

KIP-177: Consumer perf tool should count rebalance time

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

Status

Current state: *Accepted*

Discussion thread: [here](#)

JIRA: [KAFKA-5358](#)

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

Motivation

Currently ConsumerPerformance command shows the result at the end of the performance testing for the new consumer, as shown below:

```
$ bin/kafka-consumer-perf-test.sh --broker-list localhost:9092 --messages 50000000 --topic test
start.time, end.time, data.consumed.in.MB, MB.sec, data.consumed.in.nMsg, nMsg.sec
2017-07-13 09:56:00:302, 2017-07-13 09:56:31:093, 7820.1294, 253.9745, 20500000, 665778.9614
```

The information included in the command output includes:

- start time of the test
- end time of the test
- total consumed message bytes(MB)
- average consumed message bytes per second
- total consumed message count
- average consumed message count per second

This KIP suggests adding some metrics to measure rebalance time for the new consumer in this performance tool so that throughput between different versions can be compared more easily in spite of improvements such as [KIP-134: Delay initial consumer group rebalance](#). At the moment, running the perf tool on 0.11.0 or trunk for a short amount of time will present a severely skewed picture since the overall time will be dominated by the join group delay.

Public Interfaces

None.

Proposed Changes

To describe the proposed changes let us revisit the example above. This KIP proposes to count and display rebalance time for the new consumer in ConsumerPerformance tool, as shown below:

```
$ bin/kafka-consumer-perf-test.sh --broker-list localhost:9092 --messages 50000000 --topic test
start.time, end.time, data.consumed.in.MB, MB.sec, data.consumed.in.nMsg, nMsg.sec, rebalance.time.ms, fetch.
time.ms, fetch.MB.sec, fetch.nMsg.sec
2017-07-13 10:38:29:445, 2017-07-13 10:39:03:249, 7820.1294, 231.3374, 20500000, 606437.1080, 35, 33769,
231.5772, 607065.6519
```

The output above adds several metrics including:

- `rebalance.time.ms`: total rebalance time for the consumer group
- `fetch.time.ms`: total fetching time for the group excluding the rebalance time
- `fetch.MB.sec`: average fetched message bytes per second (based on `fetch.time.ms`)
- `fetch.nMsg.sec`: average fetched message count per second (based on `fetch.nMsg.sec`)

Compatibility, Deprecation, and Migration Plan

The proposed changes apply to the new Java-based consumer only. Therefore, the consumer groups based on the old consumer will be unaffected.

Users who use the new-consumer based consumer groups and somehow rely on the output of the test may have to adjust their clients to understand new format of the output.

Rejected Alternatives

None.