

# KIP-29 - Add an `IsrPropagateIntervalMs` configuration to `KafkaConfig`

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

## Status

**Current state:** *["Under Discussion"]*

**Discussion thread:** [here](#)

**JIRA:** [KAFKA-2406](#)

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

## Motivation

This KIP is related to KIP-4, KAFKA-1367 and KAFKA-2406. In the future, all the Kafka clients and admin tools will not access zookeeper directly. Therefore KAFKA-1367 was introduced to propagate the ISR change to all the brokers to sync up topic metadata on brokers. The approach is to create a ZK path to notify controller update ISR and send `UpdateMetadataRequest` to brokers. The issue is during a broker bounce or failure, the amount of ISR change are huge and overwhelm the controller. This KIP is to introduce a new config to batch the ISR update and throttle the ISR propagation rate.

## Public Interfaces

```
val IsrChangePropagateIntervalMsProp = "isr.change.report.interval.ms"
val IsrChangePropagateIntervalMsDoc = "Specify the interval to propagate ISR change to the entire cluster." +
    "The ISR propagation involves zookeeper path creation. Propagating ISR" +
    " change too frequently might cause performance issue."
```

Default value for the config will be set to 5 seconds.

## Proposed Changes

Add a new config **`IsrPropagateIntervalMs`** to `KafkaConfig`. The broker will only propagate ISR change at most every `IsrPropagateIntervalMs`. The propagation will create a sequential ZK path in zookeeper `/isr_change_notification` with data of all the partition whose ISR has changed since last propagation.

## Compatibility, Deprecation, and Migration Plan

*This is a backward compatible change.*

## Rejected Alternatives

None