KIP-781: Improve MirrorMaker2's client configuration

Status

Current state: Under Discussion

Discussion thread: here

JIRA: KAFKA-13365

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

Motivation

As of present, MirrorMaker 2 (aka MM2) 's client configurtaion feature has some problems:

1. The replication-level client configuration works only to the common properties like `bootstrap.servers`, `security.protocol`, ssl, sasl, etc; that is, a configuration like `'AB.producer.batch.size'` is ignored.
   a. Also, which admin client is affected by the replication-level configuration like AB.admin.retry.backoff.ms is unclear; MM2 uses two admin clients for both upstream and downstream clusters, respectively.
2. MM2 is based on Kafka Connect framework's connector; Since MM2 Connectors (`MirrorSourceConnector`, `MirrorCheckpointConnector`, and `MirrorHeartbeatConnector`) are source connectors, they use producer instance created by Kafka Connector, which uses `'producer.override.{property-name}'` in connector configuration; But, `'target.producer.{property-name}'` are not automatically applied to `'producer.override.{property-name}'` so not actually applied to producer instance.
3. MM2 requires to define the `'bootstrap.servers'` of the clusters in cluster-level, like `'A.bootstrap.servers'` or `'B.bootstrap.servers'`; but it also allows to override them in cluster-level and replication-level configs, like `'A.producer.bootstrap.servers'` or `'AB.consumer.bootstrap.servers'`; actually these configurations are not used but, it would be better to ignore it and give a warning.

For the proof of the problem, please refer here.

Public Interfaces

This proposal does not introduce any new public interface or configurable settings; it only makes the existing ones (which are currently ignored by MirrorMaker 2) working by providing consistent, easy-to-understand configuration rules.

Proposed Changes

1. The replication-level client configuration now works; `A->B.producer.{property-name}` and `A->B.admin.{property-name}` are applied to target cluster clients and `A->B.consumer.{property-name}` is applied to source cluster clients.
2. In standalone mode, `target.producer.{property-name}` is automatically copied to `producer.override.{property-name}`. In connector mode, the user should manually configure them, and the documentation will explicitly mention it.
3. Cluster-level and replication-level client bootstrap servers configurations will be ignored and show a warning message.

Compatibility, Deprecation, and Migration Plan

Since this proposal only makes currently not working settings work, there is no deprecation or migration plan.

Rejected Alternatives

None.