Grafana Labs

The Role of Observability in Cloud Native Environments

Juraci Paixão Kröhling
Software engineer
@jpkrohling

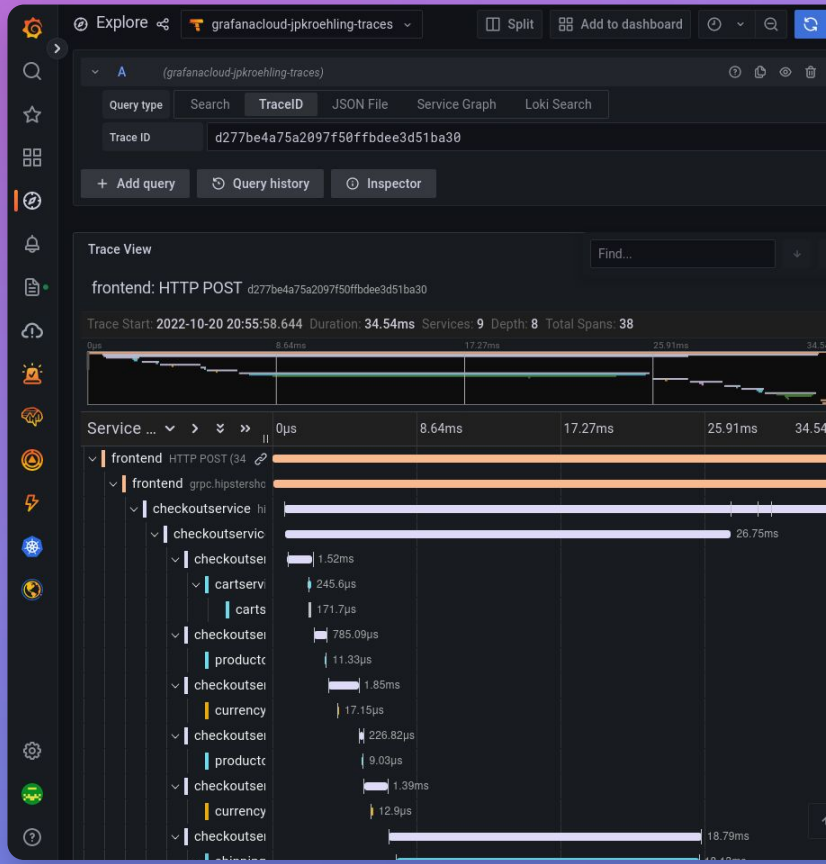

Speaker

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 (tail-sampling, load balancer, ...)
- Creator of ocb and OTel Operator
- Jaeger emeritus maintainer
- OpenTracing emeritus maintainer

Speaker

Juraci Paixão Kröhling
Software Engineer

Agenda

- Agile movement
- Microservices
- DevOps
- Reliability Engineering
- Telemetry vs. Monitoring vs. Observability
- Telemetry signals

Agile movement

Agile movement

Iterative development

Cross-functional teams

Customer feedback, deliver value

Embracing change over planning

Deliverables are functional

Microservices

Microservices

Decentralized & Independent Services

Loose Coupling & High Cohesion

Scalability & Flexibility

Faster Deployment Cycles

Resilience & Fault Isolation

DevOps Culture

DevOps Culture

Agile development

Continuous integration and delivery

Fast feedback loops

Change, evolution is constant

Reliability

Reliability Engineering

Reliability engineering

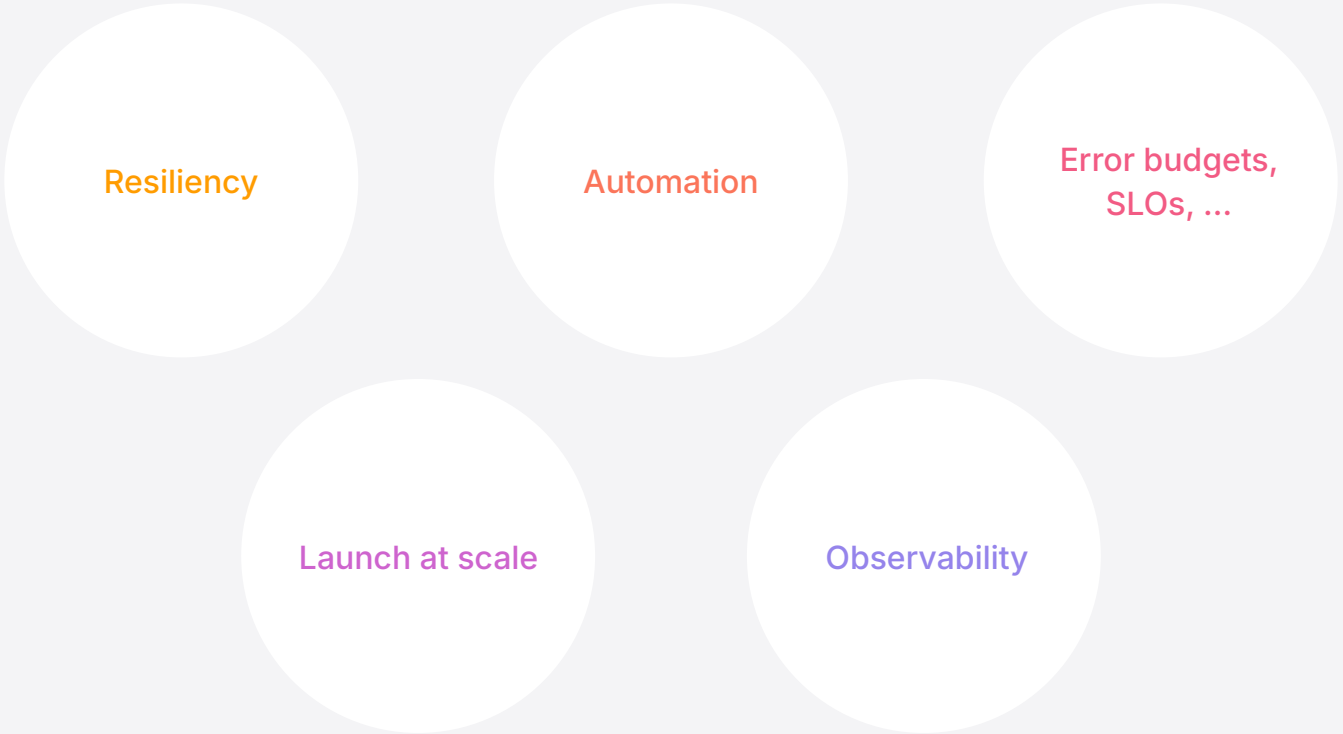

Observability


O11y

Observability is the ability to find answers to questions we don’t have yet.

Monitoring

Monitoring is the practice of obtaining quick answers to frequent questions.

Telemetry

Data that is generated by our application, informing its state while it's running.

Telemetry data types

Logs

Metrics

Traces

Profiles

Events

Cloud native

Cloud native

Containers & Orchestration

Dynamic & Ephemeral

Horizontally scalable

API-Driven

Resilience & Self-Healing

Observability for Cloud Native

Observability for Cloud Native

- What happened, and where?
  - Which pod? Which cluster? Which region? Is it a widespread event?
- Should I proceed with the rollout?
  - Provide business metrics to determine whether the individual service is healthy
- Who am I talking to?
  - Understand the call graph for specific requests
- What’s the baseline? What does “good” look like?
  - Document what’s the expected behavior

Am I able to answer questions about my services when I'm paged at 2am?

Grafana Labs

Obrigado!