



**CLOUD NATIVE**  
**COMPUTING FOUNDATION**

# Jaeger

## SIG Project Update

Pavol Loffay (Red Hat), Yuri Shkuro (Uber)

CloudNativeCon NA, Austin, Dec-7-2017

# Agenda

- Project introduction
- Demo
- Roadmap
- Discussion

# Jaeger, a Distributed Tracing System



<https://jaegertracing.io>

# Jaeger - /'yāgər/, *noun*: hunter

- Inspired by Google's Dapper and OpenZipkin
- Started at Uber in August 2015
- Open sourced in April 2017
- Official CNCF project, Sep 2017
- Built-in OpenTracing support
- <https://jaegertracing.io>



**UBER**

# Technology Stack

- Go backend
- Pluggable storage
  - Cassandra, Elasticsearch, memory, ...
- React/Javascript frontend
- OpenTracing Instrumentation libraries



Go



python™



# Community

- 10 full time engineers at Uber and Red Hat
- 30+ contributors on GitHub
- Already used by many organizations
  - including Symantec, Red Hat, Base CRM, Massachusetts Open Cloud, Nets, FarmersEdge, GrafanaLabs, Northwestern Mutual, Zenly





Let's look at some traces

demo time: <http://bit.do/jaeger-hotrod>



# Distributed Tracing Systems

distributed  
transaction  
monitoring

performance  
and latency  
optimization

root cause  
analysis

service  
dependency  
analysis

distributed context propagation

# Announcing Release 1.0

- UI performance and usability improvements to view large traces
- Storage backends: Cassandra and Elasticsearch
- Spark job for building service dependencies diagram
- Client libraries: Go, Java, Python, Node.js, C++ (new)
- Integration with other CNCF projects
  - Templates for deploying Jaeger on Kubernetes
  - All Jaeger components expose Prometheus metrics by default
  - Integration with Envoy/Istio
- Drop-in replacement for Zipkin backend



# Roadmap

Things we are working on



# Adaptive Sampling

- APIs have endpoints with different QPS
- Service owners do not know the full impact of sampling probability

Adaptive Sampling is per service + endpoint,  
decided by Jaeger backend based on traffic

# Data Pipeline

- Based on Kafka and Apache Flink
- Support aggregations and data mining
- Examples:
  - Pairwise dependencies diagram
  - Path-based dependencies diagram
  - Latency histograms





Q & A

Open Discussion

# Jaeger at CloudNativeCon

- SIG Jaeger Deep Dive Session
  - Thursday, December 7 • 2:00pm - 3:20pm
- Jaeger Project Salon
  - Friday, December 8 • 2:00pm - 3:20pm
- OpenTracing Salon
  - Thursday, December 7 • 3:50pm - 4:50pm

