

KIP-175: Additional '--describe' views for ConsumerGroupCommand


- [Status](#)
- [Motivation](#)
- [Public Interfaces](#)
- [Proposed Changes](#)
 - [--describe](#)
 - [--describe --members](#) and [--describe --members --verbose](#)
 - [--describe --offsets](#)
 - [--describe --state](#)
- [Compatibility, Deprecation, and Migration Plan](#)
- [Rejected Alternatives](#)

Status

Current state: *Adopted*

Discussion thread: [here](#)

JIRA:

 Unable to render Jira issues macro, execution error.

Please keep the discussion on the mailing list rather than commenting on the wiki (wiki discussions get unwieldy fast).

Motivation

As far as describing a consumer group goes, the `ConsumerGroupCommand` currently lists all members of the group along with associated topic partitions and their committed offsets all in one view. For example:

```
$ bin/kafka-consumer-groups.sh --bootstrap-server localhost:9092 --describe --group my-group
Note: This will only show information about consumers that use the Java consumer API (non-ZooKeeper-based
consumers).
TOPIC          PARTITION  CURRENT-OFFSET  LOG-END-OFFSET  LAG      CONSUMER-
ID          HOST          CLIENT-ID
my-topic       0          16577           33154           16577    consumer1-81fe8a2a-0b29-
4964-ald4-4741582132bb /127.0.0.1    consumer1
my-topic       1          16578           33155           16577    consumer1-81fe8a2a-0b29-
4964-ald4-4741582132bb /127.0.0.1    consumer1
my-topic       2          16577           33155           16578    consumer2-db6f12f2-e1de-
4bb7-93ec-6170fef8830 /127.0.0.1    consumer2
my-second-topic 0          33023           44543           11520    consumer7-c1d056fa-eabc-
42ed-b557-c848f747ec49 /127.0.0.1    consumer7
-              -          -              -              -        consumer8-6078de6a-af63-
4c7a-8bad-f03cb3606075 /127.0.0.1    consumer8
my-other-topic 1          7901           8652           751      -
-              -          -              -              -        -
my-other-topic 0          7902           8654           752      -
-              -          -              -              -        -
```

The information included in the command output includes:

- all topic partitions the group is consuming or has consumed from (including those without an active consumer)
- current offset, log end offset, and the lag for each topic partition in the group
- all active members of the group (including those without an assigned topic partition)

This KIP suggests breaking down the command to cover some of the information pieces above separately using additional switches. There are a couple of reasons behind this improvement:

- The above output, with all the missing column values, looks weird and could be difficult to interpret. For a big group, where there are many consumers without assigned partition, the command will output a long list, while the user may only be interested in, for example, active members only.

- A user is typically interested in one of the information pieces above when issuing the command: either offset information, or members, or topic partitions associated with the group.

As part of this KIP, additional information about the group will also be generated that is currently missing. For example, generation id of the group, or the assignment policy used by each member.

Public Interfaces

The change required for this includes updating the `ConsumerGroupCommand` tool:

1. The output of `--describe` will change to return one row per topic partition in the group. This means that group members with no assigned partition will not be present in the default output.
2. A new switch `--members` will be introduced, that can only be provided along with `--describe`, to return a list of all active members of the group. For each member the number of assigned topic partitions (without any offset information) and the assignment strategy used by the member are also returned. Additional information (such as member subscription and assignment) can also be provided with a `--members --verbose` option.
3. A new switch `--offsets` will be introduced, that can only be provided along with `--describe`, that returns the same default output of `--describe`. i.e., this will be the default switch that can be additionally used with `--describe`.
4. A new switch `--state` will be introduced, that can only be provided along with `--describe`, to return group-level status information that is returned from the `DescribeGroups` API call.

Proposed Changes

To describe the proposed changes let us revisit the example above. This KIP proposes to change the `--describe` option of the `ConsumerGroupCommand` according to the following:

`--describe`

This option of the command returns offset and consumer information for each topic partition in the group. This is how the output of the example above would look like:

```
$ bin/kafka-consumer-groups.sh --bootstrap-server localhost:9092 --describe --group my-group
Note: This will only show information about consumers that use the Java consumer API (non-ZooKeeper-based consumers).
```

TOPIC ID	PARTITION	CURRENT-OFFSET	LOG-END-OFFSET	LAG	CONSUMER-CLIENT-ID
my-topic	0	16577	33154	16577	consumer1-81fe8a2a-0b29-4964-ald4-4741582132bb
my-topic	1	16578	33155	16577	consumer1-81fe8a2a-0b29-4964-ald4-4741582132bb
my-topic	2	16577	33155	16578	consumer2-db6f12f2-e1de-4bb7-93ec-6170fefa8830
my-second-topic	0	33023	44543	11520	consumer7-c1d056fa-eabc-42ed-b557-c848f747ec49
my-other-topic	1	7901	8652	751	-
my-other-topic	0	7902	8654	752	-

Note that consumers without any assigned partitions will no longer appear in the listing.

`--describe --members` and `--describe --members --verbose`

A new option `--members` is introduced that is allowed with `--describe` only and returns info about active members of the group along with their assigned topic partitions:

```
$ bin/kafka-consumer-groups.sh --bootstrap-server localhost:9092 --describe --group my-group --members
Note: This will only show information about consumers that use the Java consumer API (non-ZooKeeper-based consumers).
```

CONSUMER-ID	HOST	CLIENT-ID	#PARTITIONS
consumer1-81fe8a2a-0b29-4964-ald4-4741582132bb	/127.0.0.1	consumer1	2
consumer2-db6f12f2-e1de-4bb7-93ec-6170fefa8830	/127.0.0.1	consumer2	1
consumer7-c1d056fa-eabc-42ed-b557-c848f747ec49	/127.0.0.1	consumer7	1

There is an additional `--verbose` switch that can be used along with this option to include additional details about each group member:

```
$ bin/kafka-consumer-groups.sh --bootstrap-server localhost:9092 --describe --group my-group --members --verbose
Note: This will only show information about consumers that use the Java consumer API (non-ZooKeeper-based consumers).
```

CONSUMER-ID	HOST	CLIENT-ID
#PARTITIONS	ASSIGNMENT	
consumer1-81fe8a2a-0b29-4964-ald4-4741582132bb	/127.0.0.1	consumer1
2	my-topic(0,1)	
consumer2-db6f12f2-elde-4bb7-93ec-6170fefa8830	/127.0.0.1	consumer2
1	my-topic(2)	
consumer7-c1d056fa-eabc-42ed-b557-c848f747ec49	/127.0.0.1	consumer7
1	my-second-topic(0)	

`--describe --offsets`

The new `--offsets` option, that is allowed with `--describe` only, produces the same result as the `--describe` option itself. In other words, `--offsets` is the default sub-option of `--describe`:

```
$ bin/kafka-consumer-groups.sh --bootstrap-server localhost:9092 --describe --group my-group --offsets
Note: This will only show information about consumers that use the Java consumer API (non-ZooKeeper-based consumers).
```

TOPIC	PARTITION	CURRENT-OFFSET	LOG-END-OFFSET	LAG	CONSUMER-
ID		HOST		CLIENT-ID	
my-topic	0	16577	33154	16577	consumer1-81fe8a2a-0b29-4964-ald4-4741582132bb
	/127.0.0.1		consumer1		
my-topic	1	16578	33155	16577	consumer1-81fe8a2a-0b29-4964-ald4-4741582132bb
	/127.0.0.1		consumer1		
my-topic	2	16577	33155	16578	consumer2-db6f12f2-elde-4bb7-93ec-6170fefa8830
	/127.0.0.1		consumer2		
my-second-topic	0	33023	44543	11520	consumer7-c1d056fa-eabc-42ed-b557-c848f747ec49
	/127.0.0.1		consumer7		
my-other-topic	1	7901	8652	751	-
-		-		-	
my-other-topic	0	7902	8654	752	-
-		-		-	

`--describe --state`

The `--state` option, that can also be used with `--describe` only, reports the group level information that comes back in the response of `DescribeGroups` API call along with the coordinator id of the group:

```
$ bin/kafka-consumer-groups.sh --bootstrap-server localhost:9092 --describe --group my-group --state
Note: This will only show information about consumers that use the Java consumer API (non-ZooKeeper-based consumers).
```

COORDINATOR (ID)	ASSIGNMENT-STRATEGY	STATE	#MEMBERS
localhost:9092 (0)	range	Stable	3

In the above outputs, the width of the following columns will be dynamically determined (as opposed to the fixed widths that have been applied so far):

- TOPIC
- CONSUMER-ID
- HOST
- CLIENT-ID
- COORDINATOR (ID)

Compatibility, Deprecation, and Migration Plan

The proposed changes apply to the new Java-based consumer only. Therefore, the consumer groups based on the old consumer will be unaffected.

Users who use the new-consumer based consumer groups and somehow rely on the output of the `--describe` option of the `ConsumerGroupCommand` may have to adjust their clients to use one of the new options provided; for example, in case they rely on the fact the command currently returns members with no assignment. Hopefully, how they use the command output will be simplified with these new options.

Rejected Alternatives