Exactly Once vs At Least Once (Dataflow)

Ihaffa Murtopo
Data Engineer
Google

BEAM SUMMIT

September 4-5, 2024
Sunnyvale, CA. USA

What Is?

Feature	Exactly-Once (Default)	At-Least-Once (New)
Definition	Each record is processed exactly one time, even in the presence of failures. No duplicates or data loss.	Each record is processed at **minimum** one time, even if there are failures. Duplicates possible.
Advantages	• Ensures data accuracy	• Lower Cost • Higher Throughput

BEAM SUMMIT

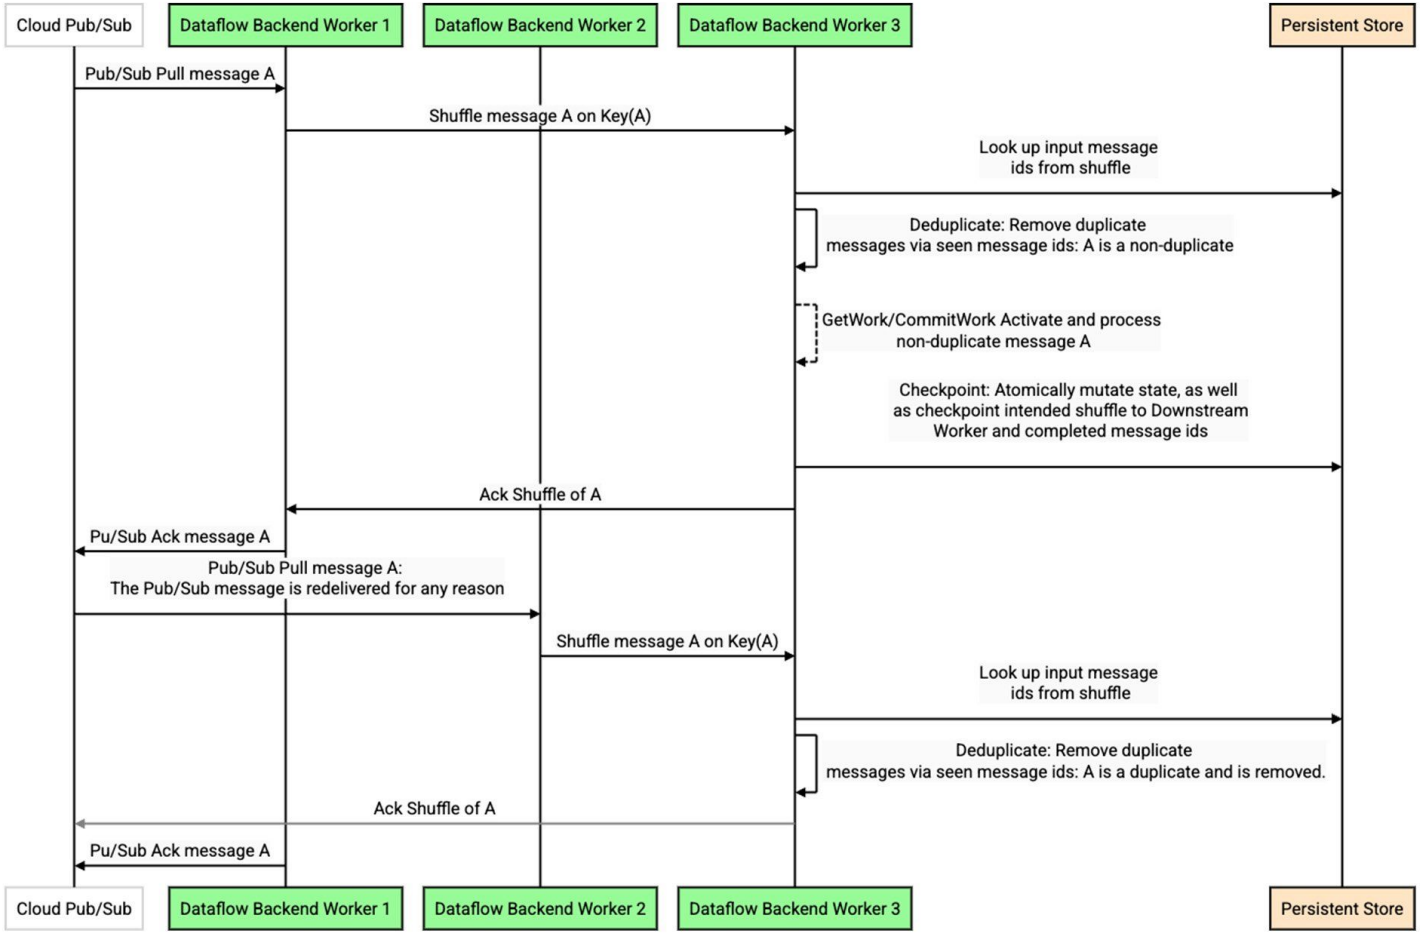

Summary:

- Deterministically assign key to Pubsub Message

- Require to constant interaction with Persistent Store

- If a particular key or worker is slow, impacts unrelated key

Setup Benchmark

- Consider using Dataflow template: Streaming Data Generator
- BQ or Table to count number of duplicates

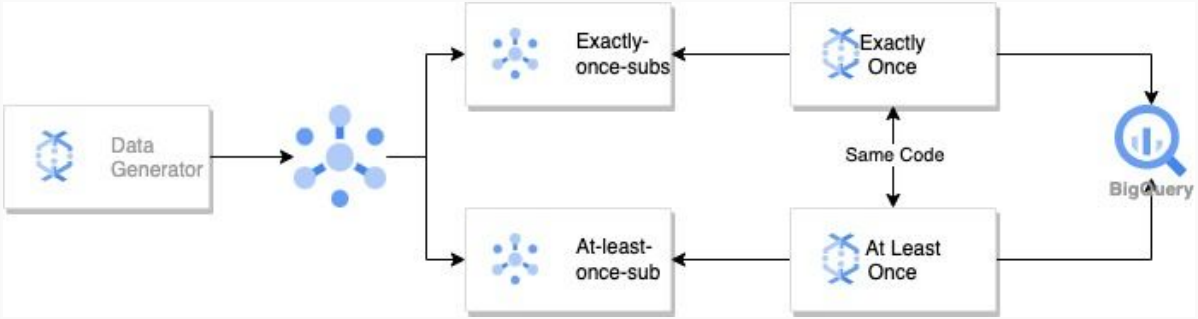

Run At-Least-Once

- Add “streaming_mode_at_least_once”
- Exactly-Once is default configuration

```
mvn compile exec:java -Dexec.mainClass=org.apache.beam.examples.Demo  \
    -Dexec.args="--runner=DataflowRunner --gcpTempLocation=gs://ihaffa-sandbox
    --usePublicIps=false --subnetwork=https://www.googleapis.com/compute/v1/pr
    --project=ihaffa-sandbox --region=us-central1 \
    --jobName=demo-exactly-once \
    --maxNumWorkers=30 \
    --streaming --enableStreamingEngine=true \
    --subsName=projects/ihaffa-sandbox/subscriptions/exactly-once-subs  \
    --outputTable=ihaffa-sandbox.beam_demo.exactly_once" \
    -Pdataflow-runner
```

```
mvn compile exec:java -Dexec.mainClass=org.apache.beam.examples.Demo  \
    -Dexec.args="--runner=DataflowRunner --gcpTempLocation=gs://ihaffa-sandbox
    --usePublicIps=false --subnetwork=https://www.googleapis.com/compute/v1/pr
    --project=ihaffa-sandbox --region=us-central1 \
    --jobName=demo-at-least-once \
    --maxNumWorkers=30 \
    --streaming --enableStreamingEngine=true \
    --subsName=projects/ihaffa-sandbox/subscriptions/at-least-once-subs  \
    --outputTable=ihaffa-sandbox.beam_demo.at_least_once  \
    --dataflowServiceOptions=streaming_mode_at_least_once" \
    -Pdataflow-runner
```

BEAM SUMMIT

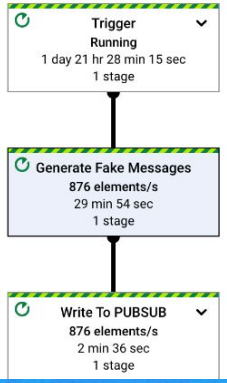
Trigger
Running
1 day 21 hr 28 min 15 sec
1 stage
Generate Fake Messages
876 elements/s
29 min 54 sec
1 stage
Write To PUBSUB
876 elements/s
2 min 36 sec
1 stage

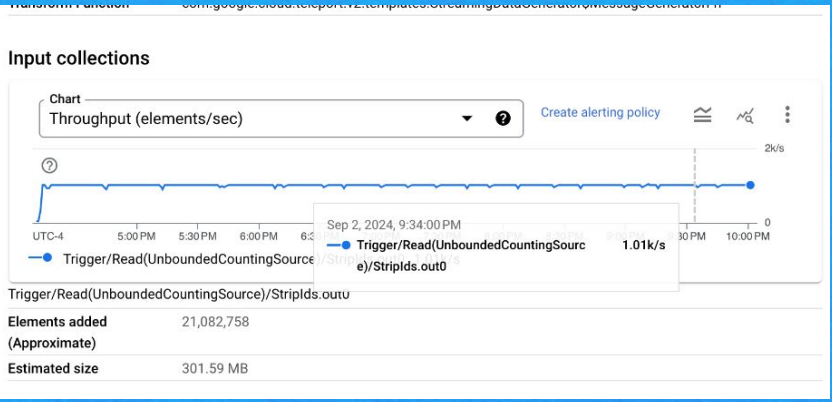
Input collections
Chart
Throughput (elements/sec)
Create alerting policy
2k/s
0
UTC-4
5:00 PM
5:30 PM
6:00 PM
10:00 PM
Sep 2, 2024, 9:34:00 PM
Trigger/Read(UnboundedCountingSource)/StripIds.out0
1.01k/s
Trigger/Read(UnboundedCountingSource)/StripIds.out0
Elements added (Approximate) 21,082,758
Estimated size 301.59 MB

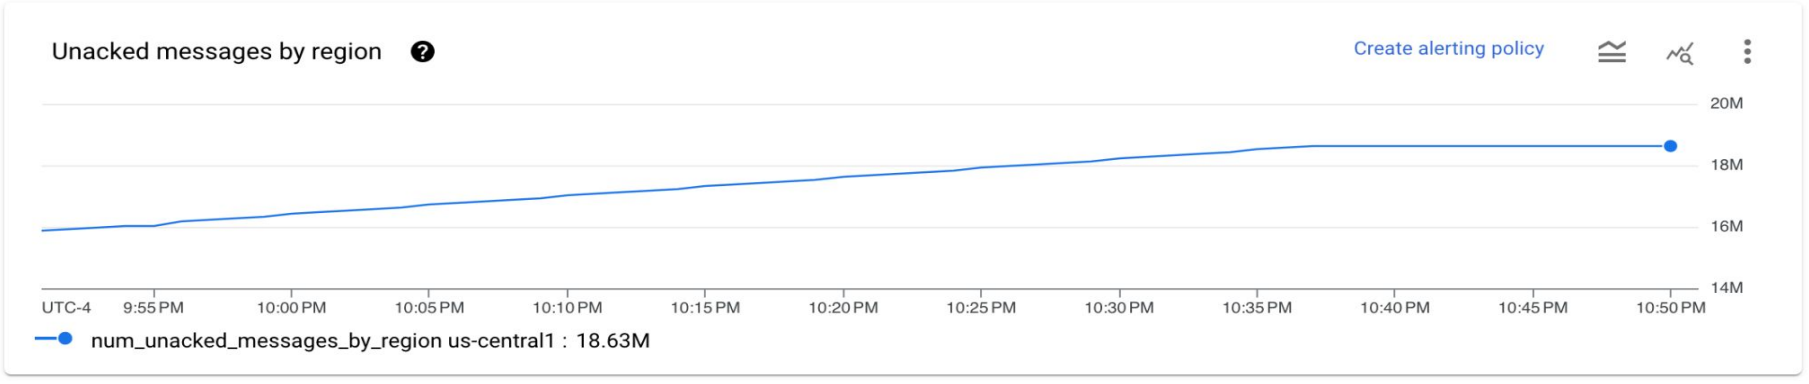
Unacked messages by region
Create alerting policy
20M
18M
16M
14M
UTC-4
9:55 PM
10:00 PM
10:05 PM
10:10 PM
10:15 PM
10:20 PM
10:25 PM
10:30 PM
10:35 PM
10:40 PM
10:45 PM
10:50 PM
num_unacked_messages_by_region us-central1 : 18.63M

Duplicates Metric

- Duplication will indicate 0 when using At-Least-Once

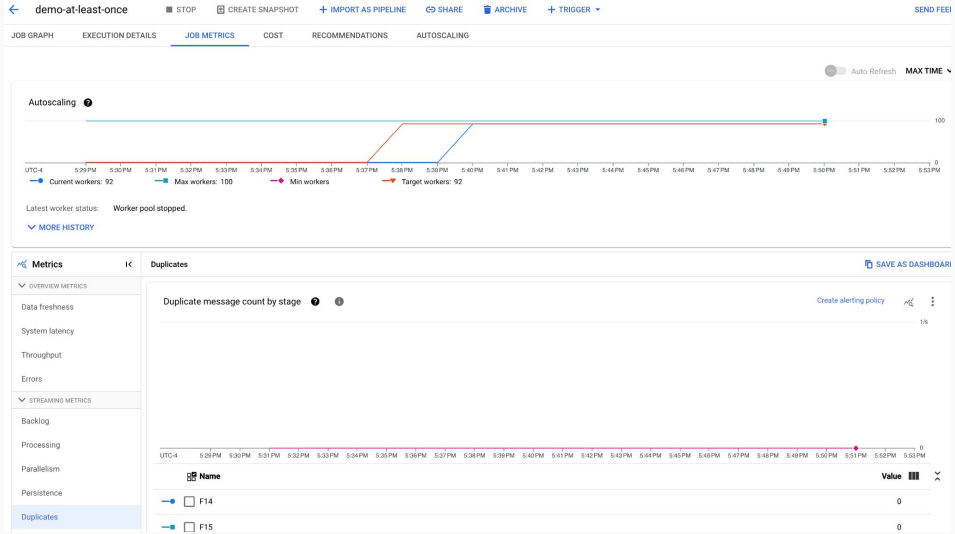

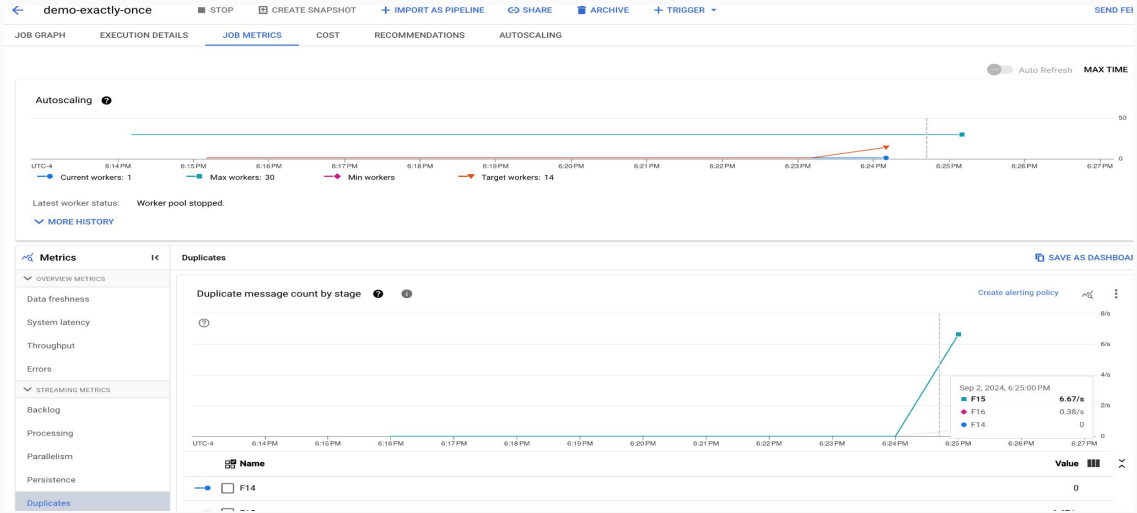

Queries On BQ

- Count exact number of duplicates

```
SELECT uuid, count(*) as duplicate FROM `ihaffa-sandbox.beam_demo.at_least_once`
group by uuid
having duplicate > 1);
```

ry results

INFORMATION | RESULTS | CHART | JSON | EXECUTION DETAILS | E

uuid	duplicate
57a47b0f-fb55-493e-95c7-d48f...	2
a5cbbe2e-c1ff-4272-9631-19b...	2
918d6685-0ad9-4713-bdfa-ee0...	2
30e2dc58-2084-4917-91e7-7ca...	2
76ee7883-4bb4-49a2-8206-13...	2

```
SELECT COUNT(*) AS `Total Duplicate` from(
SELECT uuid, count(*) as duplicate FROM `ihaffa-sandbox.beam_demo.at_least_once`
group by uuid
having duplicate > 1);
```

ery results

INFORMATION | RESULTS | CHART | JSON | EXECUTION DETAILS | E

Total Duplicate
838332

Cost Savings (At-Least-Once)

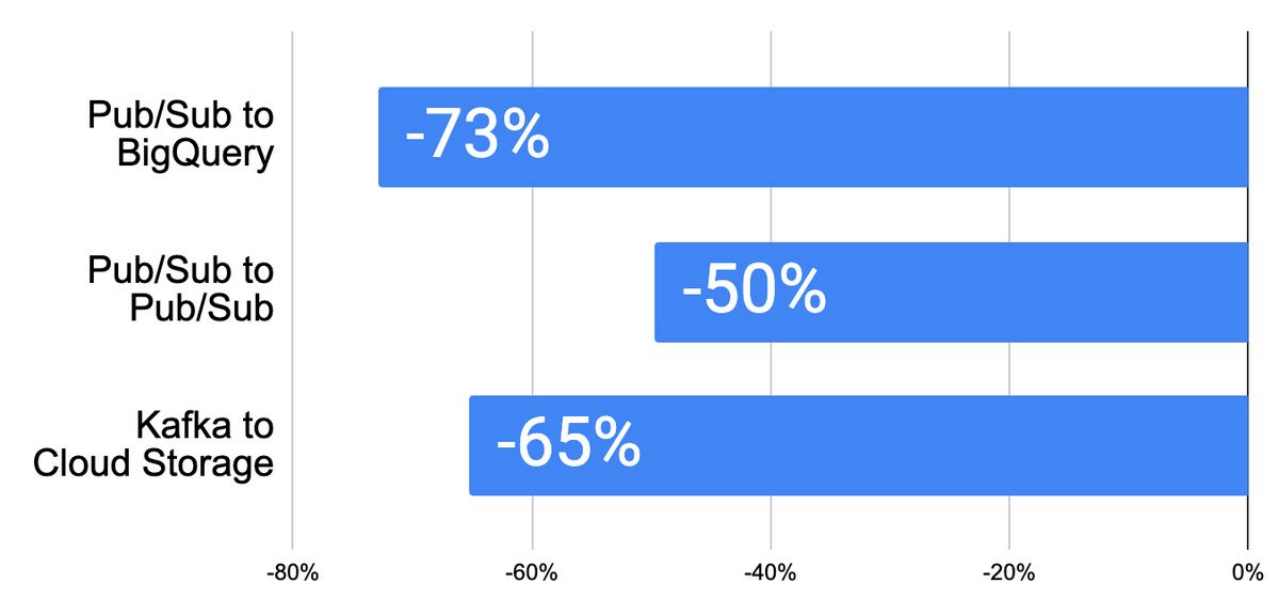

When To Use?

Case	At-Least-Once (New)	Exactly-Once (Default)
Pipeline	• Map-only Pipeline • Downstream app perform deduplication • Log processing, CDC	• Aggregation, e.g count, sum, mean • Business-critical: fraud, financing, etc
Other Consideration	• Cautious of external call • Cautious of high number of duplicates in pipeline	• Pubsub read is less performant

BEAM SUMMIT

Thank you!

Questions?

Blog: Dataflow At Least Once vs Exactly Once

Metric Exporter: Git Repo, Run Dataflow cost & performance benchmarking (Streaming & batch)

Email: lhaffa@google.com

BEAM SUMMIT