

# Kafka Streams Architecture

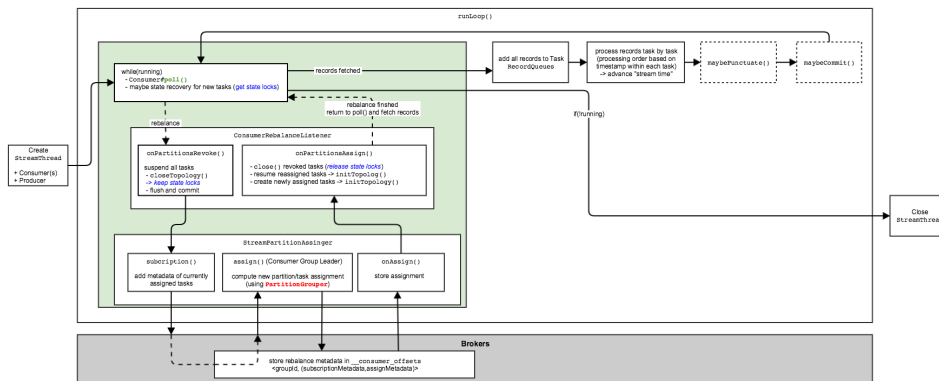
- Lifecycle of a StreamThread
- Lifecycle of a StreamTask and StandbyTask
- Exception Handling
  - Types of Exceptions:
  - Coding implications:



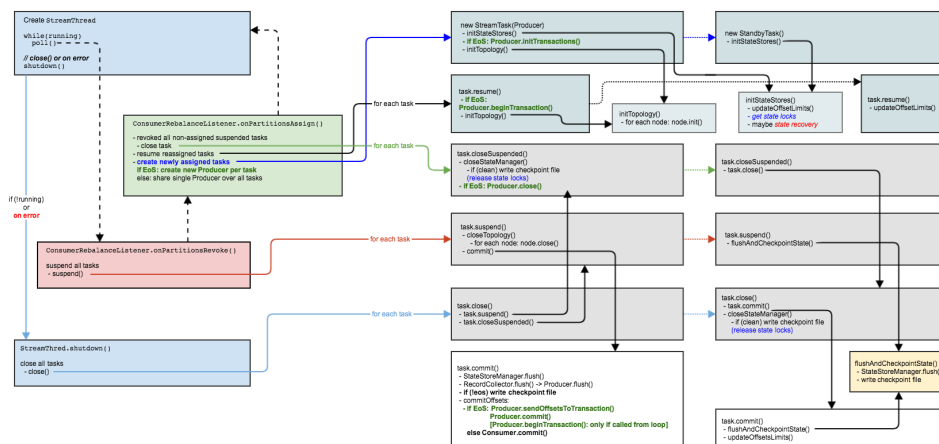
## Warning

We try to keep this doc up to date, however, as it describes internals that might change at any point in time, there is no guarantee that this doc reflects the latest state of the code base.

## Lifecycle of a StreamThread



## Lifecycle of a StreamTask and StandbyTask



# Exception Handling

A Kafka Streams client need to handle multiple different types of exceptions. We try to summarize what kind of exceptions are there, and how Kafka Streams should handle those. In general, Kafka Streams should be resilient to exceptions and keep processing even if some internal exceptions occur.

## Types of Exceptions:

There are different categories how exceptions can be categorized.

First, we can distinguish between recoverable and fatal exceptions. Recoverable exception should be handled internally and never bubble out to the user. For fatal exceptions, Kafka Streams is doomed to fail and cannot start/continue to process data.

Related to this are retrievable exception. While retrievable exception are recoverable in general, it might happen that the (configurable) retry counter is exceeded; for this case, we end up with an fatal exception.

The second category are "external" vs "internal" exception. By "external" we refer to any exception that could be returned by the brokers. "Internal" exceptions are those that are raised locally.

For "external" exceptions, we need to consider `KafkaConsumer`, `KafkaProducer`, and `KafkaAdminClient`. For internal exceptions, we have for example (de)serialization, state store, and user code exceptions as well as any other exception Kafka Streams raises itself (e.g., configuration exceptions).

Last but not least, we distinguish between exception that should never occur. If those exception do really occur, they indicate a bug and thus all those exception are fatal. All regular Java exception (eg, `NullPointerException`) are in this category.

## Coding implications:

- We should never try to handle any fatal exceptions but clean up and shutdown
  - We should catch all those exceptions for clean up only and rethrow unmodified (they will eventually bubble out of the thread and trigger uncaught exception handler if one is registered)
  - We should only log those exception (with `ERROR` level) once at the thread level before they bubble out of the thread to avoid duplicate logging
- We need to do fine grained exception handling, ie, catch exceptions individually instead of coarse grained and react accordingly
- All methods should have complete JavaDocs about exception they might throw
- All exception classes must have strictly defined semantics that are documented in their JavaDocs
- In *runtime* code, we should never throw any regular Java exception (except it's fatal) but define our own exceptions if required (this allows us to distinguish between bugs and our own exceptions)
- We should catch, wrap, and rethrow exceptions each time we can add important information to it that helps users and us to figure out the root cause of what when wrong

To be discussed:

- How to handle `Throwable` ?
  - Should we try to catch-and-rethrow in order to clean up?
    - `Throwable` is fatal, so clean up might fail anyway?
    - Furthermore, should we assume that the whole JVM is dying anyway?
  - Should we be harsh and call `System.exit` (note, we are a library – but maybe we are "special" enough to justify this?)
    - Note, if a thread dies without clean up, but other threads are still running fine, we might end up in a deadlock as locks are not released
    - Could also be configurable
    - Could also be a hybrid: try to clean up on `Throwable` but call `System.exit` if clean up fails (as we would end up in a deadlock anyway – maybe only if running with more than one thread?)
  - Should we force users to provide uncaught exception handler via `KafkaStreams` constructor to make sure they get notified about dying streams?
- Restructure exception class hierarchy:
  - Remove all sub-classed of `StreamsException` from public API (we only hand out this one to the user)
  - A `StreamsException` indicates a fatal error (we could sub-class `StreamsException` with more detailed fatal errors if required – but don't think this is necessary)
  - We sub-class `StreamsException` with (an abstract?) `RecoverableStreamsException` in internal package for any internal exception that should be handled by Streams and never bubble out
    - As an alternative (that I would prefer) we could introduce this as an independent and *checked* exception instead of inheriting from `StreamsException` (this forces us to declare and handle those exceptions in our code and makes it hard to miss – otherwise, one might bubble out due to a bug)
  - We sub-class individual recoverable exceptions in a fine grained manner from `RecoverableStreamsException` for individual errors
  - We can further group all retrievable exceptions by sub-classing them from abstract `RetriableStreamsException` extends `RecoverableStreamsException` – the more details/categories the better?

	KafkaConsumer	KafkaProducer	StreamsKafkaClient	AdminClient	Streams API
--	---------------	---------------	--------------------	-------------	-------------

fatal (should never occur)	local: - IllegalArgumentException - IllegalStateException - WakeupException - InterruptException remote: - UnknownServerException	local: - IllegalArgumentException - IllegalStateException - WakeupException - InterruptException remote: - UnknownServerException - OffsetMetadataTooLarge - SerializationException (we use < byte[] , byte[] > as types)	local: - IllegalArgumentException - IllegalStateException - WakeupException - InterruptException remote: - UnknownServerException - InvalidTopicException	local: - IllegalArgumentException - IllegalStateException - WakeupException - InterruptException remote: - UnknownServerException - InvalidTopicException	
fatal	local: - ConfigException remote: - AuthorizationException (including all subclasses) - AuthenticationException (including all subclasses) - SecurityDisabledException - InvalidTopicException	local: - ConfigException remote: - AuthorizationException (including all subclasses) - AuthenticationException (including all subclasses) - SecurityDisabledException - InvalidTopicException - UnknownTopicOrPartitionsException (retyable? refresh metadata?) - RecordBatchTooLargeException - RecordTooLargeException	local: - ConfigException remote: - AuthorizationException (including all subclasses) - AuthenticationException (including all subclasses) - SecurityDisabledException	local: - ConfigException remote: - AuthorizationException (including all subclasses) - AuthenticationException (including all subclasses) - SecurityDisabledException	local: - ConfigException - SerializationException
retriable	local: remote: - InvalidOffsetException (OffsetOutOfRangeException, NoOffsetForPartitionsException) - CommitFailedException - TimeoutException - QuotaViolationException?	local: remote: - CorruptedRecordException - NotEnoughReplicasAfterAppendException - OffsetOutOfRangeException (when can producer get this?) - TimeoutException - QuotaViolationException? - BufferExhaustedException (verify)	local: remote:	local: remote:	local:
recoverable	local: remote:	local: remote: - ProducerFencedException	local: remote:	local: remote:	local: - LockException

Having a look at all `KafkaException` there are some exception we need to double check if they could bubble out any client (or maybe we should not care, and treat all of them as fatal/remote exceptions).

-> DataException, SchemaBuilderException, SchemaProjectorException, RequestTargetException, NotAssignedException, IllegalWorkerStateException, ConnectRestException, BadRequestException, AlreadyExistsException (might be possible to occur, or only TopicExistsException), NotFoundException, ApiException, InvalidTimestampException, InvalidGroupException, InvalidReplicationFactorException (might be possible, but indicate bug), o.a.k.common.errors.InvalidOffsetException and o.a.k.common.errors.OffsetOutOfRangeException (*side note: do those need cleanup – seems to be duplicates?*), ReplicaNotAvailableException, UnknownServerException, OperationNotAttemptedException, PolicyViolationException, InvalidConfigurationException, InvalidFetchSizeException, InvalidReplicaAssignmentException, InconsistentGroupProtocolException, RebalanceInProgressException, LogDirNotFoundException, BrokerNotAvailableException, InvalidOffsetCommitSizeException, InvalidTxnTimeoutException, InvalidPartitionsException, TopicExistsException (cf. AlreadyExistsException), InvalidTxnStateException, UnsupportedForMessageFormatException, InvalidSessionTimeoutException, InvalidRequestException, IllegalGenerationException, InvalidRequiredAckException,

-> RetryableException, CoordinatorNotAvailableException, RetryableCommitException, DuplicateSequenceNumberException, NotEnoughReplicasException, NotEnoughReplicasAfterAppendException, InvalidRecordException, DisconnectException, InvalidMetadataException (NotLeaderForPartitionException, NoAvailableBrokersException, UnknownTopicOrPartitionException, KafkaStoreException, LeaderNotAvailableException), GroupCoordinatorNotAvailableException

Should never happen:

Handled by client (consumer, producer, admin) internally and should never bubble out of a client: (verify)

- ConnectionException, RebalanceNeededException, InvalidPidMappingException, ConcurrentTransactionException, NotLeaderException, TransactionalCoordinatorFencedException, ControllerMovedException, UnknownMemberIdException, OutOfOrderSequenceException, CoordinatorLoadInProgressException, GroupLoadInProgressException, NotControllerException, NotCoordinatorException, NotCoordinatorForGroupException, StaleMetadataException, NetworkException,

