

KIP-945: Update threading model for Consumer

- [Status](#)
- [Motivation](#)
- [Public Interfaces](#)
- [Proposed Changes](#)
 - [Terminology](#)
 - [Threading Model](#)
 - [Background thread](#)
 - [Providing Data to the Background Thread](#)
 - [Getting Data from the Background Thread](#)
 - [Network I/O](#)
- [Compatibility, Deprecation, and Migration Plan](#)
- [Test Plan](#)
- [Rejected Alternatives](#)

Status

Current state: *WIP*

Discussion thread: [here](#)

JIRA: [here](#)

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

Note this is a joint proposal by [Philip Nee](#), [Kirk True](#), and [Lianet Magrans](#).

Motivation

This KIP documents the updated threading model of the `Consumer` implementation of the client.

The complexity of the consumer has increased and with it the code to support and fix bugs. Patches and hotfixes in the past years have heavily impacted the readability of the code. The complex code path and intertwined logic make the code difficult to modify and comprehend. Additionally, logic is at times executed on application threads and at other times on the dedicated, internal *heartbeat* thread. The asynchronous nature of the current implementation has lead to many bugs (which are labeled with the [new-consumer-threading-should-fix](#) label). The motivation is to simplify the structure of the code by clearly defining—and removing, where possible—the asynchronous code.

The simplification will also allow us to implement the necessary primitives for [KIP-848](#).

Public Interfaces

This KIP has the explicit goal of making *no* changes to the public interfaces. The protocol, configuration, APIs, etc. will remain as they currently are. The internal behavior of the consumer is substantially changing and we want to ensure it is reviewed and vetted by the community.

Proposed Changes

Terminology

To help understand the design, we need to introduce some terminology. Terms designated with ¹ apply to the current `KafkaConsumer` implementation and terms with ² apply to the new implementation; a term may apply to both.

Term	Definition
Application event	<p>A data structure specific to each <code>Consumer</code> API call that encapsulates application-provided data. For example, the application event specific to the <code>seek</code> event would include the user-provided topic information and offset. These events are enqueued onto the <i>application event queue</i> by the <code>Consumer</code>.</p> <p>Application events can optionally include a <code>Future</code> on which the <i>application thread</i> can issue a timed block, awaiting completion by the <i>background thread</i>.</p>
Application event processor	<p>Logic which processes <i>application events</i> on the <i>background thread</i>, interacting with the <i>request managers</i>.</p>

Application event queue	A shared queue which stores <i>application events</i> enqueued by the <i>application thread</i> . These events are later dequeued by the <i>background thread</i> and given to the <i>application event processor</i> for execution.
Application thread	The thread that is executing the user's code that interacts with the <code>Consumer</code> API. Per the current implementation in <code>KafkaConsumer</code> , only one thread may call APIs at a time.
Background event	
Background event queue	A shared queue which stores <i>background events</i> enqueued by the <i>background thread</i> . The events are later dequeued by the <i>application thread</i> inside the <code>Consumer</code> and handled appropriately.
Background thread	An internal thread created for each <code>Consumer</code> instance on which the following operations are performed: <ul style="list-style-type: none"> • Execution of <i>application events</i> • Group membership • Managing network I/O requests and responses • Forwarding results to <i>application events</i> • Submitting <i>background events</i> for processing by the <i>application thread</i>
Event handler	Logic that pulls <i>events</i> from the <i>event queue</i> for processing on the <i>background thread</i> .
Heartbeat	Logic related to communicating liveness, group membership, etc. as introduced in KIP-62 .
Network client delegate	
Request manager	An internal interface that is used by the background thread to handle the management of requested, inflight, and responded network I/O.

Threading Model

<TBD>

Background thread

<TBD>

Providing Data to the Background Thread

<TBD>

Getting Data from the Background Thread

<TBD>

Network I/O

<TBD>

Compatibility, Deprecation, and Migration Plan

- *What impact (if any) will there be on existing users?*
- *If we are changing behavior how will we phase out the older behavior?*
- *If we need special migration tools, describe them here.*
- *When will we remove the existing behavior?*

Test Plan

Describe in few sentences how the KIP will be tested. We are mostly interested in system tests (since unit-tests are specific to implementation details). How will we know that the implementation works as expected? How will we know nothing broke?

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.