

KIP-804: OfflinePartitionsCount Tagged by Topic

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

Status

Current state: *Under Discussion*

Discussion thread: [thread](#)

JIRA:



Unable to render Jira issues macro, execution error.

Motivation

When a partition becomes offline it is important to quickly determine what customers/service are being impacted - possibly reducing overall disruption. Additionally, while many sources recommend instrumenting alerting on the current `OfflinePartitionCount` metric as reported by the active controller, this metric is not tagged in any way. However, as an example, it may be beneficial to only have alerts on a subset of topics or disable alerting for certain test topics. Further as topics are often synonymous with differing use cases (e.g. metrics, logs, etc), having the ability to associate an offline partition with a topic would provide new granularity for debugging, alerting, and SLO/SLA reporting. This KIP proposes to tag the current `OfflinePartitionsCount` metric by the topic name associated with the offline partition(s) for these reasons.

Public Interfaces

The current `OfflinePartitionsCount` metric will be tagged by the topic name of the offline partition(s) *if at least one partition is offline for that given topic*. The current (untagged) metric will still remain, reporting the sum of all offline partitions across all topics. No monitoring changes will be needed to maintain the current functionality.

Monitoring

Current MBean	Proposed MBeans
<code>kafka.controller:type=KafkaController</code>	<ul style="list-style-type: none">• <code>kafka.controller:type=KafkaController</code>• <code>kafka.controller:type=KafkaController,topic=[-.\w]+)</code>

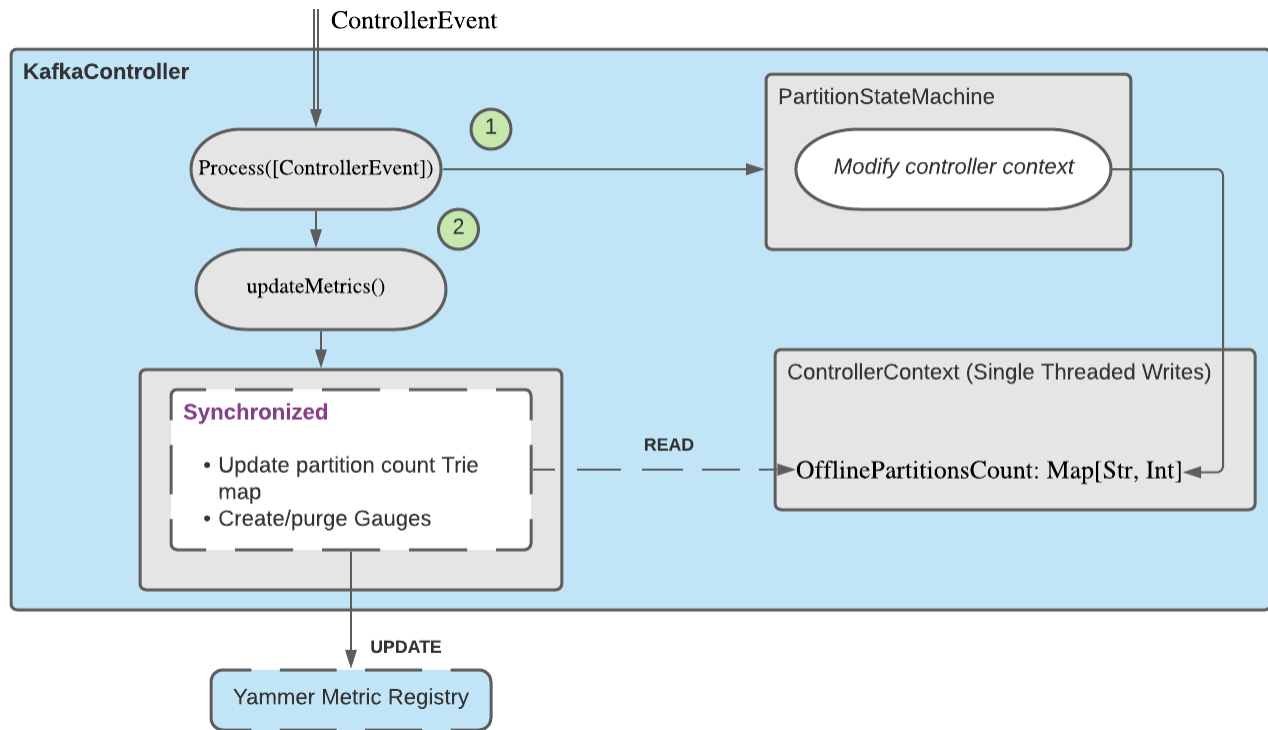
Proposed Changes

The proposed changes are based on the requirements:

- When a topic: `${TOPIC}` has offline partitions, a Gauge exists reporting the number of offline partitions under the MBean name `kafka.controller:type=KafkaController,topic=${TOPIC}`
- Gauges (metrics) do not exist in the metrics registry for topics that have zero offline partitions.
 - When a topic no longer has any offline partitions its corresponding Gauge should be purged from the Yammer metric registry.

The only changes required for this is the replacement of the current atomic `OfflinePartitionsCount` to a concurrent map within the `KafkaController`. Additionally, some logic will needed to ensure gauges are created/purged as needed.

In my ad-hoc process notation, this flow can largely be define in the diagram below.



Compatibility, Deprecation, and Migration Plan

No impact will be seen on existing users and no migration plan will be necessary.

Rejected Alternatives

- Emitting on topics with zero offline partitions
 - A given cluster may have many topics leading – all of which should only have offline partitions a small portion of the time - leading to a spammy metric
 - The original, untagged, `OfflinePartitionsCount` metric, representing the sum of all offline partitions will remain as before as a constant indicator.
 - Reporting on each topic would increase the overhead of the controller in the current design (i.e. referencing state within the `ControllerContext` after each `ControllerEvent`)