC2 SELECT * FROM TABLE2
```

```
connect.rethink.source.host
```

Specifies the rethink server.

- Data type : string
- Importance: high
- Optional : no

```
connect.rethink.source.port
```

Specifies the rethink server port number.

- Data type : int
- Importance: high
- Optional : yes

Example

```
name=rethink-source
connect.rethink.source.db=localhost
connect.rethink.source.port=28015
connector.class=com.datamountaineer.streamreactor.connect.rethink.source.
↪ReThinkSourceConnector
tasks.max=1
connect.rethink.sink.kcql=INSERT INTO rethink-topic SELECT * FROM source-test
```

Schema Evolution

The schema is fixed. The following schema is used:

Name	Type	Optional
state	string	yes
new_val	string	yes
old_val	string	yes
type	string	yes

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect Yahoo

A Connector and Source to write events from the Yahoo Finance API to Kafka.

DOCS WIP!

Kafka Connect Mqtt Source

A Connector to read events from Mqtt and push them to Kafka. The connector subscribes to the specified topics and streams the records to Kafka.

The Source supports:

1. Pluggable converters of MQTT payloads.
2. *Out of the box converters for Json/Avro and Binary*
3. *The KCQL routing querying* - Topic to Topic mapping and Field selection.

Prerequisites

- Confluent 3.0.1
- Mqtt server
- Java 1.8
- Scala 2.11

Setup

Confluent Setup

Follow the instructions [here](#).

Mqtt Setup

For testing we will use a simple application spinning up an mqtt server using Moquette. Download and unzip [this](#) Also download the [schema](#) required to read the avro records

Once you have unpacked the archiver you should start the server .. sourcecode:: bash

```
bin/mqtt-server
```

You should see the following outcome:

```
log4j:WARN No appenders could be found for logger (io.moquette.server.Server).
log4j:WARN Please initialize the log4j system properly.
log4j:WARN See http://logging.apache.org/log4j/1.2/faq.html#noconfig for more info.
Starting mqtt service on port 11883
Hit Enter to start publishing messages on topic: /mjson and /mavro
```

The server has started but no records have been published yet. More on this later once we start the source.

Source Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools` cli to post in our distributed properties file for MQTT. If you are using the [dockers](#) you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create mqtt-source < conf/source.kcql/mqtt-source.properties

#Connector name=`mqtt-source`
name=mqtt-source
tasks.max=1
connect.mqtt.connection.clean=true
connect.mqtt.connection.timeout=1000
connect.mqtt.source.kcql=INSERT INTO kjson SELECT * FROM /mjson;INSERT INTO kavro_
↪SELECT * FROM /mavro
connect.mqtt.connection.keep.alive=1000
connect.mqtt.source.converters=/mjson=com.datamountaineer.streamreactor.connect.
↪converters.source.JsonSimpleConverter;/mavro=com.datamountaineer.streamreactor.
↪connect.converters.source.AvroConverter
connect.source.converter.avro.schemas=/mavro=$PATH_TO/temperaturemeasure.avro
connect.mqtt.client.id=dm_source_id,
connect.mqtt.converter.throw.on.error=true
connect.mqtt.hosts=tcp://127.0.0.1:11883
connect.mqtt.service.quality=1
connector.class=com.datamountaineer.streamreactor.connect.mqtt.source.
↪MqttSourceConnector
#task ids: 0
```

The `mqtt-source.properties` file defines:

1. The name of the source.
2. The name number of tasks.
3. Clean the mqtt connection.
4. The Kafka Connect Query statements to read from json and avro topics and insert into Kafka kjson and kavro topics.
5. Setting the time window to emit keep alive pings

6. Set the converters for each of the Mqtt topics. If a source doesn't get a converter set it will default to BytesConverter
7. Set the avro schema for the 'avro' Mqtt topic.
8. The mqtt client identifier.
9. If a conversion can't happen it will throw an exception.
10. The connection to the Mqtt server.
11. The quality of service for the messages.
12. Set the connector source class.

If you switch back to the terminal you started Kafka Connect in you should see the MQTT Source being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
mqtt-source
```

```
[2016-12-20 16:51:08,058] INFO  
by Stefan Bocutiu  
(com.datamountaineer.streamreactor.connect.mqtt.source.MqttSourceTask:37)  
[2016-12-20 16:51:08,090] INFO MqttSourceConfig values:  
connect.mqtt.source.kcql = INSERT INTO kjson SELECT * FROM /mjson;INSERT INTO_  
←kavro SELECT * FROM /mavro  
connect.mqtt.service.quality = 1  
connect.mqtt.connection.ssl.cert = null  
connect.mqtt.source.converters = /mjson=com.datamountaineer.streamreactor.connect.  
←converters.source.JsonSimpleConverter;/mavro=com.datamountaineer.streamreactor.  
←connect.converters.source.AvroConverter  
connect.mqtt.connection.keep.alive = 1000  
connect.mqtt.hosts = tcp://127.0.0.1:11883  
connect.mqtt.converter.throw.on.error = true  
connect.mqtt.connection.timeout = 1000  
connect.mqtt.user = null  
connect.mqtt.connection.clean = true  
connect.mqtt.connection.ssl.ca.cert = null  
connect.mqtt.connection.ssl.key = null  
connect.mqtt.password = null  
connect.mqtt.client.id = dm_source_id  
(com.datamountaineer.streamreactor.connect.mqtt.config.MqttSourceConfig:178)
```

Test Records

Go to the mqtt-server application you downloaded and unzipped and execute:

```
./bin/mqtt-server
```

This will put the following records into the avro and json Mqtt topic:

```
TemperatureMeasure(1, 31.1, "EMEA", System.currentTimeMillis())
TemperatureMeasure(2, 30.91, "EMEA", System.currentTimeMillis())
TemperatureMeasure(3, 30.991, "EMEA", System.currentTimeMillis())
TemperatureMeasure(4, 31.061, "EMEA", System.currentTimeMillis())
TemperatureMeasure(101, 27.001, "AMER", System.currentTimeMillis())
TemperatureMeasure(102, 38.001, "AMER", System.currentTimeMillis())
TemperatureMeasure(103, 26.991, "AMER", System.currentTimeMillis())
TemperatureMeasure(104, 34.17, "AMER", System.currentTimeMillis())
```

Check for records in Kafka

Check for records in Kafka with the console consumer. the topic for kjson (the Mqtt payload was a json and we translated that into a Kafka Connect Struct)

```
bin/kafka-avro-console-consumer --zookeeper localhost:2181 --topic kjson --from-
↪beginning
```

You should see the following output

```
SLF4J: Actual binding is of type [org.slf4j.impl.Log4jLoggerFactory]
{"deviceId":1,"value":31.1,"region":"EMEA","timestamp":1482236627236}
{"deviceId":2,"value":30.91,"region":"EMEA","timestamp":1482236627236}
{"deviceId":3,"value":30.991,"region":"EMEA","timestamp":1482236627236}
{"deviceId":4,"value":31.061,"region":"EMEA","timestamp":1482236627236}
{"deviceId":101,"value":27.001,"region":"AMER","timestamp":1482236627236}
{"deviceId":102,"value":38.001,"region":"AMER","timestamp":1482236627236}
{"deviceId":103,"value":26.991,"region":"AMER","timestamp":1482236627236}
{"deviceId":104,"value":34.17,"region":"AMER","timestamp":1482236627236}
```

Check for records in Kafka with the console consumer. the topic for kavro (the Mqtt payload was a avro and we translated that into a Kafka Connect Struct)

```
bin/kafka-avro-console-consumer --zookeeper localhost:2181 --topic kavro --from-
↪beginning
```

You should see the following output

```
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.slf4j.impl.Log4jLoggerFactory]
{"deviceId":1,"value":31.1,"region":"EMEA","timestamp":1482236627236}
{"deviceId":2,"value":30.91,"region":"EMEA","timestamp":1482236627236}
{"deviceId":3,"value":30.991,"region":"EMEA","timestamp":1482236627236}
{"deviceId":4,"value":31.061,"region":"EMEA","timestamp":1482236627236}
{"deviceId":101,"value":27.001,"region":"AMER","timestamp":1482236627236}
{"deviceId":102,"value":38.001,"region":"AMER","timestamp":1482236627236}
{"deviceId":103,"value":26.991,"region":"AMER","timestamp":1482236627236}
{"deviceId":104,"value":34.17,"region":"AMER","timestamp":1482236627236}
```

Features

The Mqtt source allows you to plugin your own converter. Say you receive protobuf data, all you have to do is to write your own very specific converter that knows how to convert from protobuf to SourceRecord. All you have to do is set the `connect.mqtt.source.converters` for the topic containing the protobuf data.

Converters

We provide four converters out of the box but you can plug your own. See an example [here](#).

AvroConverter

```
com.datamountaineer.streamreactor.connect.mqtt.source.converters.AvroConverter
```

The payload for the Mqtt message is an Avro message. In this case you need to provide a path for the Avro schema file to be able to decode it.

JsonSimpleConverter

```
com.datamountaineer.streamreactor.connect.mqtt.source.converters.
JsonSimpleConverter
```

The payload for the Mqtt message is a Json message. This converter will parse the json and create an Avro record for it which will be sent over to Kafka.

JsonConverterWithSchemaEvolution

An experimental converter for converting Json messages to Avro. The resulting Avro schema is fully compatible as new fields are added as the MQTT json payload evolves.

BytesConverter

```
com.datamountaineer.streamreactor.connect.mqtt.source.converters.
BytesConverter
```

This is the default implementation. The Mqtt payload is taken as is: an array of bytes and sent over Kafka as an avro record with `Schema.BYTES`. You don't have to provide a mapping for the source to get this converter!!

Kafka Connect Query Language

Kafka **C**onnect **Q**uery **L**anguage found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The Mqtt Source supports the following:

```
INSERT INTO <target topic> SELECT * FROM <mqtt source topic>
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA
INSERT INTO kafkaTopic1 SELECT * FROM mqttTopicA
```

Configurations

```
connect.mqtt.source.kcql
```

Kafka connect query language expression. Allows for expressive Mqtt topic to Kafka topic routing. Currently there is no support for filtering the fields from the incoming payload.

- Data type : string
- Importance: high
- Optional : no

Examples:

```
INSERT INTO KAFKA_TOPIC1 SELECT * FROM MQTT_TOPIC1; INSERT INTO KAFKA_TOPIC2 SELECT *  
FROM MQTT_TOPIC2
```

`connect.mqtt.hosts`

Specifies the mqtt connection endpoints.

- Data type : string
- Importance: high
- Optional : no

Example:

```
tcp://broker.datamountaineer.com:1883
```

`connect.mqtt.service.quality`

The Quality of Service (QoS) level is an agreement between sender and receiver of a message regarding the guarantees of delivering a message. There are 3 QoS levels in MQTT: At most once (0); At least once (1); Exactly once (2).

- Data type : int
- Importance: high
- Optional : yes
- Default: 1

`connect.mqtt.user`

Contains the Mqtt connection user name

- Data type : string
- Importance: medium
- Optional : yes
- Default: null

`connect.mqtt.password`

Contains the Mqtt connection password

- Data type : string
- Importance: medium
- Optional : yes
- Default: null

`connect.mqtt.client.id`

Provides the client connection identifier. If is not provided the framework will generate one.

- Data type: string
- Importance: medium
- Optional: yes
- Default: generated

`connect.mqtt.connection.timeout`

Sets the timeout to wait for the broker connection to be established

- Data type: int
- Importance: medium
- Optional: yes
- Default: 3000 (ms)

`connect.mqtt.connection.clean`

The clean session flag indicates the broker, whether the client wants to establish a persistent session or not. A persistent session (the flag is false) means, that the broker will store all subscriptions for the client and also all missed messages, when subscribing with Quality of Service (QoS) 1 or 2. If clean session is set to true, the broker won't store anything for the client and will also purge all information from a previous persistent session.

- Data type: boolean
- Importance: medium
- Optional: yes
- Default: true

`connect.mqtt.connection.keep.alive`

The keep alive functionality assures that the connection is still open and both broker and client are connected to one another. Therefore the client specifies a time interval in seconds and communicates it to the broker during the establishment of the connection. The interval is the longest possible period of time, which broker and client can endure without sending a message.”

- Data type: int
- Importance: medium
- Optional: yes
- Default: 5000

`connect.mqtt.connection.ssl.ca.cert`

Provides the path to the CA certificate file to use with the Mqtt connection”

- Data type: string
- Importance: medium
- Optional: yes
- Default: null

`connect.mqtt.connection.ssl.cert`

Provides the path to the certificate file to use with the Mqtt connection

- Data type: string
- Importance: medium

- Optional: yes
- Default: null

`connect.mqtt.connection.ssl.key`

Certificate private key file path.

- Data type: string
- Importance: medium
- Optional: yes
- Default: null

`connect.mqtt.source.converters`

Contains a tuple (mqtt source topic and the canonical class name for the converter of a raw Mqtt message bytes to a SourceRecord). If the source topic is not matched it will default to the BytesConverter. This will send an avro message over Kafka using Schema.BYTES

- Data type: string
- Importance: medium
- Optional: yes
- Default: null

```
mqtt_source1=com.datamountaineer.streamreactor.connect.mqtt.source.converters.  
↳AvroConverter;mqtt_source2=com.datamountaineer.streamreactor.connect.mqtt.source.  
↳converters.JsonSimpleConverter
```

`connect.mqtt.converter.throw.on.error`

If set to false the conversion exception will be swallowed and everything carries on BUT the message is lost!!; true will throw the exception.Default is false.”

- Data type: bool
- Importance: medium
- Optional: yes
- Default: false

`connect.source.converter.avro.schemas`

If the AvroConverter is used you need to provide an avro Schema to be able to read and translate the raw bytes to an avro record. The format is \$MQTT_TOPIC=\$PATH_TO_AVRO_SCHEMA_FILE

- Data type: bool
- Importance: medium
- Optional: yes
- Default: null

Example

```

name=mqtt-source
tasks.max=1
connect.mqtt.connection.clean=true
connect.mqtt.connection.timeout=1000
connect.mqtt.source.kcql=INSERT INTO kjson SELECT * FROM /mjson;INSERT INTO kavro_
↪SELECT * FROM /mavro
connect.mqtt.connection.keep.alive=1000
connect.mqtt.source.converters=/mjson=com.datamountaineer.streamreactor.connect.
↪converters.source.JsonSimpleConverter;/mavro=com.datamountaineer.streamreactor.
↪connect.converters.source.AvroConverter
connect.source.converter.avro.schemas=/mavro=$PATH_TO/temperaturemeasure.avro
connect.mqtt.client.id=dm_source_id,
connect.mqtt.converter.throw.on.error=true
connect.mqtt.hosts=tcp://127.0.0.1:11883
connect.mqtt.service.quality=1
connector.class=com.datamountaineer.streamreactor.connect.mqtt.source.
↪MqttSourceConnector

```

Provide your own Converter

You can always provide your own logic for converting the raw Mqtt message bytes to your an avro record. If you have messages coming in Protobuf format you can deserialize the message based on the schema and create the avro record. All you have to do is create a new project and add our dependency:

Gradle:

```
compile "com.datamountaineer:kafka-connect-common:0.6.1"
```

Maven:

```

<dependency>
  <groupId>com.datamountaineer</groupId>
  <artifactId>kafka-connect-common</artifactId>
  <version>0.6.1</version>
</dependency>

```

Then all you have to do is implement `com.datamountaineer.streamreactor.connect.converters.source.Converter`.

Here is our BytesConverter class code:

```

class BytesConverter extends Converter {
  override def convert(kafkaTopic: String, sourceTopic: String, messageId: Int,
↪bytes: Array[Byte]): SourceRecord = {
    new SourceRecord(Collections.singletonMap(Converter.TopicKey, sourceTopic),
      null,
      kafkaTopic,
      MsgKey.schema,
      MsgKey.getStruct(sourceTopic, messageId),
      Schema.BYTES_SCHEMA,
      bytes)
  }
}

```

All our implementation will send a a MsgKey object as the Kafka message key. It contains the Mqtt source topic and the Mqtt message id

Deployment Guidelines

TODO

TroubleShooting

TODO

Sink Connectors

Sink connectors stream data from Kafka into external systems.

Kafka Connect Cassandra Sink

The Cassandra Sink allows you to write events from Kafka to Cassandra. The connector converts the value from the Kafka Connect SinkRecords to Json and uses Cassandra's JSON insert functionality to insert the rows. The task expects pre-created tables in Cassandra.

Note: The table and keyspace must be created before hand!

Note: If the target table has TimeUUID fields the payload string the corresponding field in Kafka must be a UUID.

The Sink supports:

1. *The KCQL routing querying* - Kafka topic payload field selection is supported, allowing you to have choose selection of fields or all fields written to Cassandra.
2. Topic to table routing via KCQL.
3. Error policies for handling failures.
4. Schema.Struct and payload Struct, Schema.String and Json payload and Json payload with no schema

The Sink supports three Kafka payloads type:

Connect entry with Schema.Struct and payload Struct. If you follow the best practice while producing the events, each message should carry its schema information. Best option is to send Avro. Your connect configurations should be set to `value.converter=io.confluent.connect.avro.AvroConverter`. You can find an example [here](#). To see how easy is to have your producer serialize to Avro have a look at [this](#). This requires SchemaRegistry which is open source thanks to Confluent! Alternatively you can send Json + Schema. In this case your connect configuration should read `value.converter=org.apache.kafka.connect.json.JsonConverter`. The difference would be to point your serialization to `org.apache.kafka.connect.json.JsonSerializer`. This doesn't require the SchemaRegistry.

Connect entry with Schema.String and payload json String. Sometimes the producer would find it easier, despite sending Avro to produce a GenericRecord, to just send a message with Schema.String and the json string.

Connect entry without a schema and the payload json String. There are many existing systems which are publishing json over Kafka and bringing them in line with best practices is quite a challenge. Hence we added the support

Prerequisites

- Cassandra 2.2.4
- Confluent 3.0.0
- Java 1.8
- Scala 2.11

Setup

Before we can do anything, including the QuickStart we need to install Cassandra and the Confluent platform.

Cassandra Setup

First download and install Cassandra if you don't have a compatible cluster available.

```
#make a folder for cassandra
mkdir cassandra

#Download Cassandra
wget http://apache.cs.uu.nl/cassandra/3.5/apache-cassandra-3.5-bin.tar.gz

#extract archive to cassandra folder
tar -xvf apache-cassandra-3.5-bin.tar.gz -C cassandra

#Set up environment variables
export CASSANDRA_HOME=~/.cassandra/apache-cassandra-3.5-bin
export PATH=$PATH:$CASSANDRA_HOME/bin

#Start Cassandra
sudo sh ~/.cassandra/bin/cassandra
```

Confluent Setup

Follow the instructions [here](#).

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Test data

The Sink currently expects precreated tables and keyspaces. So lets create a keyspace and table in Cassandra via the CQL shell first.

Once you have installed and started Cassandra create a table to write records to. This snippet creates a table called orders to hold fictional orders on a trading platform.

Start the Cassandra cql shell

```
bin ./cqlsh
Connected to Test Cluster at 127.0.0.1:9042.
[cqlsh 5.0.1 | Cassandra 3.0.2 | CQL spec 3.3.1 | Native protocol v4]
Use HELP for help.
cqlsh>
```

Execute the following to create the keyspace and table:

```
CREATE KEYSPACE demo WITH REPLICATION = {'class' : 'SimpleStrategy', 'replication_
↪factor' : 3};
use demo;

create table orders (id int, created varchar, product varchar, qty int, price float,
↪PRIMARY KEY (id, created))
WITH CLUSTERING ORDER BY (created asc);
```

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools` cli to post in our distributed properties file for Cassandra. If you are using the *docker*s you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create cassandra-sink-orders < conf/cassandra-sink.properties

#Connector `cassandra-sink-orders`:
name=cassandra-sink-orders
connector.class=com.datamountaineer.streamreactor.connect.cassandra.sink.
↪CassandraSinkConnector
tasks.max=1
topics=orders-topic
connect.cassandra.sink.kcql=INSERT INTO orders SELECT * FROM orders-topic
connect.cassandra.contact.points=localhost
connect.cassandra.port=9042
connect.cassandra.key.space=demo
connect.cassandra.contact.points=localhost
connect.cassandra.username=cassandra
connect.cassandra.password=cassandra
#task ids: 0
```

The `cassandra-sink.properties` file defines:

1. The name of the sink.
2. The Sink class.
3. The max number of tasks the connector is allowed to created (1 task only).

4. The topics to read from.
5. *The `KCQL` routing querying.*
6. The Cassandra host.
7. The Cassandra port.
8. The Cassandra Keyspace.
9. The username.
10. The password.

If you switch back to the terminal you started the Connector in you should see the Cassandra Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
cassandra-sink
```

```
[2016-05-06 13:52:28,178] INFO

By Andrew Stevenson. (com.datamountaineer.streamreactor.connect.cassandra.sink.
↳CassandraSinkTask:50)
[2016-05-06 13:52:28,179] INFO Attempting to connect to Cassandra cluster at
↳localhost and create keyspace demo. (com.datamountaineer.streamreactor.connect.
↳cassandra.CassandraConnection$:49)
[2016-05-06 13:52:28,187] WARN You listed localhost/0:0:0:0:0:0:0:1:9042 in your
↳contact points, but it wasn't found in the control host's system.peers at startup
↳(com.datastax.driver.core.Cluster:2105)
[2016-05-06 13:52:28,211] INFO Using data-center name 'datacenter1' for
↳DCAwareRoundRobinPolicy (if this is incorrect, please provide the correct
↳datacenter name with DCAwareRoundRobinPolicy constructor) (com.datastax.driver.core.
↳policies.DCAwareRoundRobinPolicy:95)
[2016-05-06 13:52:28,211] INFO New Cassandra host localhost/127.0.0.1:9042 added (com.
↳datastax.driver.core.Cluster:1475)
[2016-05-06 13:52:28,290] INFO Initialising Cassandra writer. (com.datamountaineer.
↳streamreactor.connect.cassandra.sink.CassandraJsonWriter:40)
[2016-05-06 13:52:28,295] INFO Preparing statements for orders-topic. (com.
↳datamountaineer.streamreactor.connect.cassandra.sink.CassandraJsonWriter:62)
[2016-05-06 13:52:28,305] INFO Sink task org.apache.kafka.connect.runtime.
↳WorkerSinkTask@37e65d57 finished initialization and start (org.apache.kafka.connect.
↳runtime.WorkerSinkTask:155)
[2016-05-06 13:52:28,331] INFO Source task Thread[WorkerSourceTask-cassandra-source-
↳orders-0,5,main] finished initialization and start (org.apache.kafka.connect.
↳runtime.WorkerSourceTask:342)
```

Test Records

Now we need to put some records it to the orders-topic. We can use the `kafka-avro-console-producer` to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema matches the table created earlier.

Hint: If your input topic doesn't match the target use Kafka Streams to transform in realtime the input. Also checkout the [Plumber](#), which allows you to inject a Lua script into [Kafka Streams](#) to do this, no Java or Scala required!

```
${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic orders-topic \
--property value.schema='{ "type": "record", "name": "myrecord", "fields": [{ "name": "id",
↪ "type": "int"}, { "name": "created", "type": "string"}, { "name": "product", "type": "string"},
↪ { "name": "price", "type": "double"} ] }'
```

Now the producer is waiting for input. Paste in the following (each on a line separately):

```
{ "id": 1, "created": "2016-05-06 13:53:00", "product": "OP-DAX-P-20150201-95.7",
↪ "price": 94.2 }
{ "id": 2, "created": "2016-05-06 13:54:00", "product": "OP-DAX-C-20150201-100", "price
↪ ": 99.5 }
{ "id": 3, "created": "2016-05-06 13:55:00", "product": "FU-DATAMOUNTAINEER-20150201-
↪ 100", "price": 10000 }
{ "id": 4, "created": "2016-05-06 13:56:00", "product": "FU-KOSPI-C-20150201-100",
↪ "price": 150 }
```

Now if we check the logs of the connector we should see 2 records being inserted to Cassandra:

```
[2016-05-06 13:55:10,368] INFO Setting newly assigned partitions [orders-topic-0] for
↪ group connect-cassandra-sink-1 (org.apache.kafka.clients.consumer.internals.
↪ ConsumerCoordinator:219)
[2016-05-06 13:55:16,423] INFO Received 4 records. (com.datamountaineer.streamreactor.
↪ connect.cassandra.sink.CassandraJsonWriter:96)
[2016-05-06 13:55:16,484] INFO Processed 4 records. (com.datamountaineer.
↪ streamreactor.connect.cassandra.sink.CassandraJsonWriter:138)
```

```
use demo;
SELECT * FROM orders;
```

id	created	price	product	qty
1	2016-05-06 13:53:00	94.2	OP-DAX-P-20150201-95.7	100
2	2016-05-06 13:54:00	99.5	OP-DAX-C-20150201-100	100
3	2016-05-06 13:55:00	10000	FU-DATAMOUNTAINEER-20150201-100	500
4	2016-05-06 13:56:00	150	FU-KOSPI-C-20150201-100	200

```
(4 rows)
```

Bingo, our 4 rows!

Features

The Sink connector uses Cassandra's [JSON](#) insert functionality. The SinkRecord from Kafka Connect is converted to JSON and feed into the prepared statements for inserting into Cassandra.

See Cassandra's [documentation](#) for type mapping.

Kafka Connect Query Language

Kafka Connect Query Language found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The Cassandra Sink supports the following:

```
INSERT INTO <target table> SELECT <fields> FROM <source topic>
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA
INSERT INTO tableA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to tableB
INSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB
```

Error Policies

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers..

Retry

Any error on write to the target database causes the RetryIterable exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.cassandra.sink.max.retries` and the `connect.cassandra.sink.retry.interval`.

Topic Routing

The Sink supports topic routing that allows mapping the messages from topics to a specific table. For example map a topic called “bloomberg_prices” to a table called “prices”. This mapping is set in the `connect.cassandra.sink.kcql` option.

Field Selection

The Sink supports selecting fields from the Source topic or selecting all fields and mapping of these fields to columns in the target table. For example, map a field called “qty” in a topic to a column called “quantity” in the target table.

All fields can be selected by using “*” in the field part of `connect.cassandra.sink.kcql`.

Leaving the column name empty means trying to map to a column in the target table with the same name as the field in the source topic.

Legacy topics (plain text payload with a json string)

We have found some of the clients have already an infrastructure where they publish pure json on the topic and obviously the jump to follow the best practice and use schema registry is quite an ask. So we offer support for them as well.

This time we need to start the connect with a different set of settings.

```
#create a new configuration for connect
cp etc/schema-registry/connect-avro-distributed.properties etc/schema-registry/
↪connect-distributed-json.properties
vi etc/schema-registry/connect-distributed-json.properties
```

Replace the following 4 entries in the config

```
key.converter=io.confluent.connect.avro.AvroConverter
key.converter.schema.registry.url=http://localhost:8081
value.converter=io.confluent.connect.avro.AvroConverter
value.converter.schema.registry.url=http://localhost:8081
```

with the following

```
key.converter=org.apache.kafka.connect.json.JsonConverter
key.converter.schemas.enable=false
value.converter=org.apache.kafka.connect.json.JsonConverter
value.converter.schemas.enable=false
```

Now let’s restart the connect instance: .. sourcecode:: bash

```
#start a new instance of connect $CONFLUENT_HOME/bin/connect-distributed etc/schema-
registry/connect-distributed-json.properties
```

Configurations

Configurations common to both Sink and Source are:

```
connect.cassandra.contact.points
```

Contact points (hosts) in Cassandra cluster.

- Data type: string
- Optional : no

`connect.cassandra.key.space`

Key space the tables to write belong to.

- Data type: string
- Optional : no

`connect.cassandra.port`

Port for the native Java driver.

- Data type: int
- Optional : yes
- Default : 9042

`connect.cassandra.username`

Username to connect to Cassandra with if `connect.cassandra.authentication.mode` is set to *username_password*.

- Data type: string
- Optional : yes

`connect.cassandra.password`

Password to connect to Cassandra with if `connect.cassandra.authentication.mode` is set to *username_password*.

- Data type: string
- Optional : yes

`connect.cassandra.ssl.enabled`

Enables SSL communication against SSL enable Cassandra cluster.

- Data type: boolean
- Optional : yes
- Default : false

`connect.cassandra.trust.store.password`

Password for truststore.

- Data type: string
- Optional : yes

`connect.cassandra.key.store.path`

Path to truststore.

- Data type: string
- Optional : yes

`connect.cassandra.key.store.password`

Password for key store.

- Data type: string

- Optional : yes

`connect.cassandra.ssl.client.cert.auth`

Path to keystore.

- Data type: string
- Optional : yes

`connect.cassandra.sink.kcql`

Kafka connect query language expression. Allows for expressive topic to table routing, field selection and renaming.

Examples:

```
INSERT INTO TABLE1 SELECT * FROM TOPIC1;INSERT INTO TABLE2 SELECT field1, field2, ↵
↵field3 as renamedField FROM TOPIC2
```

- Data Type: string
- Optional : no

`connect.cassandra.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.cassandra.sink.max.retries` option. The `connect.cassandra.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: high
- Default: throw

`connect.cassandra.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.cassandra.sink.error.policy` is set to **retry**.

- Type: string
- Importance: high
- Default: 10

`connect.cassandra.sink.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.cassandra.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: medium
- Default : 60000 (1 minute)

Example

```

name=cassandra-sink-orders
connector.class=com.datamountaineer.streamreactor.connect.cassandra.sink.
↳CassandraSinkConnector
tasks.max=1
topics=orders-topic
connect.cassandra.sink.kcql = INSERT INTO TABLE1 SELECT * FROM TOPIC1;INSERT INTO_
↳TABLE2 SELECT field1,
field2, field3 as renamedField FROM TOPIC2
connect.cassandra.contact.points=localhost
connect.cassandra.port=9042
connect.cassandra.key.space=demo
connect.cassandra.contact.points=localhost
connect.cassandra.username=cassandra
connect.cassandra.password=cassandra

```

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

For the Sink connector, if columns are added to the target Cassandra table and not present in the Source topic they will be set to null by Cassandra's JSON insert functionality. Columns which are omitted from the JSON value map are treated as a null insert (which results in an existing value being deleted, if one is present), if a record with the same key is inserted again.

Future releases will support auto creation of tables and adding columns on changes to the topic schema.

Deployment Guidelines

TODO

Troubleshooting

TODO

Kafka Connect CoAP Sink

A Connector and Sink to stream messages from Kafka to a CoAP server.

The Sink supports:

1. DTLS secure clients.
2. *The KCQL routing querying* - Topic to measure mapping and Field selection.
3. Schema registry support for Connect/Avro with a schema.
4. Schema registry support for Connect and no schema (schema set to Schema.String)
5. JSON payload support, no Schema Registry.
6. Error policies.

7. Schema.Struct and payload Struct, Schema.String and Json payload and Json payload with no schema.

The payload of the CoAP request sent to the CoAP server is sent as json.

Prerequisites

- Confluent 3.0.1
- Java 1.8
- Scala 2.11

Setup

Confluent Setup

Follow the instructions [here](#).

CoAP Setup

The connector uses [Californium](#) Java API under the hood. Copper, a [FireFox](#) browser addon is available so you can browse the server and resources.

We will use a simple CoAP test server we have developed for testing. Download the CoAP test server from our github release page and start the server in a new terminal tab.

```
mkdir coap_server
cd coap_server
wget https://github.com/datamountaineer/coap-test-server/releases/download/v1.0/start-
server.sh
chmod +x start-server.sh
./start-server.sh
```

You will see the server start listening on port 5864 for secure DTLS connections and on port 5633 for insecure connections.

```
m.DTLSConnector$Worker.java:-1) run() in thread DTLS-Receiver-0.0.0.0/0.0.0.0:5634 at
(2017-01-10 15:41:08)
1 INFO [CoapEndpoint]: Starting endpoint at localhost/127.0.0.1:5633 - (org.eclipse.
californium.core.network.CoapEndpoint.java:192) start() in thread main at (2017-01-
10 15:41:08)
1 CONFIG [UDPConnector]: UDPConnector starts up 1 sender threads and 1 receiver
threads - (org.eclipse.californium.elements.UDPConnector.java:261) start() in
thread main at (2017-01-10 15:41:08)
1 CONFIG [UDPConnector]: UDPConnector listening on /127.0.0.1:5633, recv buf = 65507,
send buf = 65507, recv packet size = 2048 - (org.eclipse.californium.elements.
UDPConnector.java:261) start() in thread main at (2017-01-10 15:41:08)
Secure CoAP server powered by Scandium (Sc) is listening on port 5634
UnSecure CoAP server powered by Scandium (Sc) is listening on port 5633
```

Source Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools cli` to post in our distributed properties file for MQTT. If you are using the [docker](#)s you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create coap-source < conf/coap-source.properties

#Connector name=`coap-sink`
name = coap-sink
tasks = 1
connector.class = com.datamountaineer.streamreactor.connect.coap.sink.
↳CoapSinkConnector
connect.coap.uri = coap://localhost:5683
connect.coap.kcql = INSERT INTO unsecure SELECT * FROM coap_topic
topics = coap_topic
#task ids: 0
```

The `coap-source.properties` file defines:

1. The name of the sink.
2. The name number of tasks.
3. The class containing the connector.
4. The uri of the CoAP Server and port to connect to.
5. *The KCQL routing querying..* This specifies the target resources on the CoAP server and the source topic.
6. The topics to source (Required by Connect Framework).

If you switch back to the terminal you started Kafka Connect in you should see the CoAP Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
coap-sink
```

INFO

[illegible]

Test Records

Now we need to put some records it to the `coap_topic` topics. We can use the `kafka-avro-console-producer` to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a `firstname` field of type string, a `lastname` field of type string, an `age` field of type int and a `salary` field of type double.

```

${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic coap-topic \
--property value.schema='{ "type": "record", "name": "User", "namespace": "com.
datamountaineer.streamreactor.connect.rethink"
, "fields": [{ "name": "firstName", "type": "string" }, { "name": "lastName", "type": "string" },
, "name": "age", "type": "int" }, { "name": "salary", "type": "double" } ] }'

```

Now the producer is waiting for input. Paste in the following:

```
{ "firstName": "John", "lastName": "Smith", "age": 30, "salary": 4830 }
```

Check for Records in the CoAP server via Copper

Now check the logs of the connector you should see this:

```
[2017-01-10 13:47:36,525] INFO Delivered 1 records for coap-topic. (com.
↳ datamountaineer.streamreactor.connect.coap.sink.CoapSinkTask:47)
```


In Firefox go the following url. If you have not installed Copper do so [here](#).

```
coap://127.0.0.1:5633/insecure
```

Hit the get button and the records will be displayed in the bottom panel.

The screenshot shows the StreamReactor web interface. The top navigation bar includes buttons for GET, POST, PUT, and DELETE, along with a search bar and a menu icon. The main content area displays a CoAP message log with the following entries:

Time	CoAP Message	MID	Token	Options	Payload
14:36:10	NON-GET	199	empty	Uri-Path: insecure, Block2: 0/0/512	
14:36:10	NON-2.05 Content	29712	empty	Content-Format: 0, Max-Age: 2, Block2: 0/0/512	{"firstName":"John","lastName":"Smith","age":30,"salary":4830.0}

The details panel for the selected message shows the following information:

- Header:** Type: NON, Code: 2.05 Content, MID: 29712, Token: empty
- Option:** Content-Format: text/plain, Max-Age: 2, Block2: 0 (512 B/block)
- Payload (64):** {"firstName":"John","lastName":"Smith","age":30,"salary":4830.0}

The right sidebar contains a 'Debug Control' section with a 'Reset' button and a 'Token' field. Below this is a 'Request Options' section with a 'Content-Format' dropdown and a 'Block2 (Res.)' checkbox.

Configurations

```
connect.coap.uri
```

Uri of the CoAP server.

- Data Type : string
- Importance: high
- Optional : no

```
connect.coap.kcql
```

The KCQL statement to select and route resources to topics.

- Data Type : string
- Importance: high
- Optional : no

```
connect.coap.bind.port
```

The port the DTLS connector will bind to on the Connector host.

- Data Type : int
- Importance: medium

- Optional : yes
- Default : 0

`connect.coap.bind.host`

The hostname the DTLS connector will bind to on the Connector host.

- Data Type : string
- Importance: medium
- Optional : yes
- Default : localhost

`connect.coap.keystore.pass`

The password of the key store

- Data Type : string
- Importance: medium
- Optional : yes
- Default : rootPass

`connect.coap.keystore.path`

The path to the keystore.

- Data Type : string
- Importance: medium
- Optional : yes
- Default :

`connect.coap.truststore.pass`

The password of the trust store

- Data Type : string
- Importance: medium
- Optional : yes
- Default : rootPass

`connect.coap.truststore.path`

The path to the truststore.

- Data Type : string
- Importance: medium
- Optional : yes
- Default :

`connect.coap.certs`

The certificates to load from the trust store.

- Data Type : list
- Importance: medium

- Optional : yes
- Default :

`connect.coap.cert.chain.key`

The key to use to get the certificate chain.

- Data Type : string
- Importance: medium
- Optional : yes
- Default : client

`connect.coap.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.coap.max.retries` option. The `connect.coap.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: medium
- Optional: yes
- Default: RETRY

`connect.coap.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.coap.error.policy` is set to `retry`.

- Type: string
- Importance: high
- Optional: yes
- Default: 10

`connect.coap.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.coap.error.policy` set to **RETRY**.

- Type: int
- Importance: medium
- Optional: yes
- Default : 60000 (1 minute)

Kafka Connect Druid

WORK IN PROGRESS NOT COMPLETE!

Kafka Connect Elastic

A Connector and Sink to write events from Kafka to Elastic Search using [Elastic4s](#) client. The connector converts the value from the Kafka Connect SinkRecords to Json and uses Elastic4s's JSON insert functionality to index.

The Sink creates an Index and Type corresponding to the topic name and uses the JSON insert functionality from Elastic4s.

The Sink supports:

1. Auto index creation at start up.
2. *The KCQL routing querying* - Topic to index mapping and Field selection.
3. Auto mapping of the Kafka topic schema to the index.

Prerequisites

- Confluent 3.0.1
- Elastic Search 2.2
- Java 1.8
- Scala 2.11

Setup

Confluent Setup

Follow the instructions [here](#).

Elastic Setup

Download and start Elastic search.

```
curl -L -O https://download.elastic.co/elasticsearch/release/org/elasticsearch/
↪distribution/tar/elasticsearch/2.2.0/elasticsearch-2.2.0.tar.gz
tar -xvf elasticsearch-2.2.0.tar.gz
cd elasticsearch-2.2.0/bin
./elasticsearch --cluster.name elasticsearch
```

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.


```
/ / _ / / / _ ( _ ) / _ / / _ _ / / / / / , <
/ _ _ / _ / \ , _ / _ _ / \ _ / _ / \ _ / / _ _ / _ / / _ / _ | _ |
```

by Andrew Stevenson

```
(com.datamountaineer.streamreactor.connect.elastic.ElasticSinkTask:33)
```

```
[2016-05-08 20:56:52,327] INFO [Hebe] loaded [], sites [] (org.elasticsearch.
↳plugins:149)
[2016-05-08 20:56:52,765] INFO Initialising Elastic Json writer (com.datamountaineer.
↳streamreactor.connect.elastic.ElasticJsonWriter:31)
[2016-05-08 20:56:52,777] INFO Assigned List(test_table) topics. (com.datamountaineer.
↳streamreactor.connect.elastic.ElasticJsonWriter:33)
[2016-05-08 20:56:52,836] INFO Sink task org.apache.kafka.connect.runtime.
↳WorkerSinkTask@69b6b39 finished initialization and start (org.apache.kafka.connect.
↳runtime.WorkerSinkTask:155)
```

Test Records

Now we need to put some records it to the test_table topics. We can use the kafka-avro-console-producer to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a id field of type int and a random_field of type string.

```
/${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic TOPIC1 \
--property value.schema='{ "type": "record", "name": "myrecord", "fields": [{ "name": "id",
↳ "type": "int" },
{ "name": "random_field", "type": "string" } ] }'
```

Now the producer is waiting for input. Paste in the following:

```
{ "id": 999, "random_field": "foo" }
{ "id": 888, "random_field": "bar" }
```

Check for records in Elastic Search

Now if we check the logs of the connector we should see 2 records being inserted to Elastic Search:

```
[2016-05-08 21:02:52,095] INFO Flushing Elastic Sink (com.datamountaineer.
↳streamreactor.connect.elastic.ElasticSinkTask:73)
[2016-05-08 21:03:52,097] INFO No records received. (com.datamountaineer.
↳streamreactor.connect.elastic.ElasticJsonWriter:63)
[2016-05-08 21:03:52,097] INFO org.apache.kafka.connect.runtime.
↳WorkerSinkTask@69b6b39 Committing offsets (org.apache.kafka.connect.runtime.
↳WorkerSinkTask:187)
[2016-05-08 21:03:52,097] INFO Flushing Elastic Sink (com.datamountaineer.
↳streamreactor.connect.elastic.ElasticSinkTask:73)
[2016-05-08 21:04:20,613] INFO Elastic write successful for 2 records! (com.
↳datamountaineer.streamreactor.connect.elastic.ElasticJsonWriter:77)
```

If we query Elastic Search for id 999:

```
curl -XGET 'http://localhost:9200/INDEX_1/_search?q=id:999'

{
  "took": 45,
  "timed_out": false,
  "_shards": {
    "total": 5,
    "successful": 5,
    "failed": 0
  },
  "hits": {
    "total": 1,
    "max_score": 1.2231436,
    "hits": [{
      "_index": "test_table",
      "_type": "test_table",
      "_id": "AVMY4eZXFguf2uMZyxjU",
      "_score": 1.2231436,
      "_source": {
        "id": 999,
        "random_field": "foo"
      }
    }]
  }
}
```

Features

1. Auto index creation at start up.
2. Topic to index mapping.
3. Auto mapping of the Kafka topic schema to the index.
4. Field selection

Kafka Connect Query Language

Kafka **C**onnect **Q**uery **L**anguage found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The Elastic Sink supports the following:

```
INSERT INTO <index> SELECT <fields> FROM <source topic>
```

Example:

```
#Insert mode, select all fields from topicA and write to indexA
INSERT INTO indexA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to indexB
INSERT INTO indexB SELECT x AS a, y AS b and z AS c FROM topicB PK y
```

This is set in the `connect.elastic.sink.kcql` option.

Auto Index Creation

The Sink will automatically create missing indexes at startup. The Sink use elastic4s, more details can be found [here](#)

Configurations

`connect.elastic.url`

Url of the Elastic cluster.

- Data Type : string
- Importance: high
- Optional : no

`connect.elastic.port`

Port of the Elastic cluster.

- Data Type : string
- Importance: high
- Optional : no

`connect.elastic.sink.kcql`

Kafka connect query language expression. Allows for expressive table to topic routing, field selection and renaming.

Examples:

```
INSERT INTO INDEX_1 SELECT field1, field2 FROM TOPIC1
```

- Data type : string
- Importance: high
- Optional : no

Example

```
name=elastic-sink
connector.class=com.datamountaineer.streamreactor.connect.elastic.ElasticSinkConnector
connect.elastic.url=localhost:9300
connect.elastic.cluster.name=elasticsearch
tasks.max=1
topics=test_table
connect.elastic.sink.kcql=INSERT INTO INDEX_1 SELECT field1, field2 FROM TOPIC1
```

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

Elastic Search is very flexible about what is inserted. All documents in Elasticsearch are stored in an index. We do not need to tell Elasticsearch in advance what an index will look like (eg what fields it will contain) as Elasticsearch

will adapt the index dynamically as more documents are added, but we must at least create the index first. The Sink connector automatically creates the index at start up if it doesn't exist.

The Elastic Search Sink will automatically index if new fields are added to the Source topic, if fields are removed the Kafka Connect framework will return the default value for this field, dependent of the compatibility settings of the Schema registry.

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect HazelCast

A Connector and Sink to write events from Kafka to HazelCast. The connector takes the value from the Kafka Connect SinkRecords and inserts a new entry to a HazelCast reliable topic. The Sink only supports writing to reliable topics.

The Sink supports:

1. *The KCQL routing querying* - Kafka topic payload field selection is supported, allowing you to have choose selection of fields or all fields written to Hazelcast.
2. Topic to table routing via KCQL.
3. Error policies for handling failures.
4. Storing as JSON or Avro in Hazelcast via KCQL.

Prerequisites

- Confluent 3.0.1
- Hazelcast 3.6.4
- Java 1.8
- Scala 2.11

Setup

Confluent Setup

Follow the instructions [here](#).

HazelCast Setup

Download and install HazelCast from [here](#)

When you download and extract the Hazelcast ZIP or TAR.GZ package, you will see 3 scripts under the `/bin` folder which provide basic functionality for member and cluster management.

The following are the names and descriptions of each script:

- `start.sh` - Starts a Hazelcast member with default configuration in the working directory.
- `stop.sh` - Stops the Hazelcast member that was started in the current working directory.

Start HazelCast:

```
bin/start.sh

INFO: [10.128.137.102]:5701 [dev] [3.6.4] Address[10.128.137.102]:5701 is STARTING
Aug 16, 2016 2:43:04 PM com.hazelcast.nio.tcp.nonblocking.NonBlockingIOThreadingModel
INFO: [10.128.137.102]:5701 [dev] [3.6.4] TcpIpConnectionManager configured with Non-
↳Blocking IO-threading model: 3 input threads and 3 output threads
Aug 16, 2016 2:43:07 PM com.hazelcast.cluster.impl.MulticastJoiner
INFO: [10.128.137.102]:5701 [dev] [3.6.4]

Members [1] {
  Member [10.128.137.102]:5701 this
}

Aug 16, 2016 2:43:07 PM com.hazelcast.core.LifecycleService
INFO: [10.128.137.102]:5701 [dev] [3.6.4] Address[10.128.137.102]:5701 is STARTED
```

This will start Hazelcast with a default group called *dev* and password *dev-pass*

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools cli` to post in our distributed properties file for HazelCast. If you are using the *dockers* you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create hazelcast-sink < conf/hazelcast-sink.properties

#Connector name=`hazelcast-sink`
name=hazelcast-sink
connector.class=com.datamountaineer.streamreactor.connect.hazelcast.sink.
↳HazelCastSinkConnector
max.tasks=1
```

```
topics = hazelcast-topic
connect.hazelcast.sink.cluster.members=localhost
connect.hazelcast.sink.group.name=dev
connect.hazelcast.sink.group.password=dev-pass
connect.hazelcast.sink.kcql=INSERT INTO sink-test SELECT * FROM hazelcast-topic_
↳WITHFORMAT JSON BATCH 100
#task ids: 0
```

The `hazelcast-sink.properties` configuration defines:

1. The name of the sink.
2. The Sink class.
3. The max number of tasks the connector is allowed to created.
4. The topics to read from (Required by framework)
5. The name of the HazelCast host to connect to.
6. The name of the group to connect to.
7. The password for the group.
8. *The KCQL routing querying.*

If you switch back to the terminal you started the Connector in you should see the Hazelcast Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
hazelcast-sink
```

```
(org.apache.kafka.clients.consumer.ConsumerConfig:178)
[2016-08-20 16:45:39,518] INFO Kafka version : 0.10.0.0 (org.apache.kafka.common.
↳utils.AppInfoParser:83)
[2016-08-20 16:45:39,518] INFO Kafka commitId : b8642491e78c5a13 (org.apache.kafka.
↳common.utils.AppInfoParser:84)
[2016-08-20 16:45:39,520] INFO Created connector hazelcast-sink (org.apache.kafka.
↳connect.cli.ConnectStandalone:91)
[2016-08-20 16:45:39,520] INFO
```

The diagram illustrates a sequence of 16 steps for a path on a grid. The path is represented by a series of blue and black line segments. The steps are labeled with numbers 1 through 16, and some steps have additional labels like '(-)' or '(<)'.

by Andrew Stevenson

```
(com.datamountaineer.streamreactor.connect.hazelcast.sink.HazelCastSinkTask:41)
[2016-08-20 16:45:39,521] INFO HazelCastSinkConfig values:
  connect.hazelcast.connection.buffer.size = 32
  connect.hazelcast.connection.keep.alive = true
```

```

connect.hazelcast.connection.tcp.no.delay = true
connect.hazelcast.sink.group.password = [hidden]
connect.hazelcast.connection.retries = 2
connect.hazelcast.connection.linger.seconds = 3
connect.hazelcast.sink.retry.interval = 60000
connect.hazelcast.max.retires = 20
connect.hazelcast.sink.batch.size = 1000
connect.hazelcast.connection.reuse.address = true
connect.hazelcast.sink.group.name = dev
connect.hazelcast.sink.cluster.members = [192.168.99.100]
connect.hazelcast.sink.error.policy = THROW
connect.hazelcast.sink.kcql = INSERT INTO sink-test SELECT * FROM hazelcast-topic
↪WITHFORMAT JSON BATCH 100
    connect.hazelcast.connection.timeout = 5000
(com.datamountaineer.streamreactor.connect.hazelcast.config.HazelCastSinkConfig:178)
Aug 20, 2016 4:45:39 PM com.hazelcast.core.LifecycleService
INFO: HazelcastClient[dev-kafka-connect-05e64989-41d9-433e-ad21-b54894486384] [3.6.4]
↪is STARTING
Aug 20, 2016 4:45:39 PM com.hazelcast.core.LifecycleService
INFO: HazelcastClient[dev-kafka-connect-05e64989-41d9-433e-ad21-b54894486384] [3.6.4]
↪is STARTED
Aug 20, 2016 4:45:39 PM com.hazelcast.client.spi.impl.ClientMembershipListener
INFO:
Members [1] {
    Member [172.17.0.2]:5701
}
Aug 20, 2016 4:45:39 PM com.hazelcast.core.LifecycleService
INFO: HazelcastClient[dev-kafka-connect-05e64989-41d9-433e-ad21-b54894486384] [3.6.4]
↪is CLIENT_CONNECTED

```

Test Records

Now we need to put some records it to the test_table topics. We can use the kafka-avro-console-producer to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a firstname field of type string a lastname field of type string, an age field of type int and a salary field of type double.

```

${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
  --broker-list localhost:9092 --topic hazelcast-topic \
  --property value.schema='{ "type": "record", "name": "User", "namespace": "com."
↪datamountaineer.streamreactor.connect.HazelCast"
  , "fields": [{ "name": "firstName", "type": "string"}, { "name": "lastName", "type": "string"},
↪{ "name": "age", "type": "int"}, { "name": "salary", "type": "double"} ] }'

```

Now the producer is waiting for input. Paste in the following:

```
{ "firstName": "John", "lastName": "Smith", "age": 30, "salary": 4830 }
```

Check for records in HazelCast

Now check the logs of the connector you should see this:

```
[2016-08-20 16:53:58,608] INFO Received 1 records. (com.datamountaineer.streamreactor.
↪connect.hazelcast.sink.HazelCastWriter:62)
[2016-08-20 16:53:58,644] INFO Written 1 (com.datamountaineer.streamreactor.connect.
↪hazelcast.sink.HazelCastWriter:71)
```

Now stop the connector.

Features

Kafka Connect Query Language

Kafka Connect Query Language found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The HazelCast Sink supports the following:

```
INSERT INTO <reliable topic> SELECT <fields> FROM <source topic> WITHFORMAT
↪<JSON|AVRO> STOREAS <RELIABLE_TOPIC|RING_BUFFER> BATCH BATCH_SIZE
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA
INSERT INTO tableA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to tableB, store as_
↪serialized avro encoded byte arrays, write in batches of 100
INSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB WITHFORMAT avro_
↪STOREAS RING_BUFFER BATCH 100
```

This is set in the `connect.hazelcast.sink.kcql` option.

Error Policies

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers..

Retry

Any error on write to the target database causes the `RetryIterable` exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it

will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.hazelcast.sink.max.retries` and the `connect.hazelcast.sink.retry.interval`.

With Format

Hazelcast requires that data stored in collections and topics is serializable. The Sink offers two modes to store data.

Avro In this mode the Sink converts the SinkRecords from Kafka to Avro encoded byte arrays. *Json* In this mode the Sink converts the SinkRecords from Kafka to Json strings and stores the resulting bytes.

This behaviour is controlled by the KCQL statement in the `connect.hazelcast.sink.kcql` option. The default is JSON.

Configurations

`connect.hazelcast.sink.kcql`

KCQL expression describing field selection and routes.

- Data type : string
- Importance: high
- Optional : no

`connect.hazelcast.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.hazelcast.sink.max.retries` option. The `connect.hazelcast.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: high
- Optional: yes
- Default: throw

`connect.hazelcast.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.hazelcast.sink.error.policy` is set to `retry`.

- Type: string
- Importance: medium
- Optional: yes
- Default: 10

```
connect.hazelcast.sink.retry.interval
```

The interval, in milliseconds between retries if the Sink is using `connect.hazelcast.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: medium
- Optional: yes
- Default : 60000 (1 minute)

```
connect.hazelcast.sink.batch.size
```

Specifies how many records to insert together at one time. If the connect framework provides less records when it is calling the Sink it won't wait to fulfill this value but rather execute it.

- Type : int
- Importance : medium
- Optional: yes
- Defaults : 1000

```
connect.hazelcast.sink.cluster.members
```

Address List is the initial list of cluster addresses to which the client will connect. The client uses this list to find an alive node. Although it may be enough to give only one address of a node in the cluster (since all nodes communicate with each other), it is recommended that you give the addresses for all the nodes.

- Data type : string
- Importance : high
- Optional: no
- Default: localhost

```
connect.hazelcast.sink.group.name
```

The group name of the connector in the target Hazelcast cluster.

- Data type : string
- Importance : high
- Optional: no
- Default: dev

```
connect.hazelcast.sink.group.password
```

The password for the group name.

- Data type : string
- Importance : high
- Optional : yes
- Default : dev-pass

```
connect.hazelcast.connection.timeout
```

Connection timeout is the timeout value in milliseconds for nodes to accept client connection requests.

- Data type : int

- Importance : low
- Optional : yes
- Default : 5000

`connect.hazelcast.connection.retries`

Number of times a client will retry the connection at startup.

- Data type : int
- Importance : low
- Optional : yes
- Default : 2

`connect.hazelcast.connection.keep.alive`

Enables/disables the SO_KEEPALIVE socket option. The default value is true.

- Data type : boolean
- Importance : low
- Optional : yes
- Default : true

`connect.hazelcast.connection.tcp.no.delay`

Enables/disables the SO_REUSEADDR socket option. The default value is true.

- Data type : boolean
- Importance : low
- Optional : yes
- Default : true

`connect.hazelcast.connection.linger.seconds`

Enables/disables SO_LINGER with the specified linger time in seconds. The default value is 3.

- Data type : int
- Importance : low
- Optional : yes
- Default : 3

`connect.hazelcast.connection.buffer.size`

Sets the SO_SNDBUF and SO_RCVBUF options to the specified value in KB for this Socket. The default value is 32.

- Data type : int
- Importance : low
- Optional : yes
- Default : 32

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

The Sink serializes either an Avro or Json representation of the Sink record to the target reliable topic in Hazelcast. Hazelcast is agnostic to the schema.

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect HBase

A Connector and Sink to write events from Kafka to HBase. The connector takes the value from the Kafka Connect SinkRecords and inserts a new entry to HBase.

The Sink supports:

1. *The [KCQL routing querying](#)* - Kafka topic payload field selection is supported, allowing you to select fields written to HBase.
2. Topic to table routing via KCQL.
3. RowKey selection - Selection of fields to use as the row key, if none specified the topic name, partition and offset are used via KCQL.
4. Error policies.

Prerequisites

- Confluent 3.0.1
- HBase 1.2.0
- Java 1.8
- Scala 2.11

Setup

HBase Setup

Download and extract HBase:

```
wget https://www.apache.org/dist/hbase/1.2.1/hbase-1.2.1-bin.tar.gz
tar -xvf hbase-1.2.1-bin.tar.gz -C hbase
```

Edit `conf/hbase-site.xml` and add the following content:

```
<?xml version="1.0"?>
<?xml-stylesheet type="text/xsl" href="configuration.xsl"?>
<configuration>
  <property>
    <name>hbase.cluster.distributed</name>
    <value>true</value>
  </property>
  <property>
    <name>hbase.rootdir</name>
    <value>file:///tmp/hbase</value>
  </property>
  <property>
    <name>hbase.zookeeper.property.dataDir</name>
    <value>/tmp/zookeeper</value>
  </property>
</configuration>
```

The `hbase.cluster.distributed` is required since when you start hbase it will try and start its own Zookeeper, but in this case we want to use Confluents.

Now start HBase and check the logs to ensure it's up:

```
bin/start-hbase.sh
```

Confluent Setup

Follow the instructions [here](#).

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

HBase Table

The Sink expects a precreated table in HBase. In the HBase shell create the test table, go to your HBase install location.

```
bin/hbase shell
hbase(main):001:0> create 'person_hbase',{NAME=>'d', VERSIONS=>1}

hbase(main):001:0> list
person
1 row(s) in 0.9530 seconds

hbase(main):002:0> describe 'person'
DESCRIPTION
'person', {NAME => 'd', BLOOMFILTER => 'ROW', VERSIONS => '1', IN_MEMORY => 'false',
↳KEEP_DELETED_CELLS => 'false', DATA_BLOCK_ENCODING => 'NONE', TTL => 'FOREVER', COMPRESSION => 'NONE', MIN_VERSIONS => '0',
↳BLOCKCACHE => 'true', BLOCKSIZE => '65536', REPLICATION
```

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

```
bin/start-connect.sh
```

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create hbase-sink < conf/hbase-sink.properties

#Connector name=`hbase-sink`
name=person-hbase-test
connector.class=com.datamountaineer.streamreactor.connect.hbase.HbaseSinkConnector
tasks.max=1
topics=hbase-topic
connect.hbase.sink.column.family=d
connect.hbase.sink.kcql=INSERT INTO person_hbase SELECT * FROM hbase-topic
#task ids: 0
```

1. The name of the sink.
2. The Sink class.
3. The max number of tasks the connector is allowed to created. Should not be greater than the number of partitions in the Source topics otherwise tasks will be idle.
4. The Source kafka topics to take events from.
5. The HBase column family to write to.
6. *The KCQL routing querying.*

```
#check for running connectors with the CLI
bin/cli.sh ps
hbase-sink
```

INFO

```

/_____/\\_,\\\\\\\\\\_,/_/ /_/\\_\\\\\\\\\\,\\_/ /_/\\_\\_\\_,/_/_/ /_/\\_\\_\\_/_/
/ / / / _ )____ / ____/ ( ) ____ / / ____
/ / / / ____ / ____ / ____ \\ / ____ \\ / / /
/ ____ / / / / / ( ) ____ / / / / / / / ,<
/ / / / ____ / ____ / ____ / / / / / / / | |

```

By Stefan Bocutiu (com.datamountaineer.streamreactor.connect.hbase.HbaseSinkTask:44)

Test Records

Now we need to put some records it to the test_table topics. We can use the kafka-avro-console-producer to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a firstname field of type string, a lastname field of type string, an age field of type int and a salary field of type double.

```

${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
  --broker-list localhost:9092 --topic hbase-topic \
  --property value.schema='{ "type": "record", "name": "User", "namespace": "com.
  ↪ datamountaineer.streamreactor.connect.hbase"
    "fields": [{ "name": "firstName", "type": "string" }, { "name": "lastName", "type": "string" }, {
  ↪ "name": "age", "type": "int" },
    { "name": "salary", "type": "double" } ] }'

```

Now the producer is waiting for input. Paste in the following:

```

{"firstName": "John", "lastName": "Smith", "age": 30, "salary": 4830}
{"firstName": "Anna", "lastName": "Jones", "age": 28, "salary": 5430}

```

Check for records in HBase

Now check the logs of the connector you should see this

```

INFO Sink task org.apache.kafka.connect.runtime.WorkerSinkTask@48ffb4dc finished_
  ↪ initialization and start (org.apache.kafka.connect.runtime.WorkerSinkTask:155)
INFO Writing 2 rows to Hbase... (com.datamountaineer.streamreactor.connect.hbase.
  ↪ writers.HbaseWriter:83)

```

In HBase:

```

hbase(main):004:0* scan 'person_hbase'
ROW                                     COLUMN+CELL
Anna\x0AJones                         column=d:age,
  ↪ timestamp=1463056888641, value=\x00\x00\x00\x1C
Anna\x0AJones                         column=d:firstName,
  ↪ timestamp=1463056888641, value=Anna
Anna\x0AJones                         column=d:income,
  ↪ timestamp=1463056888641, value=@\xB56\x00\x00\x00\x00
Anna\x0AJones                         column=d:lastName,
  ↪ timestamp=1463056888641, value=Jones
John\x0ASmith                         column=d:age,
  ↪ timestamp=1463056693877, value=\x00\x00\x00\x1E
John\x0ASmith                         column=d:firstName,
  ↪ timestamp=1463056693877, value=John
John\x0ASmith                         column=d:income,
  ↪ timestamp=1463056693877, value=@\xB2\xDE\x00\x00\x00\x00

```

```
John\0ASmith                                column=d:lastName,
↪timestamp=1463056693877, value=Smith
2 row(s) in 0.0260 seconds
```

Now stop the connector.

Features

The HBase Sink writes records from Kafka to HBase.

The Sink supports:

1. Field selection - Kafka topic payload field selection is supported, allowing you to select fields written to HBase.
2. Topic to table routing.
3. RowKey selection - Selection of fields to use as the row key, if none specified the topic name, partition and offset are used.
4. Error policies.

Kafka Connect Query Language

Kafka Connect Query Language found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The HBase Sink supports the following:

```
INSERT INTO <table> SELECT <fields> FROM <source topic> <PK> primary_key_cols
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA and use the default
↪rowkey (topic name, partition, offset)
INSERT INTO tableA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to tableB, use field y
↪from the topic as the row key
INSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB PK y
```

This is set in the `connect.hbase.sink.kcql` option.

Error Policies

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers.

Retry

Any error on write to the target database causes the `RetryIterable` exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.hbase.sink.max.retries` and the `connect.hbase.sink.retry.interval`.

Configurations

`connect.hbase.sink.column.family`

The hbase column family.

- Type: string
- Importance: high
- Optional: no

`connect.hbase.sink.kcql`

Kafka connect query language expression. Allows for expressive topic to table routing, field selection and renaming. Fields to be used as the row key can be set by specifying the PK. The below example uses field1 and field2 are the row key.

Examples:

```
INSERT INTO TABLE1 SELECT * FROM TOPIC1; INSERT INTO TABLE2 SELECT * FROM TOPIC2 PK_
↪field1, field2
```

If no primary keys are specified the topic name, partition and offset converted to bytes are used as the HBase rowkey.

- Type: string
- Importance: high
- Optional: no

`connect.hbase.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.hbase.sink.max.retries` option. The `connect.hbase.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: medium
- Optional: yes

- Default: RETRY

`connect.hbase.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.hbase.sink.error.policy` is set to `retry`.

- Type: string
- Importance: medium
- Optional: yes
- Default: 10

`connect.hbase.sink.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.hbase.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: medium
- Optional: yes
- Default : 60000 (1 minute)

Example

```
connect.hbase.sink.column.family=d
connect.hbase.sink.kcql=INSERT INTO person_hbase SELECT * FROM TOPIC1
connector.class=com.datamountaineer.streamreactor.connect.hbase.HbaseSinkConnector
tasks.max=1
topics=TOPIC1
name=hbase-test
```

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

The HBase Sink will automatically write and update the HBase table if new fields are added to the Source topic, if fields are removed the Kafka Connect framework will return the default value for this field, dependent of the compatibility settings of the Schema registry. This value will be put into the HBase column family cell based on the `connect.hbase.sink.fields.mappings`.

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect Influx

A Connector and Sink to write events from Kafka to InfluxDB. The connector takes the value from the Kafka Connect SinkRecords and inserts a new entry to InfluxDB.

The Sink supports:

1. *The KCQL routing querying* - Topic to index mapping and Field selection.
2. Auto mapping of the Kafka topic schema to the index.
3. Schema.Struct and payload Struct, Schema.String and Json payload and Json payload with no schema

The Sink supports three Kafka payloads type:

Connect entry with Schema.Struct and payload Struct. If you follow the best practice while producing the events, each message should carry its schema information. Best option is to send Avro. Your connect configurations should be set to `value.converter=io.confluent.connect.avro.AvroConverter`. You can find an example [here](#). To see how easy is to have your producer serialize to Avro have a look at [this](#). This requires SchemaRegistry which is open source thanks to Confluent! Alternatively you can send Json + Schema. In this case your connect configuration should read `value.converter=org.apache.kafka.connect.json.JsonConverter`. The difference would be to point your serialization to `org.apache.kafka.connect.json.JsonSerializer`. This doesn't require the SchemaRegistry.

Connect entry with Schema.String and payload json String. Sometimes the producer would find it easier, despite sending Avro to produce a GenericRecord, to just send a message with Schema.String and the json string.

Connect entry without a schema and the payload json String. There are many existing systems which are publishing json over Kafka and bringing them in line with best practices is quite a challenge. Hence we added the support

Prerequisites

- Confluent 3.0.1
- Java 1.8
- Scala 2.11

Setup

Confluent Setup

Follow the instructions [here](#).

InfluxDB Setup

[Download](#) and start InfluxDB. Users of OS X 10.8 and higher can install InfluxDB using the Homebrew package manager. Once brew is installed, you can install InfluxDB by running:

```
brew update
brew install influxdb
```

Note: InfluxDB starts an Admin web server listening on port 8083 by default. For this quickstart this will collide with Kafka Connects default port of 8083. Since we are running on a single node we will need to edit the InfluxDB config.


```
#create config dir
sudo mkdir /etc/influxdb
#dump the config
influxd config > /etc/influxdb/influxdb.generated.conf
```

Now change the following section to a port 8087 or any other free port.

```
[admin]
enabled = true
bind-address = ":8087"
https-enabled = false
https-certificate = "/etc/ssl/influxdb.pem"
```

Now start InfluxDB.

```
influxd
```

If you are running on a single node start InfluxDB with the new configuration file we generated.

```
influxd -config /etc/influxdb/influxdb.generated.conf
```

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Test data

The Sink expects a database to exist in InfluxDB. Use the InfluxDB CLI to create this:

```
~ influx
Visit https://enterprise.influxdata.com to register for updates, InfluxDB server
↔management, and monitoring.
Connected to http://localhost:8086 version v1.0.2
InfluxDB shell version: v1.0.2
```

```
> CREATE DATABASE mydb
```

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools cli` to post in our distributed properties file for InfluxDB. If you are using the *dockers* you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create influx-sink < conf/influxdb-sink.properties

#Connector name=`influx-sink`
name=influxdb-sink
connector.class=com.datamountaineer.streamreactor.connect.influx.InfluxSinkConnector
tasks.max=1
topics=influx-topic
connect.influx.sink.route.query=INSERT INTO influxMeasure SELECT * FROM influx-topic_
↳ WITHTIMESTAMP sys_time()
connect.influx.connection.url=http://localhost:8086
connect.influx.connection.database=mydb
#task ids: 0
```

The `elastic-sink.properties` file defines:

1. The name of the connector.
2. The class containing the connector.
3. The max number of task allowed for this connector.
4. The Source topic to get records from.
5. *The KCQL routing querying.*
6. The InfluxDB connection URL.
7. The InfluxDB database.

If you switch back to the terminal you started Kafka Connect in you should see the InfluxDB Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
influxdb-sink
```

[illegible]

```

connect.influx.connection.password = [hidden]
connect.influx.connection.url = http://localhost:8086
connect.influx.retry.interval = 60000
connect.influx.sink.route.query = INSERT INTO influxMeasure SELECT * FROM influx-
↳topic WITHTIMESTAMP sys_time()
connect.influx.max.retires = 20
(com.datamountaineer.streamreactor.connect.influx.config.InfluxSinkConfig:178)

```

Test Records

Now we need to put some records it to the test_table topics. We can use the kafka-avro-console-producer to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a company field of type string a address field of type string, an latitude field of type int and a longitude field of type int.

```

${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic influx-topic \
--property value.schema='{ "type": "record", "name": "User", "namespace": "com.
↳datamountaineer.streamreactor.connect.influx", "fields": [{ "name": "company", "type":
↳"string"}, { "name": "address", "type": "string"}, { "name": "latitude", "type": "float"}, {
↳"name": "longitude", "type": "float"}] }'

```

Now the producer is waiting for input. Paste in the following:

```

{"company": "DataMountaineer", "address": "MontainTop", "latitude": -49.817964,
↳"longitude": -141.645812}

```

Check for records in InfluxDB

Now check the logs of the connector you should see this:

```

INFO Setting newly assigned partitions [influx-topic-0] for group connect-influx-sink_
↳(org.apache.kafka.clients.consumer.internals.ConsumerCoordinator:231)
INFO Received 1 record(-s) (com.datamountaineer.streamreactor.connect.influx.
↳InfluxSinkTask:81)
INFO Writing 1 points to the database... (com.datamountaineer.streamreactor.connect.
↳influx.writers.InfluxDbWriter:45)
INFO Records handled (com.datamountaineer.streamreactor.connect.influx.
↳InfluxSinkTask:83)

```

Check in InfluxDB.

```

influx
Visit https://enterprise.influxdata.com to register for updates, InfluxDB server_
↳management, and monitoring.
Connected to http://localhost:8086 version v1.0.2
InfluxDB shell version: v1.0.2
> use mydb;
Using database mydb
> show measurements;
name: measurements
-----
name

```

```
influxMeasure

> select * from influxMeasure;
name: influxMeasure
-----
time                address      async    company      latitude
↳      longitude
1478269679104000000 MontainTop    true     DataMountaineer -49.817962646484375
↳141.64581298828125
```

Features

1. Topic to index mapping. 3. Auto mapping of the Kafka topic schema to the index. 4. Field selection

Kafka Connect Query Language

Kafka **C**onnect **Q**uery **L**anguage found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The Influx Sink supports the following:

```
INSERT INTO <measure> SELECT <fields> FROM <source topic> WITHTIMESTAMP <field_name>
↳|sys_time()
```

Example:

```
#Insert mode, select all fields from topicA and write to indexA
INSERT INTO measureA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to indexB, use field Y
↳as the point measurement
INSERT INTO measureB SELECT x AS a, y AS b and z AS c FROM topicB WITHTIMESTAMP y

#Insert mode, select 3 fields and rename from topicB and write to indexB, use field Y
↳as the current system time for
#Point measurement
INSERT INTO measureB SELECT x AS a, y AS b and z AS c FROM topicB WITHTIMESTAMP sys_
↳time()
```

This is set in the `connect.influx.sink.kcql` option.

Error Policies

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers.

Retry

Any error on write to the target database causes the `RetryIterable` exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.influx.sink.max.retries` and the `connect.influx.sink.retry.interval`.

Configurations

`connect.influx.sink.kcql`

Kafka connect query language expression. Allows for expressive topic to table routing, field selection and renaming. For InfluxDB it allows either setting a default or selecting a field from the topic as the Point measurement.

- Data type : string
- Importance: high
- Optional : no

`connect.influx.connection.url`

The InfluxDB database url.

- Data type : string
- Importance: high
- Optional : no

`connect.influx.connection.database`

The InfluxDB database.

- Data type : string
- Importance: high
- Optional : no

`connect.influx.connection.username`

The InfluxDB username.

- Data type : string
- Importance: high
- Optional : yes

`connect.influx.connection.password`

The InfluxDB password.

- Data type : string
- Importance: high

- Optional : yes

`connect.influx.retention.policy`

Determines how long InfluxDB keeps the data - the options for specifying the duration of the retention policy are listed below. Note that the minimum retention period is one hour. DURATION determines how long InfluxDB keeps the data - the options for specifying the duration of the retention policy are listed below. Note that the minimum retention period is one hour.

m minutes h hours d days w weeks INF infinite

Default retention is *autogen* from 1.0 onwards or *default* for any previous version

- Data type : string
- Importance: medium
- Optional : yes

`connect.influx.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.influx.sink.max.retries` option. The `connect.influx.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: medium
- Optional: yes
- Default: RETRY

`connect.influx.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.influx.sink.error.policy` is set to **retry**.

- Type: string
- Importance: medium
- Optional: yes
- Default: 10

`connect.influx.sink.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.influx.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: high
- Optional: no
- Default : 60000 (1 minute)

Example

```
name=elastic-sink
connector.class=com.datamountaineer.streamreactor.connect.elastic.ElasticSinkConnector
connect.elastic.url=localhost:9300
connect.elastic.cluster.name=elasticsearch
tasks.max=1
topics=test_table
connect.elastic.sink.kcql=INSERT INTO INDEX_1 SELECT field1, field2 FROM TOPIC1
```

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

Elastic Search is very flexible about what is inserted. All documents in Elasticsearch are stored in an index. We do not need to tell Elasticsearch in advance what an index will look like (eg what fields it will contain) as Elasticsearch will adapt the index dynamically as more documents are added, but we must at least create the index first. The Sink connector automatically creates the index at start up if it doesn't exist.

The Elastic Search Sink will automatically index if new fields are added to the Source topic, if fields are removed the Kafka Connect framework will return the default value for this field, dependent of the compatibility settings of the Schema registry.

Deployment Guidelines

TODO

Troubleshooting

TODO

Kafka Connect JMS

The JMS Sink connector allows you to extract entries from a Kafka topic with the CQL driver and pass them to a JMS topic/queue. The connector allows you to specify the payload type sent to the JMS target:

1. JSON
2. AVRO
3. MAP
4. OBJECT

The Sink supports:

1. *The KCQL routing querying*. Kafka topic payload field selection is supported, allowing you to select fields written to the queue or topic in JMS.
2. Topic to topic routing via KCQL.
3. Payload format selection via KCQL.

4. Error policies for handling failures.

Prerequisites

- Confluent 3.0.0
- Java 1.8
- Scala 2.11
- A JMS framework (ActiveMQ for example)

Setup

Before we can do anything, including the QuickStart we need to install the Confluent platform. For ActiveMQ follow <http://activemq.apache.org/getting-started.html> for the instruction of setting it up.

Confluent Setup

Follow the instructions [here](#).

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools` cli to post in our distributed properties file for JMS. If you are using the [dockers](#) you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create jms-sink < conf/jms-sink.properties
#Connector name=`jms-sink`
name=jms-sink
connector.class=com.datamountaineer.streamreactor.connect.jms.sink.JMSSinkConnector
tasks.max=1
topics=jms-topic
connect.jms.sink.url=tcp://somehost:61616
connect.jms.sink.connection.factory=org.apache.activemq.ActiveMQConnectionFactory
connect.jms.sink.sink.kcql=INSERT INTO topic_1 SELECT * FROM jms-topic
```



```
connect.jms.sink.message.type=AVRO
connect.jms.error.policy=THROW
connect.jms.sink.export.route.topics=topic_1
#task ids: 0
```

The `jms-sink.properties` file defines:

1. The name of the sink.
2. The Sink class.
3. The max number of tasks the connector is allowed to created.
4. The Kafka topics to take events from.
5. The JMS url.
6. The factory class for the JSM endpoint.
7. *The KCQL routing querying.*
8. The message type storage format.
9. The error policy.
10. The list of target topics (must match the targets set in `connect.jms.sink.sink.kcql`

If you switch back to the terminal you started the Connector in you should see the JMS Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
jms-sink
```

Test Records

Now we need to put some records it to the `test_table` topics. We can use the `kafka-avro-console-producer` to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a `id` field of type `int` and a `random_field` of type `string`.

```
${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic jms_test \
--property value.schema='{ "type": "record", "name": "User", "namespace": "com.
↳ datamountaineer.streamreactor.connect.jms",
"fields": [{ "name": "firstName", "type": "string" }, { "name": "lastName", "type": "string" }, {
↳ "name": "age", "type": "int" }, { "name": "salary", "type": "double" } ] }'
```

Now the producer is waiting for input. Paste in the following:

```
{ "firstName": "John", "lastName": "Smith", "age": 30, "salary": 4830 }
{ "firstName": "Anna", "lastName": "Jones", "age": 28, "salary": 5430 }
```

Now check for records in ActiveMQ.

Now stop the connector.

Features

The Sink supports:

1. Field selection - Kafka topic payload field selection is supported, allowing you to select fields written to the queue or topic in JMS.
2. Topic to topic routing.
3. Payload format selection.
4. Error policies for handling failures.

Kafka Connect Query Language

Kafka Connect Query Language found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The JMS Sink supports the following:

```
INSERT INTO <jms target> SELECT <fields> FROM <source topic>
```

Example:

```
#select all fields from topicA and write to jmsA
INSERT INTO jmsA SELECT * FROM topicA

#select 3 fields and rename from topicB and write to jmsB
INSERT INTO jmsB SELECT x AS a, y AS b and z AS c FROM topicB
```

JMS Payload

When a message is sent to a JMS target it can be one of the following:

1. JSON - Send a TextMessage;
2. AVRO - Send a BytesMessage;
3. MAP - Send a MapMessage;
4. OBJECT - Send an ObjectMessage

Topic Routing

The Sink supports topic routing that allows mapping the messages from topics to a specific jms target. For example, map a topic called “bloomberg_prices” to a jms target named “prices”. This mapping is set in the `connect.jms.sink.sink.kcql` option.

Example:

```
//Select all
INSERT INTO jms1 SELECT * FROM topic1; INSERT INTO jms3 SELECT * FROM
topicCConfigurations
```

Configurations

`connect.jms.sink.url`

Provides the JMS broker url

- Data Type: string
- Importance: high
- Optional : no

`connect.jms.sink.user`

Provides the user for the JMS connection.

- Data Type: string
- Importance: high
- Optional : no

`connect.jms.sink.password`

Provides the password for the JMS connection.

- Data Type: string
- Importance: high
- Optional : no

`connect.jms.sink.connection.factory`

Provides the full class name for the ConnectionFactory implementation to use.

- Data Type: string
- Importance: high
- Optional : no

`connect.jms.sink.sink.kcql`

KCQL expression describing field selection and routes.

- Data Type: string
- Importance: high
- Optional : no

`connect.jms.sink.export.route.topics`

Lists all the jms target topics.

- Data Type: list (comma separated strings)
- Importance: medium
- Optional : yes

`connect.jms.sink.export.route.queue`

Lists all the jms target queues.

- Data Type: list (comma separated strings)
- Importance: medium
- Optional : yes

`connect.jms.sink.message.type`

Specifies the JMS payload. If JSON is chosen it will send a `TextMessage`.

- Data Type: string
- Importance: medium
- Optional : yes
- Default : AVRO

`connect.jms.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.jms.sink.max.retries` option. The `connect.jms.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: medium
- Optional: yes
- Default: RETRY

`connect.jms.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.jms.sink.error.policy` is set to `retry`.

- Type: string
- Importance: medium
- Optional: yes
- Default: 10

`connect.jms.sink.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.jms.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: medium
- Optional: yes
- Default : 60000 (1 minute)

Schema Evolution

Not applicable.

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect Kudu

A Connector and Sink to write events from Kafka to kudu. The connector takes the value from the Kafka Connect SinkRecords and inserts a new entry to Kudu.

The Sink supports:

1. *The KCQL routing querying* - Kafka topic payload field selection is supported, allowing you to select fields written to Kudu.
2. Topic to table routing via KCQL.
3. Auto table create with DISTRIBUTE BY partition strategy via KCQL.
4. Auto evolution of tables via KCQL.
5. Error policies for handling failures.

Prerequisites

- Confluent 3.0.1
- Kudu 0.8
- Java 1.8
- Scala 2.11

Setup

Kudu Setup

Download and check Kudu QuickStart VM starts up.

```
curl -s https://raw.githubusercontent.com/cloudera/kudu-examples/master/demo-vm-setup/  
↪bootstrap.sh | bash
```

Confluent Setup

Follow the instructions [here](#).

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Kudu Table

Lets create a table in Kudu via Impala. The Sink does support auto creation of tables but they are not sync'd yet with Impala.

```
#demo/demo
ssh demo@quickstart -t impala-shell

CREATE TABLE default.kudu_test (id INT,random_field STRING )
TBLPROPERTIES ('kudu.master_addresses'='127.0.0.1', 'kudu.key_columns'='id',
'kudu.table_name'='kudu_test', 'transient_lastDdlTime'='1456744118',
'storage_handler'='com.cloudera.kudu.hive.KuduStorageHandler')
exit;
```

Note: The Sink will fail to start if the tables matching the topics do not already exist and the Sink is not in auto create mode.

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools` cli to post in our distributed properties file for Kudu. If you are using the *dockers* you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create kudu-sink < conf/kudu-sink.properties
#Connector name=kudu-sink
connector.class=com.datamountaineer.streamreactor.connect.kudu.KuduSinkConnector
tasks.max=1
connect.kudu.master=quickstart
connect.kudu.sink.kcql = INSERT INTO kudu_test SELECT * FROM kudu-test
topics=kudu_test
#task ids: 0
```

The `kudu-sink.properties` file defines:

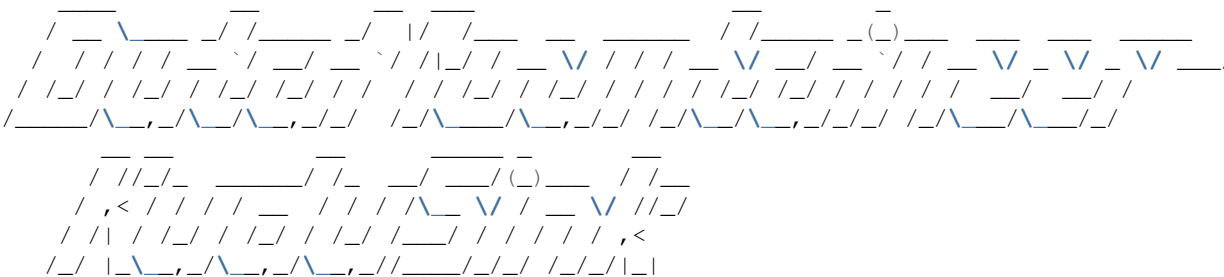
1. The name of the sink.
2. The Sink class.
3. The max number of tasks the connector is allowed to created. Should not be greater than the number of partitions in the Source topics otherwise tasks will be idle.
4. The Kudu master host.
5. *The KCQL routing querying.*
6. The Source kafka topics to take events from.

If you switch back to the terminal you started the Connector in you should see the Kudu Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
kudu-sink
```

```
[2016-05-08 22:00:20,823] INFO
```



```

by Andrew Stevenson
    (com.datamountaineer.streamreactor.connect.kudu.KuduSinkTask:37)
[2016-05-08 22:00:20,823] INFO KuduSinkConfig values:
    connect.kudu.master = quickstart
    (com.datamountaineer.streamreactor.connect.kudu.KuduSinkConfig:165)
[2016-05-08 22:00:20,824] INFO Connecting to Kudu Master at quickstart (com.
↳ datamountaineer.streamreactor.connect.kudu.KuduWriter$:33)
[2016-05-08 22:00:20,875] INFO Initialising Kudu writer (com.datamountaineer.
↳ streamreactor.connect.kudu.KuduWriter:40)
[2016-05-08 22:00:20,892] INFO Assigned topics (com.datamountaineer.streamreactor.
↳ connect.kudu.KuduWriter:42)
[2016-05-08 22:00:20,904] INFO Sink task org.apache.kafka.connect.runtime.
↳ WorkerSinkTask@b60ba7b finished initialization and start (org.apache.kafka.connect.
↳ runtime.WorkerSinkTask:155)

```

Test Records

Now we need to put some records it to the `test_table` topics. We can use the `kafka-avro-console-producer` to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a `id` field of type `int` and a `random_field` of type `string`.

```

${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic kudu-test \
--property value.schema='{ "type": "record", "name": "myrecord", "fields": [{ "name": "id",
↪ "type": "int"},
{ "name": "random_field", "type": "string" } ] }'

```

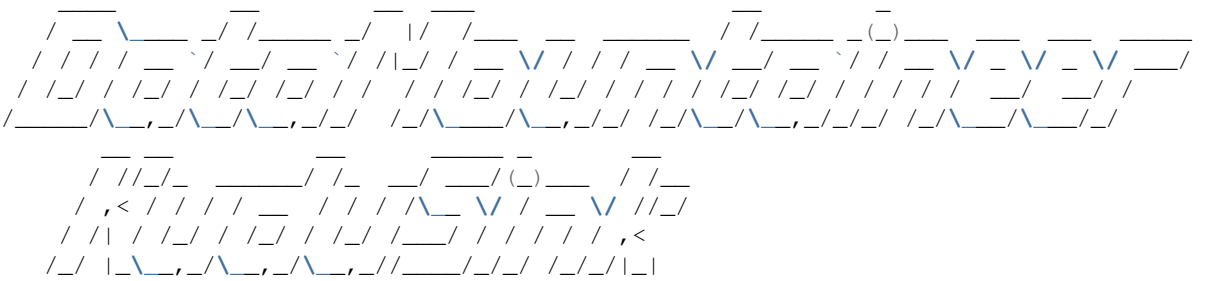
Now the producer is waiting for input. Paste in the following:

```
{ "id": 999, "random_field": "foo" }
{ "id": 888, "random_field": "bar" }
```

Check for records in Kudu

Now check the logs of the connector you should see this:

```
[2016-05-08 22:09:22,065] INFO
```



```
by Andrew Stevenson
    (com.datamountaineer.streamreactor.connect.kudu.KuduSinkTask:37)
[2016-05-08 22:09:22,065] INFO KuduSinkConfig values:
    connect.kudu.master = quickstart
    (com.datamountaineer.streamreactor.connect.kudu.KuduSinkConfig:165)
[2016-05-08 22:09:22,066] INFO Connecting to Kudu Master at quickstart (com.
↳ datamountaineer.streamreactor.connect.kudu.KuduWriter$:33)
[2016-05-08 22:09:22,116] INFO Initialising Kudu writer (com.datamountaineer.
↳ streamreactor.connect.kudu.KuduWriter:40)
[2016-05-08 22:09:22,134] INFO Assigned topics kudu_test (com.datamountaineer.
↳ streamreactor.connect.kudu.KuduWriter:42)
[2016-05-08 22:09:22,148] INFO Sink task org.apache.kafka.connect.runtime.
↳ WorkerSinkTask@68496440 finished initialization and start (org.apache.kafka.connect.
↳ runtime.WorkerSinkTask:155)
[2016-05-08 22:09:22,476] INFO Written 2 for kudu_test (com.datamountaineer.
↳ streamreactor.connect.kudu.KuduWriter:90)
```

In Kudu:

```
#demo/demo
ssh demo@quickstart -t impala-shell

SELECT * FROM kudu_test;

Query: select * FROM kudu_test
+-----+-----+
| id | random_field |
+-----+-----+
| 888 | bar          |
| 999 | foo          |
+-----+-----+
Fetched 2 row(s) in 0.14s
```

Now stop the connector.

Features

The Kudu Sink writes records from Kafka to Kudu.

The Sink supports:

1. Field selection - Kafka topic payload field selection is supported, allowing you to select fields written to Kudu.
2. Topic to table routing.
3. Auto table create with DISTRIBUTE BY partition strategy.
4. Auto evolution of tables.
5. Error policies for handling failures.

Kafka Connect Query Language

Kafka Connect Query Language found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The Kudu Sink supports the following:

```
<write mode> INTO <target table> SELECT <fields> FROM <source topic> <AUTOCREATE>
↳<DISTRIBUTEBY> <PK_FIELDS> INTO <NBR_OF_BUCKETS> BUCKETS <AUTOEVOLVE>
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA
INSERT INTO tableA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to tableB
INSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB

#Upsert mode, select all fields from topicC, auto create tableC and auto evolve, use
↳field1 and field2 as the primary keys
UPSERT INTO tableC SELECT * FROM topicC AUTOCREATE DISTRIBUTEBY field1, field2 INTO
↳10 BUCKETS AUTOEVOLVE
```

Error Polices

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the Sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers.

Retry

Any error on write to the target database causes the RetryIterable exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.kudu.sink.max.retries` and the `connect.kudu.sink.retry.interval`.

Auto conversion of Connect records to Kudu

The Sink automatically converts incoming Connect records to Kudu inserts or upserts.

Topic Routing

The Sink supports topic routing that allows mapping the messages from topics to a specific table. For example, map a topic called “bloomberg_prices” to a table called “prices”. This mapping is set in the `connect.kudu.sink.kcql` option.

Example:

```
//Select all
INSERT INTO table1 SELECT * FROM topic1; INSERT INTO tableA SELECT * FROM topicC
```

Field Selection

The Kudu Sink supports field selection and mapping. This mapping is set in the `connect.kudu.sink.kcql` option.

Examples:

```
//Rename or map columns
INSERT INTO table1 SELECT lst_price AS price, qty AS quantity FROM topicA

//Select all
INSERT INTO table1 SELECT * FROM topic1
```

Tip: Check you mappings to ensure the target columns exist.

Warning: Field selection disables evolving the target table if the upstream schema in the Kafka topic changes. By specifying field mappings it is assumed the user is not interested in new upstream fields. For example they may be tapping into a pipeline for a Kafka stream job and not be intended as the final recipient of the stream.

If you chose field selection you must include the primary key fields otherwise the insert will fail.

Write Modes

The Sink supports both **insert** and **upsert** modes. This mapping is set in the `connect.kudu.sink.export.mappings` option.

Insert

Insert is the default write mode of the sink.

Insert Idempotency

Kafka currently provides at least once delivery semantics. Therefore, this mode may produce errors if unique constraints have been implemented on the target tables. If the error policy has been set to NOOP then the Sink will discard the error and continue to process, however, it currently makes no attempt to distinguish violation of integrity constraints from other exceptions such as casting issues.

Upsert

The Sink support Kudu upserts which replaces the existing row if a match is found on the primary keys.

Upsert Idempotency

Kafka currently provides at least once delivery semantics and order is a guaranteed within partitions.

This mode will, if the same record is delivered twice to the sink, result in an idempotent write. The existing record will be updated with the values of the second which are the same.

If records are delivered with the same field or group of fields that are used as the primary key on the target table, but different values, the existing record in the target table will be updated.

Since records are delivered in the order they were written per partition the write is idempotent on failure or restart. Redelivery produces the same result.

Auto Create Tables

The Sink supports auto creation of tables for each topic. This mapping is set in the `connect.kudu.sink.kcql` option.

Primary keys are set in the `DISTRIBUTE BY` clause of the `connect.kudu.sink.kcql`.

Tables are created with the Kudu hash partition strategy. The number of buckets must be specified in the `kcql` statement.

```
#AutoCreate the target table
INSERT INTO table1 SELECT * FROM topic AUTOCREATE DISTRIBUTE BY field1, field2 INTO 10
↪BUCKETS
```

Note: The fields specified as the primary keys (distributeby) must be in the `SELECT` clause or all fields must be selected

The Sink will try and create the table at start up if a schema for the topic is found in the Schema Registry. If no schema is found the table is created when the first record is received for the topic.

Auto Evolve Tables

The Sink supports auto evolution of tables for each topic. This mapping is set in the `connect.kudu.sink.kcql` option. When set the Sink will identify new schemas for each topic based on the schema version from the Schema registry. New columns will be identified and an alter table DDL statement issued against Kudu.

Schema evolution can occur upstream, for example any new fields or change in data type in the schema of the topic, or downstream DDLs on the database.

Upstream changes must follow the schema evolution rules laid out in the Schema Registry. This Sink only supports BACKWARD and FULLY compatible schemas. If new fields are added the Sink will attempt to perform a `ALTER` table DDL statement against the target table to add columns. All columns added to the target table are set as nullable.

Fields cannot be deleted upstream. Fields should be of Avro union type [null, <dataType>] with a default set. This allows the Sink to either retrieve the default value or null. The Sink is not aware that the field has been deleted as a value is always supplied to it.

Warning: If a upstream field is removed and the topic is not following the Schema Registry's evolution rules, i.e. not full or backwards compatible, any errors will default to the error policy.

Downstream changes are handled by the sink. If columns are removed, the mapped fields from the topic are ignored. If columns are added, we attempt to find a matching field by name in the topic.

Data Type Mappings

Connect Type	Kudu Data Type
INT8	INT8
INT16	INT16
INT32	INT32
INT64	INT64
BOOLEAN	BOOLEAN
FLOAT32	FLOAT
FLOAT64	FLOAT
BYTES	BINARY

Configurations

`connect.kudu.master`

Specifies a Kudu server.

- Data type : string
- Importance: high
- Optional : no

`connect.kudu.sink.kcql`

Kafka connect query language expression. Allows for expressive topic to table routing, field selection and renaming.

Examples:

```
INSERT INTO TABLE1 SELECT * FROM TOPIC1; INSERT INTO TABLE2 SELECT field1, field2, ↵
↵field3 as renamedField FROM TOPIC2
```

- Data Type: string
- Importance: high
- Optional : no

`connect.kudu.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.kudu.sink.max.retries` option. The `connect.kudu.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: high
- Optional : yes
- Default: RETRY

`connect.kudu.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.kudu.sink.error.policy` is set to `retry`.

- Type: string
- Importance: medium
- Optional : yes
- Default: 10

`connect.kudu.sink.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.kudu.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: medium
- Optional : yes
- Default : 60000 (1 minute)

`connect.kudu.sink.schema.registry.url`

The url for the Schema registry. This is used to retrieve the latest schema for table creation.

- Type : string
- Importance : high
- Optional : yes
- Default : <http://localhost:8081>

`connect.kudu.sink.batch.size`

Specifies how many records to insert together at one time. If the connect framework provides less records when it is calling the Sink it won't wait to fulfill this value but rather execute it.

- Type : int
- Importance : medium
- Optional : yes
- Defaults : 3000

Example

```
name=kudu-sink
connector.class=com.datamountaineer.streamreactor.connect.kudu.KuduSinkConnector
tasks.max=1
connect.kudu.master=quickstart
```

```
connect.kudu.sink.kcql=INSERT INTO kudu_test SELECT * FROM kudu_test AUTOCREATE_  
↳DISTRIBUTEBY id INTO 5 BUCKETS  
topics=kudu_test  
connect.kudu.sink.schema.registry.url=http://myhost:8081
```

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect Mongo Sink

The Mongo Sink allows you to write events from Kafka to your MongoDB instance. The connector converts the value from the Kafka Connect SinkRecords to MongoDB Document and will do an insert or upsert depending on the configuration you chose. It is expected the database is created upfront; the targeted MongoDB collections will be created if they don't exist

Note: The database needs to be created upfront!

The Sink supports:

1. *The KCQL routing querying* - Topic to measure mapping and Field selection.
2. Schema registry support for Connect/Avro with a schema.
3. Schema registry support for Connect and no schema (schema set to Schema.String)
4. Json payload support, no Schema Registry.
5. Error policies.
6. Schema.Struct and payload Struct, Schema.String and Json payload and Json payload with no schema.

The Sink supports three Kafka payloads type:

Connect entry with Schema.Struct and payload Struct. If you follow the best practice while producing the events, each message should carry its schema information. Best option is to send Avro. Your connect configurations should be set to `value.converter=io.confluent.connect.avro.AvroConverter`. You can find an example [here](#). To see how easy is to have your producer serialize to Avro have a look at [this](#). This requires SchemaRegistry which is open source thanks to Confluent! Alternatively you can send Json + Schema. In this case your connect configuration should read `value.converter=org.apache.kafka.connect.json.JsonConverter`. The difference would be to point your serialization to `org.apache.kafka.connect.json.JsonSerializer`. This doesn't require the SchemaRegistry.

Connect entry with Schema.String and payload json String. Sometimes the producer would find it easier, despite sending Avro to produce a GenericRecord, to just send a message with Schema.String and the json string.

Connect entry without a schema and the payload json String. There are many existing systems which are publishing json over Kafka and bringing them in line with best practices is quite a challenge. Hence we added the support

Prerequisites

- MongoDB 3.2.10
- Confluent 3.0.1
- Java 1.8
- Scala 2.11

Setup

Before we can do anything, including the QuickStart we need to install MongoDB and the Confluent platform.

Confluent Setup

Follow the instructions [here](#).

MongoDb Setup

If you already have an instance of Mongo running you can skip this step. First download and install MongoDB Community edition. This is the manual approach for installing on Ubuntu. You can follow the details <https://docs.mongodb.com/v3.2/administration/install-community/> for your OS.

```
#go to home folder
cd ~
#make a folder for mongo
mkdir mongodb

#Download Mongo
wget https://fastdl.mongodb.org/linux/mongodb-linux-x86_64-ubuntu1604-3.2.10.tgz
→tgz

#extract the archive
tar xvf mongodb-linux-x86_64-ubuntu1604-3.2.10.tgz -C mongodb
cd mongodb
mv mongodb-linux-x86_64-ubuntu1604-3.2.10/* .

#create the data folder
mkdir data
mkdir data/db

#Start MongoDB
bin/mongod --dbpath data/db
```

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

The important configuration for Connect is related to the key and value deserializer. In the first example we default to the best practice where the source sends Avro messages to a Kafka topic. It is not enough to just be Avro messages but also the producer must work with the Schema Registry to create the schema if it doesn't exist and set the schema id in the message. Every message sent will have a magic byte followed by the Avro schema id and then the actual Avro record in binary format.

Here are the entries in the config setting all the above. The are placed in the `connect-properties` file Kafka Connect is started with. Of course if your SchemaRegistry runs on a different machine or you have multiple instances of it you will have to amend the configuration.

```
key.converter=io.confluent.connect.avro.AvroConverter
key.converter.schema.registry.url=http://localhost:8081
value.converter=io.confluent.connect.avro.AvroConverter
value.converter.schema.registry.url=http://localhost:8081
```

Test Database

The Sink requires that a database be precreated in MongoDB.

```
#from a new terminal
cd ~/mongodb/bin

#start the cli
./mongo

#list all dbs
show dbs

#create a new database named connect
use connect
#create a dummy collection and insert one document to actually create the database
db.dummy.insert({"name": "Kafka Rulz!"})

#list all dbs
show dbs
```

Starting the Connector

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools` cli to post in our distributed properties file for Kudu. If you are using the *dockers* you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create mongo-sink < conf/source.kcql/mongo-sink.properties

#Connector `mongo-sink-orders`:
```



```
name=mongo-sink-orders
connector.class=com.datamountaineer.streamreactor.connect.mongodb.sink.
↳MongoSinkConnector
tasks.max=1
topics=orders-topic
connect.mongo.sink.kcql=INSERT INTO orders SELECT * FROM orders-topic
connect.mongo.database=connect
connect.mongo.connection=mongodb://localhost:27017
connect.mongo.sink.batch.size=10

#task ids: 0
```

If you switch back to the terminal you started Kafka Connect in you should see the Mongo Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
mongo-sink
```

[illegible]

Test Records

Hint: If your input topic doesn't match the target use [Kafka Streams](#) to transform in realtime the input. Also checkout the [Plumber](#), which allows you to inject a Lua script into [Kafka Streams](#) to do this, no Java or Scala required!

Now we need to put some records it to the orders-topic. We can use the `kafka-avro-console-producer` to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema matches the table created earlier.

```
bin/kafka-avro-console-producer \  
  --broker-list localhost:9092 --topic orders-topic \  
  --property value.schema='{ "type": "record", "name": "myrecord", "fields": [{ "name": "id",  
↪ "type": "int"},  
{ "name": "created", "type": "string"}, { "name": "product", "type": "string"}, { "name":  
↪ "price", "type": "double" } ] }'
```

Now the producer is waiting for input. Paste in the following (each on a line separately):

```
{ "id": 1, "created": "2016-05-06 13:53:00", "product": "OP-DAX-P-20150201-95.7",  
↪ "price": 94.2 }  
{ "id": 2, "created": "2016-05-06 13:54:00", "product": "OP-DAX-C-20150201-100", "price"  
↪ ": 99.5 }  
{ "id": 3, "created": "2016-05-06 13:55:00", "product": "FU-DATAMOUNTAINEER-20150201-  
↪ 100", "price": 10000 }  
{ "id": 4, "created": "2016-05-06 13:56:00", "product": "FU-KOSPI-C-20150201-100",  
↪ "price": 150 }
```

Now if we check the logs of the connector we should see 2 records being inserted to MongoDB:

```
[2016-11-06 22:30:30,473] INFO Setting newly assigned partitions [orders-topic-0] for_  
↪ group connect-mongo-sink-orders (org.apache.kafka.clients.consumer.internals.  
↪ ConsumerCoordinator:231)  
[2016-11-06 22:31:29,328] INFO WorkerSinkTask{id=mongo-sink-orders-0} Committing_  
↪ offsets (org.apache.kafka.connect.runtime.WorkerSinkTask:261)
```

```
#Open a new terminal and navigate to the mongodb instalation folder  
./bin/mongo  
  > show databases  
      connect  0.000GB  
      local    0.000GB  
  > use connect  
      switched to db connect  
  > show collections  
      dummy  
      orders  
  > db.orders.find()  
  { "_id" : ObjectId("581fb21b09690a24b63b35bd"), "id" : 1, "created" : "2016-05-06_  
↪ 13:53:00", "product" : "OP-DAX-P-20150201-95.7", "price" : 94.2 }  
  { "_id" : ObjectId("581fb2f809690a24b63b35c2"), "id" : 2, "created" : "2016-05-06_  
↪ 13:54:00", "product" : "OP-DAX-C-20150201-100", "price" : 99.5 }  
  { "_id" : ObjectId("581fb2f809690a24b63b35c3"), "id" : 3, "created" : "2016-05-06_  
↪ 13:55:00", "product" : "FU-DATAMOUNTAINEER-20150201-100", "price" : 10000 }  
  { "_id" : ObjectId("581fb2f809690a24b63b35c4"), "id" : 4, "created" : "2016-05-06_  
↪ 13:56:00", "product" : "FU-KOSPI-C-20150201-100", "price" : 150 }
```

Bingo, our 4 rows!

Legacy topics (plain text payload with a json string)

We have found some of the clients have already an infrastructure where they publish pure json on the topic and obviously the jump to follow the best practice and use schema registry is quite an ask. So we offer support for them as well.

This time we need to start the connect with a different set of settings.

```
#create a new configuration for connect
cp etc/schema-registry/connect-avro-distributed.properties etc/schema-registry/
connect-avro-distributed-json.properties
vi etc/schema-registry/connect-avro-distributed-json.properties
```

Replace the following 4 entries in the config

```
key.converter=io.confluent.connect.avro.AvroConverter
key.converter.schema.registry.url=http://localhost:8081
value.converter=io.confluent.connect.avro.AvroConverter
value.converter.schema.registry.url=http://localhost:8081
```

with the following

```
key.converter=org.apache.kafka.connect.json.JsonConverter
key.converter.schemas.enable=false
value.converter=org.apache.kafka.connect.json.JsonConverter
value.converter.schemas.enable=false
```

Now let's restart the connect instance:

```
#start a new instance of connect
$bin/start-connect.sh
```

Use the CLI to remove the old MongoDB Sink:

```
bin/cli.sh rm mongo-sink
```

and start the new Sink with the json properties files to read from the a different topic with json as the payload.

```
#start the connector for mongo
bin/cli.sh create mongo-sink-orders-json < mongo-sink-orders-json.properties
```

You should see in the terminal where you started Kafka Connect the following entries in the log:

```
[2016-11-06 23:53:09,881] INFO MongoConfig values:
  connect.mongo.retry.interval = 60000
  connect.mongo.sink.kcql = UPSERT INTO orders_json SELECT id, product as product_
  ↳ name, price as value FROM orders-topic-json PK id
  connect.mongo.connection = mongodb://localhost:27017
  connect.mongo.error.policy = THROW
  connect.mongo.database = connect
  connect.mongo.sink.batch.size = 10
  connect.mongo.max.retires = 20
  (com.datamountainer.streamreactor.connect.mongodb.config.MongoConfig:178)
[2016-11-06 23:53:09,927] INFO
```

by Stefan Bocutiu

```
. (com.datamountaineer.streamreactor.connect.mongodb.sink.MongoSinkTask:51)
[2016-11-06 23:53:10,270] INFO Initialising Mongo writer.Connection to mongodb://
↪localhost:27017 (com.datamountaineer.streamreactor.connect.mongodb.sink.MongoWriter
↪$:126)
```

Now it's time to produce some records. This time we will use the simple kafka-console-consumer to put simple json on the topic:

```
`${CONFLUENT_HOME}/bin/kafka-console-producer --broker-list localhost:9092 --topic_
↪orders-topic-json

{"id": 1, "created": "2016-05-06 13:53:00", "product": "OP-DAX-P-20150201-95.7",
↪"price": 94.2}
{"id": 2, "created": "2016-05-06 13:54:00", "product": "OP-DAX-C-20150201-100", "price
↪": 99.5}
{"id": 3, "created": "2016-05-06 13:55:00", "product": "FU-DATAMOUNTAINEER-20150201-
↪100", "price":10000}
```

Following the command you should have something similar to this in the logs for your connect:

```
[2016-11-07 00:08:30,200] INFO Setting newly assigned partitions [orders-topic-json-
↪0] for group connect-mongo-sink-orders-json (org.apache.kafka.clients.consumer.
↪internals.ConsumerCoordinator:231)
[2016-11-07 00:08:30,324] INFO Opened connection [connectionId{localValue:3,
↪serverValue:9}] to localhost:27017 (org.mongodb.driver.connection:71)
```

Let's check the mongo db database for the new records:

```
#Open a new terminal and navigate to the mongodb installation folder
./bin/mongo
> show databases
  connect  0.000GB
  local    0.000GB
> use connect
switched to db connect
> show collections
  dummy
  orders
  orders_json
> db.orders_json.find()
{ "_id" : ObjectId("581fc5fe53b2c9318a3c1004"), "created" : "2016-05-06 13:53:00",
↪ "id" : NumberLong(1), "product_name" : "OP-DAX-P-20150201-95.7", "value" : 94.2 }
{ "_id" : ObjectId("581fc5fe53b2c9318a3c1005"), "created" : "2016-05-06 13:54:00",
↪ "id" : NumberLong(2), "product_name" : "OP-DAX-C-20150201-100", "value" : 99.5 }
{ "_id" : ObjectId("581fc5fe53b2c9318a3c1006"), "created" : "2016-05-06 13:55:00",
↪ "id" : NumberLong(3), "product_name" : "FU-DATAMOUNTAINEER-20150201-100", "value"
↪: NumberLong(10000) }
```

Bingo, our 3 rows!

Features

The sink connector will translate the SinkRecords to json and will insert each one in the database. We support to insert modes: INSERT and UPSERT. All of this can be expressed via KCQL (our own SQL like syntax for configuration. Please see below the section for Kafka Connect Query Language)

The sink supports:

1. Field selection - Kafka topic payload field selection is supported, allowing you to have choose selection of fields or all fields written to MongoDB.
2. Topic to table routing. Your sink instance can be configured to handle multiple topics and collections. All you have to do is to set your configuration appropriately. Below you will find an example

```
connect.mongo.sink.kcql = INSERT INTO orders SELECT * FROM orders-topic; UPSERT INTO ↵
↵customers SELECT * FROM customer-topic PK customer_id; UPSERT INTO invoiceid as ↵
↵invoice_id, customerid as customer_id, value a SELECT invoice_id, FROM invoice-topic
```

3. Error policies for handling failures.

Kafka Connect Query Language

Kafka Connect Query Language, *KCQL* allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

MongoDb sink supports the following:

```
INSERT INTO <database>.<target collection> SELECT <fields> FROM <source topic> <PK ↵
↵field name>
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA
INSERT INTO collectionA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to tableB with primary ↵
↵key as the field id from the topic
INSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB PK id
```

Error Polices

The sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers..

Retry

Any error on write to the target database causes the `RetryIterable` exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the database is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the sink will retry can be controlled by using the `connect.mongo.max.retries` and the `connect.mongo.retry.interval`.

Topic Routing

The sink supports topic routing that maps the messages from topics to a specific collection. For example map a topic called “bloomberg_prices” to a collection called “prices”. This mapping is set in the `connect.mongo.kcql` option. You don’t need to set up multiple sinks for each topic or collection. The same sink instance can be configured to handle multiple collections. For example your configuration in this case:

```
connect.mongo.sink.kcql = INSERT INTO orders SELECT * FROM orders-topic; UPSERT INTO ↵
↵customers SELECT * FROM customer-topic PK customer_id; UPSERT INTO invoiceid as ↵
↵invoice_id, customerid as customer_id, value a SELECT invoice_id, FROM invoice-topic
```

Field Selection

The sink supports selecting fields from the source topic or selecting all. There is an option to rename a field as well. All of this can be easily expressed with KCQL:

- Select all fields from topic `fx_prices` and insert into the `fx` collection: `INSERT INTO fx SELECT * FROM fx_prices.`
- Select all fields from topic `fx_prices` and upsert into the `fx` collection, The assumption is there will be a `ticker` field in the incoming json: `UPSERT INTO fx SELECT * FROM fx_prices PK ticker.`
- Select specific fields from the topic `sample_topic` and insert into the `sample` collection: `INSERT INTO sample SELECT field1,field2,field3 FROM sample_topic.`
- Select specific fields from the topic `sample_topic` and upsert into the `sample` collection: `UPSERT INTO sample SELECT field1,field2,field3 FROM sample_fopic PK field1.`
- Rename some fields while selecting all from the topic `sample_topic` and insert into the `sample` collection: `INSERT INTO sample SELECT *, field1 as new_name1,field2 as new_name2 FROM sample_topic.`
- Rename some fields while selecting all from the topic `sample_topic` and upsert into the `sample` collection: `UPSERT INTO sample SELECT *, field1 as new_name1,field2 as new_name2 FROM sample_topic PK new_name1.`
- Select specific fields and rename some of them from the topic `sample_topic` and insert into the `sample` collection: `INSERT INTO sample SELECT field1 as new_name1,field2, field3 as new_name3 FROM sample_topic.`
- Select specific fields and rename some of them from the topic `sample_topic` and upsert into the `sample` collection: `INSERT INTO sample SELECT field1 as new_name1,field2, field3 as new_name3 FROM sample_fopic PK new_name3.`

Configurations

Configurations parameters:

`connect.mongo.database`

The target MongoDB database name.

- Data type: string

- Optional : no

`connect.mongo.connection`

The mongodb endpoints connections in the format `mongodb://[username:password@]host1[:port1][,host2[:port2],...[,hostN[:portN]]][/[c`

- Data type: string
- Optional : no

`connect.mongo.sink.batch.size`

The number of records the sink would push to mongo at once (improved performance)

- Data type: int
- Optional : yes
- Default: 100

`connect.mongo.sink.kcql`

Kafka connect query language expression. Allows for expressive topic to collection routing, field selection and renaming.

Examples:

```
INSERT INTO TABLE1 SELECT * FROM TOPIC1; INSERT INTO TABLE2 SELECT field1, field2, ↵
↵field3 as renamedField FROM TOPIC2
```

- Data Type: string
- Optional : no

`connect.mongo.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **NOOP**, the error is swallowed, **THROW**, the error is allowed to propagate and retry. For **RETRY** the Kafka message is redelivered up to a maximum number of times specified by the `connect.mongo.max.retires` option. The `connect.mongo.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: high
- Default: throw

`connect.mongo.max.retires`

The maximum number of times a message is retried. Only valid when the `connect.mongo.error.policy` is set to **TRHOW**.

- Type: string
- Importance: high
- Default: 10

`connect.mongo.retry.interval`

The interval, in milliseconds between retries if the sink is using `connect.mongo.error.policy` set to **RETRY**.

- Type: int
- Importance: medium

- Default : 60000 (1 minute)

Example

```
name=mongo-sink-orders
connector.class=com.datamountaineer.streamreactor.connect.mongodb.sink.
↳MongoSinkConnector
tasks.max=1
topics=orders-topic
connect.mongo.sink.kcql=INSERT INTO orders SELECT * FROM orders-topic
connect.mongo.database=connect
connect.mongo.connection=mongodb://localhost:27017
connect.mongo.sink.batch.size=10
```

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect Redis

A Connector and Sink to write events from Kafka to Redis. The connector takes the value from the Kafka Connect SinkRecords and inserts a new entry to Redis.

The Sink supports:

1. *The [KCQL routing querying](#)* - Kafka topic payload field selection is supported, allowing you to select fields written to Redis.
2. Topic to table routing via KCQL.
3. RowKey selection - Selection of fields to use as the row key, if none specified the topic name, partition and offset are used via KCQL.
4. Error policies for handling failures.

Prerequisites

- Confluent 3.0.1
- Jedis 2.8.1

- Java 1.8
- Scala 2.11

Setup

Redis Setup

Download and install Redis.

```
wget http://download.redis.io/redis-stable.tar.gz
tar xvzf redis-stable.tar.gz
cd redis-stable
sudo make install
```

Start Redis

```
bin/redis-server
```

Check Redis is running:

```
redis-cli ping
PONG
sudo service redis-server status
```

Confluent Setup

Follow the instructions [here](#).

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools cli` to post in our distributed properties file for Redis. If you are using the [dockers](#) you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create redis-sink < conf/redis-sink.properties
#Connector name='redis-sink'
connect.redis.connection.host=localhost
connect.redis.connection.port=6379
connector.class=com.datamountaineer.streamreactor.connect.redis.sink.
↳RedisSinkConnector
tasks.max=1
topics=redis-topic
connect.redis.sink.kcql=INSERT INTO TABLE1 SELECT * FROM redis-topic
#task ids: 0
```

The `redis-sink.properties` file defines:

1. The name of the sink.
2. The name of the redis host to connect to.
3. The redis port to connect to.
4. The Sink class.
5. The max number of tasks the connector is allowed to created. Should not be greater than the number of partitions in the Source topics otherwise tasks will be idle.
6. The Source kafka topics to take events from.
7. *The KCQL routing querying.*

Warning: If your redis server is requiring the connection to be authenticated you will need to provide an extra setting:

```
connect.redis.sink.connection.password=$REDIS_PASSWORD
```

Don't set the value to empty if no password is required.

If you switch back to the terminal you started the Connector in you should see the Redis Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
redis-sink
```

[illegible]

```
RedisSinkSettings(RedisConnectionInfo(localhost,6379,None),RedisKey(FIELDS,
↳ WrappedArray(firstName, lastName)),PayloadFields(false,Map(firstName -> firstName,
↳ lastName -> lastName, age -> age, salary -> income)))
    (com.datamountaineer.streamreactor.connect.redis.sink.RedisSinkTask:65)
[2016-05-08 22:37:05,687] INFO Sink task org.apache.kafka.connect.runtime.
↳ WorkerSinkTask@44b24eaa finished initialization and start (org.apache.kafka.connect.
↳ runtime.WorkerSinkTask:155)
```

Test Records

Now we need to put some records it to the test_table topics. We can use the kafka-avro-console-producer to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a firstname field of type string, a lastname field of type string, an age field of type int and a salary field of type double.

```
${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic redis-topic \
--property value.schema='{ "type": "record", "name": "User", "namespace": "com.
↳ datamountaineer.streamreactor.connect.redis"
, "fields": [{ "name": "firstName", "type": "string"}, { "name": "lastName", "type": "string"},
↳ { "name": "age", "type": "int"}, { "name": "salary", "type": "double"}] }'
```

Now the producer is waiting for input. Paste in the following:

```
{"firstName": "John", "lastName": "Smith", "age":30, "salary": 4830}
```

Check for records in Redis

Now check the logs of the connector you should see this:

```
INFO Received record from topic:redis-topic partition:0 and offset:0 (com.
↳ datamountaineer.streamreactor.connect.redis.sink.writer.RedisDbWriter:48)
INFO Empty list of records received. (com.datamountaineer.streamreactor.connect.redis.
↳ sink.RedisSinkTask:75)
```

Check in Redis.

```
redis-cli
127.0.0.1:6379> keys *
1) "John.Smith"
2) "11"
3) "10"
127.0.0.1:6379>
127.0.0.1:6379> get "John.Smith"
"{\"firstName\":\"John\",\"lastName\":\"Smith\",\"age\":30,\"income\":4830.0}"
```

Now stop the connector.

Features

The Redis Sink writes records from Kafka to Redis.

The Sink supports:

1. Field selection - Kafka topic payload field selection is supported, allowing you to select fields written to Redis.
2. Topic to table routing.
3. RowKey selection - Selection of fields to use as the row key, if none specified the topic name, partition and offset are used.
4. Error policies for handling failures.

Kafka Connect Query Language

Kafka Connect Query Language found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The Redis Sink supports the following:

```
INSERT INTO <table> SELECT <fields> FROM <source topic> <PK> primary_key_cols
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA and use the default_  
→rowkey (topic name, partition, offset)  
INSERT INTO tableA SELECT * FROM topicA  
  
#Insert mode, select 3 fields and rename from topicB and write to tableB, use field y_  
→from the topic as the primary key  
INSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB PK y
```

This is set in the `connect.redis.sink.kcql` option.

Error Policies

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers.

Retry

Any error on write to the target database causes the `RetryIterable` exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.redis.sink.max.retries` and the `connect.redis.sink.retry.interval`.

Configurations

`connect.redis.sink.kcql`

Kafka connect query language expression. Allows for expressive topic to table routing, field selection and renaming. Fields to be used as the row key can be set by specifying the PK. The below example uses field1 as the primary key.

- Data type : string
- Importance: high
- Optional : no

Examples:

```
INSERT INTO TABLE1 SELECT * FROM TOPIC1; INSERT INTO TABLE2 SELECT * FROM TOPIC2 PK_
↪field1
```

Examples:

```
INSERT INTO TABLE1 SELECT * FROM TOPIC1; INSERT INTO TABLE2 SELECT * FROM TOPIC2 PK_
↪field1, field2
```

`connect.redis.sink.connection.host`

Specifies the Redis server.

- Data type : string
- Importance: high
- Optional : no

`connect.redis.sink.connection.port`

Specifies the Redis server port number.

- Data type : int
- Importance: high
- Optional : no

`connect.redis.sink.connection.password`

Specifies the authorization password.

- Data type : string
- Importance: high
- Optional : yes
- Description: If you don't have a password set up on the redis server don't provide the value or you will see this error: "ERR Client sent AUTH, but no password is set"

`connect.redis.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.redis.sink.max.retries` option. The `connect.redis.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: medium
- Optional: yes
- Default: RETRY

`connect.redis.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.redis.sink.error.policy` is set to `retry`.

- Type: string
- Importance: medium
- Optional: yes
- Default: 10

`connect.redis.sink.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.redis.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: high
- Optional: no
- Default : 60000 (1 minute)

Example

```
name=redis-sink
connect.redis.connection.host=localhost
connect.redis.connection.port=6379
connector.class=com.datamountaineer.streamreactor.connect.redis.sink.
↳RedisSinkConnector
tasks.max=1
topics=redis-topic
connect.redis.sink.kcql=INSERT INTO TABLE1 SELECT * FROM redis-topic
```

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

The Redis Sink will automatically write and update the Redis table if new fields are added to the Source topic, if fields are removed the Kafka Connect framework will return the default value for this field, dependent of the compatibility

settings of the Schema registry. This value will be put into the Redis column family cell based on the `connect.redis.sink.kcql` mappings.

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect ReThink

A Connector and Sink to write events from Kafka to RethinkDb. The connector takes the value from the Kafka Connect SinkRecords and inserts a new entry to RethinkDb.

The Sink supports:

1. *The [KCQL routing querying](#)* - Kafka topic payload field selection is supported, allowing you to select fields written to RethinkDb.
2. Topic to table routing via KCQL.
3. RowKey selection - Selection of fields to use as the row key, if none specified the topic name, partition and offset are used via KCQL.
4. RethinkDB write modes via KCQL.
5. Error policies for handling failures.
6. Schema.Struct and payload Struct, Schema.String and Json payload and Json payload with no schema

The Sink supports three Kafka payloads type:

Connect entry with Schema.Struct and payload Struct. If you follow the best practice while producing the events, each message should carry its schema information. Best option is to send Avro. Your connect configurations should be set to `value.converter=io.confluent.connect.avro.AvroConverter`. You can find an example [here](#). To see how easy is to have your producer serialize to Avro have a look at [this](#). This requires SchemaRegistry which is open source thanks to Confluent! Alternatively you can send Json + Schema. In this case your connect configuration should read `value.converter=org.apache.kafka.connect.json.JsonConverter`. The difference would be to point your serialization to `org.apache.kafka.connect.json.JsonSerializer`. This doesn't require the SchemaRegistry.

Connect entry with Schema.String and payload json String. Sometimes the producer would find it easier, despite sending Avro to produce a GenericRecord, to just send a message with Schema.String and the json string.

Connect entry without a schema and the payload json String. There are many existing systems which are publishing json over Kafka and bringing them in line with best practices is quite a challenge. Hence we added the support

Prerequisites

- Confluent 3.0.1
- RethinkDb 2.3.3
- Java 1.8

- Scala 2.11

Setup

Rethink Setup

Download and install RethinkDb. Follow the instruction [here](#) dependent on your operating system.

Confluent Setup

Follow the instructions [here](#).

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools cli` to post in our distributed properties file for ReThinkDB. If you are using the [dockers](#) you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create rethink-sink < rethink-sink.properties
#Connector name=`rethink-sink`
name=rethink-sink
connect.rethink.sink.db=localhost
connect.rethink.sink.port=28015
connector.class=com.datamountaineer.streamreactor.connect.rethink.sink.
↪ReThinkSinkConnector
tasks.max=1
topics=rethink-topic
connect.rethink.sink.kcql=INSERT INTO TABLE1 SELECT * FROM rethink-topic
#task ids: 0
```

The `rethink-sink.properties` file defines:

1. The name of the sink.
2. The name of the rethink host to connect to.
3. The rethink port to connect to.

4. The Sink class.
5. The max number of tasks the connector is allowed to created. Should not be greater than the number of partitions in the Source topics otherwise tasks will be idle.
6. The Source kafka topics to take events from.
7. *The KCQL routing querying.*

If you switch back to the terminal you started the Connector in you should see the ReThinkDB Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
rethink-sink
```

```
[2016-05-08 22:37:05,616] INFO
```

```
(com.datamountaineer.streamreactor.connect.rethink.sink.config.RethinkSinkConfig)
```

Test Records

Now we need to put some records it to the `test_table` topics. We can use the `kafka-avro-console-producer` to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a `firstname` field of type string a `lastname` field of type string, an `age` field of type int and a `salary` field of type double.

```

${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic rethink-topic \
--property value.schema='{ "type": "record", "name": "User", "namespace": "com.
datamountaineer.streamreactor.connect.rethink"
, "fields": [ { "name": "firstName", "type": "string" }, { "name": "lastName", "type": "string" },
, "name": "age", "type": "int" }, { "name": "salary", "type": "double" } ] } '

```

Now the producer is waiting for input. Paste in the following:

```
{ "firstName": "John", "lastName": "Smith", "age": 30, "salary": 4830 }
```

Check for records in Rethink

Now check the logs of the connector you should see this:

```
INFO Received record from topic:person_rethink partition:0 and offset:0 (com.
↳datamountaineer.streamreactor.connect.rethink.sink.writer.rethinkDbWriter:48)
INFO Empty list of records received. (com.datamountaineer.streamreactor.connect.
↳rethink.sink.RethinkSinkTask:75)
```

Check for records in Rethink

Now stop the connector.

Features

The ReThinkDb Sink writes records from Kafka to RethinkDb.

The Sink supports:

1. Field selection - Kafka topic payload field selection is supported, allowing you to select fields written to RethinkDb.
2. Topic to table routing.
3. RowKey selection - Selection of fields to use as the row key, if none specified the topic name, partition and offset are used.
4. RethinkDB write modes.
5. Error policies for handling failures.

Kafka Connect Query Language

Kafka Connect Query Language found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The ReThink Sink supports the following:

```
<write mode> INTO <target table> SELECT <fields> FROM <source topic> <AUTOCREATE> <PK_
↳FIELD>
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA
INSERT INTO tableA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to tableB
INSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB

#Upsert mode, select all fields from topicC, auto create tableC and auto evolve, use_
↳field1 as the primary key
UPSERT INTO tableC SELECT * FROM topicC AUTOCREATE PK field1
```

Write Modes

The Sink support two write modes **insert** and **upsert** which map to RethinkDb's conflict policies, **insert** to **ERROR** and **upsert** to **REPLACE**.

Error Policies

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other expectations thrown by drivers.

Retry

Any error on write to the target database causes the `RetryIterable` exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.rethink.sink.max.retries` and the `connect.rethink.sink.retry.interval`.

Topic Routing

The Sink supports topic routing that allows mapping the messages from topics to a specific table. For example, map a topic called "bloomberg_prices" to a table called "prices". This mapping is set in the `connect.rethink.sink.kcql` option.

Example:

```
//Select all
INSERT INTO table1 SELECT * FROM topic1; INSERT INTO tableA SELECT * FROM topicC
```

Field Selection

The ReThink Sink supports field selection and mapping. This mapping is set in the `connect.rethink.sink.kcql` option.

Examples:

```
//Rename or map columns
INSERT INTO table1 SELECT lst_price AS price, qty AS quantity FROM topicA

//Select all
INSERT INTO table1 SELECT * FROM topic1
```

Tip: Check you mappings to ensure the target columns exist.

Auto Create Tables

The Sink supports auto creation of tables for each topic. This mapping is set in the `connect.rethink.sink.kcql` option.

A user specified primary can be set in the PK clause for the `connect.rethink.sink.kcql` option. Only one key is supported. If more than one is set only the first is used. If no primary keys are set the default primary key called `id` is used. The value for the default key is the topic name, partition and offset of the records.

```
#AutoCreate the target table
INSERT INTO table1 SELECT * FROM topic AUTOCREATE PK field1
```

Note: The fields specified as the primary keys must be in the SELECT clause or all fields must be selected

The Sink will try and create the table at start up if a schema for the topic is found in the Schema Registry. If no schema is found the table is created when the first record is received for the topic.

Configurations

`connect.rethink.sink.kcql`

Kafka connect query language expression. Allows for expressive topic to table routing, field selection and renaming. Fields to be used as the row key can be set by specifying the PK. The below example uses `field1` as the primary key.

- Data type : string
- Importance: high
- Optional : no

Examples:

```
INSERT INTO TABLE1 SELECT * FROM TOPIC1; INSERT INTO TABLE2 SELECT * FROM TOPIC2 PK_
↪field1
```

`connect.rethink.sink.host`

Specifies the rethink server.

- Data type : string
- Importance: high
- Optional : no

`connect.rethink.sink.port`

Specifies the rethink server port number.

- Data type : int
- Importance: high
- Optional : yes

`connect.rethink.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.rethink.sink.max.retries` option. The `connect.rethink.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: medium
- Optional: yes
- Default: RETRY

`connect.rethink.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.rethink.sink.error.policy` is set to **retry**.

- Type: string
- Importance: high
- Optional: yes
- Default: 10

`connect.rethink.sink.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.rethink.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: medium
- Optional: yes
- Default : 60000 (1 minute)

`connect.rethink.sink.batch.size`

Specifies how many records to insert together at one time. If the connect framework provides less records when it is calling the Sink it won't wait to fulfill this value but rather execute it.

- Type : int
- Importance : medium
- Optional: yes
- Defaults : 3000

Example

```
name=rethink-sink
connect.rethink.sink.db=localhost
connect.rethink.sink.port=28015
connector.class=com.datamountaineer.streamreactor.connect.rethink.sink.
↳ReThinkSinkConnector
tasks.max=1
topics=person_rethink
connect.rethink.sink.kcql=INSERT INTO TABLE1 SELECT * FROM person_rethink
```

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

The rethink Sink will automatically write and update the rethink table if new fields are added to the Source topic, if fields are removed the Kafka Connect framework will return the default value for this field, dependent of the compatibility settings of the Schema registry.

Deployment Guidelines

TODO

TroubleShooting

TODO

Kafka Connect VoltDB

A Connector and Sink to write events from Kafka to VoltDB. The connector used the built in stored procedures for inserts and upserts but requires the tables to be pre-created.

The Sink supports:

1. *The KCQL routing querying* - Kafka topic payload field selection is supported, allowing you to select fields written to VoltDB.
2. Topic to table routing via KCQL.
3. Voltdb write modes, upsert and insert via KCQL.
4. Error policies for handling failures.

Prerequisites

- Confluent 3.0.1
- VoltDB 6.4
- Java 1.8
- Scala 2.11

Setup

VoltDB Setup

Download VoltDB from [here](#)

Unzip the archive

```
tar -xzf voltdb-ent-*.tar.gz
```

Start VoltDB:

```
cd voltdb-ent-*
bin/voltdb create

Build: 6.5 voltdb-6.5-0-gdlfe3fa-local Enterprise Edition
Initializing VoltDB...

  _ _ _ _ _
| | / / _ _ / / / _ _ \ / _ _ )
| | / / _ _ \ / / _ _ / / / _ _ |
| | / / _ _ / / / _ _ / / / _ _ /
| _ _ \ _ _ / _ _ \ _ _ / _ _ /

-----

Connecting to VoltDB cluster as the leader...
Host id of this node is: 0
Starting VoltDB with trial license. License expires on Sep 11, 2016.
Initializing the database and command logs. This may take a moment...
WARN: This is not a highly available cluster. K-Safety is set to 0.
```

Confluent Setup

Follow the instructions [here](#).

Sink Connector QuickStart

We will start the connector in distributed mode. Each connector exposes a rest endpoint for stopping, starting and updating the configuration. We have developed a Command Line Interface to make interacting with the Connect Rest API easier. The CLI can be found in the Stream Reactor download under the `bin` folder. Alternatively the Jar can be pulled from our GitHub [releases](#) page.

Create Voltdb Table

At present the Sink doesn't support auto creation of tables so we need to login to VoltDb to create one. In the directory you extracted Voltdb start the `sqlcmd` shell and enter the following DDL statement. This creates a table called person.

```
create table person(firstname varchar(128), lastname varchar(128), age int, salary_
↪float, primary key (firstname, lastname));
```

```
bin ./sqlcmd
SQL Command :: localhost:21212
1> create table person(firstname varchar(128), lastname varchar(128), age int, salary_
↪float, primary key (firstname, lastname));
Command succeeded.
2>
```

Starting the Connector (Distributed)

Download, unpack and install the Stream Reactor. Follow the instructions [here](#) if you haven't already done so. All paths in the quickstart are based in the location you installed the Stream Reactor.

Start Kafka Connect in distributed mode by running the `start-connect.sh` script in the `bin` folder.

```
bin/start-connect.sh
```

Once the connector has started we can now use the `kafka-connect-tools` cli to post in our distributed properties file for VoltDB. If you are using the *dockers* you will have to set the following environment variable to for the CLI to connect to the Rest API of Kafka Connect of your container.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

```
bin/cli.sh create voltdb-sink < conf/voltdb-sink.properties

#Connector `voltdb-sink`:
name=voltdb-sink
connector.class=com.datamountaineer.streamreactor.connect.voltdb.VoltSinkConnector
max.tasks=1
topics=sink-test
connect.volt.connection.servers=localhost:21212
connect.volt.sink.kcql=INSERT INTO person SELECT * FROM sink-test
connect.volt.connection.password=
connect.volt.connection.user=
#task ids:
```

The `voltdb-sink.properties` file defines:

1. The name of the sink.
2. The Sink class.
3. The max number of tasks the connector is allowed to created.
4. The topics to read from (Required by framework)
5. The name of the voltdb host to connect to.
6. Username to connect as.
7. The password for the username.
8. *The KCQL routing querying.*

If you switch back to the terminal you started the Connector in you should see the VoltDb Sink being accepted and the task starting.

We can use the CLI to check if the connector is up but you should be able to see this in logs as-well.

```
#check for running connectors with the CLI
bin/cli.sh ps
voltdb-sink
```

```
[2016-08-21 20:31:36,398] INFO Finished starting connectors and tasks (org.apache.
↳ kafka.connect.runtime.distributed.DistributedHerder:769)
[2016-08-21 20:31:36,406] INFO
```

[illegible]


```

| | | | / _ | _ ) / _ | \ / _ \ | | | _ \ | _ ) / _ | | _ \ / _ ) _ ) / _
| | | | / ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( ( (
| | | | / \ | | | \ | | | \ | | | \ | | | \ | | | \ | | | \ | | | \ | | | \
by Stefan Bocutiu

| | | | | _ ( _ \ | | | | ( _ ) | |
| | | | | _ | _ \ \ | | | | \ \ | | | | _ \ | | | | / )
\ \ / _ \ | | | _ ) | | | | | \ \ \ \ | | | | _ \ | | | | / )
\ \ / | | | | | | | | / / | | ) _ _ ) ) | | | | | < (
\ \ \ \ / | | \ ) _ _ / | _ _ / ( _ _ / | | | | | | \ )
(com.datamountaineer.streamreactor.connect.voltdb.VoltSinkTask:44)
[2016-08-21 20:31:36,407] INFO VoltSinkConfig values:
    connect.volt.error.policy = THROW
    connect.volt.retry.interval = 60000
    connect.volt.sink.kcql = INSERT INTO person SELECT * FROM sink-test
    connect.volt.max.retires = 20
    connect.volt.connection.servers = localhost:21212
    connect.volt.connection.user =
    connect.volt.connection.password =
(com.datamountaineer.streamreactor.connect.voltdb.config.VoltSinkConfig:178)
[2016-08-21 20:31:36,501] INFO Settings:com.datamountaineer.streamreactor.connect.
↪ voltdb.config.VoltSettings$@34c34c3e (com.datamountaineer.streamreactor.connect.
↪ voltdb.VoltSinkTask:71)
[2016-08-21 20:31:36,565] INFO Connecting to VoltDB... (com.datamountaineer.
↪ streamreactor.connect.voltdb.writers.VoltConnectionConnectFn$:28)
[2016-08-21 20:31:36,636] INFO Connected to VoltDB node at: localhost:21212 (com.
↪ datamountaineer.streamreactor.connect.voltdb.writers.VoltConnectionConnectFn$:46)

```

Test Records

Now we need to put some records it to the `test_table` topics. We can use the `kafka-avro-console-producer` to do this.

Start the producer and pass in a schema to register in the Schema Registry. The schema has a `firstname` field of type string a `lastname` field of type string, an `age` field of type int and a `salary` field of type double.

```

${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
  --broker-list localhost:9092 --topic sink-test \
  --property value.schema='{ "type": "record", "name": "User", "namespace": "com.
↪ datamountaineer.streamreactor.connect.voltdb"
  , "fields": [{ "name": "firstName", "type": "string"}, { "name": "lastName", "type": "string"},
↪ { "name": "age", "type": "int"}, { "name": "salary", "type": "double"} ] }'

```

Now the producer is waiting for input. Paste in the following:

```
{ "firstName": "John", "lastName": "Smith", "age": 30, "salary": 4830 }
```

Check for records in VoltDb

Now check the logs of the connector you should see this:

```

[2016-08-21 20:41:25,361] INFO Writing complete (com.datamountaineer.streamreactor.
↪ connect.voltdb.writers.VoltDbWriter:61)
[2016-08-21 20:41:25,362] INFO Records handled (com.datamountaineer.streamreactor.
↪ connect.voltdb.VoltSinkTask:86)

```

In Voltdb sqlcmd terminal

```
SELECT * FROM PERSON;

FIRSTNAME  LASTNAME  AGE  SALARY
-----
John       Smith     30   4830.0

(Returned 1 rows in 0.01s)
```

Now stop the connector.

Features

The Sink supports:

1. Field selection - Kafka topic payload field selection is supported, allowing you to select fields written to VoltDB.
2. Topic to table routing.
3. Voltdb write modes, upsert and insert.
4. Error policies for handling failures.

Kafka Connect Query Language

Kafka Connect Query Language found here [GitHub repo](#) allows for routing and mapping using a SQL like syntax, consolidating typically features in to one configuration option.

The Voltdb Sink supports the following:

```
INSERT INTO <table> SELECT <fields> FROM <source topic>
UPSERT INTO <table> SELECT <fields> FROM <source topic>
```

Example:

```
#Insert mode, select all fields from topicA and write to tableA
INSERT INTO tableA SELECT * FROM topicA

#Insert mode, select 3 fields and rename from topicB and write to tableB
INSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB

#Upsert mode, select 3 fields and rename from topicB and write to tableB
UPSERT INTO tableB SELECT x AS a, y AS b and z AS c FROM topicB
```

This is set in the `connect.volt.sink.kcql` option.

Error Policies

The Sink has three error policies that determine how failed writes to the target database are handled. The error policies affect the behaviour of the schema evolution characteristics of the sink. See the schema evolution section for more information.

Throw

Any error on write to the target database will be propagated up and processing is stopped. This is the default behaviour.

Noop

Any error on write to the target database is ignored and processing continues.

Warning: This can lead to missed errors if you don't have adequate monitoring. Data is not lost as it's still in Kafka subject to Kafka's retention policy. The Sink currently does **not** distinguish between integrity constraint violations and or other exceptions thrown by drivers..

Retry

Any error on write to the target database causes the `RetryIterable` exception to be thrown. This causes the Kafka connect framework to pause and replay the message. Offsets are not committed. For example, if the table is offline it will cause a write failure, the message can be replayed. With the Retry policy the issue can be fixed without stopping the sink.

The length of time the Sink will retry can be controlled by using the `connect.hazelcast.sink.max.retries` and the `connect.hazelcast.sink.retry.interval`.

Topic Routing

The Sink supports topic routing that allows mapping the messages from topics to a specific table. For example, map a topic called "bloomberg_prices" to a table called "prices". This mapping is set in the `connect.volt.sink.kcql` option.

Example:

```
//Select all
INSERT INTO table1 SELECT * FROM topic1; INSERT INTO tableA SELECT * FROM topicC
```

Write Modes

The Sink supports both **insert** and **upsert** modes. This mapping is set in the `connect.volt.sink.export.mappings` option.

Insert

Insert is the default write mode of the sink.

Insert Idempotency

Kafka currently provides at least once delivery semantics. Therefore, this mode may produce errors if unique constraints have been implemented on the target tables. If the error policy has been set to NOOP then the Sink will discard the error and continue to process, however, it currently makes no attempt to distinguish violation of integrity constraints from other exceptions such as casting issues.

Upsert

The Sink support VoltDB upserts which replaces the existing row if a match is found on the primary keys.

Upsert Idempotency

Kafka currently provides at least once delivery semantics and order is a guaranteed within partitions.

This mode will, if the same record is delivered twice to the sink, result in an idempotent write. The existing record will be updated with the values of the second which are the same.

If records are delivered with the same field or group of fields that are used as the primary key on the target table, but different values, the existing record in the target table will be updated.

Since records are delivered in the order they were written per partition the write is idempotent on failure or restart. Redelivery produces the same result.

Configurations

`connect.volt.sink.kcql`

KCQL expression describing field selection and routes.

- Data type : string
- Importance : high
- Optional : no

`connect.volt.connection.servers`

Comma separated server[:port].

- Type : string
- Importance : high
- Optional : no

`connect.volt.connection.user`

The user to connect to the volt database.

- Type : string
- Importance : high
- Optional : no

`connect.volt.connection.password`

The password for the voltdb user.

- Type : string
- Importance : high
- Optional : no

`connect.volt.sink.error.policy`

Specifies the action to be taken if an error occurs while inserting the data.

There are three available options, **noop**, the error is swallowed, **throw**, the error is allowed to propagate and retry. For **retry** the Kafka message is redelivered up to a maximum number of times specified by the `connect.volt.sink.max.retries` option. The `connect.volt.sink.retry.interval` option specifies the interval between retries.

The errors will be logged automatically.

- Type: string
- Importance: high
- Default: throw

`connect.volt.sink.max.retries`

The maximum number of times a message is retried. Only valid when the `connect.volt.sink.error.policy` is set to `retry`.

- Type: string
- Importance: medium
- Optional: yes
- Default: 10

`connect.volt.sink.retry.interval`

The interval, in milliseconds between retries if the Sink is using `connect.volt.sink.error.policy` set to **RETRY**.

- Type: int
- Importance: medium
- Optional: yes
- Default : 60000 (1 minute)

`connect.volt.sink.batch.size`

Specifies how many records to insert together at one time. If the connect framework provides less records when it is calling the Sink it won't wait to fulfill this value but rather execute it.

- Type : int
- Importance : medium
- Optional: yes
- Defaults : 1000

Schema Evolution

Upstream changes to schemas are handled by Schema registry which will validate the addition and removal of fields, data type changes and if defaults are set. The Schema Registry enforces Avro schema evolution rules. More information can be found [here](#).

No schema evolution is handled by the Sink yet on changes in the upstream topics.

Deployment Guidelines

TODO

TroubleShooting

TODO

Tools



Helper tools and libraries for interacting with the components of the architecture.

Kafka Connect CLI

This is a tiny command line interface (CLI) around the [Kafka Connect REST Interface](#) to manage connectors. It is used in a git like fashion where the first program argument indicates the command: it can be one of [ps|get|rm|create|run].

The CLI is meant to behave as a good unix citizen: input from `stdin`; output to `stdout`; out of band info to `stderr` and non-zero exit status on error. Commands dealing with configuration expect or producedata in `.properties` style: `key=value` lines and comments start with a `#`.

```
kafka-connect-cli 0.5
Usage: kafka-connect-cli [ps|get|rm|create|run|status] [options] [<connector-name>]

  --help
    prints this usage text
  -e <value> | --endpoint <value>
    Kafka Connect REST URL, default is http://localhost:8083/

Command: ps
list active connectors names.

Command: get
get the configuration of the specified connector.

Command: rm
remove the specified connector.

Command: create
create the specified connector with the .properties from stdin; the connector cannot
↪ already exist.

Command: run
create or update the specified connector with the .properties from stdin.

Command: status
get connector and it's task(s) state(s).

    <connector-name>
        connector name
```

You can override the default endpoint by setting an environment variable `KAFKA_CONNECT_REST` i.e.

```
export KAFKA_CONNECT_REST="http://myserver:myport"
```

Requirements

- Java 1.8

To Build

```
gradle fatJar
```

Usage

Clone this repository, do a `mvn package` and run the jar in a way you prefer, for example with the provided `cli` shell script. The CLI can be used as follows.

Get Active Connectors

Command: `ps`

Example:

```
$ ./cli ps
twitter-source
```

Get Connector Information

Command: `get`

Example:

```
$ ./cli get twitter-source
#Connector `twitter-source`:
name=twitter-source
tasks.max=1

(snip)

track.terms=test
#task ids: 0
```

Delete a Connector

Command: `rm`

Example:

```
$ ./cli rm twitter-source
```

Create a New Connector

The connector cannot already exist.

Command: `create`

Example:

```
$ ./cli create twitter-source <twitter.properties
#Connector `twitter-source`:
name=twitter-source
tasks.max=1

(snip)

track.terms=test
#task ids: 0
```

Create or Update a Connector

Either starts a new connector if it did not exist, or update an existing connector.

Command: run

Example:

```
$ ./cli run twitter-source <twitter.properties
#Connector `twitter-source`:
name=twitter-source
tasks.max=1

(snip)

track.terms=test
#task ids: 0
```

Schema Registry CLI

This repository contains a CLI and Go client for the REST API of Confluent's Kafka Schema Registry.

CLI

To install the CLI, assuming a properly setup Go installation, do:

```
go get github.com/datamountaineer/schema-registry/schema-registry-cli
```

After that, the CLI is found in `$GOPATH/bin/schema-registry-cli`. Running `schema-registry-cli` without arguments gives:

```
A command line interface for the Confluent schema registry

Usage:
  schema-registry-cli [command]

Available Commands:
  add          registers the schema provided through stdin
  exists       checks if the schema provided through stdin exists for the subject
  get          retrieves a schema specified by id or subject
  subjects     lists all registered subjects
  versions     lists all available versions

Flags:
  -h, --help          help for schema-registry-cli
```



```
-e, --url string    schema registry url (default "http://localhost:8081")
-v, --verbose       be verbose
```

Use `"schema-registry-cli [command] --help"` **for** more information about a command.

Client

The documentation of the package can be found [\[here\]\(https://godoc.org/github.com/datamountaineer/schema-registry\)](https://godoc.org/github.com/datamountaineer/schema-registry).

Kafka Socket Streamer

Akka Http with Reactive Kafka to stream topics to clients via Web sockets and Server Send Events.

This is test and not yet intended for any serious use yet.

Prerequisites

- Confluent Platform 3.0.1
- Scala 2.11.7

Setup

Confluent Setup

Follow the instructions [here](#).

QuickStart

The socket streamer pushes events out from Kafka to clients via websockets or server send events. Two different endpoints are available. But first we need some data in Kafka. Start the console producer and send some events in:

```
${CONFLUENT_HOME}/bin/kafka-avro-console-producer \
--broker-list localhost:9092 --topic socket_streamer \
--property value.schema='{ "type": "record", "name": "User", "namespace": "com.
→ datamountaineer.streamreactor.connect.redis"
, "fields": [{ "name": "firstName", "type": "string"}, { "name": "lastName", "type": "string"},
→ { "name": "age", "type": "int"},
{ "name": "salary", "type": "double"} ] }'
```

Paste the following in at the console producer:

```
{ "firstName": "John", "lastName": "Smith", "age": 30, "salary": 4830 }
{ "firstName": "Max", "lastName": "Power", "age": 30, "salary": 1000000 }
```

Now start the socket streamer. We need to set some configurations first. The socket-streamer uses Typesafe's configuration loader so we can create a file called `application.conf` and add the following.

```
system-name = "streamreactor-socket-streamer"
port = 8080

kafka {
    bootstrap-servers = "localhost:9092"
    zookeeper-servers = "localhost:2181"
    schema-registry-url = "http://localhost:8081"
}
```

To use the `application.conf` file, set its location as a Java property when starting the application like this `-Dconfig.file=path_to_file/application.conf`. We have included a start script in the `bin` folder of the Stream Reactor install.

To start the socket streamer:

```
bin/start-socket-streamer

2016-05-12 15:57:39,712 INFO [main] [c.d.s.s.Main$] [delayedEndpoint$com
↳ $datamountaineer$streamreactor$socketstreamer$Main$1:32]

by Andrew Stevenson

2016-05-12 15:57:39,716 INFO [main] [c.d.s.s.Main$] [delayedEndpoint$com
↳ $datamountaineer$streamreactor$socketstreamer$Main$1:49]
System name      : streamreactor-socket-streamer
Kafka brokers   : localhost:9092
Zookeepers      : localhost:2181
Schema registry : http://localhost:8081
Listening on port : 8080
```

Now lets have the socket streamer push using server send event by simply calling curl:

```
curl 'http://localhost:8080/api/kafka/sse?query=SELECT+%2A+FROM+socket-
↳streamer+WITHFORMAT+JSON+WITHGROUP+test'

data:{"value":{"firstName\":"John", "lastName\":"Smith", "age\":" 30, \
↳"salary\":" 4830.0}}
data:{"value":{"firstName\":"Max", "Power\":"Jones", "age\":" 30, "salary\
↳": 1000000}}
data:{"timestamp":"Thu May 12 16:42:02 CEST 2016", "system":"streamreactor-socket-
↳streamer", "message":"heartbeat"}
```

For websockets, install a websocket client, for example [Dark WebSocket Terminal](#). Start it and connect to the websocket endpoint.

Note: Dark Terminal, for some reason, needs a extra whitespace at the end of the connection url to work.

```
11:50:58 command: /connect ws://localhost:8080/ws/topics?topic=socket_streamer&consumergroup=testcgws
11:50:58 system: connection established, ws://localhost:8080/ws/topics?topic=socket_streamer&consumergroup=testcgws
11:50:59 received: {"value":{"firstName": "John", "lastName": "Smith", "age": 30, "salary": 4830.0}}
11:50:59 received: {"value":{"firstName": "Max", "lastName": "Power", "age": 30, "salary": 1000000.0}}
```

```
command: curl 'http://localhost:8080/api/kafka/ws?query=SELECT+%2A+FROM+socket-
↳streamer+WITHFORMAT+JSON+WITHGROUP+test'
system: connection established, ws://localhost:8080/ws/topics?topic=redis-topic&
↳consumergroup=testcgws
received: {"value":{"firstName": "John", "lastName": "Smith", "age": 30,
↳"salary": 4830.0}}
```

Features

1. Web Sockets
2. Server Send Events
3. Limited SQL support
4. Consumer Group Offset control
5. Column selection
6. Sample rows
7. Sliding windows

Configurations

Endpoints

GET /api/kafka/ws?query=SELECT [*|columns] FROM [TOPIC_NAME] WITHFORMAT JSON|AVRO|BINARY [T

WebSocket example request

```
GET /api/kafka/sse?query=SELECT+%2A+FROM+socket-
↳streamer+WITHFORMAT+JSON+WITHGROUP+test
```

Stream via Web Sockets the socket-streamer topic with consumer group test with
↳format json.

GET /sse/topics?topic=<topic_name>&consumergroup=<consumergroup>
Send Server Events example request

```
GET /api/kafka/sse?query=SELECT+%2A+FROM+socket-
↳streamer+WITHFORMAT+JSON+WITHGROUP+test
```

Stream via Send Server Events the socket-streamer topic with consumer group test with format json.

Deployment Guidelines

TODO

TroubleShooting

TODO

Fast Data UI's (Landoop)

Landoop has a number of UI's available to visually data in Kafka and schemas in the Schema Registry.

You can either build from their GitHub repo or install and run the docker images.

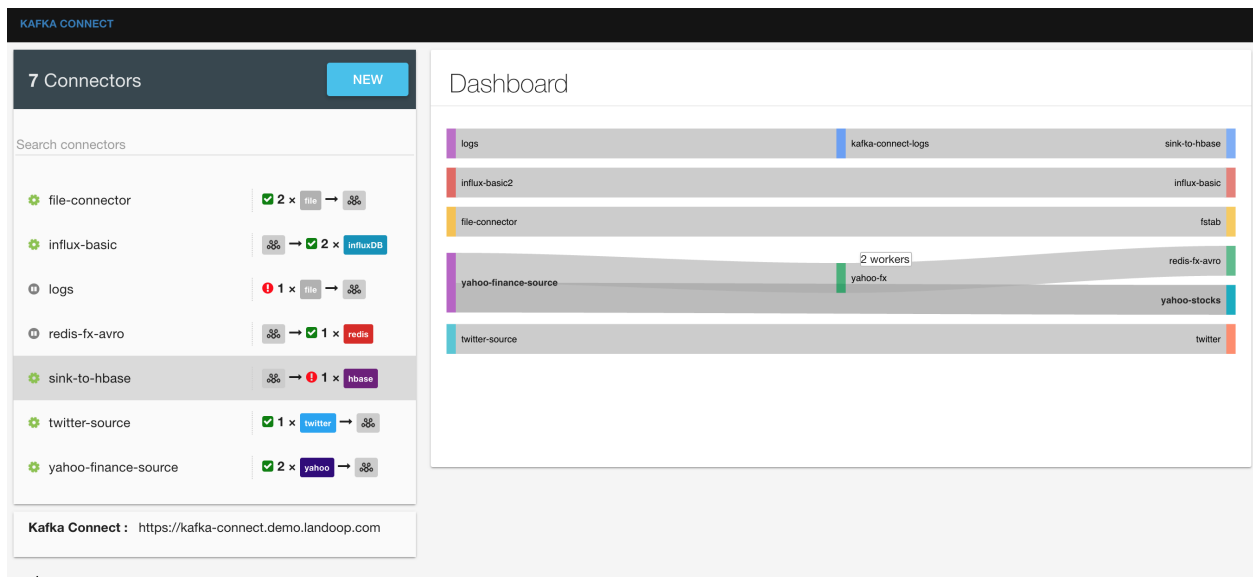
- [Kafka Topics Browser](#)
- [Schema Registry](#)
- [Kafka Connect](#)

Kafka Connect UI

The Kafka Connect UI lets you:

- Visualise your connect cluster sink & sources.
- Create new connectors with few clicks.
- Update & Delete connectors configuration.
- View workers tasks health & failures.

For the Connect UI please contact [Landoop](#).



Add new connectors.

7 Connectors
NEW

Search connectors

file-connector
2 x file →

influx-basic
→ 2 x influxDB

logs
1 x file →

redis-fx-avro
→ 1 x redis

sink-to-hbase
→ 1 x hbase

twitter-source
1 x twitter →

yahoo-finance-source
2 x yahoo →

Kafka Connect : <https://kafka-connect.demo.landoop.com>

New Connector

Search

Sources

Sinks

twitter
Subscribe to feeds using the Twitter API and stream data into a kafka topic

file
Read files and stream data into Kafka Topics

Blockchain
A Connector to hook into the live streaming providing a real time feed for new bitcoin blocks and transactions provided by www.blockchain.info

bloomberg
a source connector to subscribe to Bloomberg feeds via the Bloomberg labs open API and write to Kafka.

cassandra
The Cassandra source connector allows you to extract entries from Cassandra with the CQL driver and write them into a Kafka topic

yahoo
Subscribe to Yahoo Finance API and stream data into a kafka topic

jdbc

cassandra
The Cassandra sink connector allows you to write data from a Kafka topic into a Cassandra table.

influxDB
The InfluxDB sink connector allows you to write data from a Kafka topic into a influxDB table.

elastic
A Connector and Sink to write events from Kafka to Elastic Search using Elastic4s client.

hbase
The HBase sink connector allows you to write data from a Kafka topic into HBase

redis
The Redis sink connector allows you to write data from a Kafka topic into Redis

kudu
The Kudu sink connector allows you to write data from a Kafka topic into Kudu. Kudu 0.9 or newer

jms
The JMS sink connector allows you to extract entries

7 Connectors
NEW

Search connectors

file-connector
2 x file →

influx-basic
→ 2 x influxDB

logs
1 x file →

redis-fx-avro
→ 1 x redis


sink-to-hbase
→ 1 x hbase

twitter-source
1 x twitter →

yahoo-finance-source
2 x yahoo →

Kafka Connect : <https://kafka-connect.demo.landoop.com>

New Connector (Sink): kudu


kudu
The Kudu sink connector allows you to write data from a Kafka topic into Kudu. Kudu 0.9 or newer
class: `com.datamountaineer.streamreactor.connect.kudu.sink.KuduSinkConnector`

☐ Show Hints

BASIC INFO
ADVANCED OPTIONS
FINISH

Kafka-Connect Query Language (KCQL)

KCQL
INSERT INTO TABLE2 SELECT field1, field2, field3 as renamedField FROM TOPIC2

Kudu configuration

Master Server
quickstart

The schema registry url
<https://schema-registry.demo.landoop.com>

Error policy
noop

« BACK
NEXT »

Kafka Topic Browser

The Kafka Topic Browser allows you to look into topic without having to write code or use the command line console consumers.

Supported features are:

- Find topics & browse topic data (kafka messages)
- View topic metadata
- View topic configuration
- Download data

The screenshot displays the Kafka Topic Browser interface. On the left, a sidebar titled '6 Topics' lists several topics: `_schemas`, `connect-status`, `kafka-connect-logs`, `twitter`, `yahoo-fx` (selected), and `yahoo-stocks`. Each topic shows '1 Replication x 1 Partition'. Below the list, it indicates 'Kafka Rest : https://kafka-rest-proxy.demo.landoop.com' and 'Kafka Brokers : 1'. A note states 'Default configuration has been amended'. On the right, the 'yahoo-fx' topic is selected, showing a 'Filter' section and a table with columns 'TOPIC', 'TABLE', and 'RAW DATA'. The table has sub-columns 'Offset', 'Partition', 'Key', and 'Value'. The data shows a sequence of messages with offsets 0 through 9, all in partition 0. The 'Value' column displays JSON-like structures, with the second message (offset 1) expanded to show 'symbol: EURGBP=X' and 'price: 0.8581'. At the bottom, a pagination bar shows '1 2 3 4 5 6 7 ... 81'.

Schema Registry Browser

The Schema Registry is an integral part of the Kafka Streaming Platform.

Schema Registry provides a serving layer for your metadata. It provides a RESTful interface for storing and retrieving Avro schemas. It stores a versioned history of all schemas, provides multiple compatibility settings and allows evolution of schemas according to the configured compatibility setting. It provides serializers that plug into Kafka clients that handle schema storage and retrieval for Kafka messages that are sent in the Avro format.

Our web tool for schema registry allows:

- Visibility of schemas and their versions in the topics
- Schema validation
- Ability to add new schemas
- Ability to track changes with a graphical diff.

KAFKA TOPICS

6 Topics System Topics

Search topics

* _schemas
1 Replication x 1 Partition

connect-status
1 Replication x 1 Partition

kafka-connect-logs
1 Replication x 1 Partition

twitter
1 Replication x 1 Partition

* yahoo-fx
1 Replication x 1 Partition

yahoo-stocks
1 Replication x 1 Partition

json

binary

..

avro

avro

avro

Default configuration has been amended

Kafka Rest : <https://kafka-rest-proxy.demo.landoop.com>

Kafka Brokers : 1

yahoo-fx

Filter

TOPIC

TABLE

RAW DATA

offset	partition	key	value	
			symbol.string	price.double
0	0		USDGBP=X	0.7671
1	0		EURGBP=X	0.8581
2	0		USDGBP=X	0.767
3	0		EURGBP=X	0.8583
4	0		USDGBP=X	0.767
5	0		EURGBP=X	0.8583
6	0		USDGBP=X	0.767
7	0		EURGBP=X	0.8583
8	0		USDGBP=X	0.767
9	0		EURGBP=X	0.8583

<

1

2

3

4

5

6

7

...

81

>

Powered by Landoop

Install

For docker, pull the images:

```
docker pull landoop/kafka-topics-ui
docker pull landoop/schema-registry-ui
```

To run

```
docker run --rm -it -p 8000:8000 \
-e "SCHEMAREGISTRY_URL=http://confluent-schema-registry-host:port" \
landoop/schema-registry-ui

docker run --rm -it -p 8000:8000 \
-e "KAFKA_REST_PROXY_URL=http://kafka-rest-proxy-host:port" \
landoop/kafka-topics-ui

#Start both in one docker
#docker run --rm -it -p 8000:8000 \
# -e "SCHEMAREGISTRY_UI_URL=http://confluent-schema-registry-host:port" \
# -e "KAFKA_REST_PROXY_URL=http://kafka-rest-proxy-host:port" \
# landoop/kafka-topics-ui
```

Your schema-registry service will need to allow CORS (!!)

To do that, and in `/opt/confluent-3.0.0/etc/schema-registry/schema-registry.properties`

```
access.control.allow.methods=GET,POST,OPTIONS
access.control.allow.origin=*
```

Python Schema Client & Serializers and Deserializers

DataMountaineer recommends all payloads in Kafka are Avro. Schema Registry provides a serving layer for your metadata. It provides a RESTful interface for storing and retrieving Avro schemas. It stores a versioned history of all schemas, provides multiple compatibility settings and allows evolution of schemas according to the configured compatibility setting. It provides serializers that plug into Kafka clients that handle schema storage and retrieval for Kafka messages that are sent in the Avro format.

Confluent have provided a [Python](#) client built on [librdkafka](#). To integrated with the Schema Registry we need to encode/decode using the *schema_id*. This fork is from [Versign](#), we upgraded it to Python 3.5 for a client.

Install

Run `setup.py` from the source root

```
python setup.py install
```

or via `pip`

```
pip3 install datamountaineer-schemaregistry
```

Usage

```
from datamountaineer.schemaregistry.client import SchemaRegistryClient
from datamountaineer.schemaregistry.serializers import MessageSerializer, Util

# Initialize the client
client = SchemaRegistryClient(url='http://registry.host')
# encode/decode a record to put onto kafka
serializer = MessageSerializer(client)
```

Schema Operations

```
# register a schema for a subject
schema_id = client.register('my_subject', avro_schema)

# fetch a schema by ID
avro_schema = client.get_by_id(schema_id)

# get the latest schema info for a subject
schema_id, avro_schema, schema_version = client.get_latest_schema('my_subject')

# get the version of a schema
schema_version = client.get_version('my_subject', avro_schema)

# Compatibility tests
is_compatible = client.test_compatibility('my_subject', another_schema)

# One of NONE, FULL, FORWARD, BACKWARD
new_level = client.update_compatibility('NONE', 'my_subject')
current_level = client.get_compatibility('my_subject')
```


Encoding to write back to Kafka. Encoding by id is the most efficient as it avoids an extra trip to the Schema Registry to lookup the schema id.

Writing Messages

Encode by schema only.

```
# use an existing schema and topic
# this will register the schema to the right subject based
# on the topic name and then serialize
encoded = serializer.encode_record_with_schema('my_topic', avro_schema, record)

from confluent_kafka import Producer

p = Producer({'bootstrap.servers': 'mybroker,mybroker2'})
p.produce('my_topic', encoded)
p.flush()
```

Reading Messages

```
# decode a message from kafka

from confluent_kafka import Consumer, KafkaError

c = Consumer({'bootstrap.servers': 'mybroker', 'group.id': 'mygroup',
    'default.topic.config': {'auto.offset.reset': 'smallest'}})
c.subscribe(['my_topic'])
running = True
while running:
    msg = c.poll()
    if not msg.error():
        decoded = serializer.decode_message(message)
    elif msg.error().code() != KafkaError._PARTITION_EOF:
        print(msg.error())
        running = False
c.close()
```

FAQS

I have JSON. Can I still use DataMountaineer Sink Connectors?

Kafka Connect has two converters for both the key and payload from Kafka. These are Json and Avro, the Json converter is part of the Kafka distribution and the Avro converter from Confluents schema registry. These converters convert the records in Kafka to SinkRecords, most of our Sinks rely of the data in Kafka being Avro and written with the Confluent Avro Serializers. This is *best practice*. This allows the Connectors to receive SinkRecords with schemas for the payloads to mapping, filtering can take place based on the [The KCQL routing querying](#) provided.

However, it is possible sink Json messages from Kafka with some of our Sinks by using the JsonConverter from Kafka. If your Json messages have a *schema* field the converter will deliver the records to the Sink with a schema. If no *schema* tag is present the records will be delivered with a schema of type `SCHEMA.String`.

We are currently working on support for schemaless json records.

Can I run on multiple nodes?

Yes, Kafka Connect has two modes, standalone and distributed. Both allow for scaling by setting the *max.tasks* property.

In distributed mode each worker joins a Connect cluster defined in the *etc/schema-registry/connect-avro-distributed.properties* file that is part of the Confluent distribution. Within this file a property called `group.id` controls this.

```
# The group ID is a unique identifier for the set of workers that form a single Kafka_
↪Connect
# cluster
group.id=connect-cluster
```

ClassNotFoundException

The `start-connect.sh` in the Stream Reactor download adds all the jars from the *libs* folder to the CLASSPATH automatically. If you are not using this start script it is more than likely you have not add the relevant Connector jar to the CLASSPATH properly.

Explicit add the relevant jar to the classpath and restart Kafka Connect

```
export CLASSPATH=my_connector.jar
```

You can ask a running instance of Kafka Connect what Connector classes are on the classpath with the *CLI* `<cli>`

```
bin/cli.sh loaded
```

Guava version

The Elastic Search and HBase use different versions of guava, as does the Hive libraries supplied by Confluent with the HDFS Connector. This can cause version clashes.

Explicit add the relevant jar to the classpath and restart Kafka Connect. This sets our jars first and should solve the issues.

```
#For Elastic
#export CLASSPATH=kafka-connect-elastic-0.2.3-3.0.1.jar

#For HBASE
#export CLASSPATH=kafka-connect-hbase-0.2.3-3.0.1.jar
```

Redis authentication

If your redis server is requiring the connection to be authenticated you will need to provide an extra setting:

```
connect.redis.sink.connection.password=$REDIS_PASSWORD
```

Don't set the value to empty if no password is required.

InfluxDb Port already in use

InfluxDB starts an Admin web server listening on port 8083 by default. For this quickstart this will collide with Kafka Connects default port of 8083. Since we are running on a single node we will need to edit the InfluxDB config.

```
#create config dir
sudo mkdir /etc/influxdb
#dump the config
influxd config > /etc/influxdb/influxdb.generated.conf
```

Now change the following section to a port 8087 or any other free port.

```
[admin]
enabled = true
bind-address = ":8087"
https-enabled = false
https-certificate = "/etc/ssl/influxdb.pem"
```

How get multiple worker on different hosts to for a Connect Cluster

For workers to join a Connect cluster, set the *group.id* in the *\$CONFLUENT_HOME/etc/schema-registry/connect-avro-distributed.properties* file.

```
# The group ID is a unique identifier for the set of workers that form a single Kafka_
↪Connect
# cluster
group.id=connect-cluster
```


/api

```
GET /api/kafka/ws?query=SELECT
    [*|columns] FROM
    [TOPIC_NAME] WITHFORMAT
    JSON|AVRO|BINARY [WITHGROUP
    $YOUR_CONSUMER_GROUP]
    [WITHPARTITION
    (partition), [(partition,
    offset)] [SAMPLE
    $RECORDS_NUMBER EVERY
    $SLIDE_WINDOW], 143
```

/sse

```
GET /sse/topics?topic=<topic_name>&consumergroup=<consumergroup>,
    143
```