

KIP-503: Add metric for number of topics marked for deletion

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

Status

Current state: *Under Discussion*

Discussion thread: [here](#)

JIRA: [KAFKA-8753](#)

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

Motivation

Since topic deletion cannot always be performed immediately (due to offline replicas, partitions are being reassigned, etc), the Controller marks a topic for deletion and enqueues it for later processing. When a large number of topics (replicas, really) are deleted at once, it can take significant time for the Controller to process everything. During this time, it is not unusual for the Controller to get bogged down. During these times, it would be useful to know how many topics and replicas still remain to be deleted. Currently, an operator can only check on the progress of topic deletion by looking directly in ZooKeeper at the `/admin/delete_topics` znode. In a production environment this is rather cumbersome and is somewhat ill-advised (poking around in ZK on a running Kafka cluster).

Proposed Changes

The following new JMX gauges are proposed for KafkaController:

- `kafka.controller:type=KafkaController,name=TopicsToDeleteCount`
- `kafka.controller:type=KafkaController,name=ReplicasToDeleteCount`
- `kafka.controller:type=KafkaController,name=IneligibleTopicsToDeleteCount`
- `kafka.controller:type=KafkaController,name=IneligibleReplicasToDeleteCount`

These return an integral value for the number of topics and number of replicas known to the Controller that are enqueued for deletion as well as the count of topics/replicas that are not eligible for deletion. Note that the ineligible topics/replicas are exclusive from the pending topics/replicas marked for deletion.

Rather than listing the children of the znode directly, these metrics will read be determined using the internal state of the Controller. During initialization and controller re-elections, these values will be zero as the Controller has not yet read in the list of topics from ZK and computed deletion eligibility.

Compatibility, Deprecation, and Migration Plan

Since this is only adding a new metric, it should not affect any metrics gathering clients.

Rejected Alternatives

If there are alternative ways of accomplishing the same thing, what were they? The purpose of this section is to motivate why the design is the way it is and not some other way.