

Kafka Streams Application Reset Tool

The application reset tool allows to reset a Kafka Streams application's internal state, such that it can reprocess its input data from scratch.

Available since: Apache Kafka 0.10.0.1

Table of Contents

- [Usage](#)
 - [Step 1 of 2: Local reset \(for each application instance\)](#)
 - [Step 2 of 2: Global reset \(for the overall application\)](#)
- [What the application reset tool does](#)

Usage



Use this tool only if the application is fully stopped:

You must use this script only if *no instances of your application are running*. Otherwise the application may enter an invalid state, crash, or produce incorrect results.

Use this tool with care and double-check its parameters:

If you provide wrong parameter values (e.g. typos in `application.id`) or specify parameters inconsistently (e.g. specifying the wrong input topics for the application), this tool might invalidate the application's state or even impact other applications, consumer groups, or Kafka topics of your Kafka cluster.

Step 1 of 2: Local reset (for each application instance)



This local reset must be done on each machine on which an instance of the application has been run before.

Call `KafkaStreams#cleanup()` in your application code before `KafkaStreams#start()` to perform a local reset.

Step 2 of 2: Global reset (for the overall application)



Unlike step 1 this step needs to be performed only once, and it does not matter on which machine the application reset tool is being executed.

Run the application reset tool:

```
$ bin/kafka-streams-application-reset.sh

Option (* = required)      Description
-----
* --application-id <id>    The Kafka Streams application ID (application.id)
--bootstrap-servers <urls> Comma-separated list of broker urls with format: HOST1:PORT1,HOST2:PORT2
                           (default: localhost:9092)
--intermediate-topics <list> Comma-separated list of intermediate user topics
--input-topics <list>       Comma-separated list of user input topics
--zookeeper <url>          Format: HOST:PORT
                           (default: localhost:2181)
```

Note: *Intermediate topics* are all user-created topics that are used both as input and as output topics within a single Kafka Streams application (e.g. topics used in `through()`).

What the application reset tool does

- For any specified *input topics*:

- Reset the application's committed consumer offsets to "beginning of the topic" for all partitions (for consumer group `application.id`).
- For any specified *intermediate topics*:
 - Skip to the end of the topic, i.e. set the application's committed consumer offsets for all partitions to each partition's `logSize` (for consumer group `application.id`).
- For any *internal topics*:
 - Delete the internal topic (this will also delete committed the corresponding committed offsets).