Traces Are the Fuel, Not the Car

Making Distributed Tracing Valuable

March 4, 2019

Ben Sigelman, CEO and Co-founder, LightStep

LIGHTSTEP

Part I

Observability Dogma: A Critique

The Conventional Wisdom

Observing microservices is hard

Google and Facebook solved this (right???)

They used **Metrics, Logging, and Distributed Tracing**…

So we should, too.

The Three Pillars of Observability

Fatal Flaws

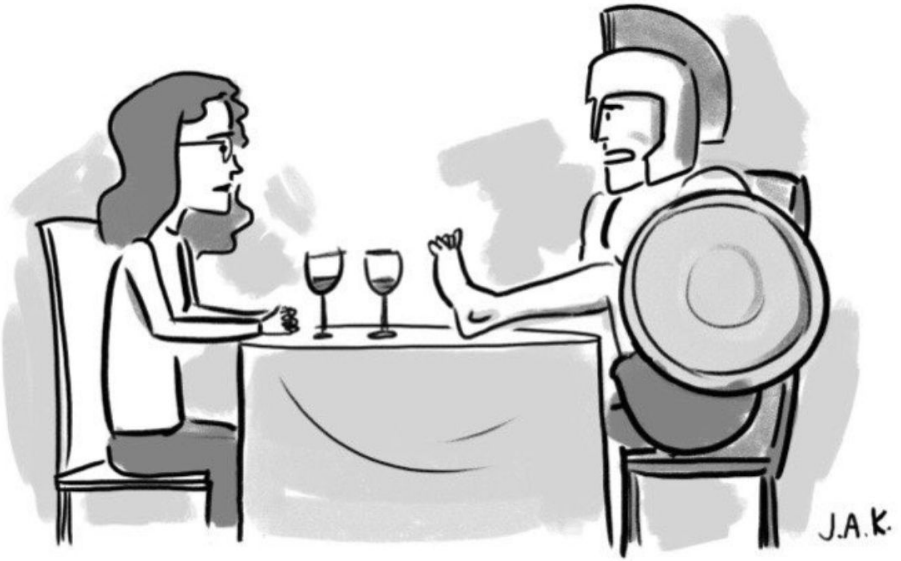

So Many Flaws, So Little Time…

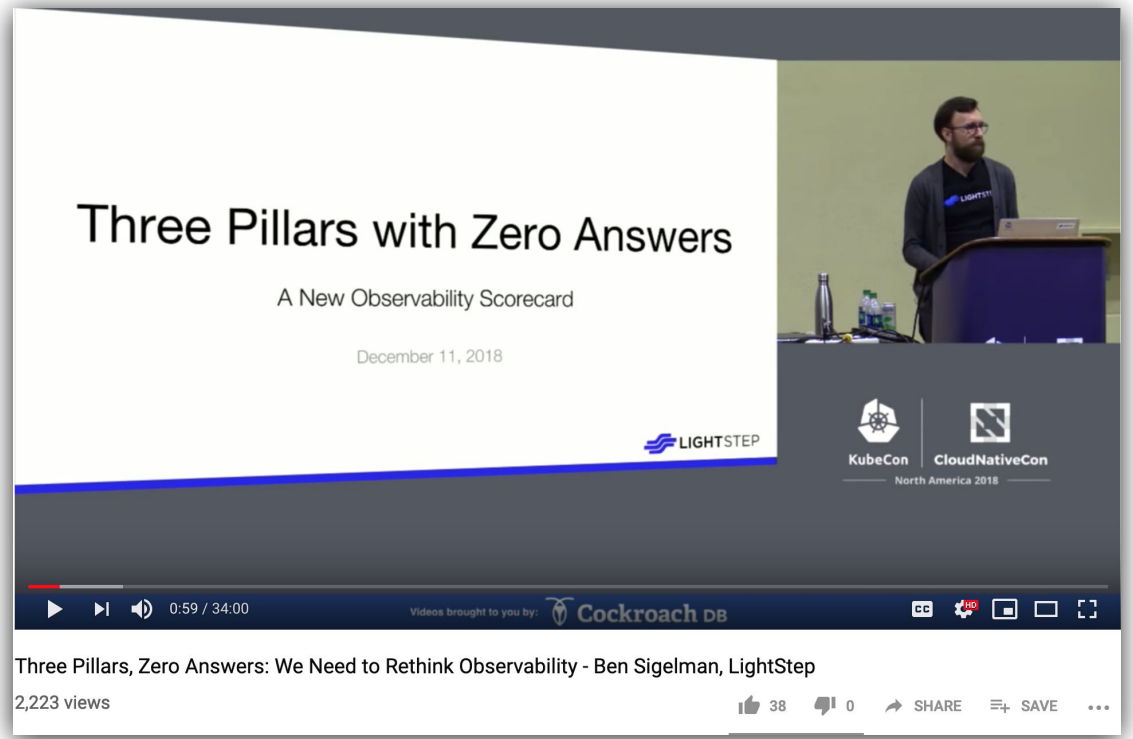

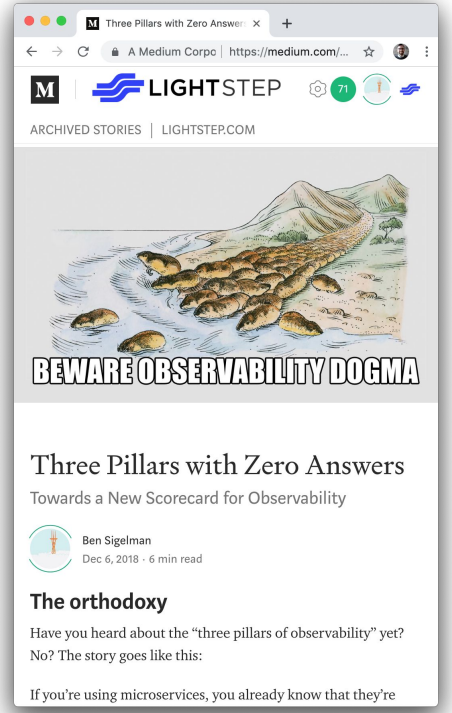

Fatal Flaws: “TL;DR” edition

	Logs	Metrics	Dist. Traces
TCO scales gracefully	—	✓	✓
Accounts for all data (i.e., unsampled)	✓	✓	—
Immune to cardinality	✓	—	✓

A fun game!

Design your own (positive-ROI) observability system:

- ☐ High-throughput
- ☐ High-cardinality
- ☐ Unsampled
- ☐ Lengthy retention window

Choose three.

Metrics, Logs, and Traces are **Just Data**,

… not a feature or use case.

The Three Pillars of Observability

The Three ~~Pillars~~ **Pipes** of Observability

Part II

Service-Centric Observability

A microservices architecture

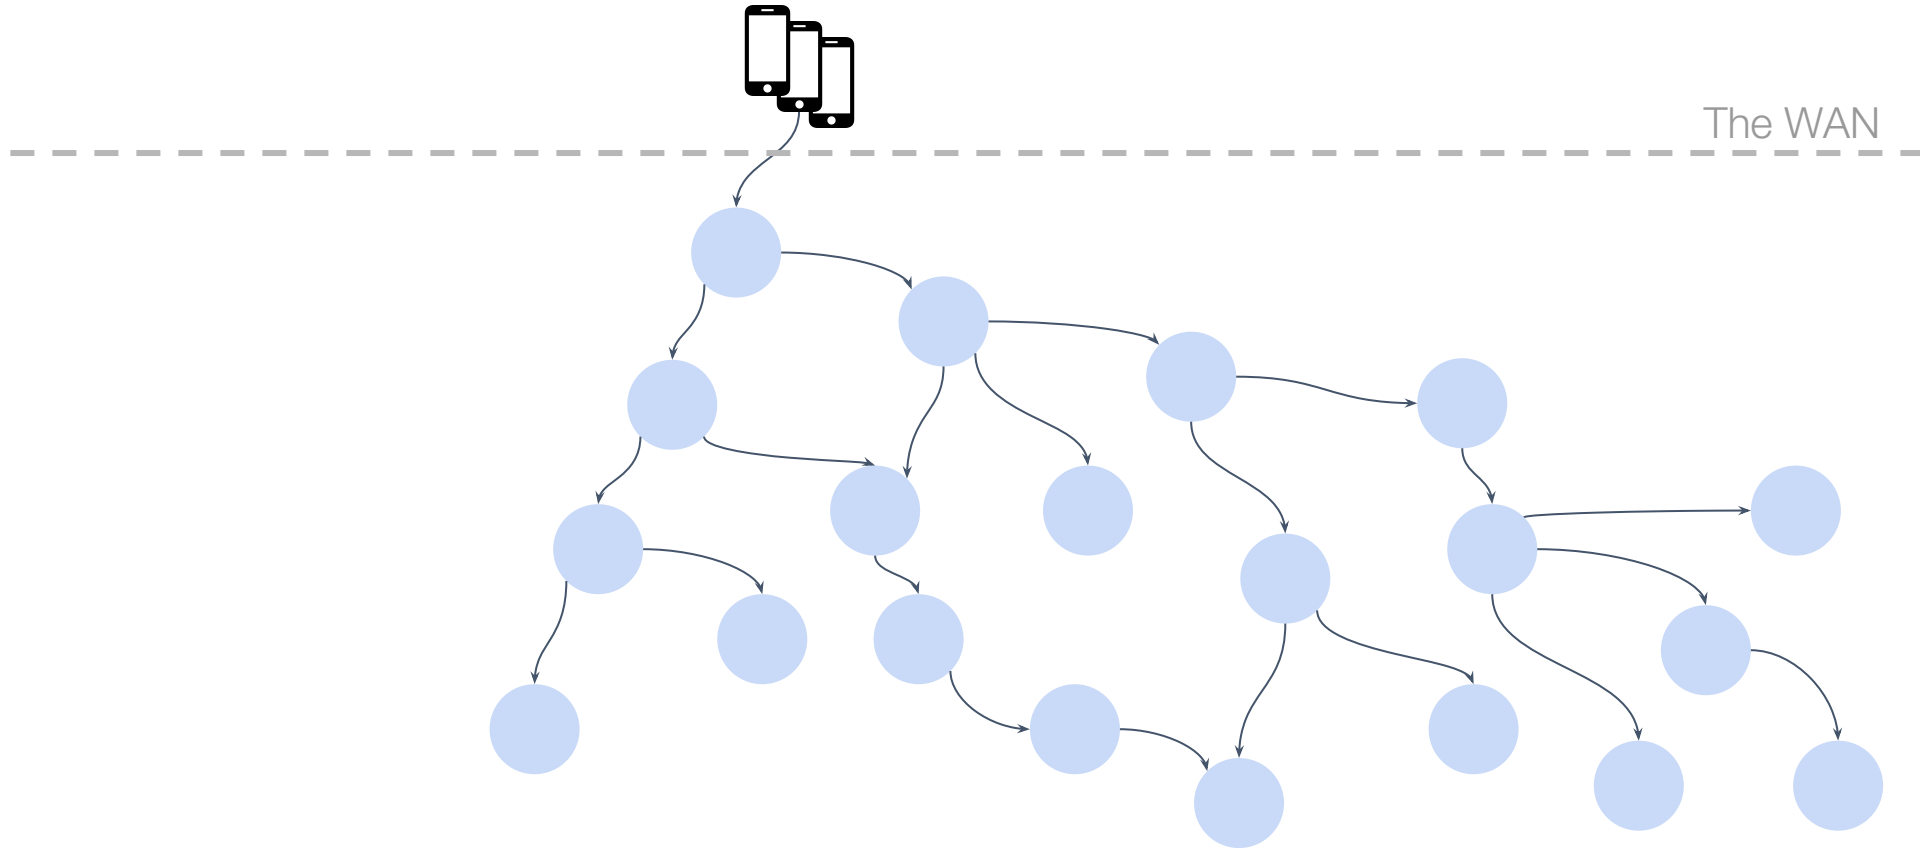

A microservices architecture

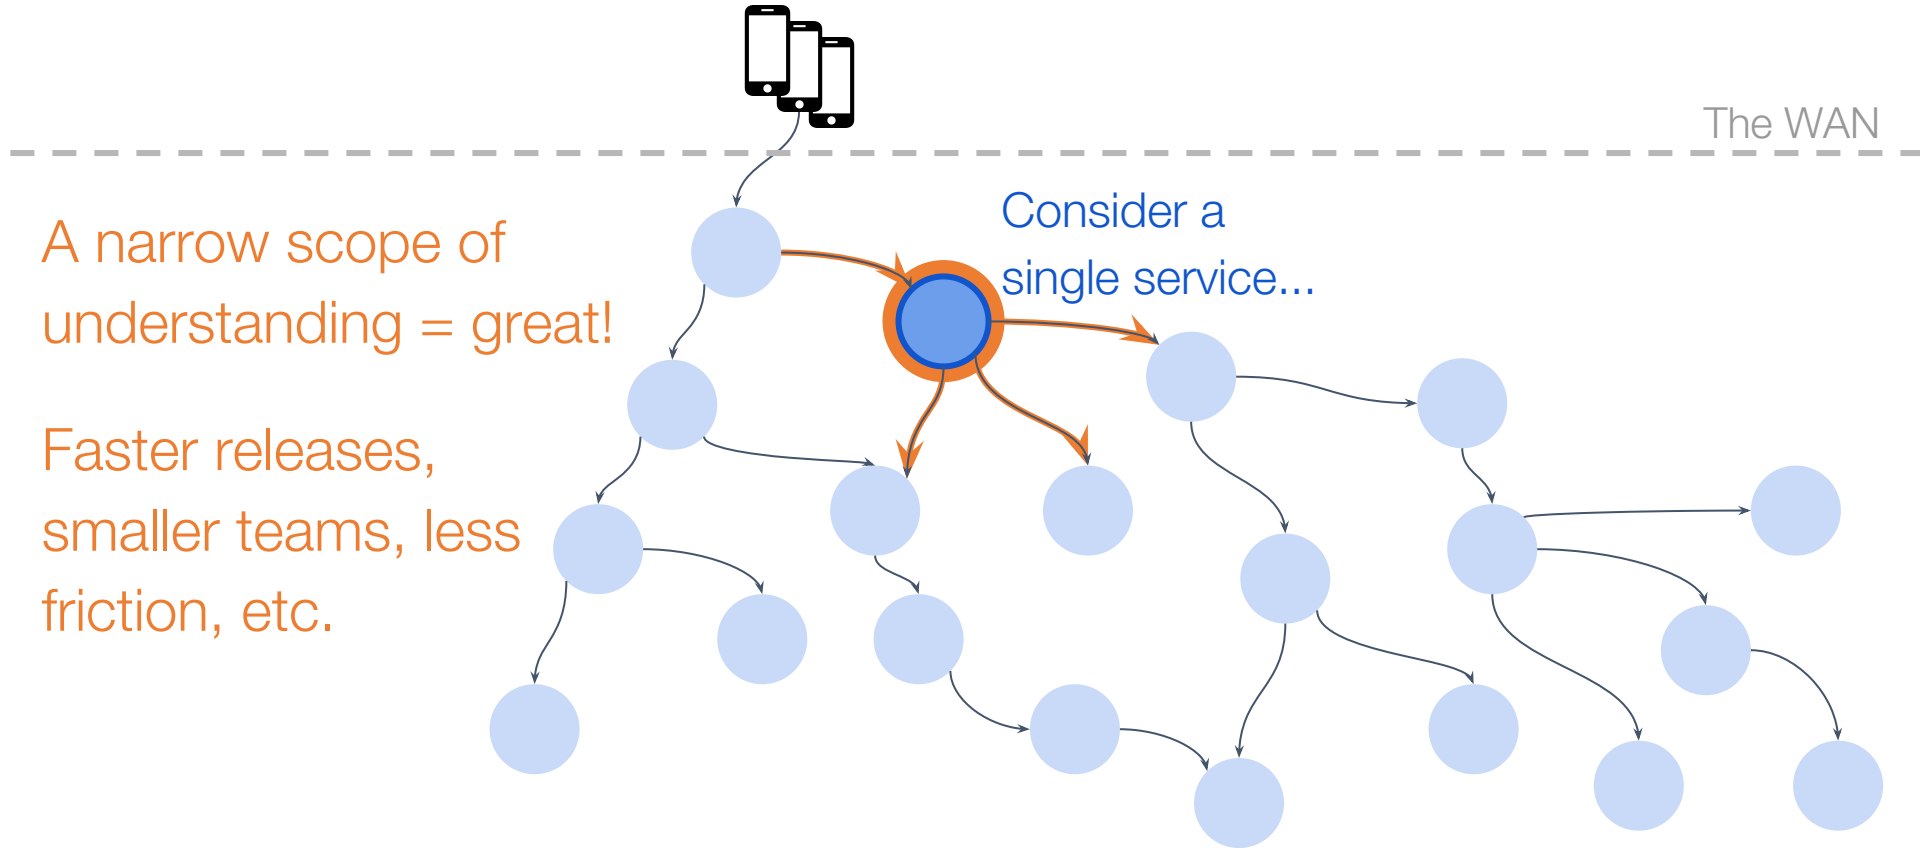

A narrow scope of understanding = great!

Faster releases, smaller teams, less friction, etc.

A microservices architecture (with a slowdown)

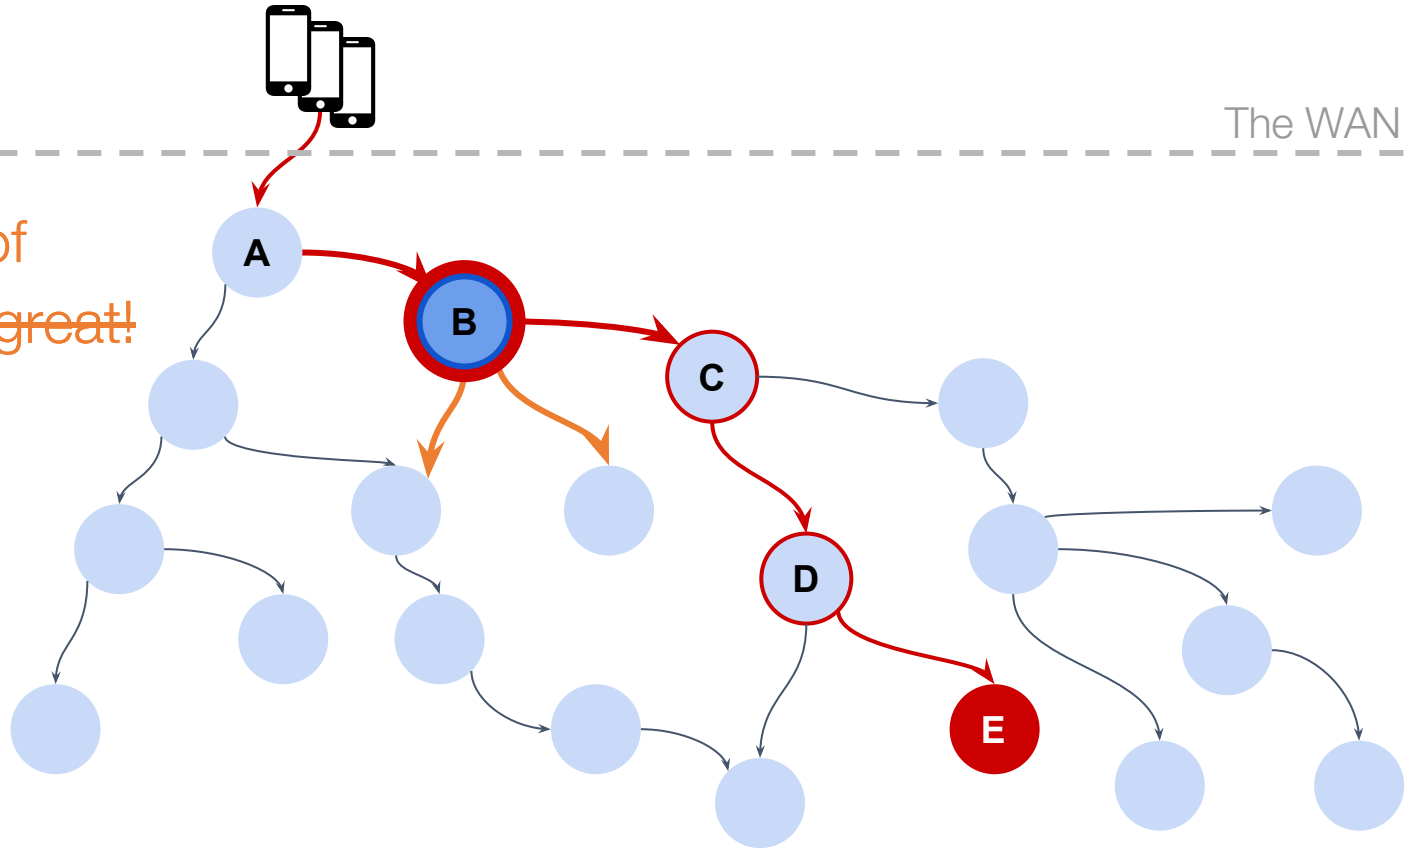

A narrow scope of understanding = ~~great!~~

… problematic.

Decoupling cuts both ways.

Hands-on with a single distributed trace

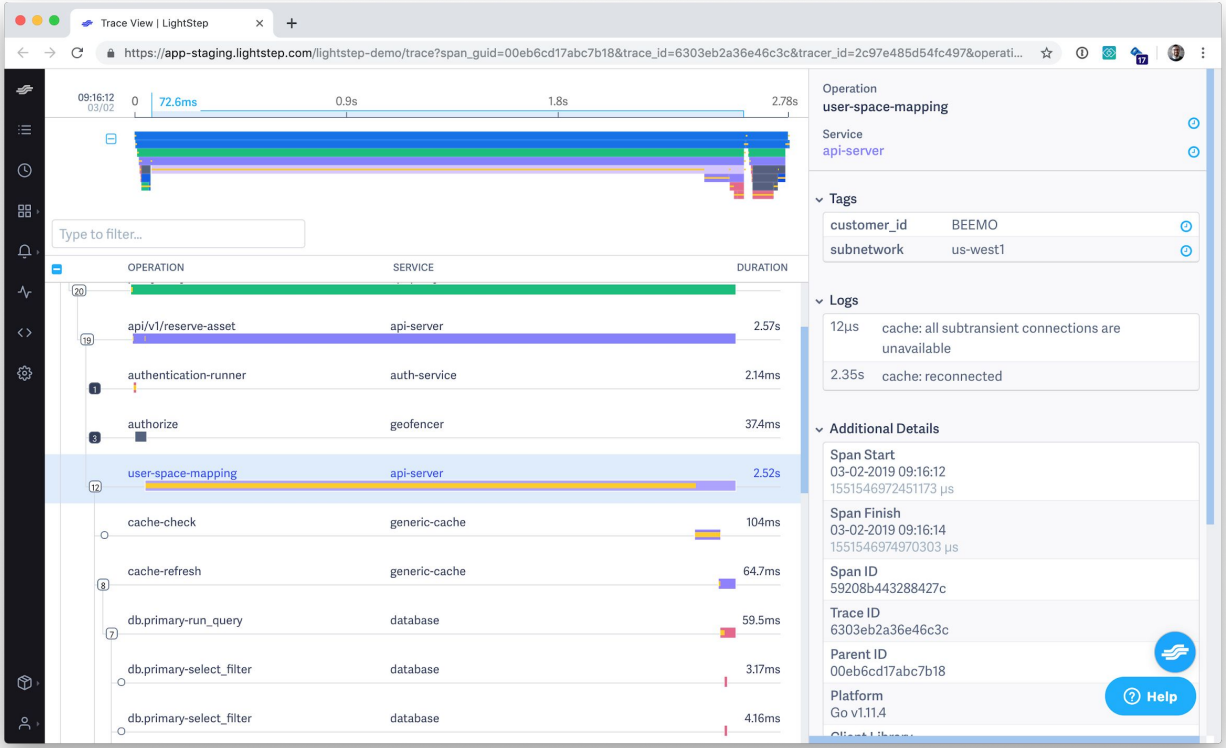

Distributed traces, in summary

- One distributed trace per transaction
- Crosses microservice boundaries
- They are **necessary** if we want to understand the relationships between distant actors in our architecture …

… and yet:

- too numerous to centralize in “standard” ways
- too data-dense for our brains to process without help

“Distributed Tracing” != “Distributed Traces”

Distributed *traces*: basically just `structs`

Distributed *tracing*: **the art and science of making distributed traces valuable**

So… how *do* we make distributed traces valuable?

Quick Vocab Refresher: **SLIs**

"SLI" = "Service Level Indicator"

TL;DR: An SLI is **an indicator of health** that a service's **consumers** would care about.

… *not* an indicator of its inner workings

Two Fundamental **Goals**

- Gradually **improving** an SLI
- Rapidly **restoring** an SLI

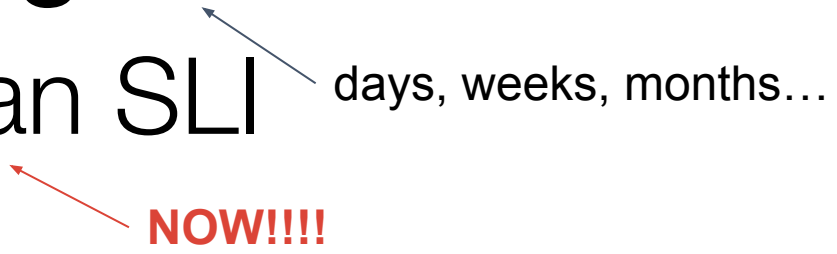

Reminder: “SLI” = “Service Level Indicator”

Two Fundamental **Activities**

1. **Detection:** measuring SLIs precisely
2. **Explaining variance:** recognizing and explaining variance, often iteratively

The Refinement Process

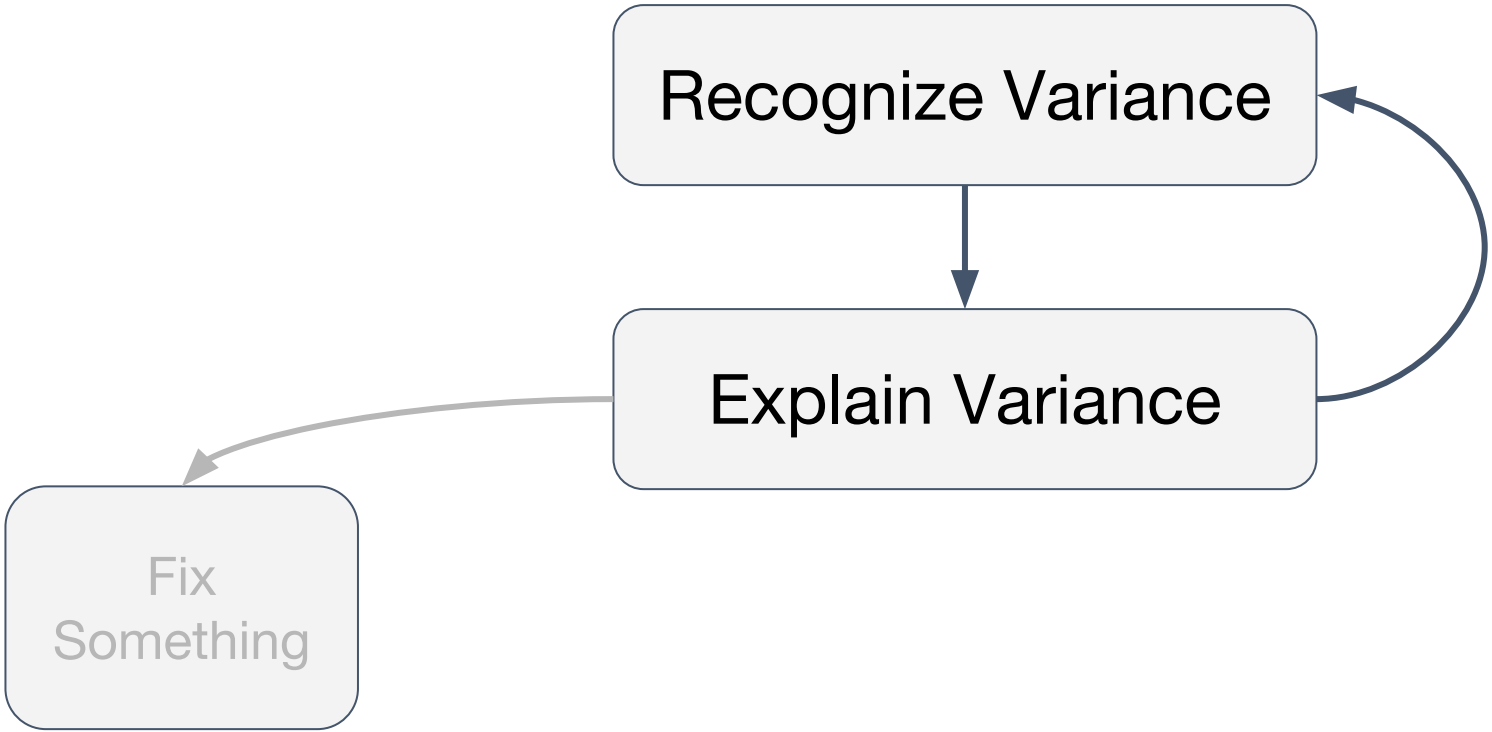

A Service-Centric Approach

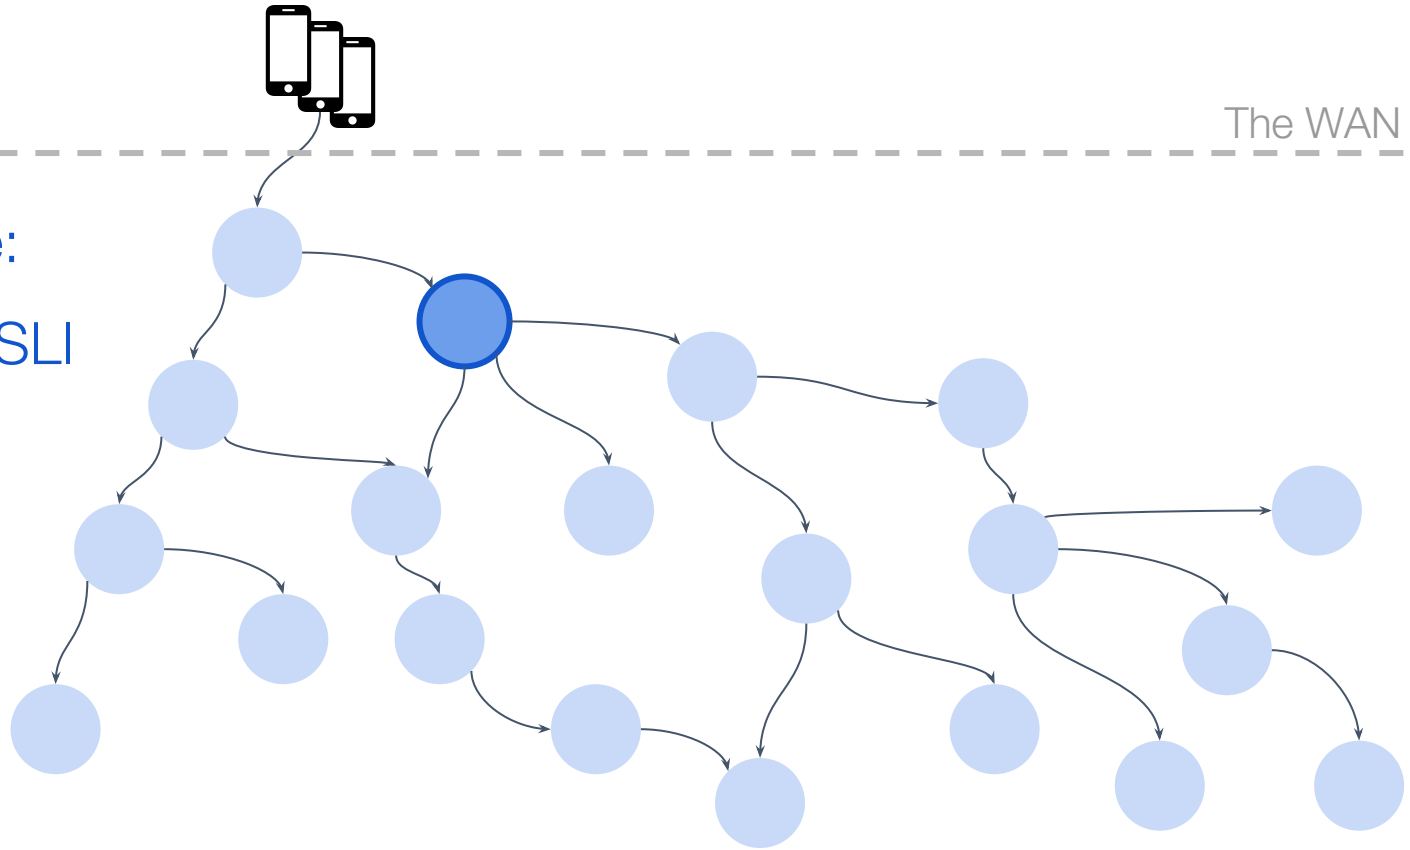

Given any service:

1. Start with an SLI
2. Find variance
3. Explain it

Part III

“Show & Tell”

A simple microservices architecture

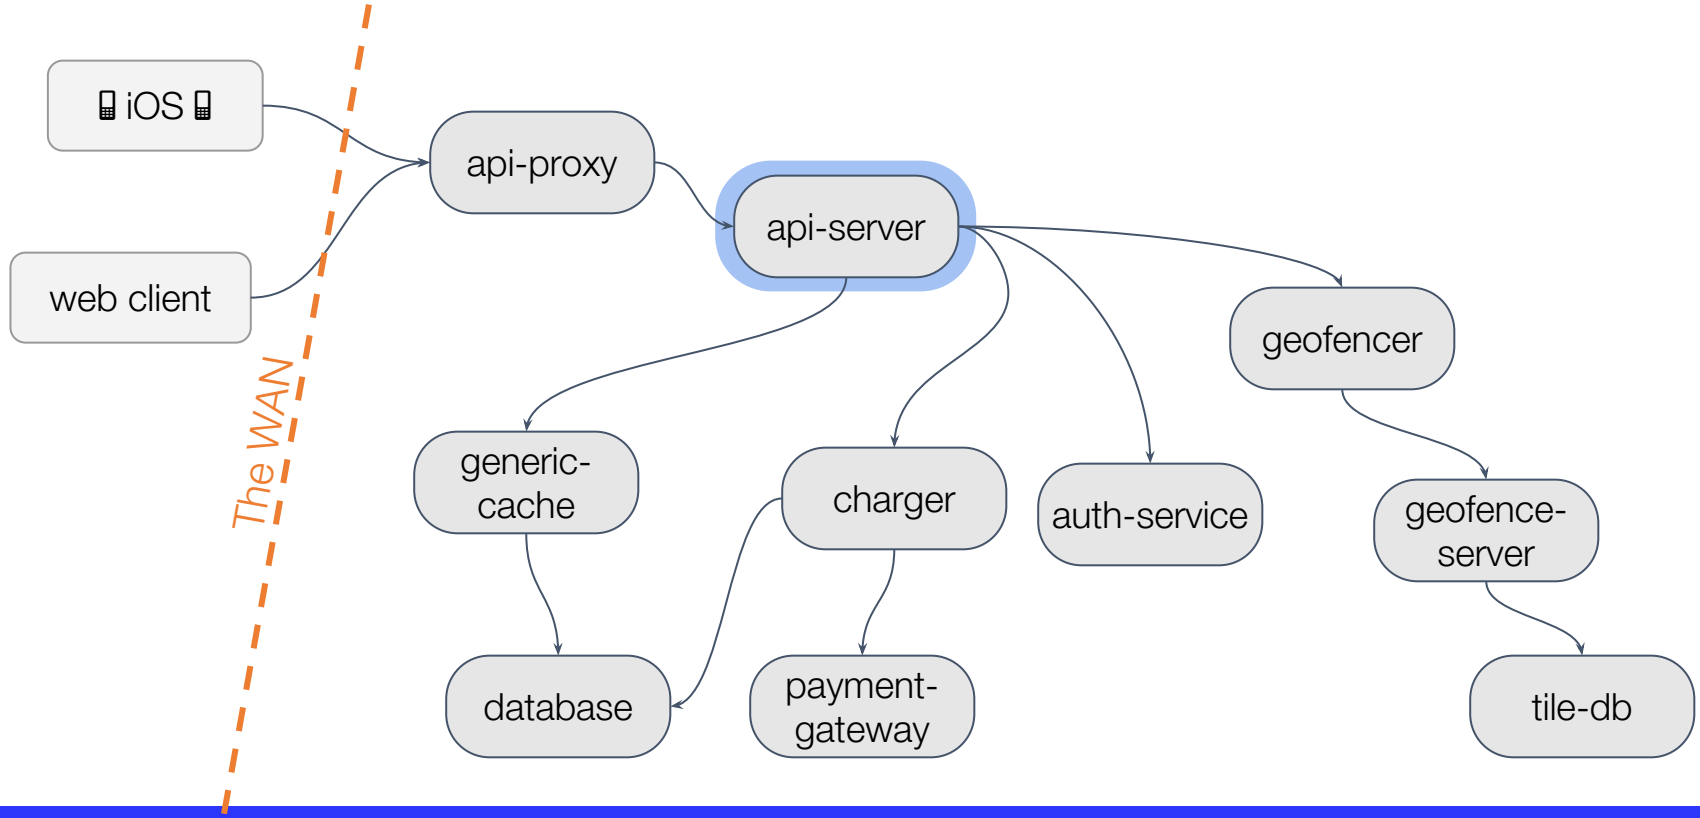

Recognizing Variance

1. Discovering SLIs (slide)
2. High-percentile latency measurement
3. “Performance is a shape” (and knowing what’s normal)
4. Examining individual traces

(link)

A blast from the past…

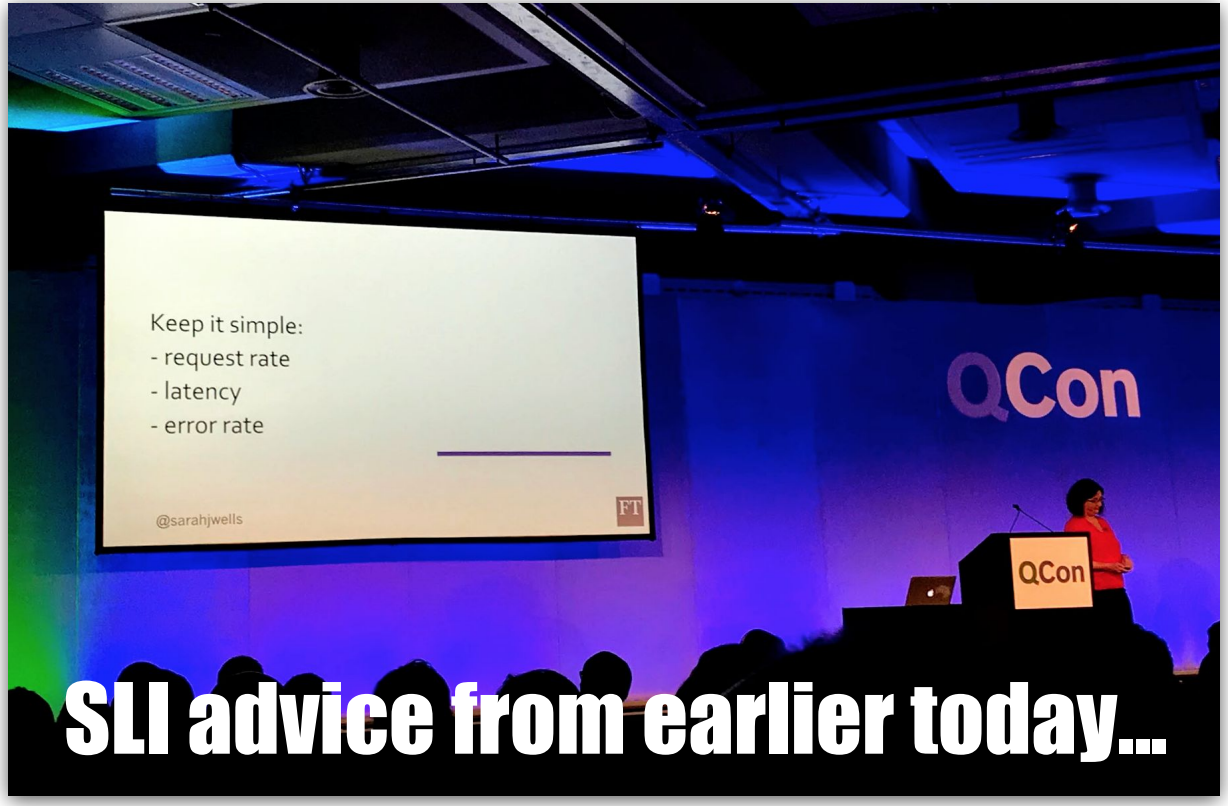

Service Diagrams

1. “Where’s Waldo” antipatterns (next slide)
2. Finding the common-case bottleneck
3. Finding the latency-outlier bottleneck

(link)

Service Diagrams and “Actionability”

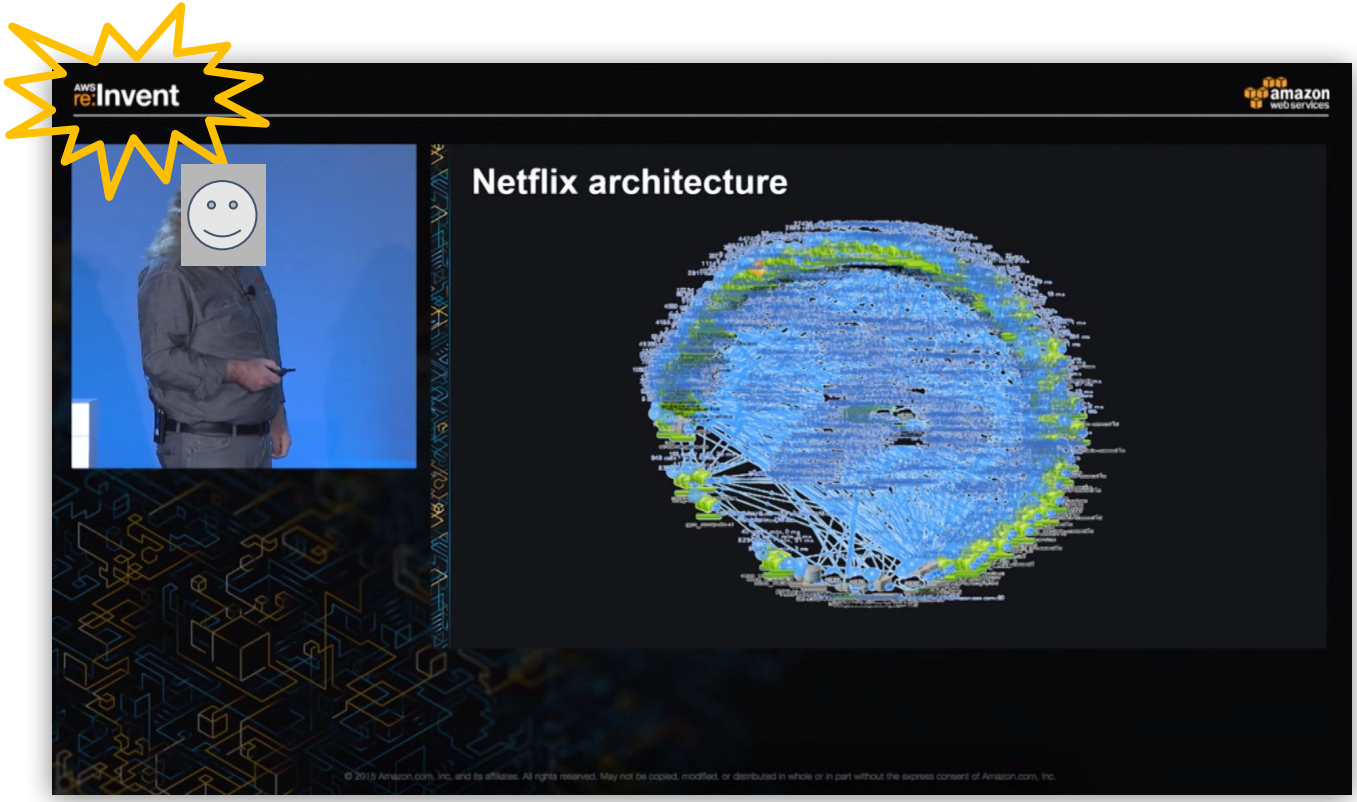

Explaining Variance With Many Dimensions

1. A “cardinality refresher” (next slide)
2. Exploring data with no cardinality limits
3. Explaining variance across the stack

(link)

A word nobody knew in 2015…

Dimensions (aka “tags”) can explain variance in timeseries data (aka “metrics”) …

… but **cardinality**

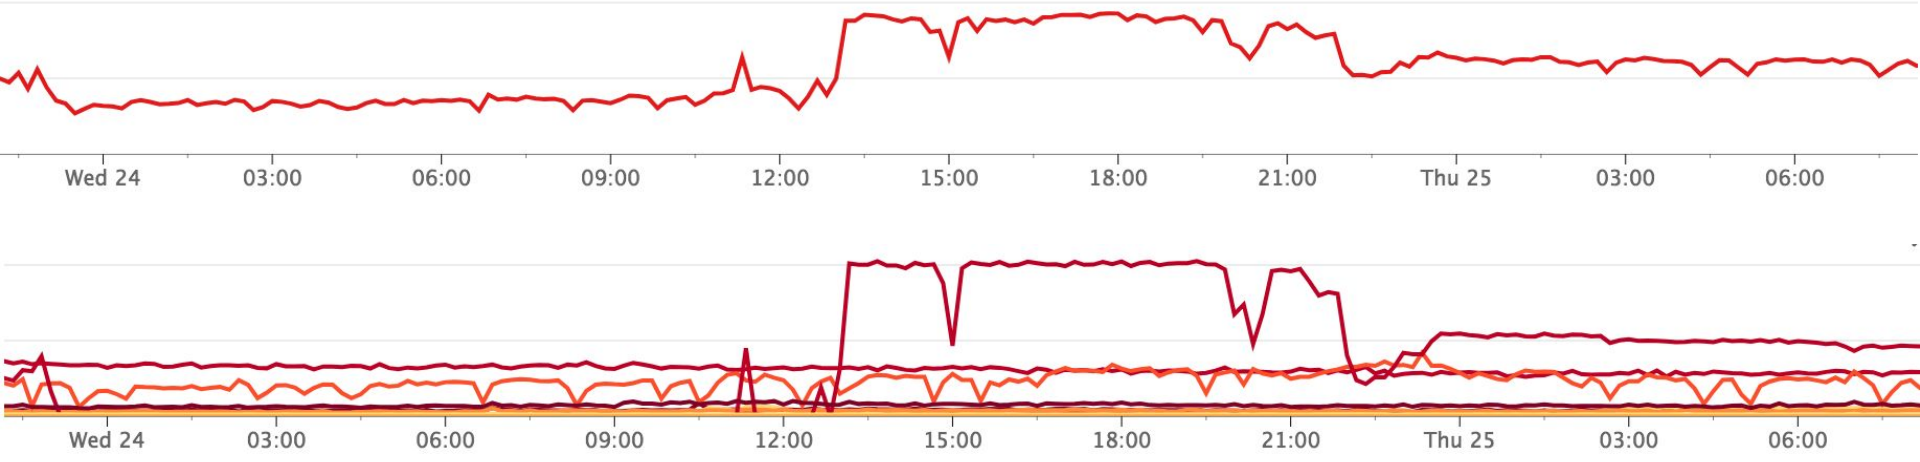

Wrapping up…

What we've learned

- Microservices helped us reduce human comms overhead
  - … and that created huge problems for observability
- Distributed traces are *necessary but not sufficient*
- Distributed tracing is much more than distributed traces
- A service-centric approach with a modern, sophisticated distributed tracing system can do amazing things

Thank you!

Ben Sigelman, Co-founder and CEO
twitter: @el_bhs
email: bhs@lightstep.com

PS: LightStep announced something cool today! Stop by Booth #3 to learn more.

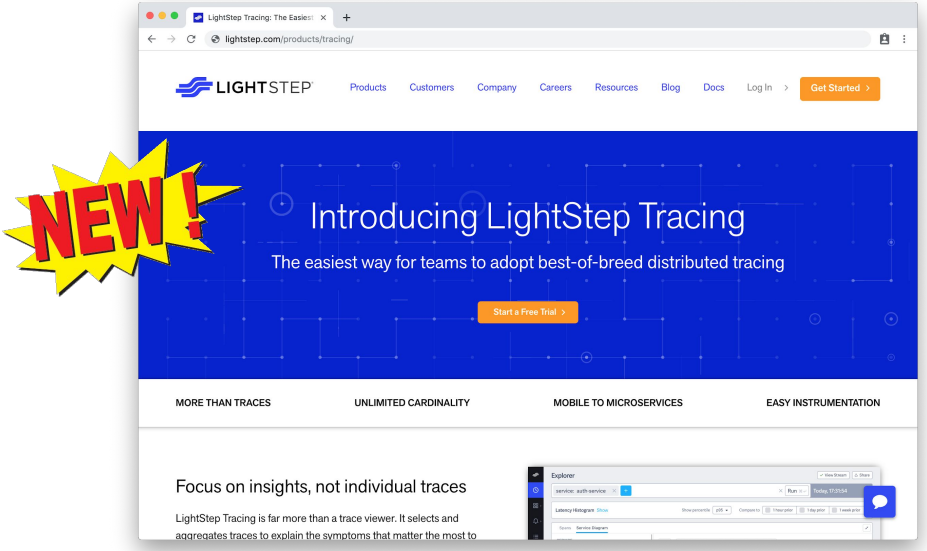

I am friendly and would love to chat… please say hello, I don't make it to Europe often!

Extra slides