

KIP-434: Add Replica Fetcher and Log Cleaner Count Metrics

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

Status

Current state: *Accepted*

Discussion thread: [here](#)

JIRA:

 Unable to render Jira issues macro, execution error.

Released: AK 2.4.0

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

Motivation

In some occasions we detected errors where replica fetcher threads or log cleaners died because of an unrecoverable error and caused more serious issues in the brokers (from lagging to offline replicas, filling up disks, etc.). It would often help if the monitoring systems attached to Kafka could detect these problems early on as it would allow a prompt response from the user and the greater possibility of capturing the root cause.

Replica Fetchers

In this case the first thing users usually notice is that replicas start lagging and it takes a considerable time by they get to the conclusion that it was because of *something* killing the replica fetchers. The problems can span from bugs to rare log divergence issues. Having a metric for dead replica fetchers would speed up the investigation of any such issues as monitoring systems attached to Kafka could record the exact time the issue raised. In some (but relatively small) of these cases it is a problem that application logs roll over too often, therefore the real root cause of the issue remains unknown. This metric would allow users to trigger alerts based on the change of this metric.

Log Directory Fetchers

Similarly to the replica fetchers when altering log directories, some log dir fetcher might get unexpectedly interrupted and in this case it won't be finished and users might have to do some digging to figure out what could have happened. Introducing a metric named `DeadLogDirFetcherThreadCount` for tracking the dead fetcher threads would speed up diagnostics.

Log Cleaners

The motivation for the dead log cleaner thread count metric is very similar. Sometimes a problem with the log cleaner threads get noticed when a disk gets full - a log cleaner died earlier because of some issue that prevented cleanup and it never got restarted. A metric for the dead thread count would help because alerts could be triggered based on this in monitoring systems and also it would be easier to find out the exact time of this issue.

Public Interfaces

I propose to add two gauge: `DeadReplicaFetcherThreadCount` for the fetcher threads, `log-cleaner-dead-thread-count` for the log cleaner. All of these are broker level metrics.

DeadFetcherThreadCount: this basically exposes the count of non-alive threads in the internal thread map in `AbstractFetcherManager`. Its `clientId` tag could either be `Fetcher` or `ReplicaAlterLogDirs` so the two metrics are distinguishable.

log-cleaner-dead-thread-count: this would expose the number of dead threads inside the `ArrayBuffer` that maintains the `CleanerThread` instances in `LogCleaner`.

Proposed Changes

There would be no changes beside the changes listed in the previous section.

Compatibility, Deprecation, and Migration Plan

No metrics will be removed or deprecated and no migration would be required.

Test Plan

It is possible to write an automated test which sets up a few brokers and injects failures in the observed threads or just interrupts them and observes the changes in the metrics.

Rejected Alternatives

No rejected alternatives so far.