

KIP-838 Simulate batching and compression

Status

Current state: *WIP*

Discussion thread:

Discussing thread: [here](#)

JIRA:

Motivation

The idea is to keep in the way of saving money getting accurate information about batching and compression.

In our company we have hundreds of clusters and hundreds of marketplaces (producers and consumers) so basically several clients we don't know.

Currently thanks to this [KIP-824](#) now we have a way to inspect the produced requests without affecting the cluster performance. For now we have a rudimentary way which is reading the STDOUT and parsing it in order to get the batching information and the compression type.

This is good enough to get information about potential benefit of batching and compression sending the STDOUT to another script, but still we are missing an important

check, SIMULATE the compression payload.

I would like to create a script which read a sample of the segment log, and output it into a json if the topic/s will benefit of the batching and compression printing the compression ratio.

Then we can reach the clients with more accurate information, for example the reduction traffic and disk saving. In this way they could see in numbers the cost saving even before applying it.

Also with this script we can monitor in real time the topics and automate a way to reach the clients.

Public Interfaces

- Similar to kafka-dump-log.sh the new script would be called (of course open to discuss) kafka-simulate-batching-and-compression.sh
- The script will accept parameters :

- Max bytes to read from the segment log
- Window time in ms to use as a potential group of batching
- Min amount of records in the ms time window mentioned above to consider batching
- Group them by producer ID (if present)
- Compression algorithm

- Skip not active topics and internal ones

Executing command

```
$ bin/kafka-simulate-batching-and-compression.sh --topic topic_test --max-bytes 5000000 --batching-window-time 300 --min-records-for-batching 30 --batching-only-by-producer-id --compression-algorithm lz4 --topic-considered-active-last-produce-request-hours 1 --output json
```

Output:

Executing command

```
{
{
  "topic": "topic_test" {
    "already_batching": "false",
    "already_compressing": "false",
    "candidate_for_batching": "true",
    "candidate_for_compression": "true",
    "compression_ratio_percent": 400,
    "average_records_per_batch": 25
  }
}
}
```

Proposed Changes

For reading the segment logs we can reuse the classes for reading the segments called FileRecords, similar to what kafka-dump-log.sh does.

The idea is to read the partition 0 of a topic (it always exists) make sure the topic is "alive" so it must be checked the mtime of the current active segment.

Then the samples are taken respecting the amount of bytes and from the end of the active segment.

Rejected alternatives

WIP

Compatibility, Deprecation, and Migration Plan

- This is a new script so neither creates any compatibility issue nor migration plan is needed.

- WIP