Grafana Labs

Demystifying Sampling

Juraci Paixão Kröhling
Software engineer
@jpkrohling

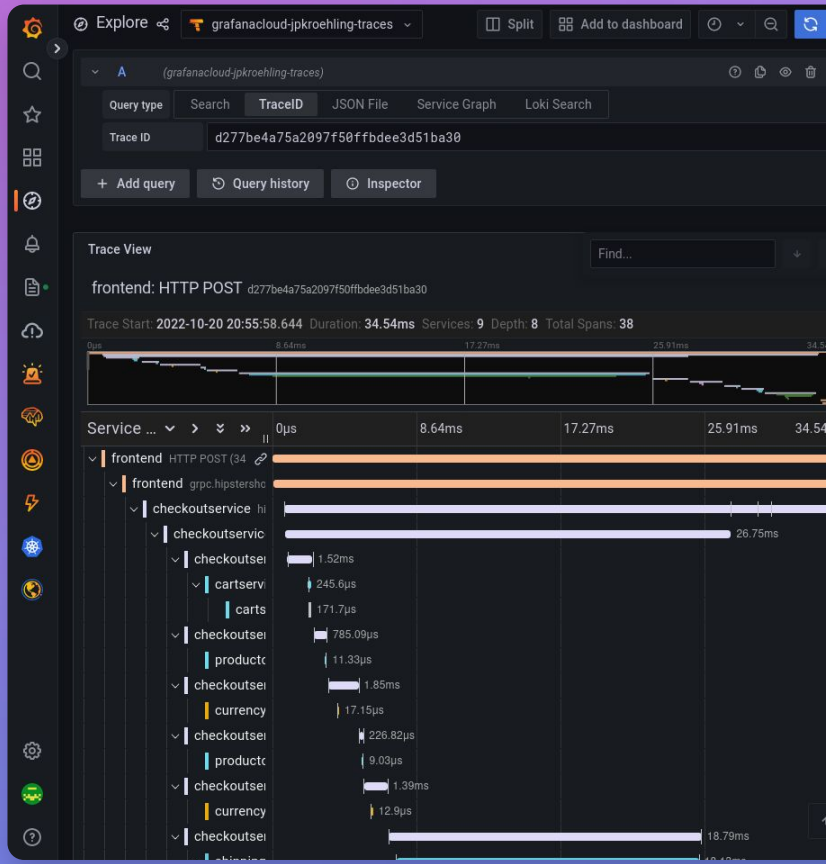

Speaker

Juraci Paixão Kröhling
Software Engineer

Agenda

1. Sampling?
2. Logs and sampling
3. Metrics and sampling
4. Traces and sampling
5. Bonus: Profiles and sampling
6. Questions and answers

Speaker

Juraci Paixão Kröhling
Software Engineer

About me

- Software engineer at Grafana Labs
- Governance Committee member for the OpenTelemetry project
- Cloud Native Computing Foundation (CNCF) Ambassador
- Maintainer of modules for OpenTelemetry Collector
- Jaeger emeritus maintainer
- OpenTracing emeritus maintainer

Sampling?

In statistics, quality assurance, and survey methodology, sampling is the selection of a subset or a statistical sample (termed sample for short) of individuals from within a statistical population to estimate characteristics of the whole population.

Source: https://en.wikipedia.org/wiki/Sampling_(statistics)

Sampling is the technique used to reduce the amount of telemetry data that is generated or stored, while retaining representativity.

Sampling and Logs

Sampling for logs

- Log levels!

Sampling for logs

- Log levels!
- Rate-limiting is the mostly used technique
- Usually done by the instrumentation library
- Prevents log flooding, especially under error conditions
    - Such as flooding the logs with stack traces that are currently happening for all incoming requests

Sampling for logs

- Go with Zap

```go
package main

import (
    "go.uber.org/zap"
)

func main() {
    logger, _ := zap.NewProduction()

    for i := 0; i < 100_000; i++ {
        logger.Info("Hello, world.")
    }
}
```

Sampling for logs

- Go with Zap

```
> go run . 2>&1 | wc -l
1099
```

Sampling for logs

- Java with Log4j

```
<BurstFilter level="INFO"
rate="16" maxBurst="100"/>
```

Sampling and Metrics

Sampling for metrics

- Downsampling
  - Typically on older data, to have longer retention in exchange for granularity
- Decreasing cardinality
  - Reduce the number of dimensions, have fewer time-series

Sampling for metrics - Downsampling

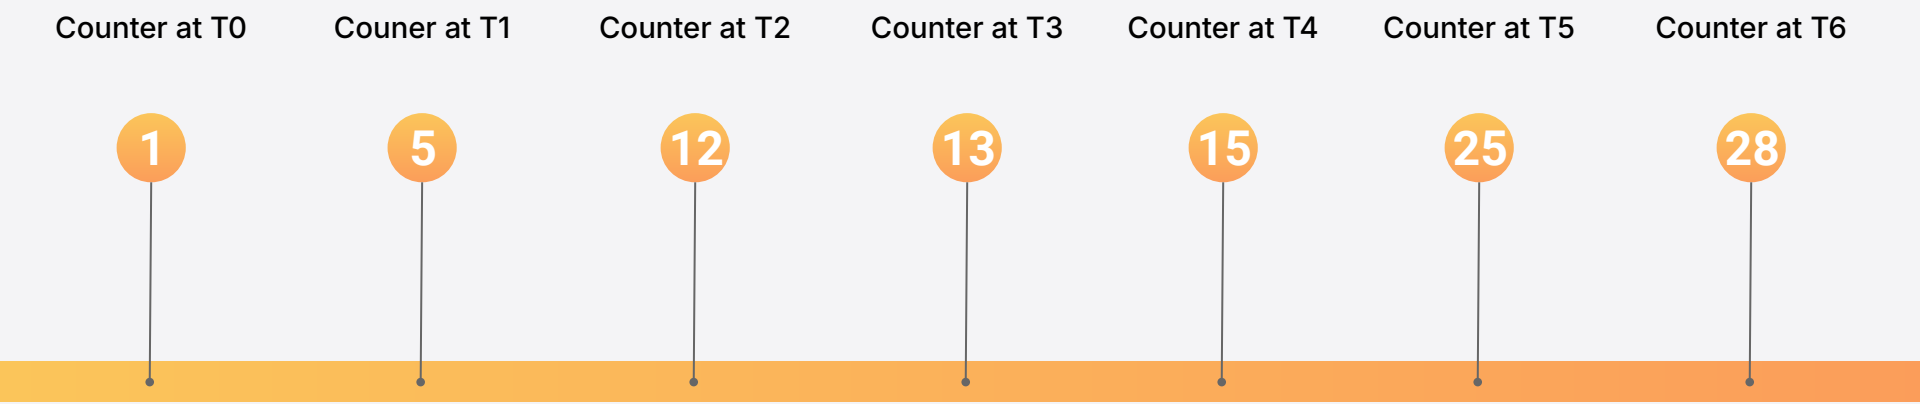

Sampling for metrics - Downsampling

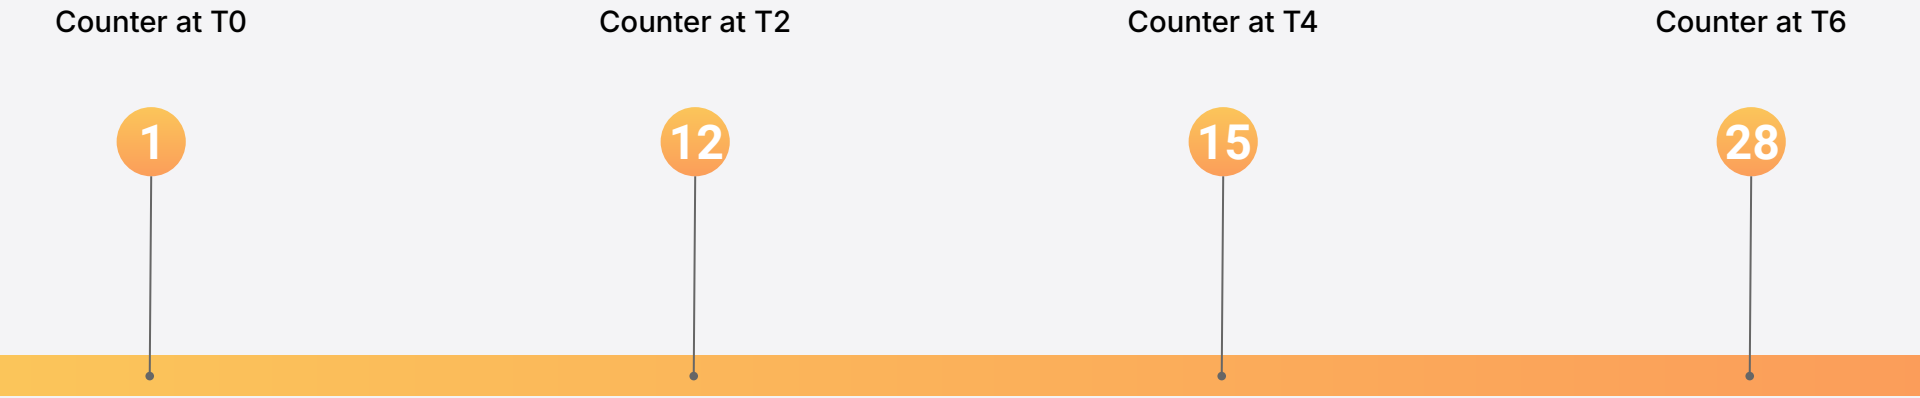

Sampling for metrics - Reduce cardinality

order_count{cluster="prod", userid=123}

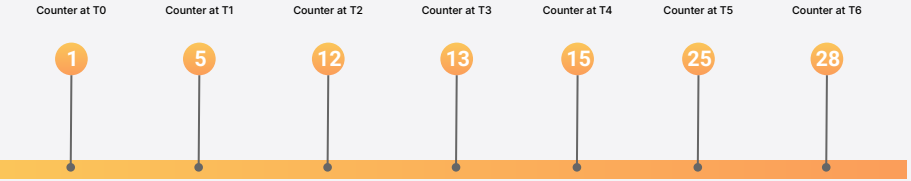

order_count{cluster="prod", userid=456}

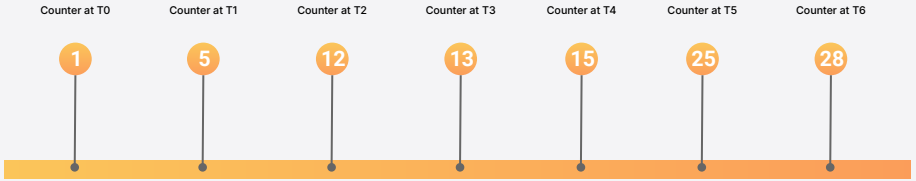

order_count{cluster="test", userid=789}

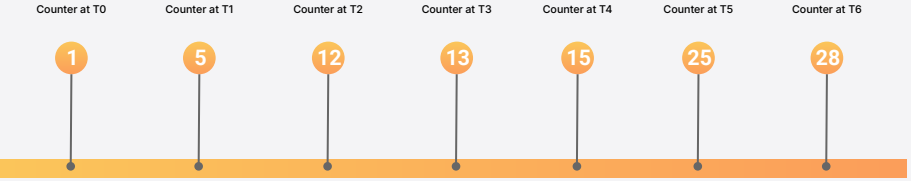

order_count{cluster="test", userid=321}

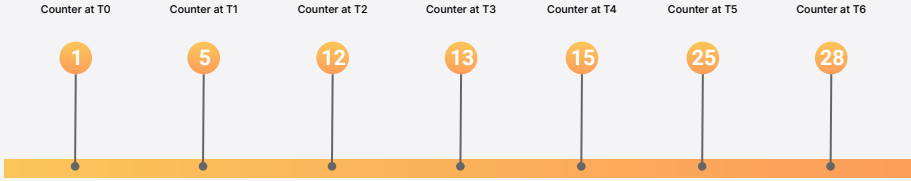

Sampling for metrics - Reduce cardinality

order_count{cluster="prod"}

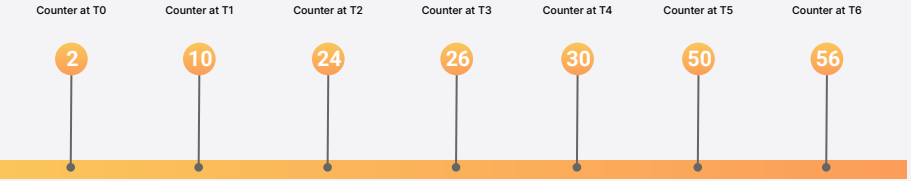

order_count{cluster="test"}

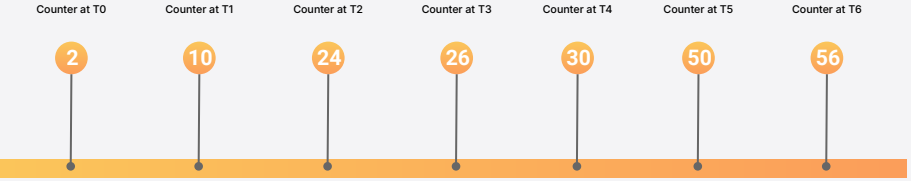

Sampling and Traces

Sampling for traces

- Different sampling strategies
  - Head sampling, at the beginning of the transaction
  - Consistent sampling, based on characteristics shared by all spans in the same trace (like the traceID)
  - Tail sampling, once the transaction has been finished
  - Out-of-band sampling, after the data has been persisted
- Different sampling policies
  - Probabilistic
  - Rate-limiting
  - Based on attributes
  - Adaptive
  - Based on trace characteristics
    - Was there a span in error?
    - Did the trace take longer than a threshold?

Head sampling

- Usually probabilistic
- Low network usage
- Low memory usage
- Simple to use
- Hard to do it consistently among services
- Hard to extract statistics out of it
  - Probability of 10% != 10%
  - Hard to know what was the probability at the time the span was created
  - How to extrapolate from highly different data?

Head sampling

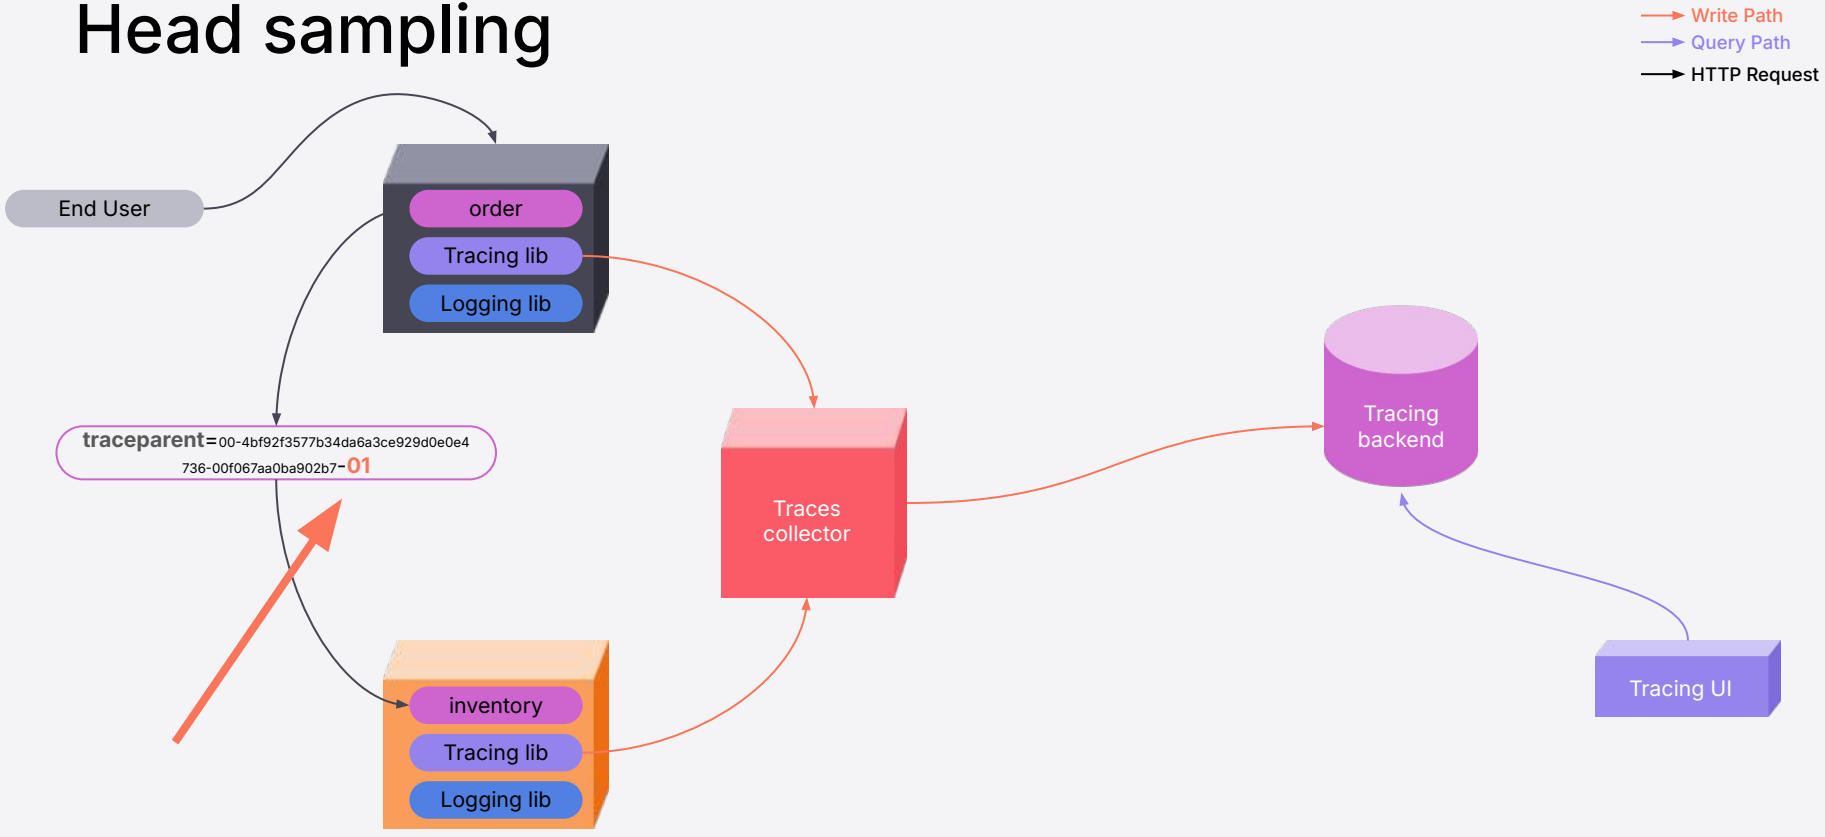

Head sampling

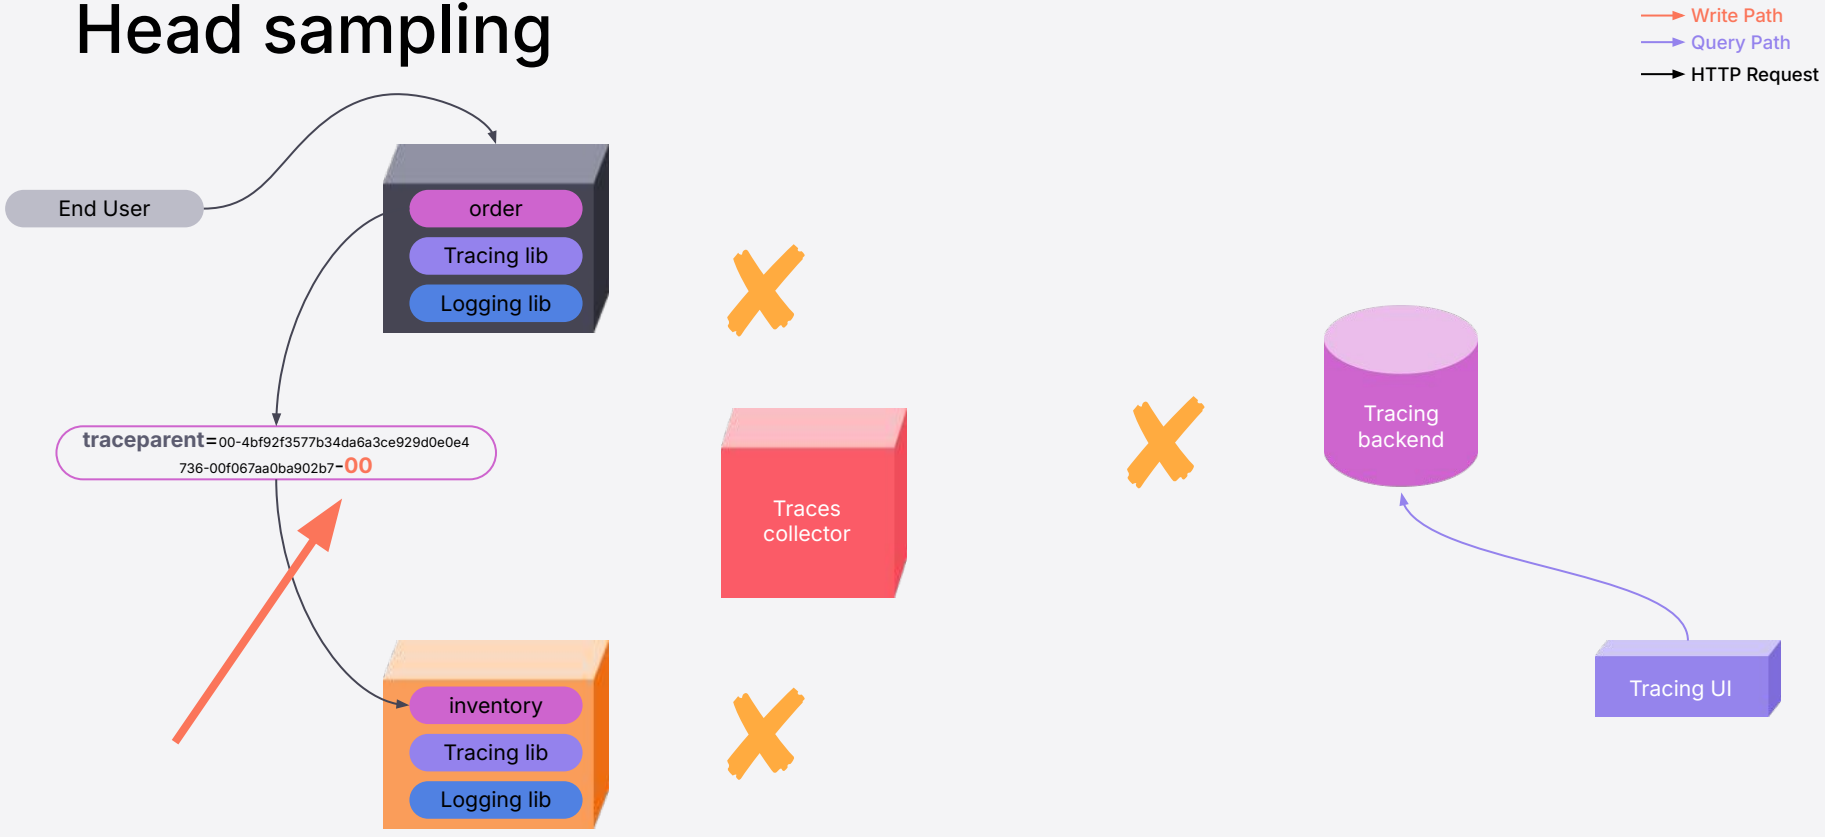

Consistent sampling

- Same decision for the spans on the same trace
  - Trace ID Ratio on OTel SDKs, for instance
  - Probabilistic sampling processor on OTel Collector, for instance
- Easier to do it consistently (one place to configure)
- More network usage
- Requires a central collection point
  - Comes with an ownership cost: engineering, computing, monitoring, ...
- Hard to extract statistics out of it

Consistent sampling SDK

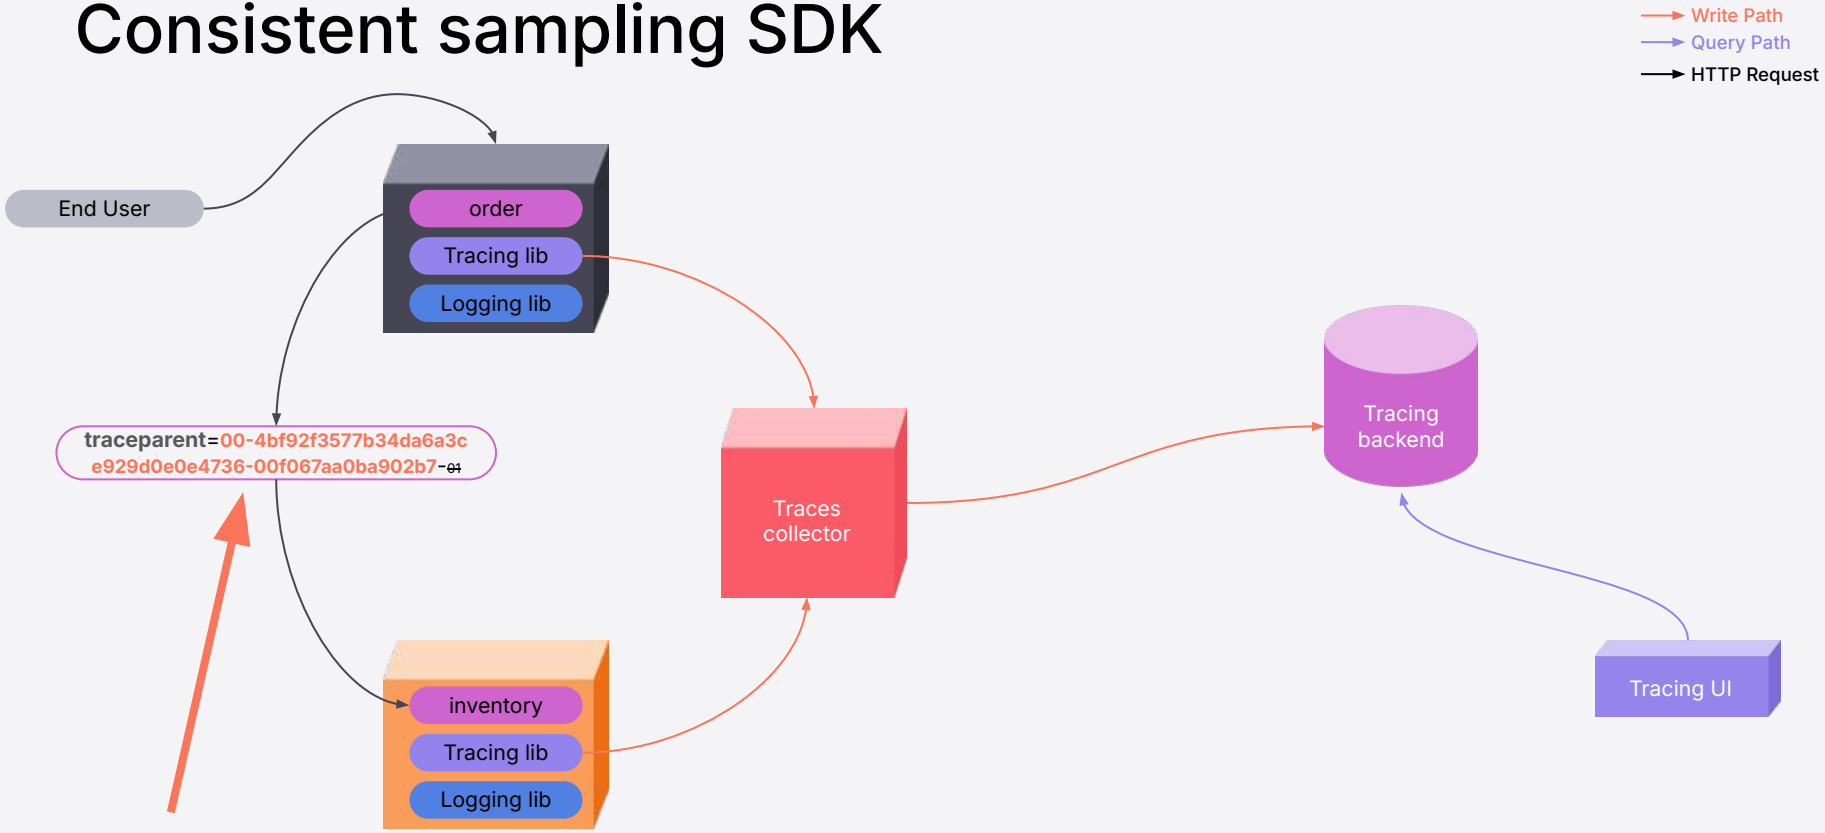

Consistent sampling SDK

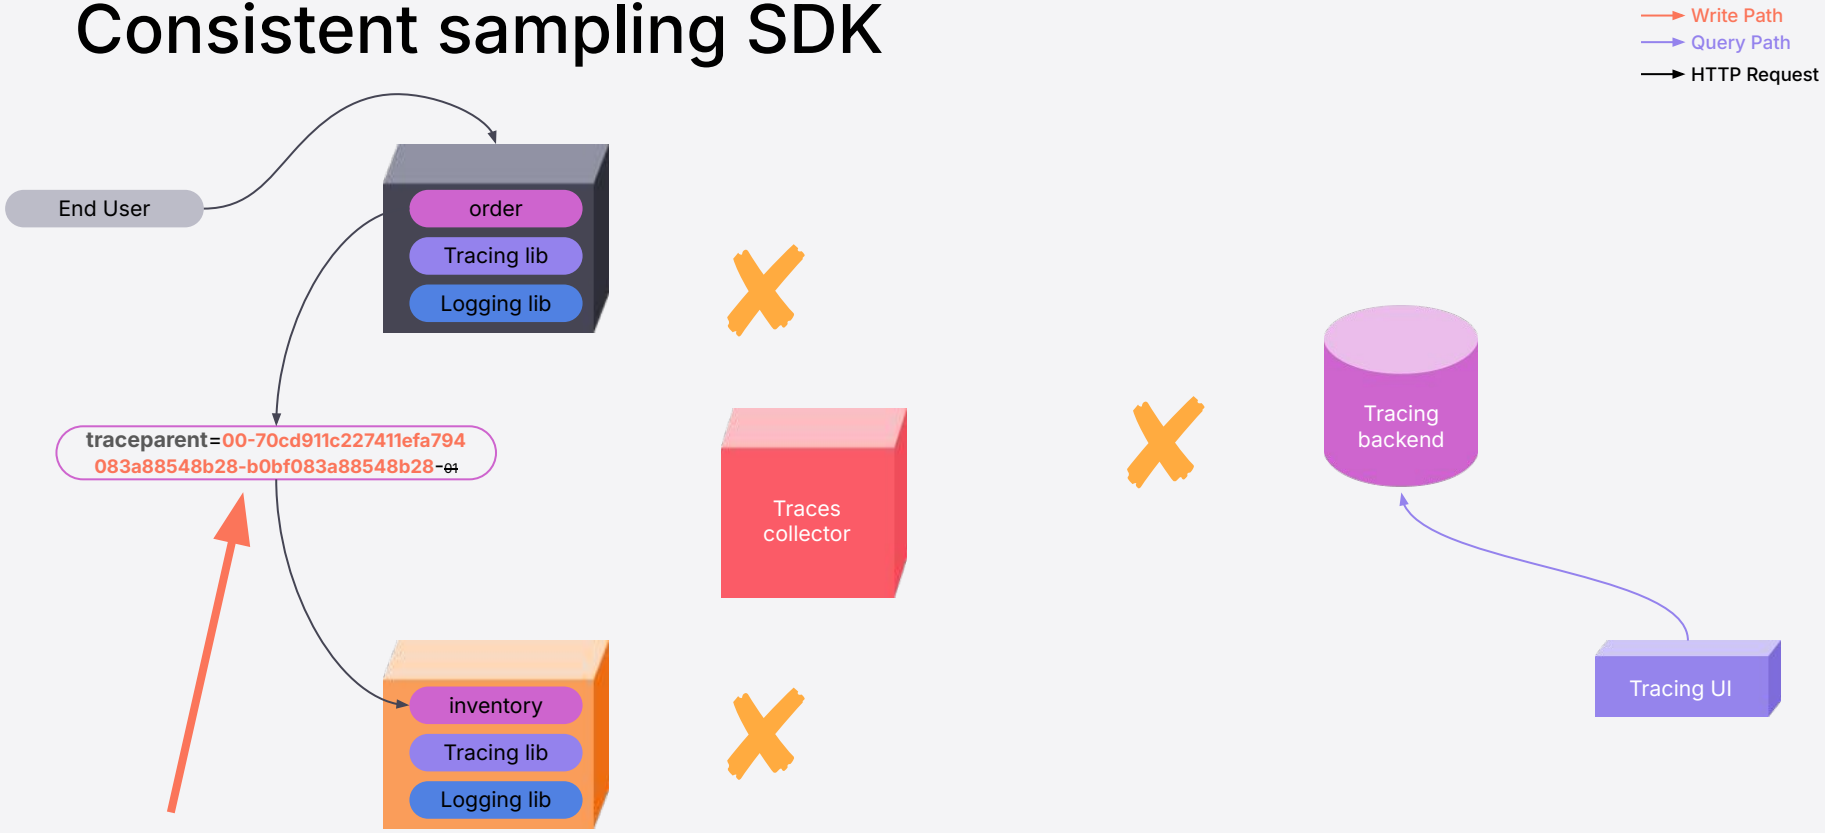

Consistent sampling - Collector

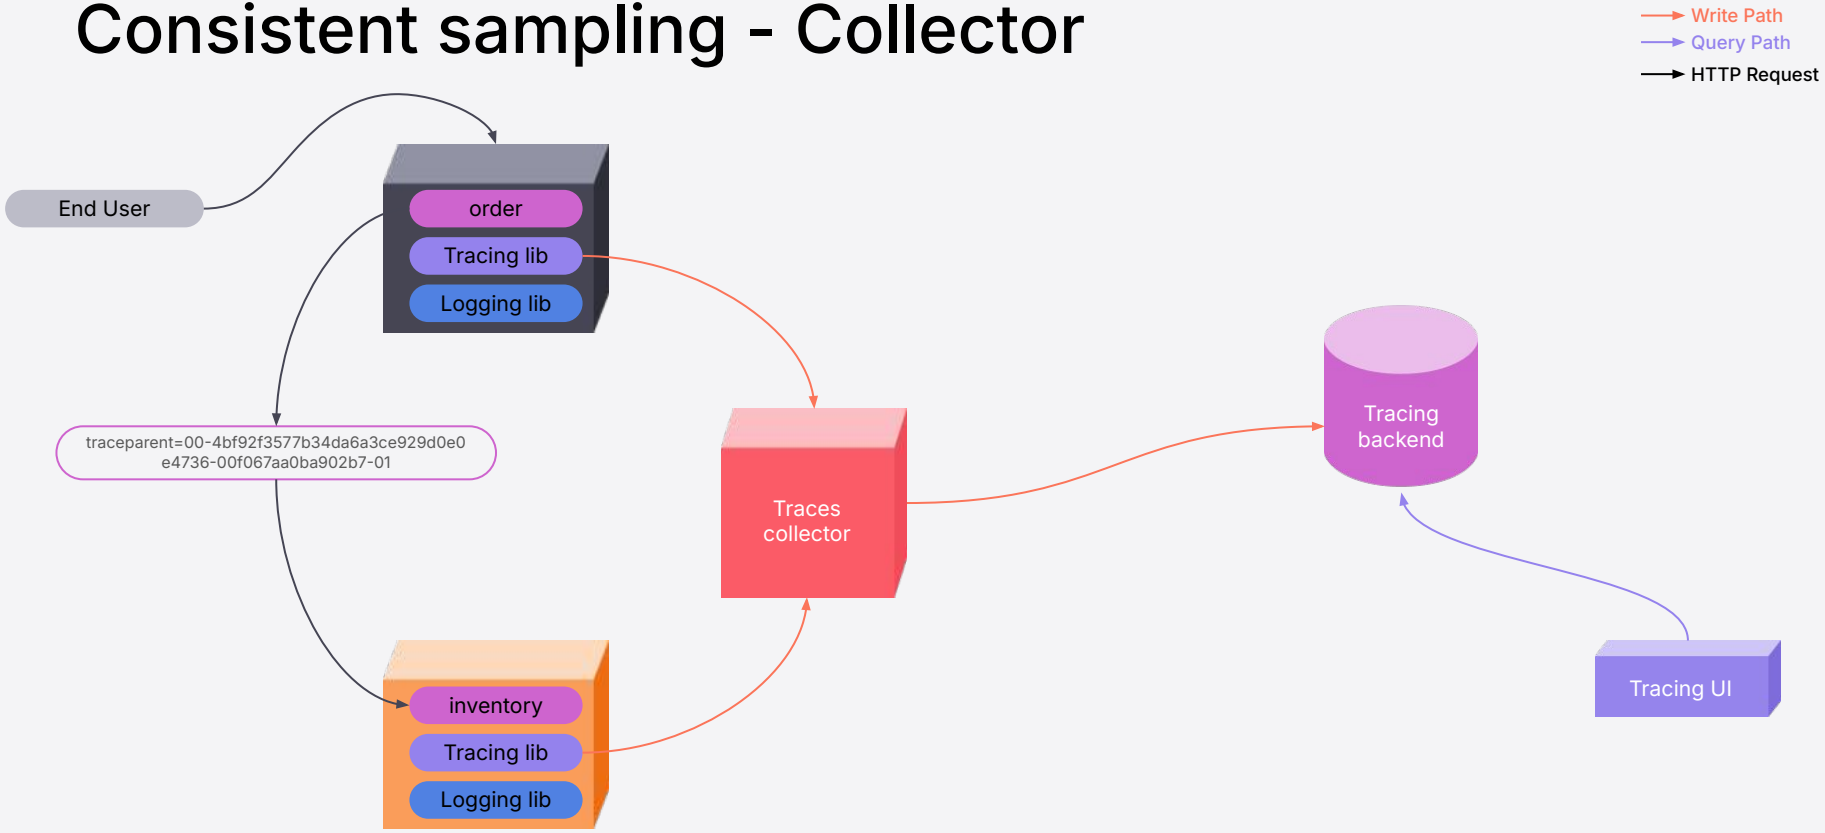

Consistent sampling - Collector

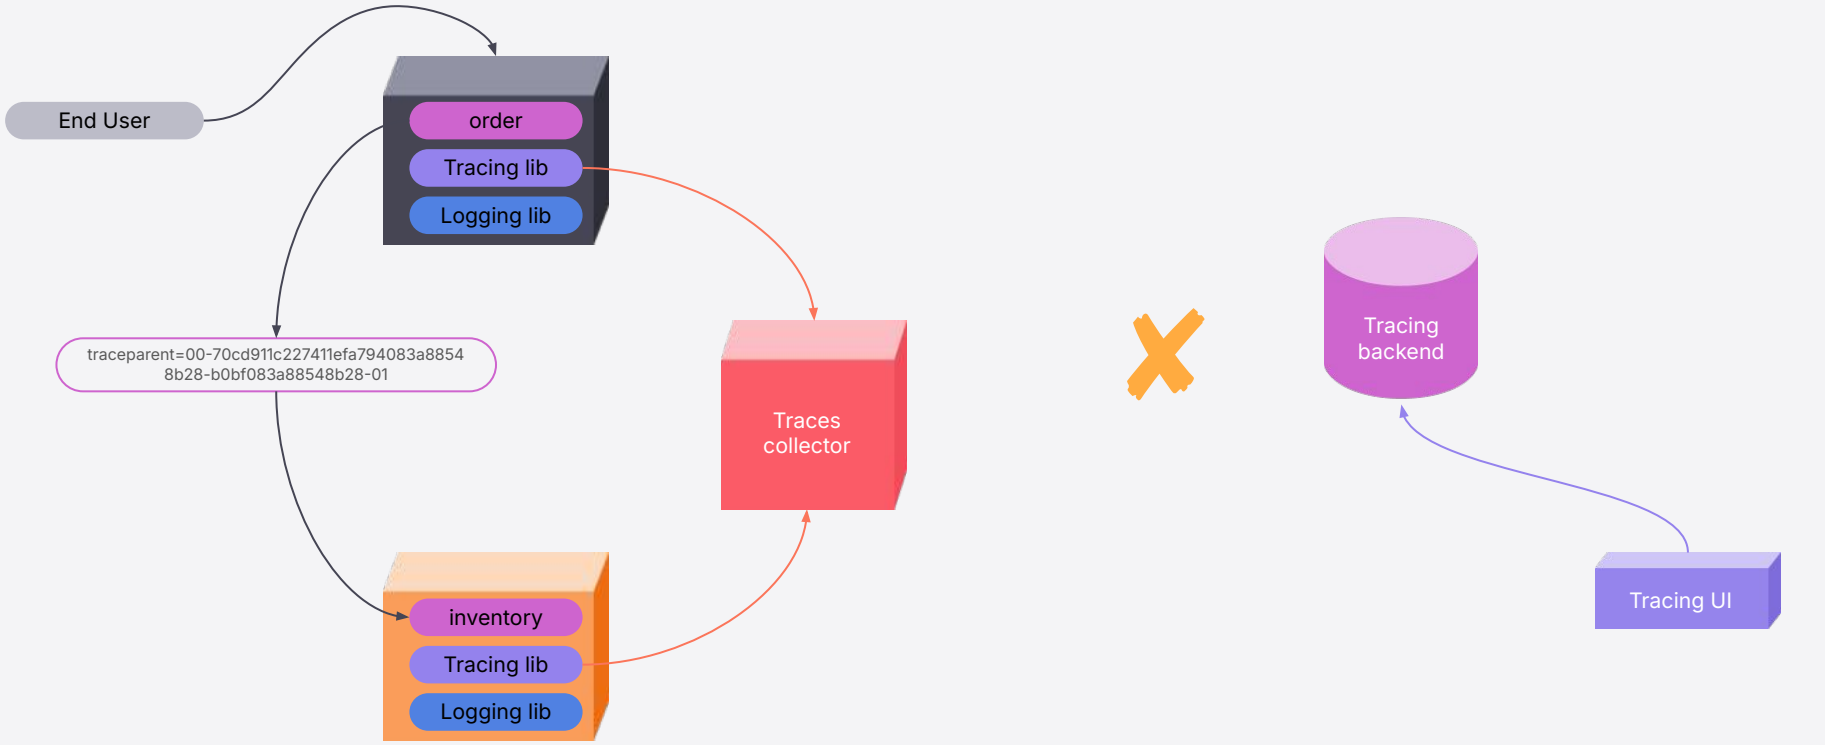

Tail sampling

- Highly interesting traces
- Allows for complex use-cases
- Requires some effort to scale
  - A load-balancer that can use the traceID to make the decision
- Higher cost of ownership
  - Higher resource consumption
  - Likely to have multiple instances of it in production
  - Likely to have another layer of collector for load-balancing
- Even harder to extract statistics out of it
  - What was the policy that caused a trace to be sampled? How many were discarded?

Tail sampling

```
tail_sampling:
  decision_wait: 1s
  num_traces: 50_000
  expected_new_traces_per_sec: 500
  policies:
    [
      {
        type: and,
        and:
          {
            and_sub_policy:
              [
                {
                  name: only-10-percent,
                  type: probabilistic,
                  probabilistic: { sampling_percentage: 10 },
                },
                {
                  name: vip,
                  type: string_attribute,
                  string_attribute: { key: vip, values: ["true"] },
                },
              ],
          },
      },
    ]
```

Tail sampling

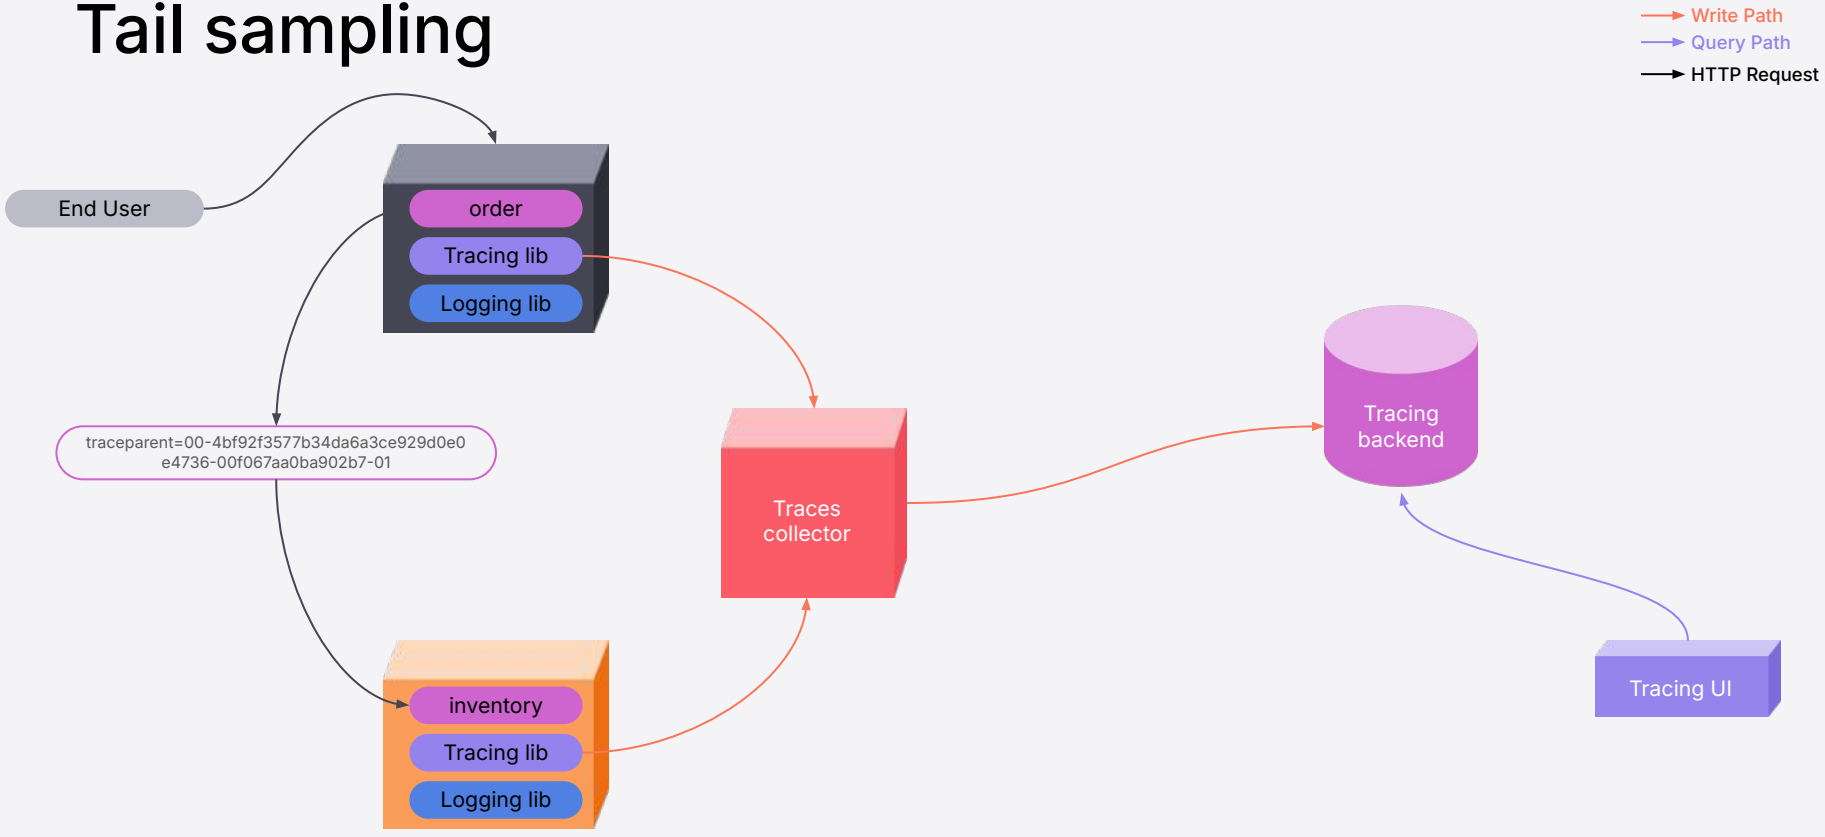

Tail sampling

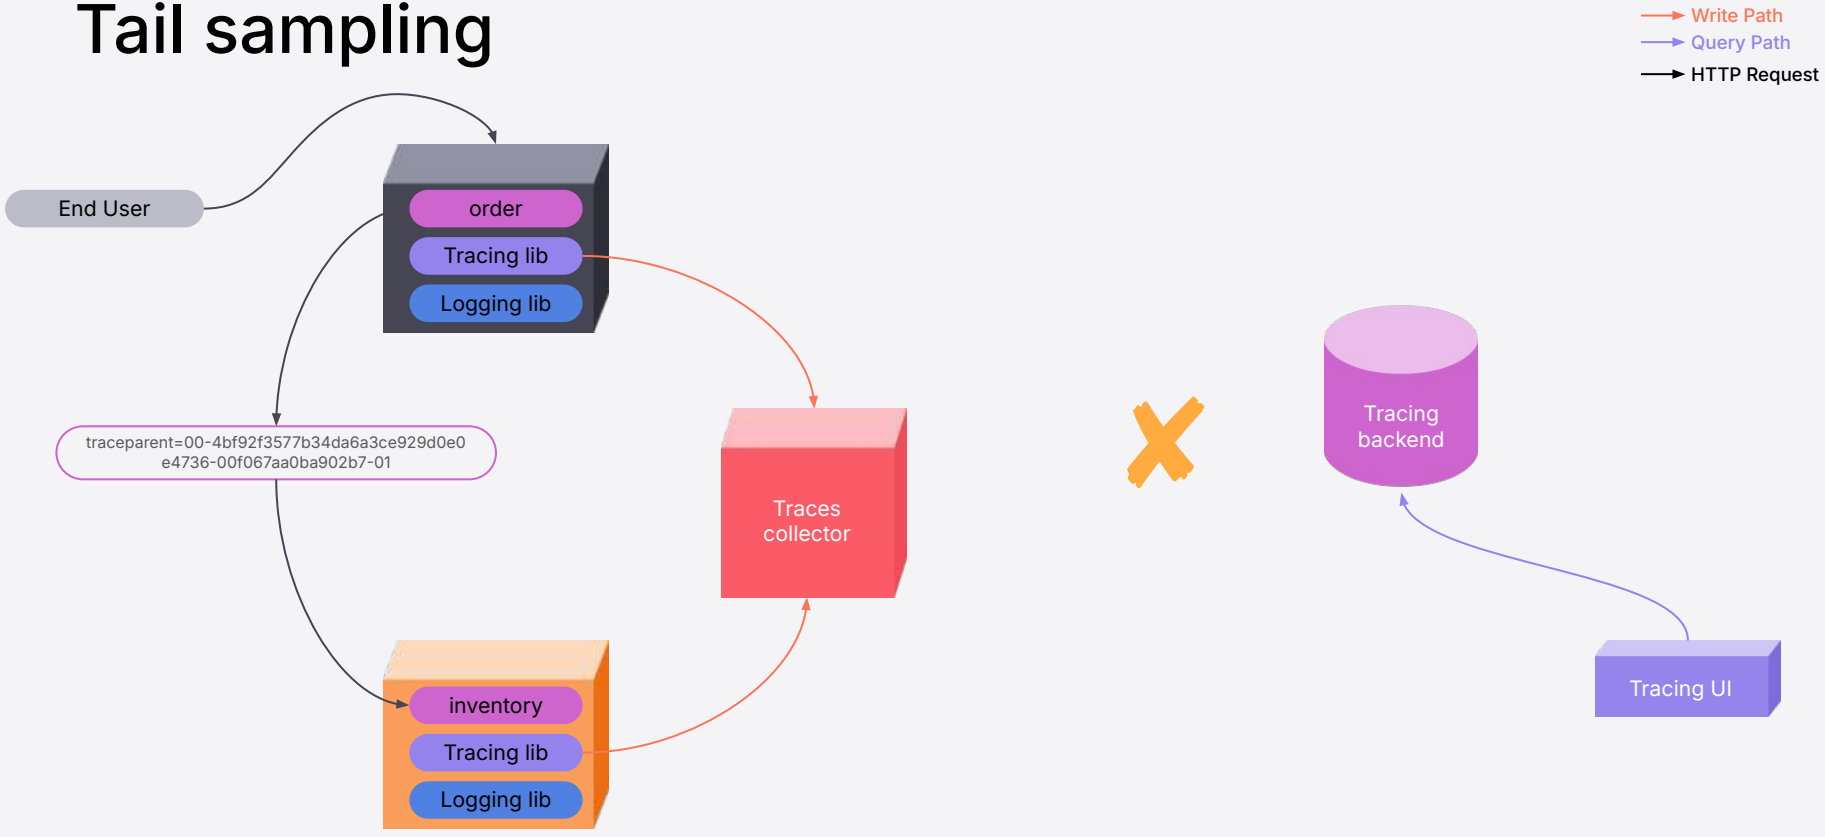

Hint: jpkrohling/otelcol-cookbook

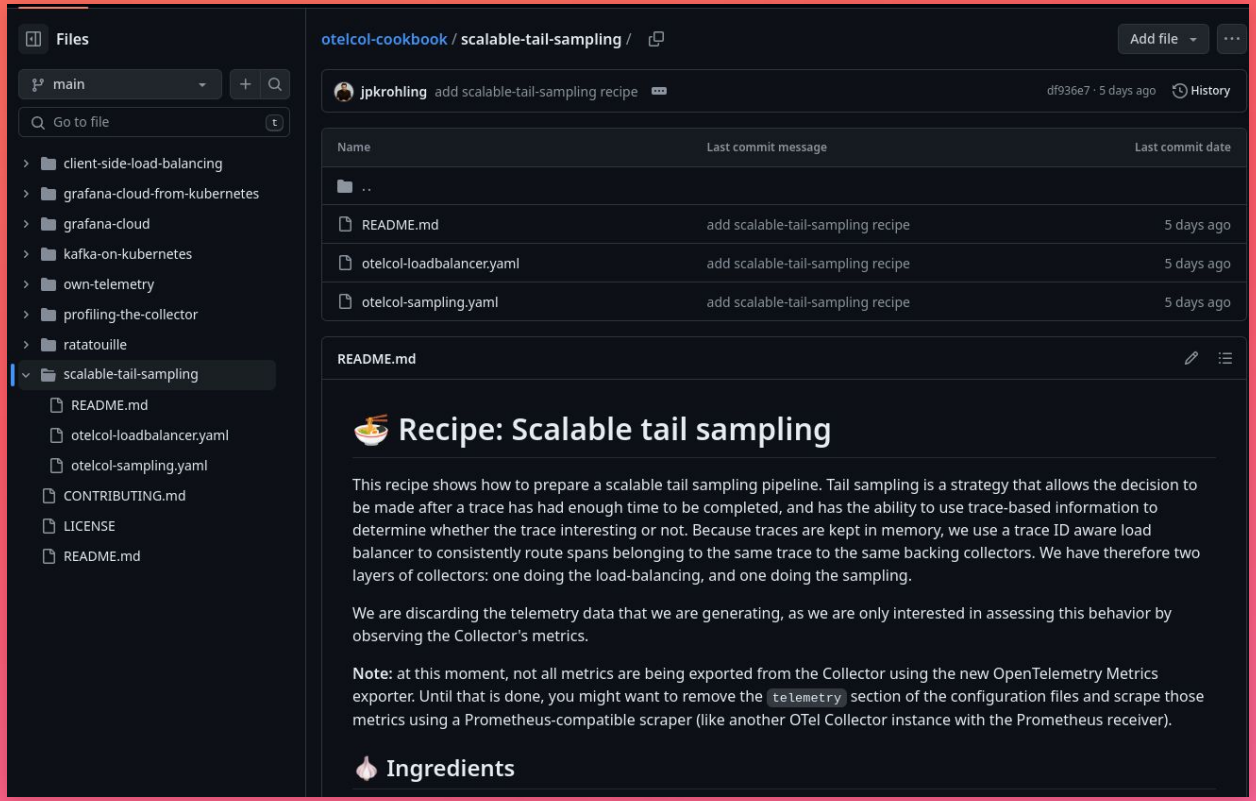

Out-of-band sampling

- Ingests everything, removes later
  - Perhaps with two stores
    - full data for a short period of time
    - sampled data for higher retention
- High-throughput systems
- Simplified architecture
- Requires development of custom components
- Statistics are possible!

Out-of-band sampling

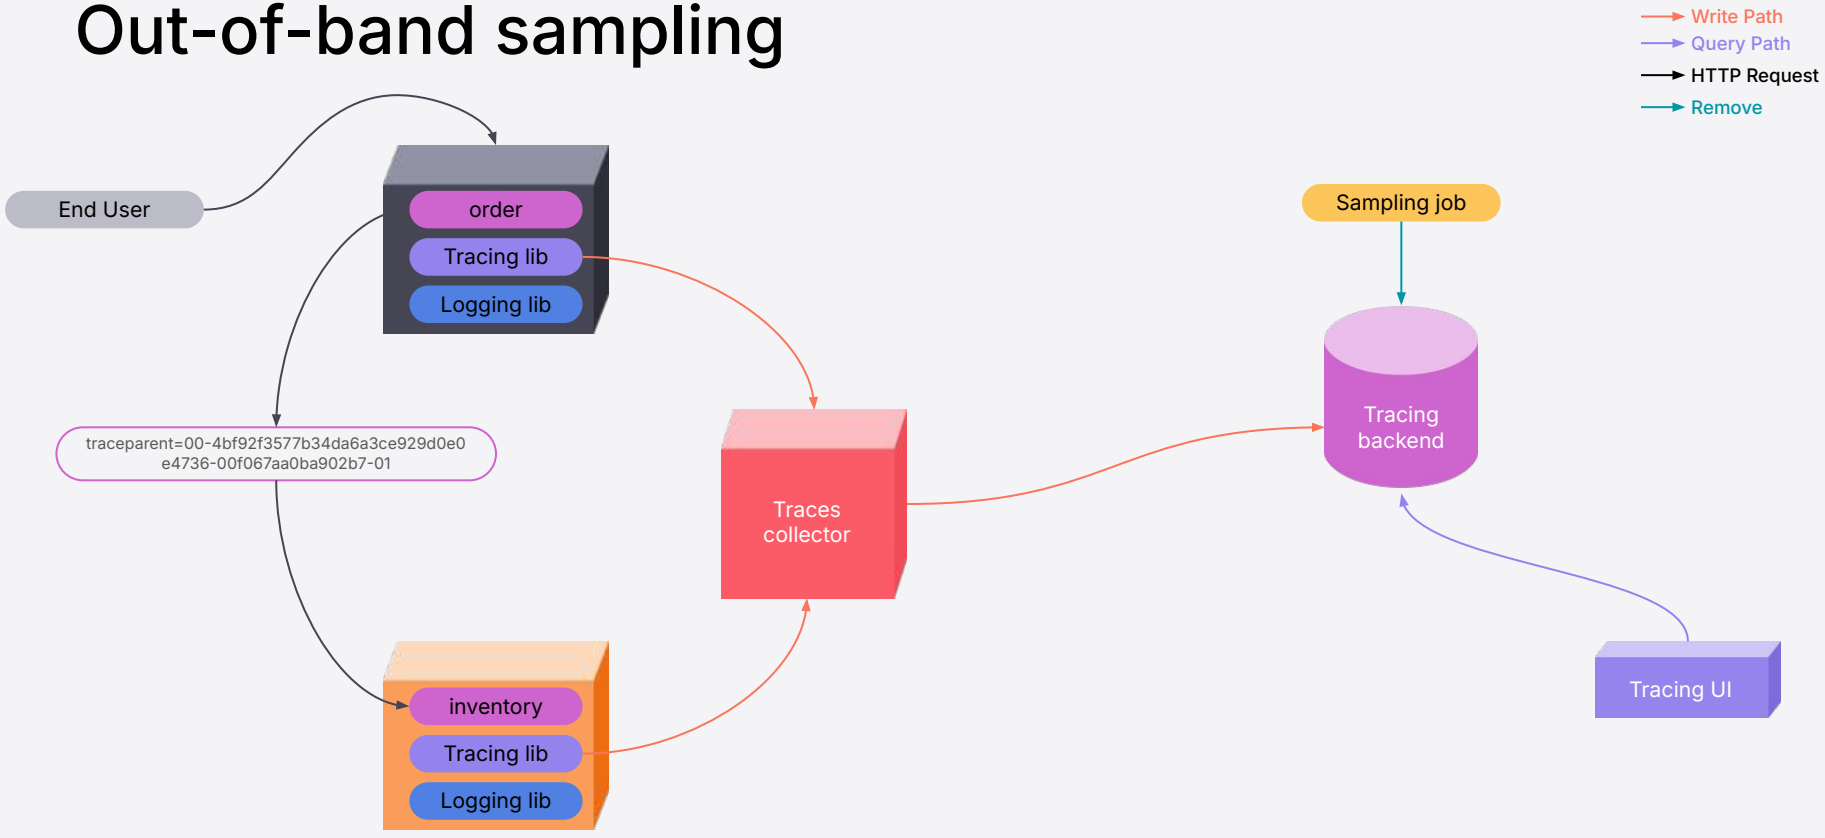

Tip:

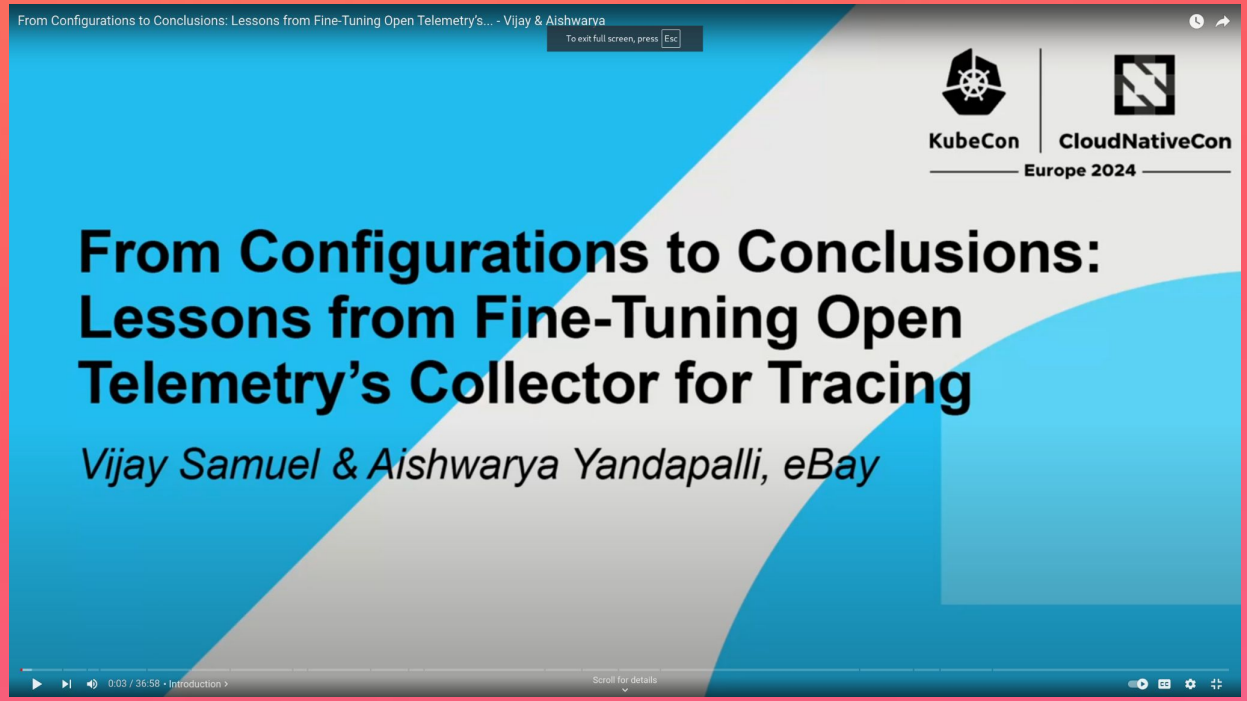

https://youtu.be/RSJwv1jOdTg

Bonus: sampling and Profiles

Sampling for profiles

- Profiling frequency
- Profiling time window
- A fraction of the instances of a service

Q&A time!

Grafana Labs

Obrigado!

@jpkrohling