

KIP-760: Minimum value for segment.ms and segment.bytes

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

Status

Current state: "Under Discussion"

Discussion thread: [here](#)

Voting thread: [here](#)

JIRA: [KAFKA-7760](#)

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

Motivation

Many times, Kafka brokers in production crash with "Too many open files" error or "Out of memory" errors because some Kafka topics have a lot of segment files as a result of small `segment.ms` or `segment.bytes`. These two configuration can be set by any user who is authorized to create topic or modify topic configuration.

To prevent these two configuration from causing Kafka broker crash, there should be a minimum value for these configuration.

Public Interfaces

This KIP proposes to add these two broker configuration:

- `min.topic.segment.ms`
- `min.topic.segment.bytes`

And also, this KIP proposes changes to the validation logic of the following topic configuration:

- `segment.ms`
- `segment.bytes`

Proposed Changes

This KIP proposes to add the following dynamic broker configuration:

Configuration	Type	Default value	Documentation
<code>min.topic.segment.ms</code>	long	3600000	The minimum value for topic configuration <code>segment.ms</code> .
<code>min.topic.segment.bytes</code>	int	1048576	The minimum value for topic configuration <code>segment.bytes</code> .

When provided via `CreateTopics` or `AlterConfigs` API, the validation of the following topic configuration changes as follow:

Topic configuration	Validation
<code>segment.ms</code>	If the value of <code>segment.ms</code> is less than the value of <code>min.topic.segment.ms</code> , Kafka broker will reject it by responding with <code>InvalidConfigurationException</code>
<code>segment.bytes</code>	If the value of <code>segment.bytes</code> is less than the value of <code>min.topic.segment.bytes</code> , Kafka broker will reject it by responding with <code>InvalidConfigurationException</code>

Compatibility, Deprecation, and Migration Plan

- If a topic has `segment.ms` that is less than `min.topic.segment.ms`, Kafka broker will log a warning on startup "Topic X has segment.ms (Y) that is less than min.topic.segment.ms (Z)."
- If a topic has `segment.bytes` that is less than `min.topic.segment.bytes`, Kafka broker will log a warning on startup "Topic X has segment.bytes (Y) that is less than min.topic.segment.bytes (Z)."
- This way, the proposed changes will not cause issue on the existing topic, but still raise attention to topics with small segment sizes.

Rejected Alternatives

- Set a fixed minimum value for `segment.ms` and `segment.bytes`. The reason this approach is rejected is because there are use cases that requires setting `segment.ms` or `segment.bytes` to very small value. For example: deleting all records in a topic by setting `segment.bytes` to 0 temporarily.