

KIP-379: Multiple Consumer Group Management

- Status
- Motivation
- Public Interfaces
- Proposed Changes
- Compatibility, Deprecation, and Migration Plan
- Rejected Alternatives

Status

Current state: Accepted

Discussion thread: [DISCUSS] KIP-379: Multiple Consumer Group Management

Voting thread: [VOTE] KIP-379: Multiple Consumer Group Management

JIRA:  Unable to render Jira issues macro, execution error.

PULL REQUEST: #5726

votes

Non-binding votes: 4 (Gwen Shapira, Rianne Dolan, Jason Gustafson, Vahid Hashemian)
Binding votes: 1 (Damian Guy)

Motivation

- 1) Running a set of commands like `--describe`, `--delete` or `--reset-offsets` for every single consumer group appears to be quite a tedious, time and memory consuming task. Each time you have to start a separate JVM to perform a query for a single group or spend some time writing your own client code.
- 2) Quite often one would like to perform operations on all consumer groups, thus, it would be great to have something like an `--all-groups` option, similar to `--all-topics`.
- 3) In other cases, you would like to perform operations on a certain subset of consumer groups, thus, it would be great to have something like a repeatable `--group` option for each group.
- 4) Performing a Describe query on all consumer groups with a combination of options `--describe` and `--state` would allow us to grep consumer groups by their state (like `Stable, Dead, Empty, PreparingRebalance, NonExisting`). The only thing we have to do is to extend the "results table" with a GROUP column.
- 5) New CSV file format to export consumer offsets for multiple groups.

Public Interfaces

- 1) Providing the command line tool **kafka-consumer-groups.sh** with support of repeating **--group** parameter for operations --describe, --delete, --reset-offsets:

Usage example:

```
bin/kafka-consumer-groups.sh --bootstrap-server :9092 --describe      --group group-1 --group group-2 --group  
group-3  
bin/kafka-consumer-groups.sh --bootstrap-server :9092 --delete        --group group-1 --group group-2 --group  
group-3  
bin/kafka-consumer-groups.sh --bootstrap-server :9092 --reset-offsets --group group-1 --group group-2 --group  
group-3 --topic test
```

2) Providing the command line tool **kafka-consumer-groups.sh** with an optional **--all-groups** parameter for operations **--describe**, **--delete**, **--reset-offsets**:

parameter name	defaults to	type	usage
--all-groups	inactive	flag	optional

Usage example:

```
bin/kafka-consumer-groups.sh --bootstrap-server :9092 --describe      --all-groups
bin/kafka-consumer-groups.sh --bootstrap-server :9092 --delete       --all-groups
bin/kafka-consumer-groups.sh --bootstrap-server :9092 --reset-offsets --all-groups --topic test
```

3) Extending data output table format by adding "**GROUP**" column for Describe queries:

--describe:

GROUP	TOPIC	PARTITION	CURRENT-OFFSET	LOG-END-OFFSET	LAG	CONSUMER-ID	HOST	CLIENT-ID
-------	-------	-----------	----------------	----------------	-----	-------------	------	-----------

--describe --members:

GROUP	CONSUMER-ID	HOST	CLIENT-ID	#PARTITIONS
-------	-------------	------	-----------	-------------

--describe --members --verbose

GROUP	CONSUMER-ID	HOST	CLIENT-ID	#PARTITIONS	ASSIGNMENT
-------	-------------	------	-----------	-------------	------------

--describe --state

GROUP	COORDINATOR (ID)	ASSIGNMENT-STRATEGY	STATE	#MEMBERS
-------	------------------	---------------------	-------	----------

4) Add new CSV file to support multiple consumer groups. Thus, we will have to update table format by adding a consumer group column. Backwards compatibility with old CSV file format should be preserved.

New CSV file format:

GROUP	TOPIC	PARTITION	OFFSET

Old CSV file format for a single group (leaving as is):

TOPIC	PARTITION	OFFSET

Proposed Changes

- Providing the command line tool **kafka-consumer-groups.sh** with support of repeating **--group** parameter for operations **--describe**, **--delete**, **--reset-offsets**
- Providing the command line tool **kafka-consumer-groups.sh** with an optional **--all-groups** parameter for operations **--describe**, **--delete**, **--reset-offsets**
- Extending data output table format by adding "**GROUP**" column for all **--describe** queries
- Update CSV file to support multiple consumer groups
- We will also have to make sure CSV generation is done properly since there are no restrictions on consumer group names and symbols like commas and quotes are allowed. That is, CSV file generation with export/import of consumer offsets needs some rework

Compatibility, Deprecation, and Migration Plan

- Positive impact on existing users
- We are not changing behavior but adding new features
- Support backwards compatibility with old CSV file format for a single group case

Rejected Alternatives

None.