Tracing is more than traces: the insights in trace aggregates

Daniela Miao
LightStep

@dymxzvf

Quick review

Quick review

Microservices may be here to stay

… but they broke our old tools.

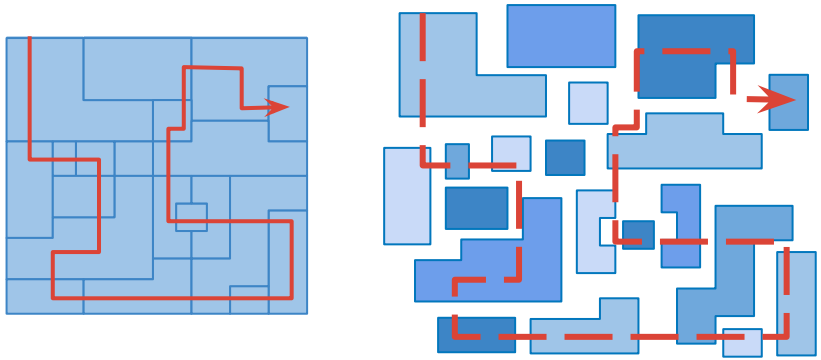

A trace tells a story in a distributed system: end-to-end

A single trace

Everyone knows this, so let's move on quickly

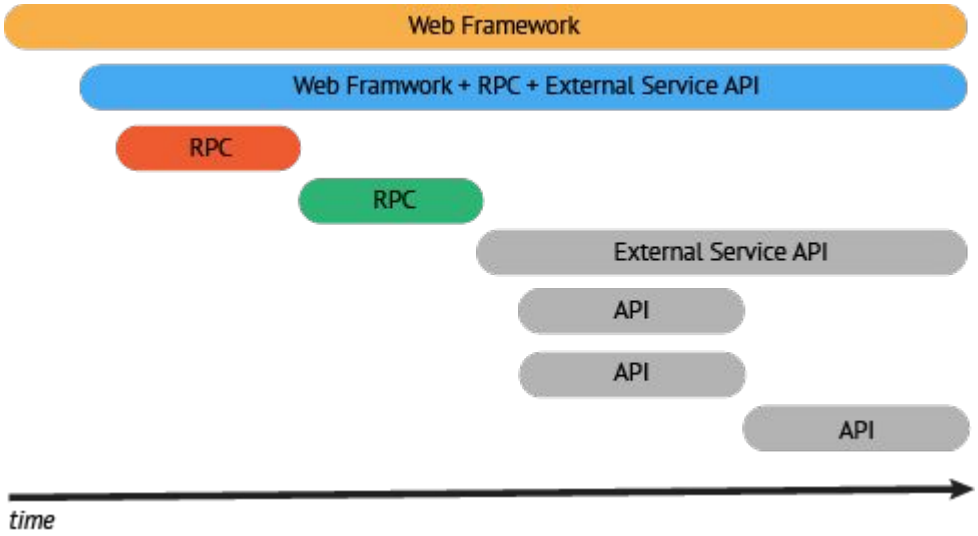

Image source: https://opentracing.io/docs/best-practices/instrumenting-your-application/

It’s only one trace though

Tracing as a cost-effective solution always involves sampling

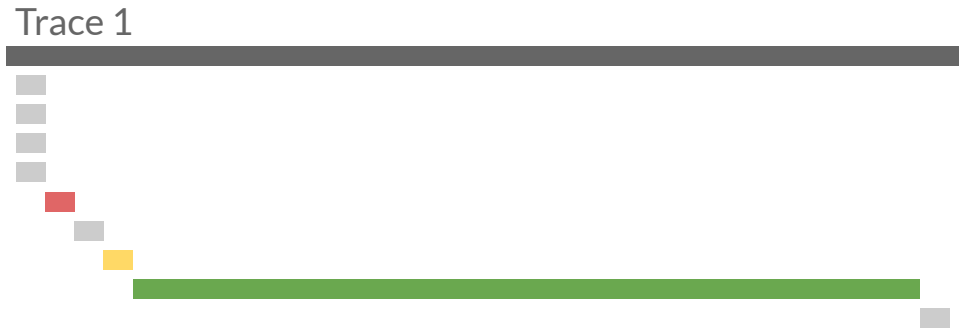

It’s only one trace though

Tracing as a cost-effective solution always involves sampling

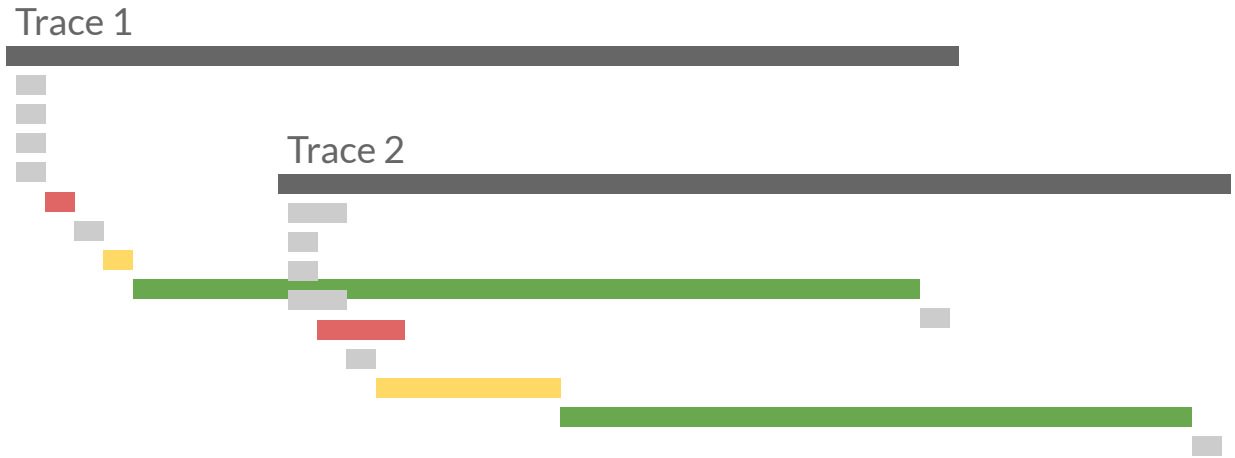

It’s only one trace though

Tracing as a cost-effective solution always involves sampling

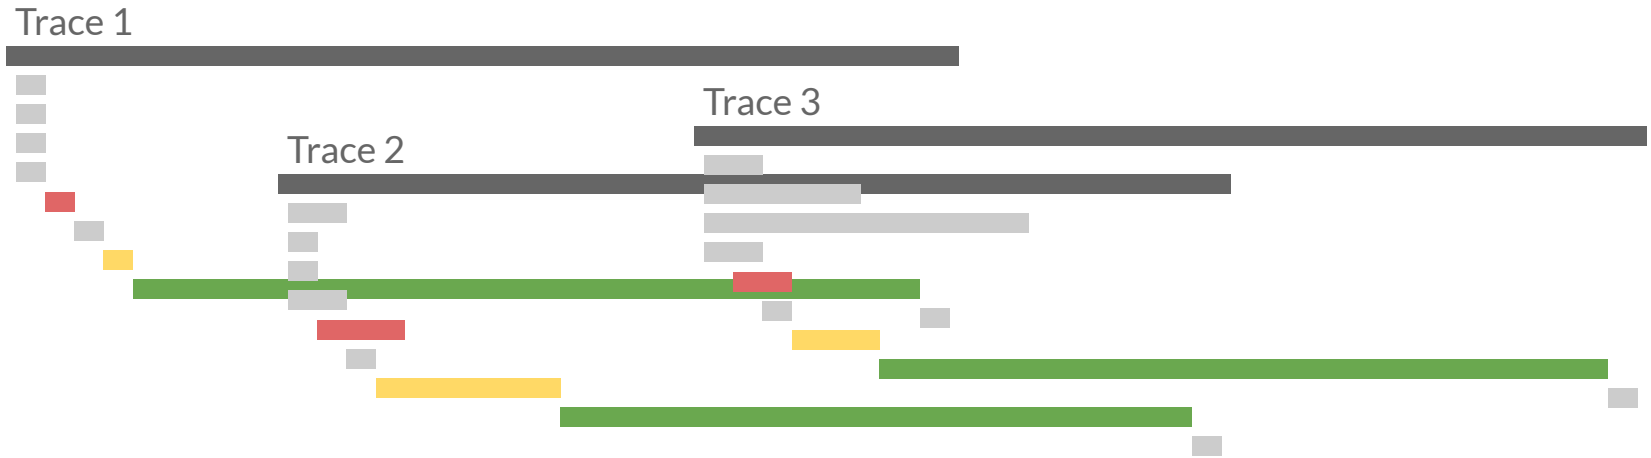

It’s only one trace though

Tracing as a cost-effective solution always involves sampling

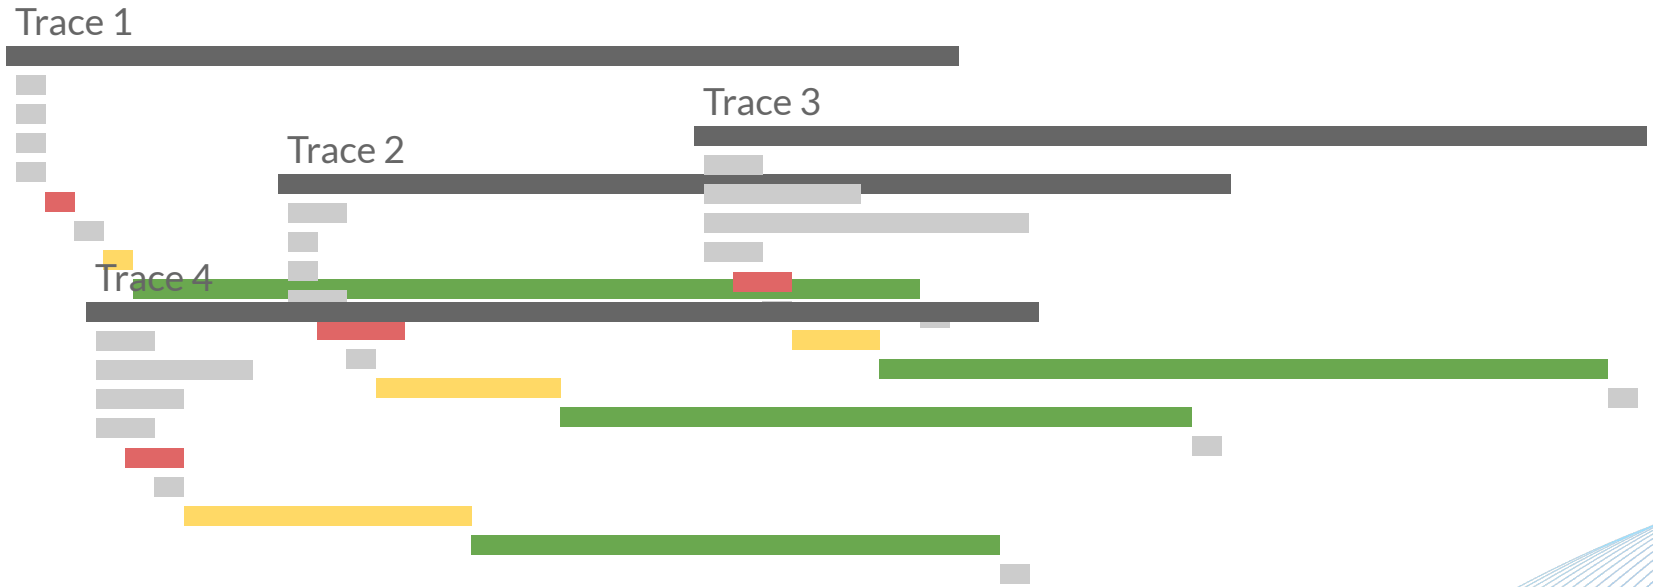

It’s only one trace though

Tracing as a cost-effective solution always involves sampling

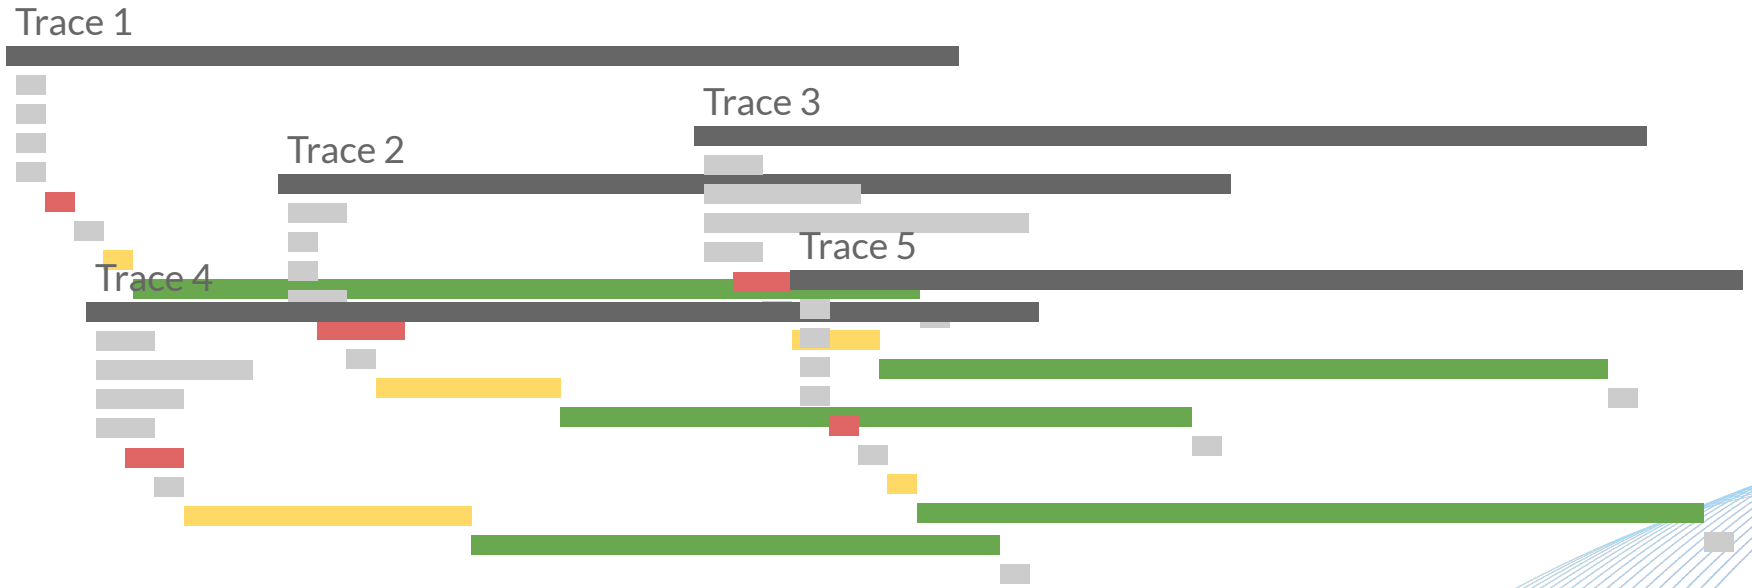

It’s only one trace though

Tracing as a cost-effective solution always involves sampling

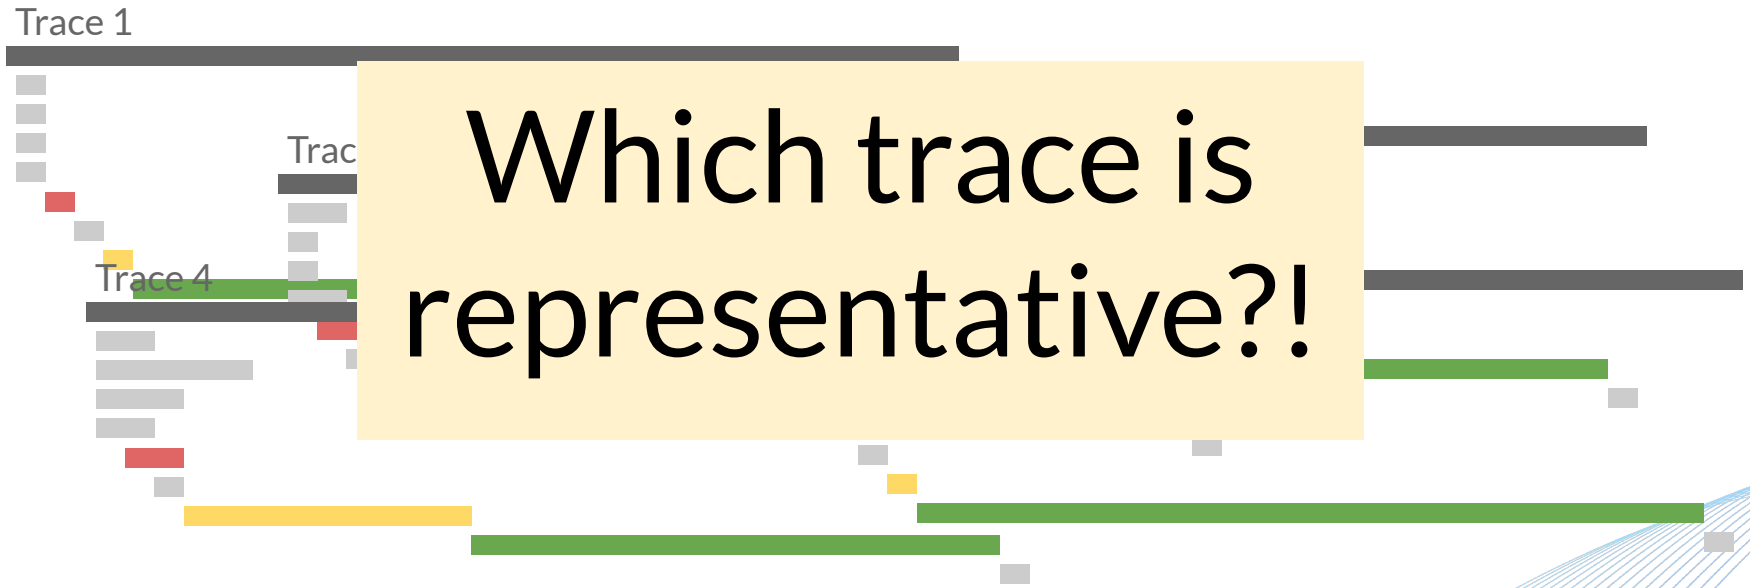

Trace aggregate analysis

- Correlating ANY characteristic of the system with metrics
- Latency anomalies in context of service infrastructure
- Critical path analysis => resource contention

Trace aggregate analysis

- Correlating ANY characteristic of the system with metrics
- Latency anomalies in context of service infrastructure
- Critical path analysis => resource contention

Correlation based on statistical analysis

Symptom: metrics out of whack

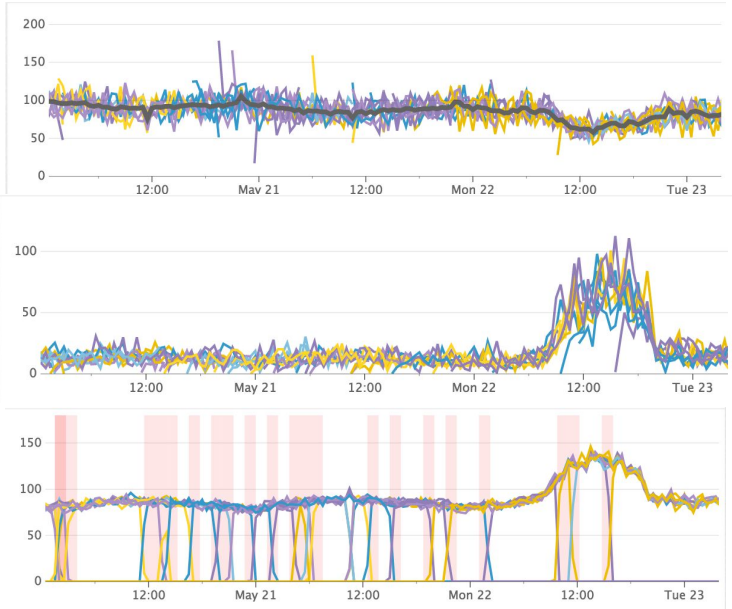

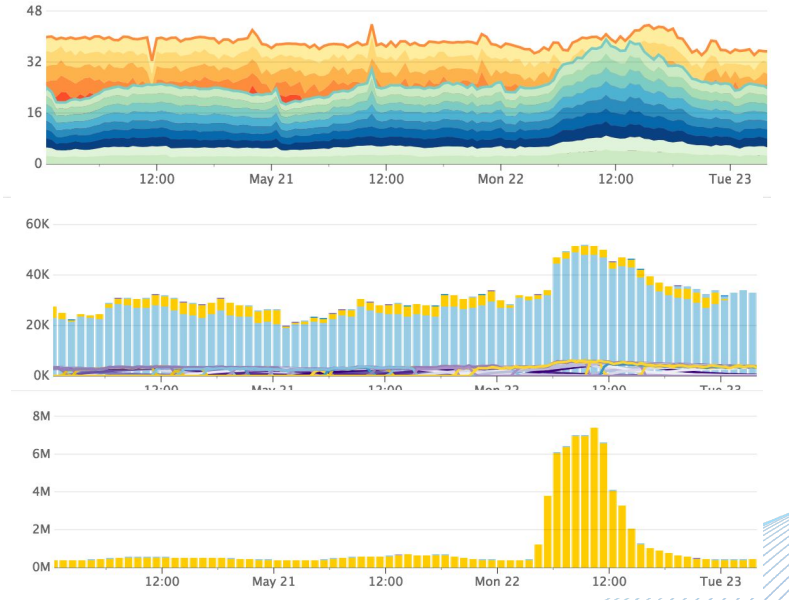

@dymxzvf

Correlation based on statistical analysis

Root cause: 1 customer inadvertently DDOS-ing the system

How long does it take to get from Symptom to Root Cause?

Too long: knowing what to segment by is hard

Correlation based on statistical analysis

Low Cardinality = Easy

High Cardinality (thousands of customers) = Hard, and Expensive

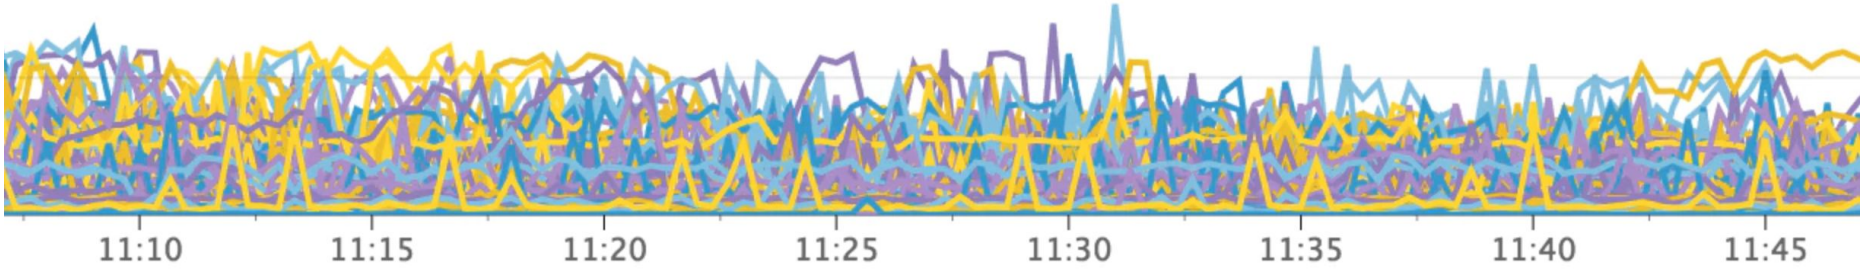

Correlation based on statistical analysis

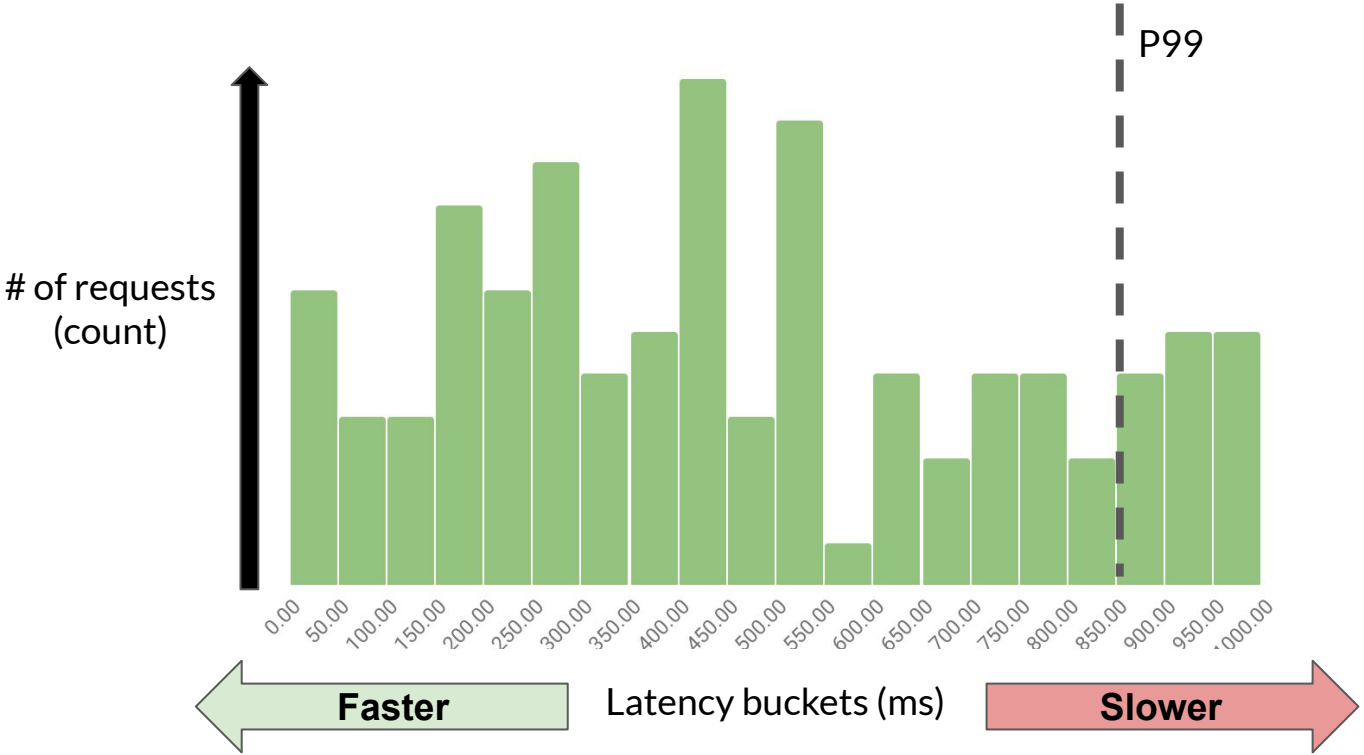

Correlation based on statistical analysis

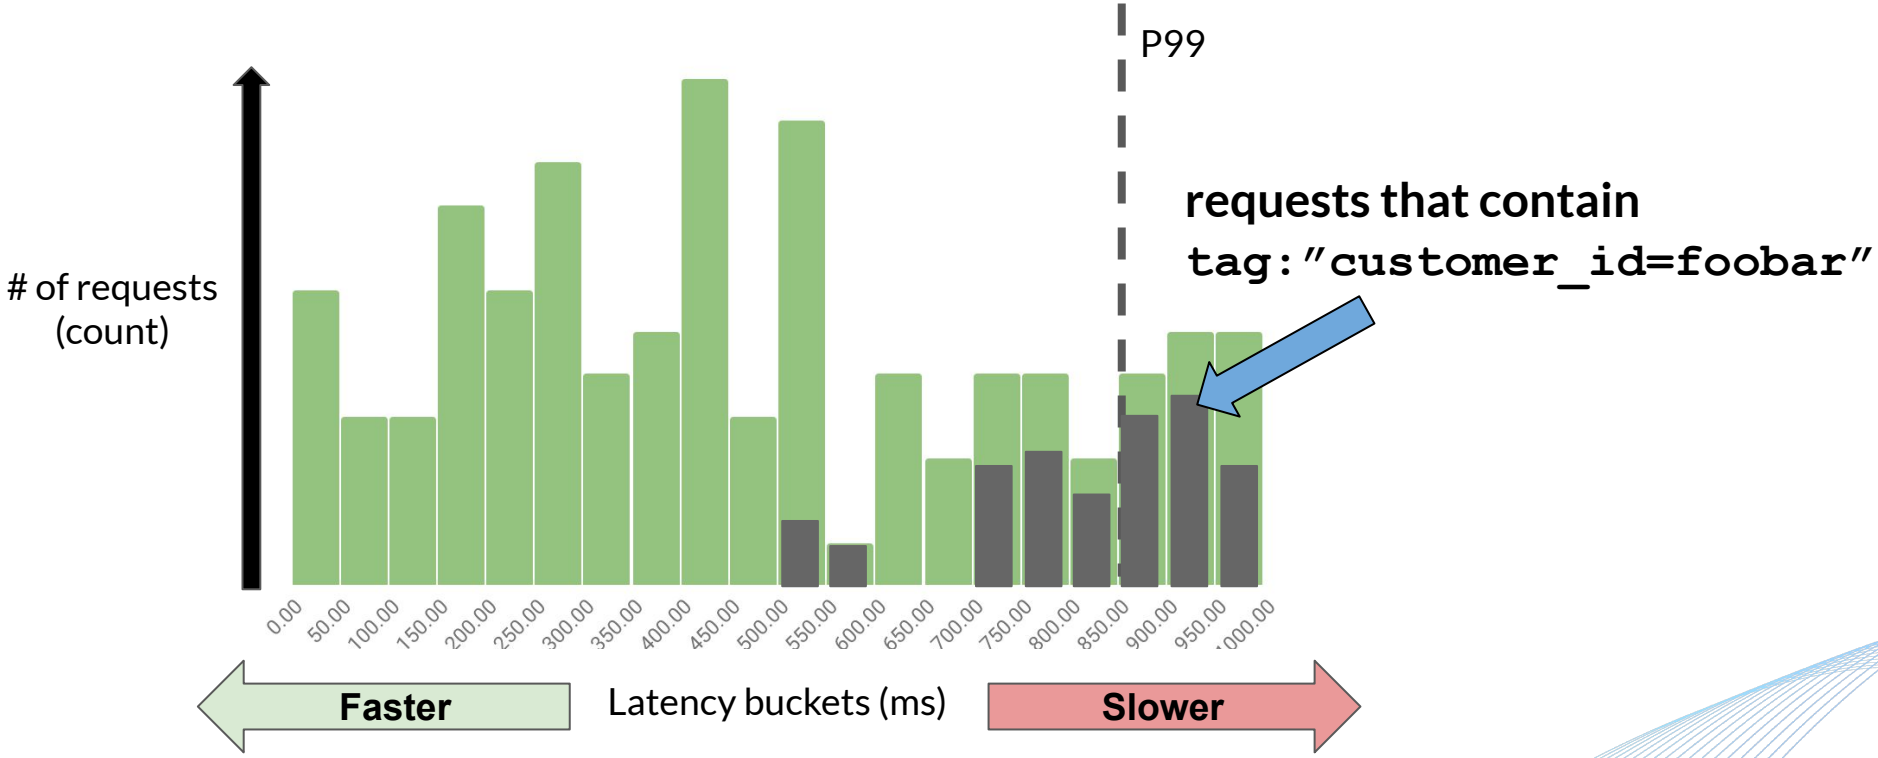

Latency anomalies in context of service infrastructure

One faulty network card: explain your p99.9!

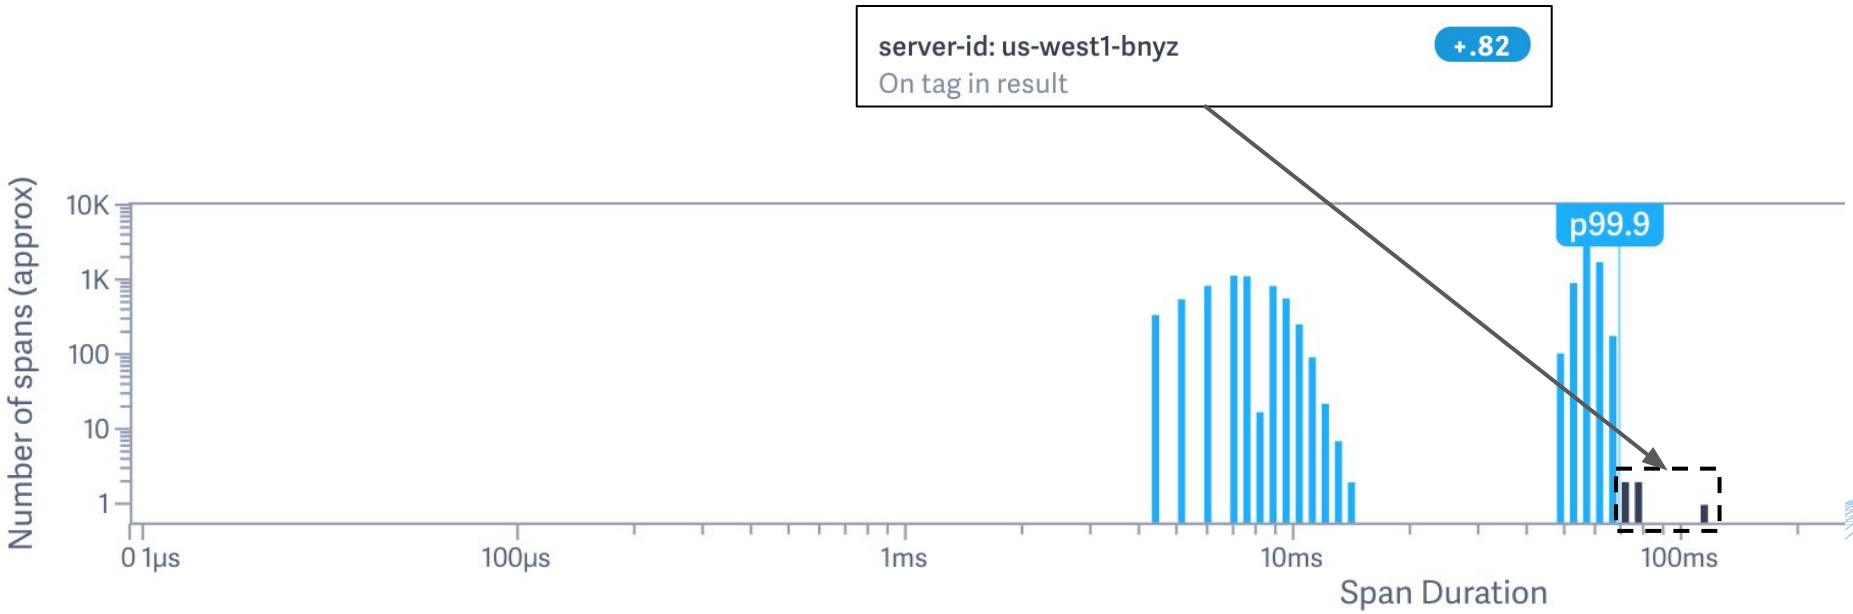

Critical path analysis

Critical path analysis

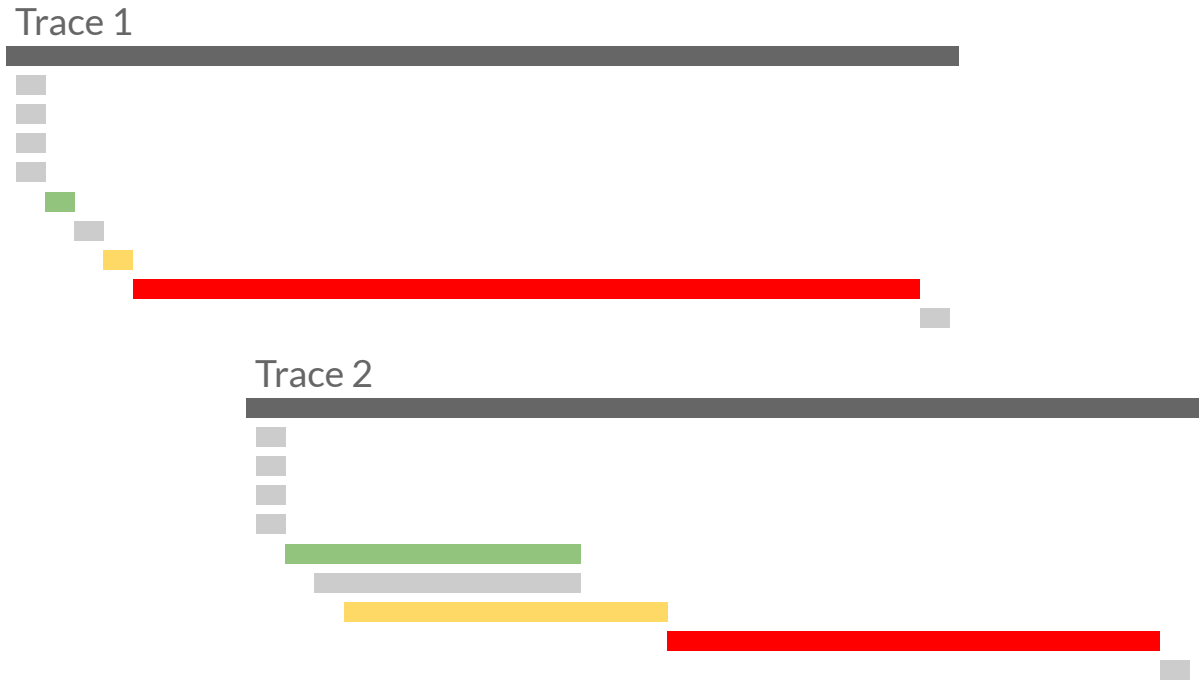

Critical path analysis

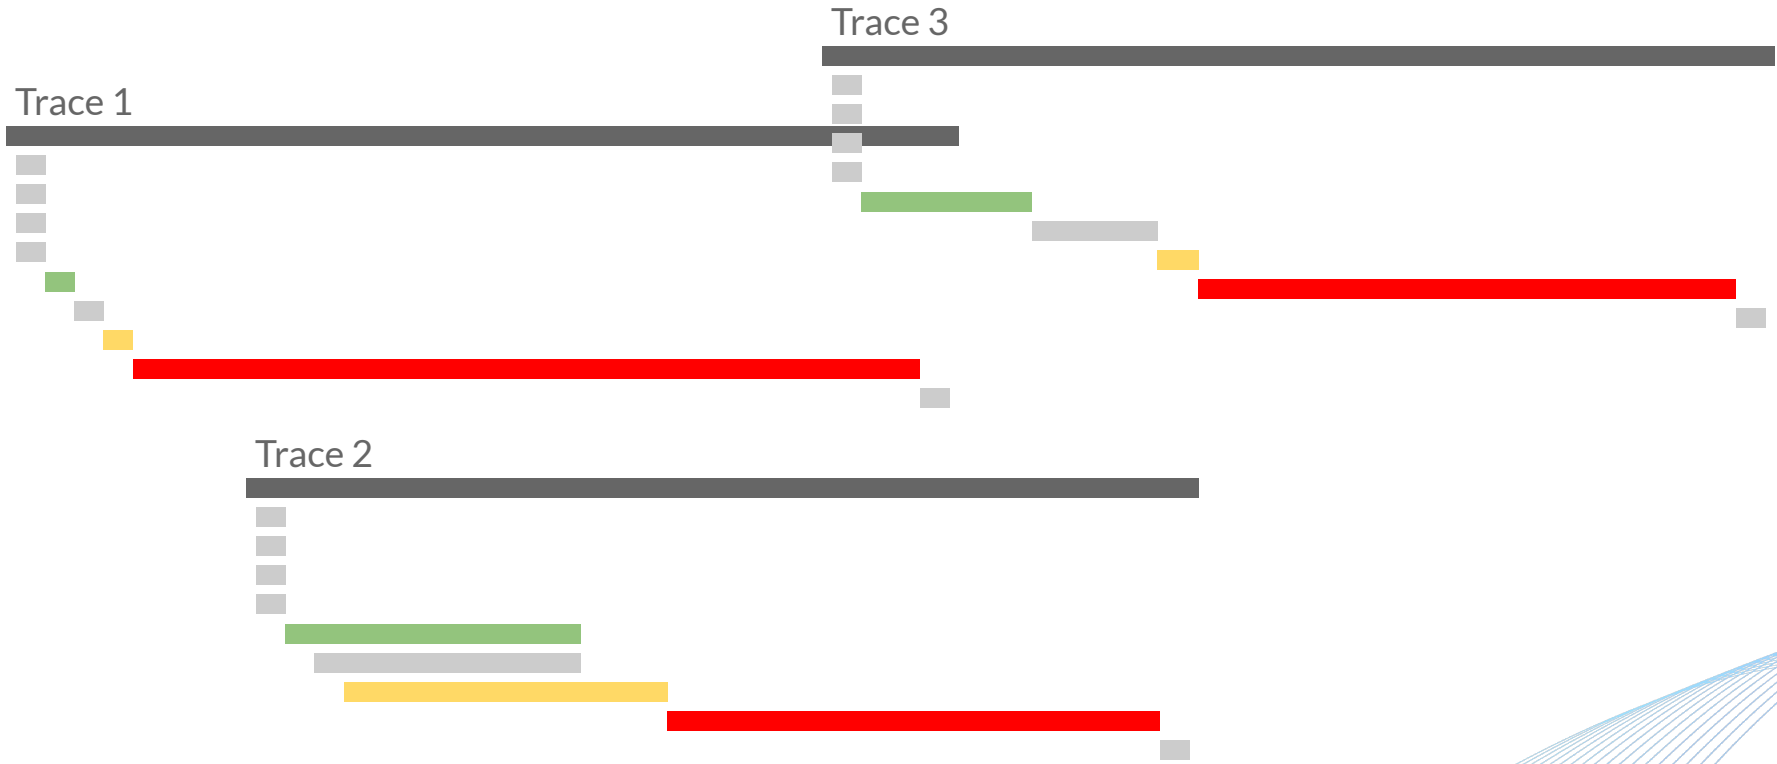

Critical path analysis => resource contention

Explain your critical path

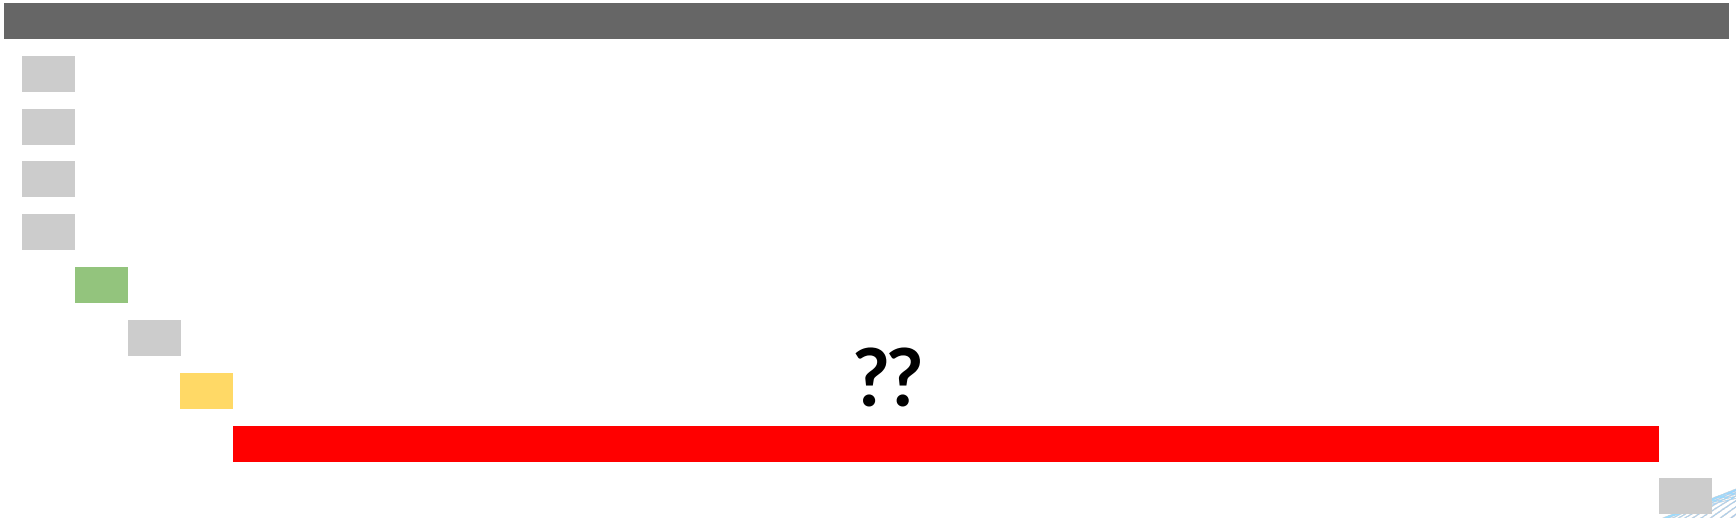

Resource contention

- The critical path is the “where” but rarely the “why”
- “Why”: nearly all latency issues are due to contention

Example: mutex contention

Explain your critical path

Resource contention analysis

Traffic jam, you want to know what's holding up all that traffic!!

Resource contention analysis

Thankfully there is a solution ...

Resource contention

Thankfully there is a solution ...

...

...

Use aggregate analysis, of course!

Example: mutex contention

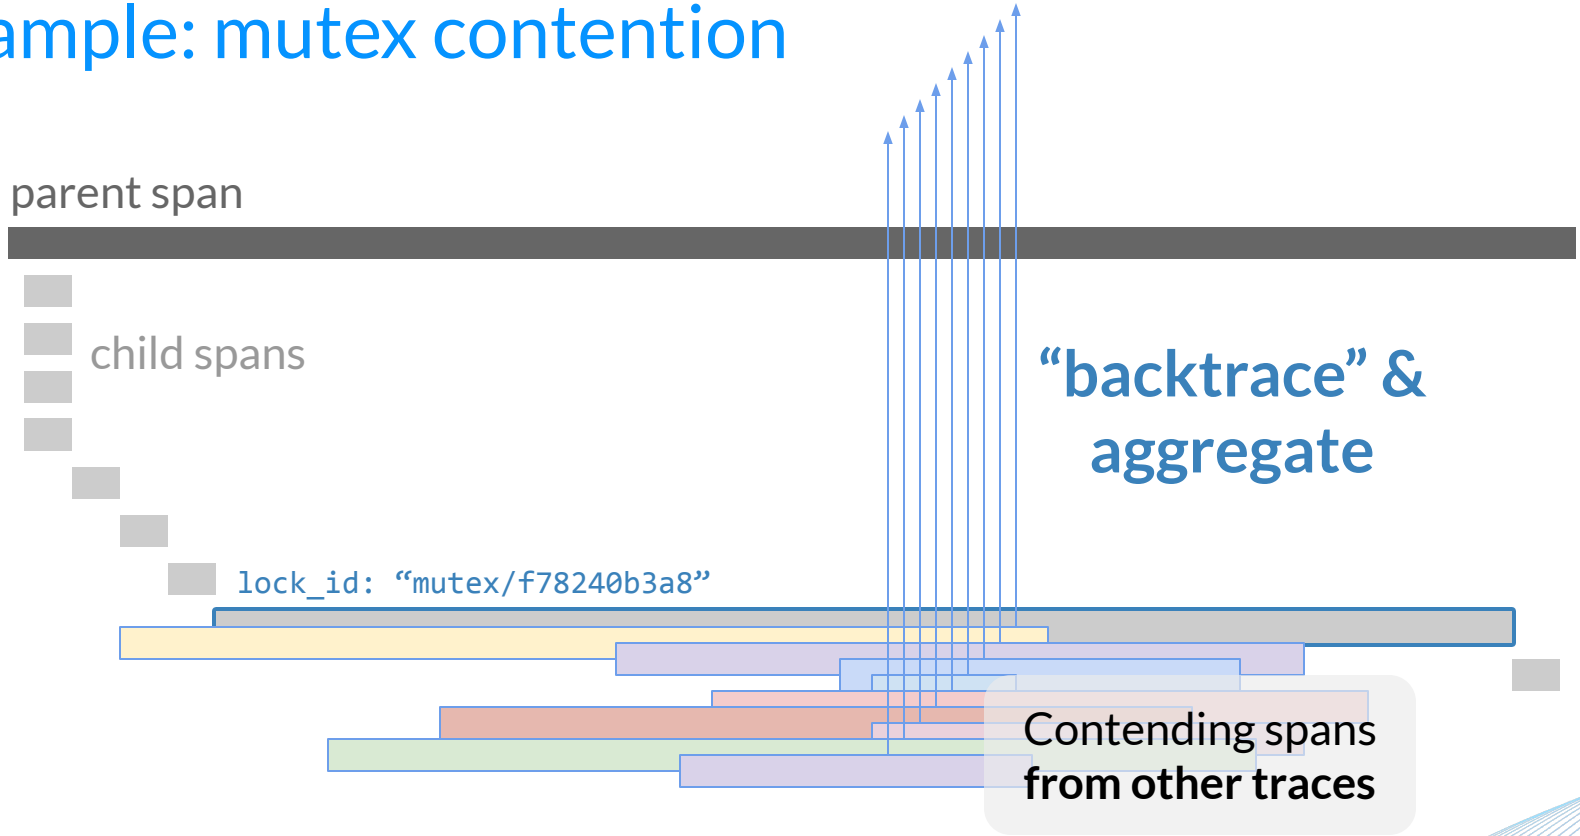

Example: mutex contention - terminology

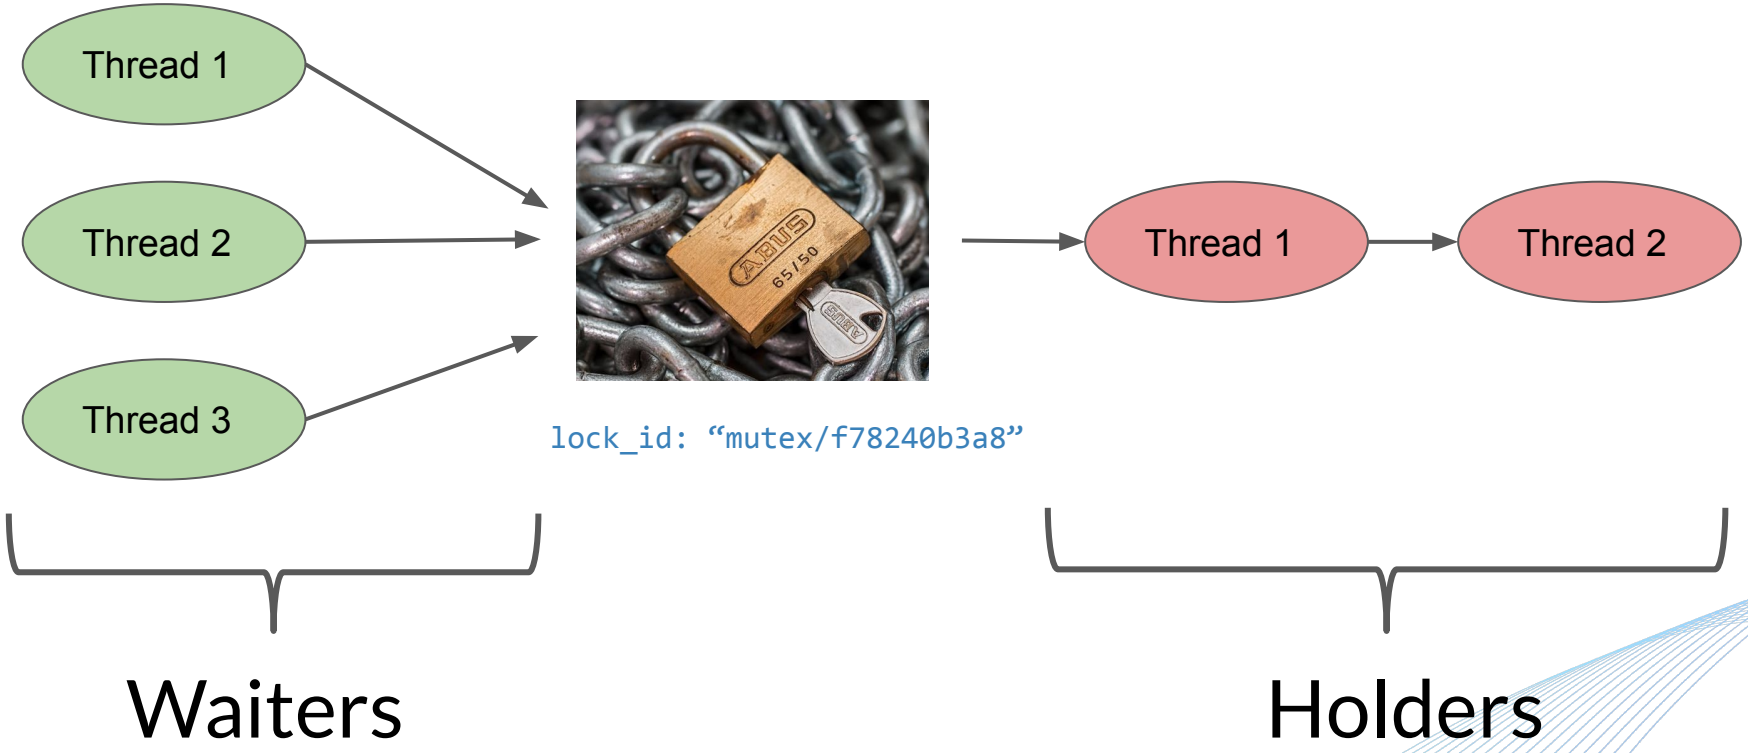

Demo time

🍩aaS Architecture: Move Fast and Bake Things

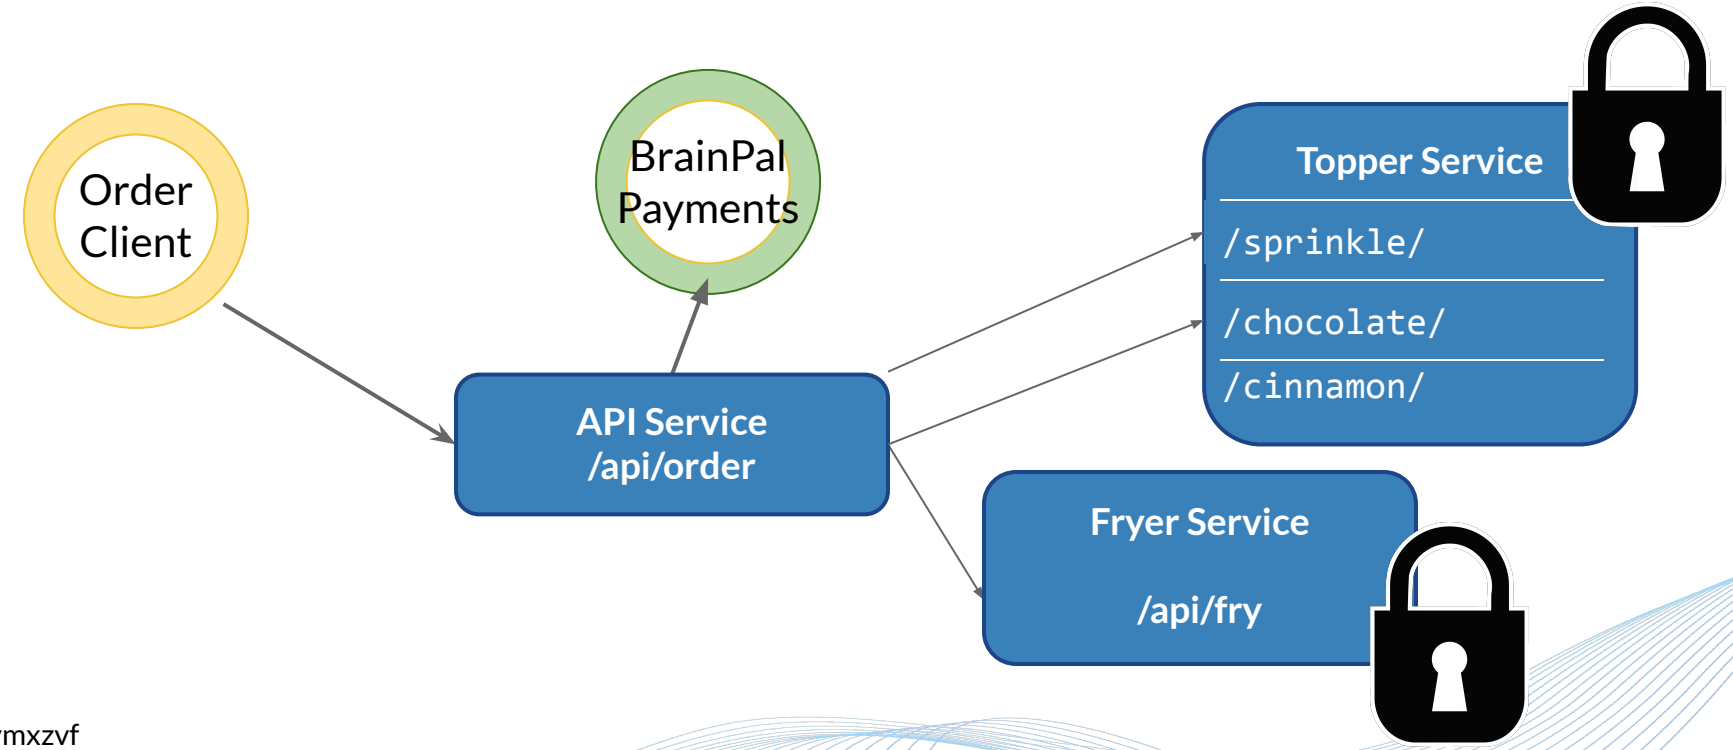

Demo review

Cinnamon (largest order count) spends a lot of time waiting on chocolate orders

Consider splitting out the lock

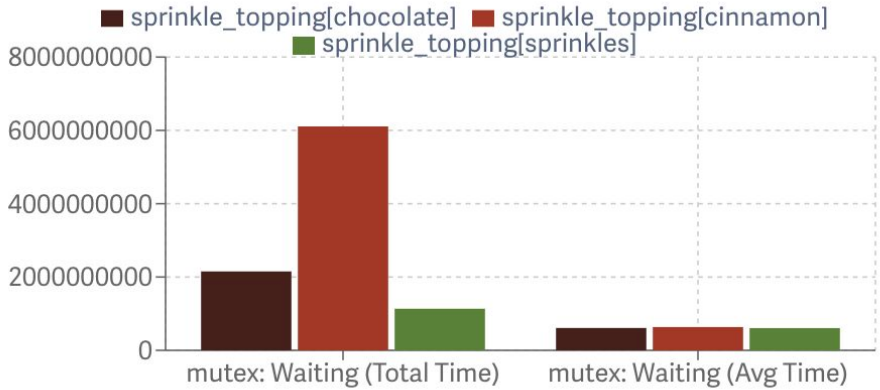

Waiters

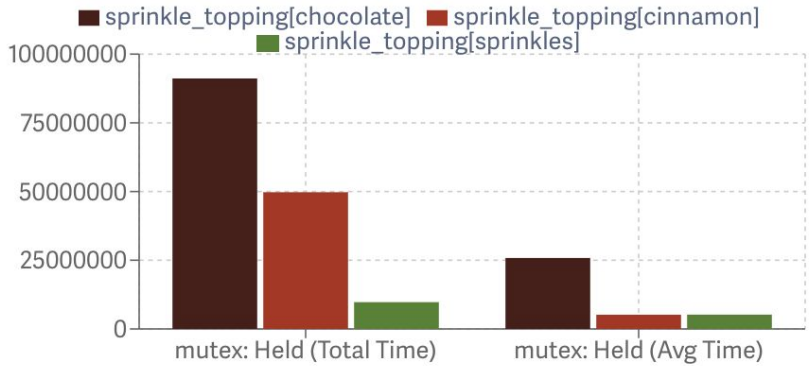

Holders

What was this talk about?

Trace Aggregates

- Aggregates => Insights
- You need lots and lots and lots of traces, without being locked in

- Enables robust, portable telemetry to be a built-in feature of cloud-native software
- It is the next major version of both OpenTracing and OpenCensus => backwards compatible with both

Thank you!

Questions?

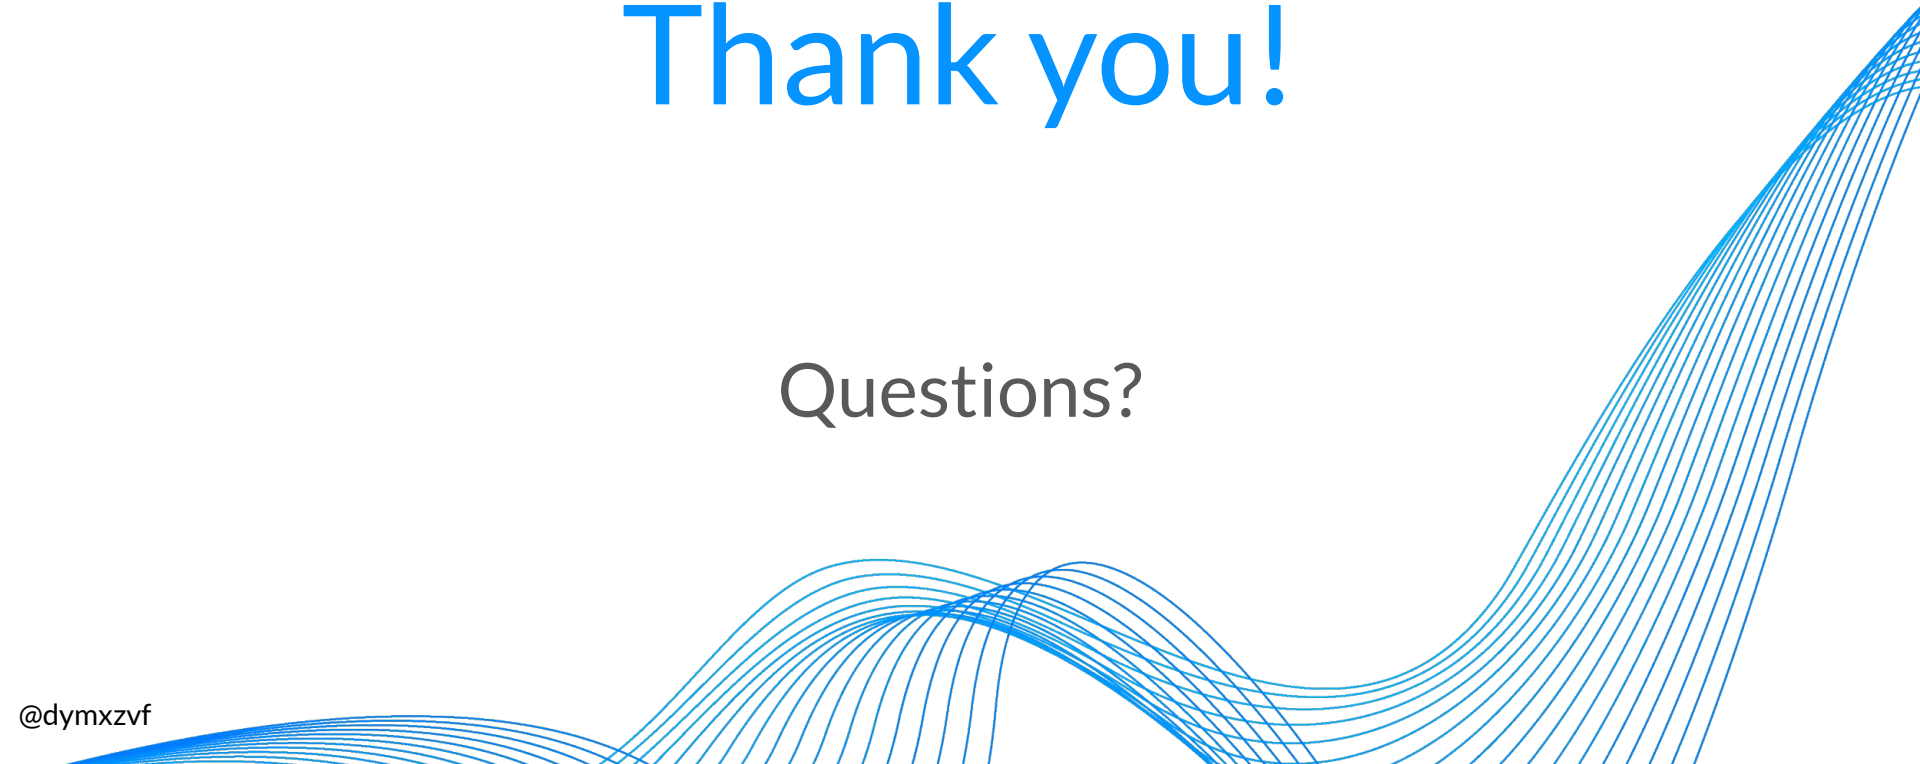

@dymxzvf