

KIP-635: GetOffsetShell: support for multiple topics and consumer configuration override

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

Status

Current state: *Accepted*

Discussion thread: [here](#)

JIRA: [here](#)

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

Motivation

1. User gets offsets for one topic only. User needs to invoke GetOffsetShell for each topic separately. In order to get offsets of all topics, user first has to obtain the list of topics (by means of another tool) and then query each topic.
2. GetOffsetShell does not have a convenient shell launcher as other console tools. It can be run only with a generic tool: `kafka-run-class.sh`.
3. GetOffsetShell is inconsistent with other command-line tools in terms of command-line arguments. For example, `--broker-list` instead of `--bootstrap-servers`. (This was meant to be handled by [KIP-499](#) and [KAFKA-8507](#).)
4. Users cannot provide additional consumer configurations for the tool, which makes it impossible to use with secure configurations.
5. The tool still accepts deprecated and ignored arguments, which can be removed (`offsets` and `max-wait-ms`).

Public Interfaces

This KIP introduces a new command line tool: `kafka-get-offsets.sh`. The tool provides the following arguments:

- `--bootstrap-server vm1:9092` - Comma-separated list of Kafka servers
- `--broker-list vm1:9092` - Comma-separated list of Kafka servers - DEPRECATED
- `--topic-partitions topic1:1,topic2*:0-1,topic3:1-,topic4:-4` - Comma-separated list of topic-partition patterns to query offsets for. A pattern consists of a topic name or regex, and a partition filter, which can be a partition id, a range (inclusive lower, exclusive upper), an inclusive lower bound or an exclusive upper bound. The topic part can be omitted to only specify the pattern filter. The partition part can be omitted to fetch all partitions of the specified topic(s). (if omitted, query all topics)
- `--topic topic*` - Topic name or regex to specify topics to query for offsets (if omitted, query all topics - if `--topic-partitions` is specified, ignored)
- `--partitions 1,2` - Comma-separated list of topic partitions to query (if omitted, query all partitions of matched topics - if `--topic-partitions` is specified, ignored)
- `--exclude-internal-topics` - Exclude Kafka-internal topics, like `__consumer_offsets` (include internal topics by default)
- `--command-config` - Consumer config properties file to be used by the consumer.

For backward compatibility, the GetOffsetShell tool does not need to change anything as:

- `--broker-list` - Deprecated, the tool will print a warning, but will keep working as before.
- `--topic` - Can still be used as a simple topic name.
- `--exclude-internal-topics` - If omitted, the old behavior is achieved.
- Deprecated arguments were ignored and printing warnings for longer than the expected deprecation period.
- Other arguments did not change, or they are newly added and optional.

Proposed Changes

New implementation of GetOffsetShell does not change the interactions between the tool and the brokers, only adds additional filtering options in the tool.

The `--broker-list` argument will be deprecated. Usage of the argument will print a warning and suggest usage of `--bootstrap-servers`. The commonly used `--bootstrap-servers` will be supported.

Now the user can get offsets for many topic-partitions at once. No need to retrieve the list of existing topics and then query them one by one.

Moreover, now the user is able to retrieve offsets for all topics - this is the default when no topics specified.

Additionally, users can provide configuration overrides for the consumer with a properties file.

Compatibility, Deprecation, and Migration Plan

Any client tools depending on deprecated command-line arguments will stop working. This is acceptable, since the deprecated arguments are ignored for a long time now, and the tool was clearly indicating this for users.

Rejected Alternatives

For the public interface, a different flag was considered for controlling the exclusion of internal topics: *--include-internal-topics*. This would have also meant that the default behavior would exclude internal topics from querying. This would have broken backward-compatibility, as the old tool does not need any additional configuration to query internal topics.