

KIP-24 - Remove ISR information from TopicMetadataRequest and add broker level metadata request

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

ISR information returned from metadata requests can be inconsistent with that stored/ returned from Zookeeper. Metadata is only propagated from controller to the brokers when the controller changes the leader or the ISR. However, when a follower catches up, the leader (not the controller) adds it back to ISR and updates ZK. This information is not propagated to all brokers.

Status

Current state: *Discarded*

Discussion thread: [here](#)

JIRA: [KAFKA-1367](#)

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

Motivation

In current design, ISR information in topic metadata response from brokers can be out of sync. The incorrect information can be a problem for the clients relying on this information.

Public Interfaces

Will update this section once we have decided on the solution we want to go ahead with.

Briefly list any new interfaces that will be introduced as part of this proposal or any existing interfaces that will be removed or changed. The purpose of this section is to concisely call out the public contract that will come along with this feature.

A public interface is any change to the following:

- *Binary log format*
- *The network protocol and api behavior*
- *Any class in the public packages under clientsConfiguration, especially client configuration*
 - *org/apache/kafka/common/serialization*
 - *org/apache/kafka/common*
 - *org/apache/kafka/common/errors*
 - *org/apache/kafka/clients/producer*
 - *org/apache/kafka/clients/consumer (eventually, once stable)*
- *Monitoring*
- *Command line tools and arguments*
- *Anything else that will likely break existing users in some way when they upgrade*

Proposed Changes

TopicMetadataRequest's ISR information, which gets out of sync for a broker that gets bounced/ restarted, does not have any good known use case. It is proposed to remove ISR information from the TopicMetadataRequest. However, a new request, BrokerMetadataRequest, is proposed to be added. The new request will include ISR information for all the partitions that the broker leads.

Compatibility, Deprecation, and Migration Plan

- *What impact (if any) will there be on existing users?*
 - *ISR information will be removed from TopicMetadataRequest. Clients using the information will have to start using new BrokerMetadataRequest.*
- *If we are changing behavior how will we phase out the older behavior?*
 - *It will be a breaking change for anyone who is using the unreliable ISR information from TopicMetadataRequest.*
- *If we need special migration tools, describe them here.*
 - *No*
- *When will we remove the existing behavior?*
 - *Kafka 0.9 (probably)*

Rejected Alternatives

- Remove ISR field from topic metadata response and enforce people to use the admin tool (with ZK dependency) for such usages, which would also require a protocol change between client / server
- Let the controller to also watch for ISR changes and propagate that information to brokers, this will not introduce protocol change to clients. This will likely add a lot of burden on controllers since ISR change is more frequent than leader migrations.