

ideas for kafka message format v.3

1. [KIP-931: Flag to ignore unused message attribute field](#)
2. Changing the batch format proposal

Problem: Currently, we perform validation of data schema (such as verifying monotonically increasing offsets etc.). To do that, we have to read the headers of each message in a batch. For a compressed batch, this means that we have to decompress the entire batch including the payload (key/value) even though we don't really require any information out of them.

Solution: If we change the ordering of messages in a batch to prefix with headers of all messages first and followed by the payload of all messages, it will lead to an optimization where we have to "partially" decompress the batch (i.e. the prefix).

3. v3 of the record format should make it cheaper to make changes in the future (perhaps it could support tagged fields or similar)
4. We'd want to fix other known issues at the same time (eg log append time should always be available, there may be others)
5. We should consider whether we would want to introduce a user header that is at the batch level vs record level for efficiency reasons
6. With the message format, v2, we have a single Timestamp field that is qualified by the timestamp type (0 for CreateTime, 1 for LogAppendTime). This forces us to choose either the timestamp of the Client or the Broker's append time and use the selected timestamp type for log retention and rotation logic. However both the client and the broker timestamps are relevant as they capture different information and can be used log retention and rotation logic. The proposal here is to include both the CreateTime(Event Time) and the LogAppendTime. If #4 above is the same as this proposal then please ignore