

KIP-988: Streams Standby Update Listener

- [Status](#)
- [Motivation](#)
- [Public Interfaces](#)
- [Proposed Changes](#)
- [Compatibility, Deprecation, and Migration Plan](#)
- [Test Plan](#)
- [Rejected Alternatives](#)
 - [Do Nothing](#)
 - [Use StateRestoreListener](#)

Status

Current state: *"Adopted"*

Discussion thread: [here](#)

JIRA: [here](#)

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

Co-Authored by Eduwer Camacaro.

Motivation

Currently, Kafka Streams allows you to provide a `StateRestoreListener` which executes callbacks upon the start and end of a state restoration, and also after each batch is restored in the restoration. This is useful for monitoring any Active Task that is undergoing restoration.

KIP-869 adds metrics which further improve visibility of Active Task Restorations. However, as of now it is difficult to get real-time updates (via callbacks) which tell you about the state of Standby Tasks on your Kafka Streams instance. This would be most useful from the operational perspective, for example when implementing a rolling restart or a smooth scale down. In these situations, I as an operator of Kafka Streams want to be able to know where each Active and Standby Task lives and also how "caught up" they are so that I know when it's safe to remove or bounce a certain Streams instance.

In addition to the operational use-case described above, knowing when a Standby Task is created and destroyed (either by promotion or migration) would help operators test various configurations of `session.timeout.ms` during rolling upgrade scenarios with regards to how noisy the shuffling of Standby Tasks is.

Public Interfaces

We will add the `StandbyUpdateListener` interface as follows:

We will add a method to the `KafkaStreams` object:

```
public void setStandbyUpdateListener(StandbyUpdateListener otterStandbyUpdateListener);
```

Proposed Changes

We propose to create a `StandbyUpdateListener` interface, and allow users to supply one to their Kafka Streams Topology via `KafkaStreams#setStandbyUpdateListener(...)` in a manner similar to the `StateRestoreListener`.

Compatibility, Deprecation, and Migration Plan

We are adding a new method without changing any existing API's. Existing users will not need to know about the `StandbyUpdateListener` functionality, and existing code will be unaffected.

Test Plan

This is a small change that can be tested sufficiently via unit tests. We at LittleHorse have implemented a draft PR and will test internally as well.

Rejected Alternatives

Do Nothing

The first alternative is “Do Nothing.” In theory, we can use the `KafkaStreams#metrics()` method to get a handle on the Restore Consumer, and look at the lag metrics of that consumer to determine how “caught up” the Standby Task is. This has a few problems:

1. It feels “clunky” to go through JMX metrics within an application to change things that we want to handle in code path. Additionally, the updates are not guaranteed to be as precise, and it is a polling mechanism rather than a push-based (callback) mechanism.
2. We have no easy way of determining why the Standby Task was migrated away to a different instance (whether it was `PROMOTED` to an Active Task, or `MIGRATED` to run on another Streams Instance).

Use StateRestoreListener

The second alternative is to use the `StateRestoreListener` for Standby Tasks as well. However, that quickly falls apart upon further examination because the API of the `StateRestoreListener` is semantically different from the `StandbyUpdateListener`. Most crucially, a State Restoration has a definitive “finish line”, which is the last offset of the Changelog TopicPartition at the time that the State Restoration begins. This is because that offset will not increase during restoration, since the Active Task is down (we know this because the Active Task itself is undergoing restoration!). In contrast, with Standby Tasks, the finish line is constantly advancing.