

KIP-592: MirrorMaker 1.0 should replicate topics from earliest

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

Status

Current state: Under Discussion

Discussion thread: [here](#)

JIRA: [KAFKA-4668](#) - Getting issue details...

STATUS

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

Motivation

MirrorMaker 1.0 currently inherits the default value for `auto.offset.reset`, which is `latest`.

While for most consumers this is a sensible default, MirrorMakers are specifically designed for replication, so they should default to replicating topics from the beginning.

A specific scenario where the current behavior causes data "loss" is when a MirrorMaker is subscribed to a regex pattern. If auto-topic creation is enabled on the cluster, and you start producing to a non-existent topic that matches the regex, then there will be a period of time where the producer is producing before the new topic's partitions have been picked up by the MirrorMaker. Those messages will never be consumed by the MirrorMaker because it will start from latest, ignoring those just-produced messages.

In fact, the new MirrorMaker 2.0 sets exactly this config: <https://github.com/apache/kafka/blob/d63eaaaa0181bb7b9b4f5ed088abc00d7b32aeb0/connect/mirror/src/main/java/org/apache/kafka/connect/mirror/MirrorConnectorConfig.java#L233>

So this change will simply bring the old MirrorMaker 1.0 into compliance with the behavior of the new MirrorMaker 2.0 which already behaves this way.

Proposed Changes

This would add a MirrorMaker 1.0 default consumer property of `auto.offset.reset==earliest`. Users can still override this in the MirrorMaker consumer config file.

Compatibility, Deprecation, and Migration Plan

This will be a silent breaking change since it flips the behavior around.

Existing mirrormakers will be unaffected for any topics they are currently consuming since they already have a saved offset.

Mirrormakers that start consuming topics for which they don't have a saved offset will start replicating the partitions from the beginning, rather than from the partition's current highwater mark. If the mirrormaker starts consuming a very large partition/topic, it will replicate far more data than previously. This has relatively low probability since most topics that get replicated are newly-created, in which case starting from the earliest simply prevents skipping the first few seconds/minutes of data written to the topic.

Since MirrorMaker 2.0 already behaves this way, this change will make future migrations from MM1 to MM2 easier for folks since the behavior will stop changing between them.

Rejected Alternatives

Leaving it as-is. As noted in the description, the existing state of affairs produces data gaps for anyone replicating topics using a regex pattern.