



redhat®

JAEGER PROJECT DEEP DIVE

KubeCon + CloudNativeCon Europe 2018

2018-05-03

JURACI PAIXÃO KRÖHLING

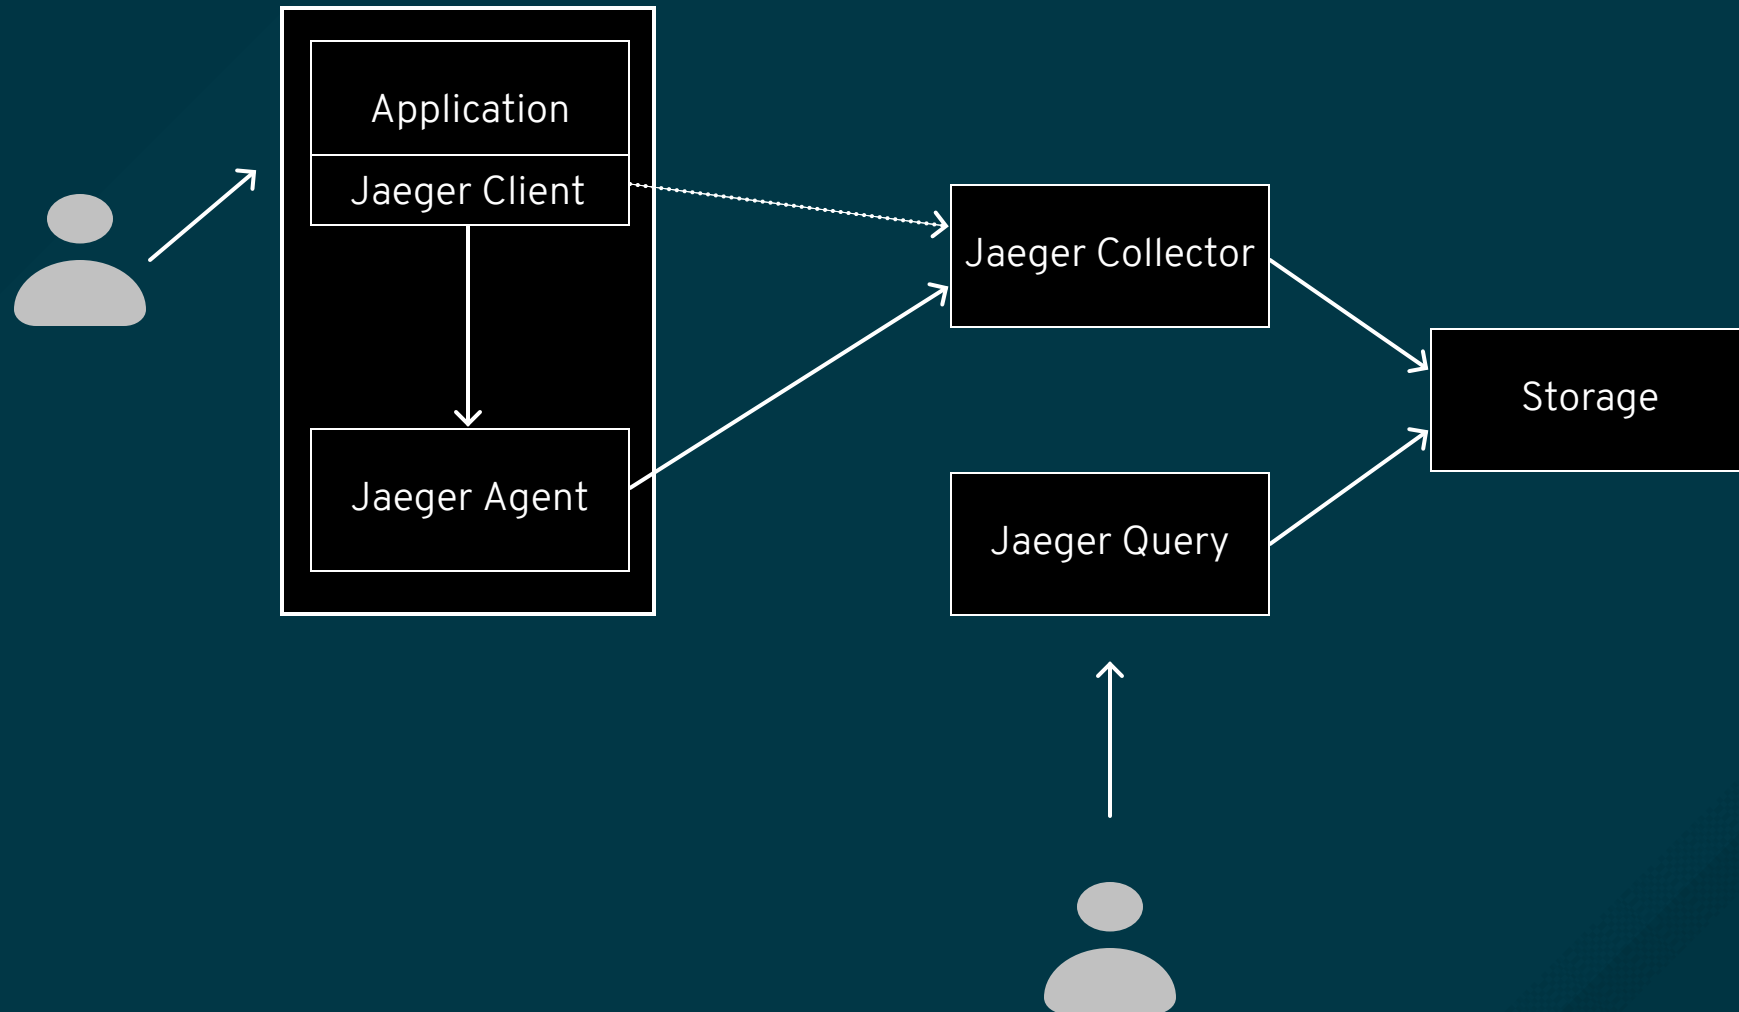
Software Engineer

jpkroehling on #kiali at irc.freenode.net

CONTEXT

- * MICROSERVICES
- * OBSERVABILITY PROBLEM
- * DISTRIBUTED TRACING
- * OPENTRACING
- * JAEGER ARCHITECTURE

JAEGER ARCHITECTURE



JAEGER CLIENT

- * BUILDS SPANS BASED ON INSTRUMENTATION
- * SENDS SPANS TO AGENT OR COLLECTOR
- * SUPPORTS CUSTOM "REPORTERS"
- * SUPPORTS CUSTOM "SENDERS"
- * EMITS INTERNAL METRICS

JAEGER AGENT

- * SITS CLOSE TO THE CLIENT
- * ACTS AS A "BUFFER"
- * RECEIVES SPANS VIA UDP
- * SENDS SPANS TO A COLLECTOR

JAEGER COLLECTOR

* WRITES SPAN TO STORAGE

JAEGER QUERY

- * READS SPAN FROM STORAGE

- * EMBEDS THE JAEGER UI

CONFIGURATION

* CONFIGURATION FILE

```
data:  
  span-storage-type: cassandra
```

* CLI OPTIONS

```
--span-storage.type=memory
```

* ENVIRONMENT VARIABLES

```
SPAN_STORAGE_TYPE=memory
```


PERFORMANCE TUNING

* QUEUE SIZES

`--collector.queue-size`

`--processor.<type>.server-queue-size`

* WORKERS

`--collector.num-workers`

`--processor.<type>.workers`

* SAMPLING STRATEGY

`--sampling.strategies-file`

MONITORING

* EXP VARS

```
--metrics-backend=expvar
```

* PROMETHEUS

```
--metrics-backend=prometheus (default)
```

DEPLOYMENT - SIMPLE

* ALL-IN-ONE (AKA STANDALONE)

DEPLOYMENT - PRODUCTION

- * EACH COMPONENT ON ITS OWN CONTAINER/POD
- * OPENSIFT AND KUBERNETES TEMPLATES AVAILABLE
- * COMMUNITY-CONTRIBUTED HELM CHART TOO!

DEPLOYMENT - AGENT

- * BARE METAL - HOST-LEVEL PROCESS

- * KUBERNETES - DAEMONSET

- * KUBERNETES - SIDECAR

DEPLOYMENT - MULTI TENANCY

* WHAT COULD BE MULTI TENANCY IN OUR CONTEXT?

- ** Multi tenant spans, with tenant being part of the span?
- ** Multi tenant traces, with spans coming from different tenants?
- ** Multi tenant client, ie, talking to different agents/collectors?
- ** Multi tenant agents, talking to different collectors based on span data?
- ** Multi tenant collector, receiving data from different agents (each agent being a tenant)?

* EVERYONE SEEMS TO HAVE A DIFFERENT REQUIREMENT

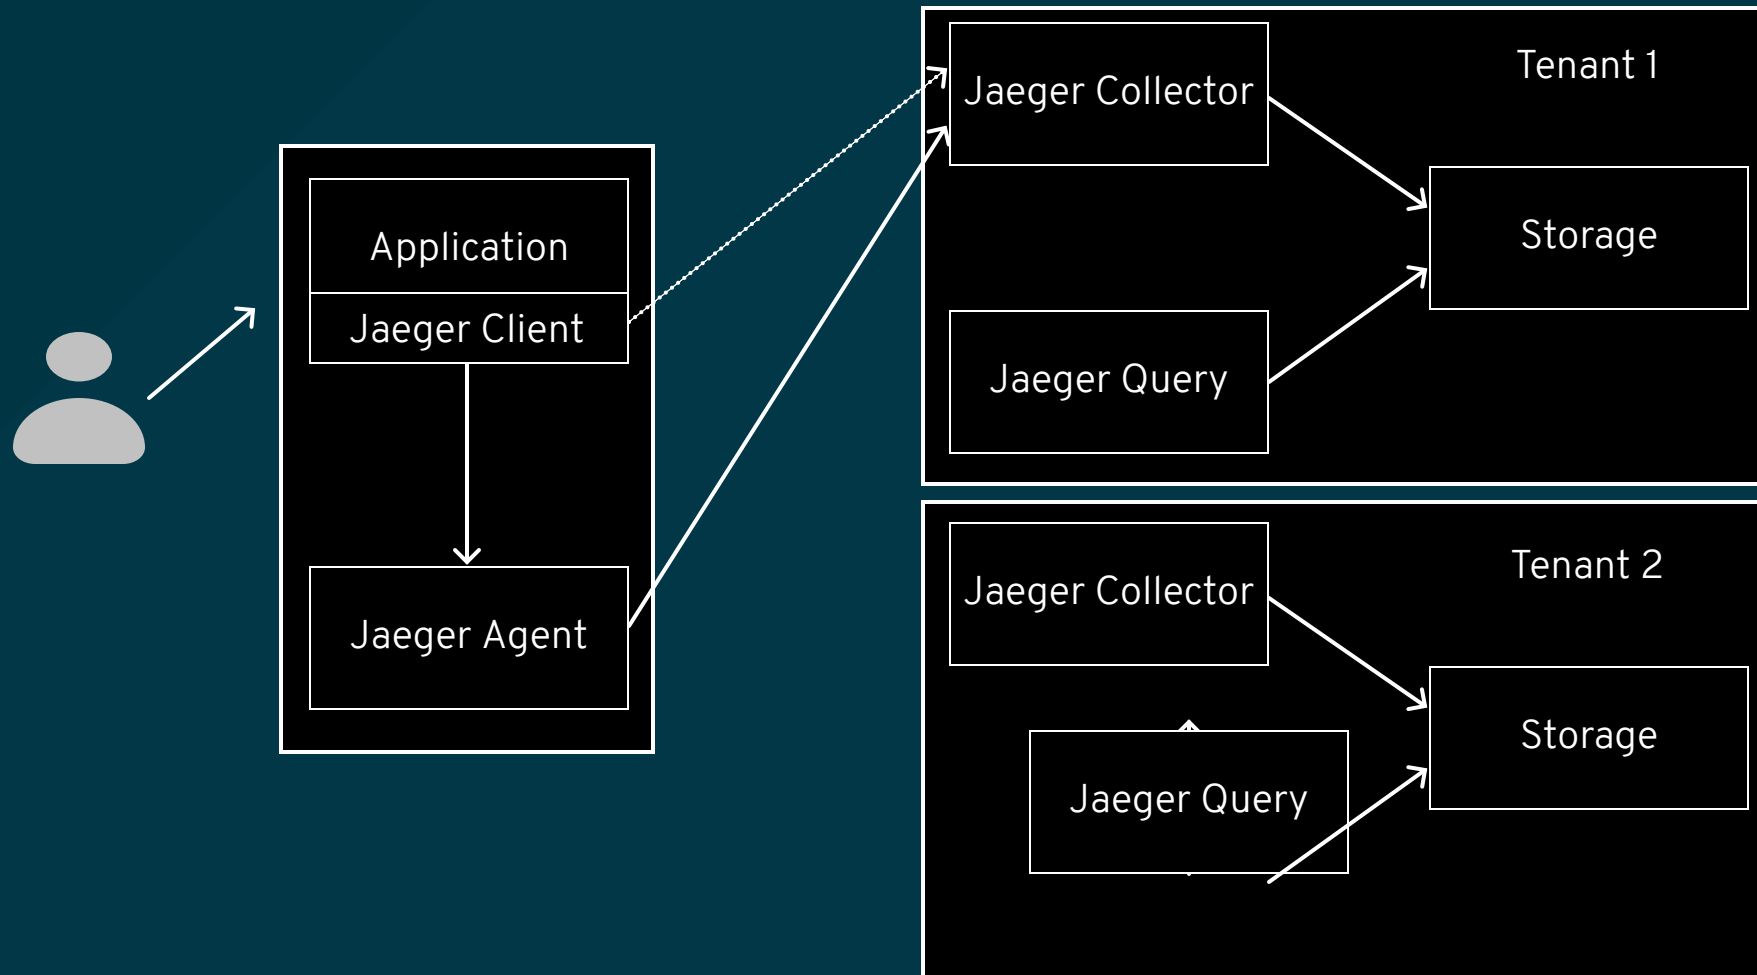
* REMEMBER: COMPONENTS ARE LIGHTWEIGHT!

DEPLOYMENT - MULTI TENANCY

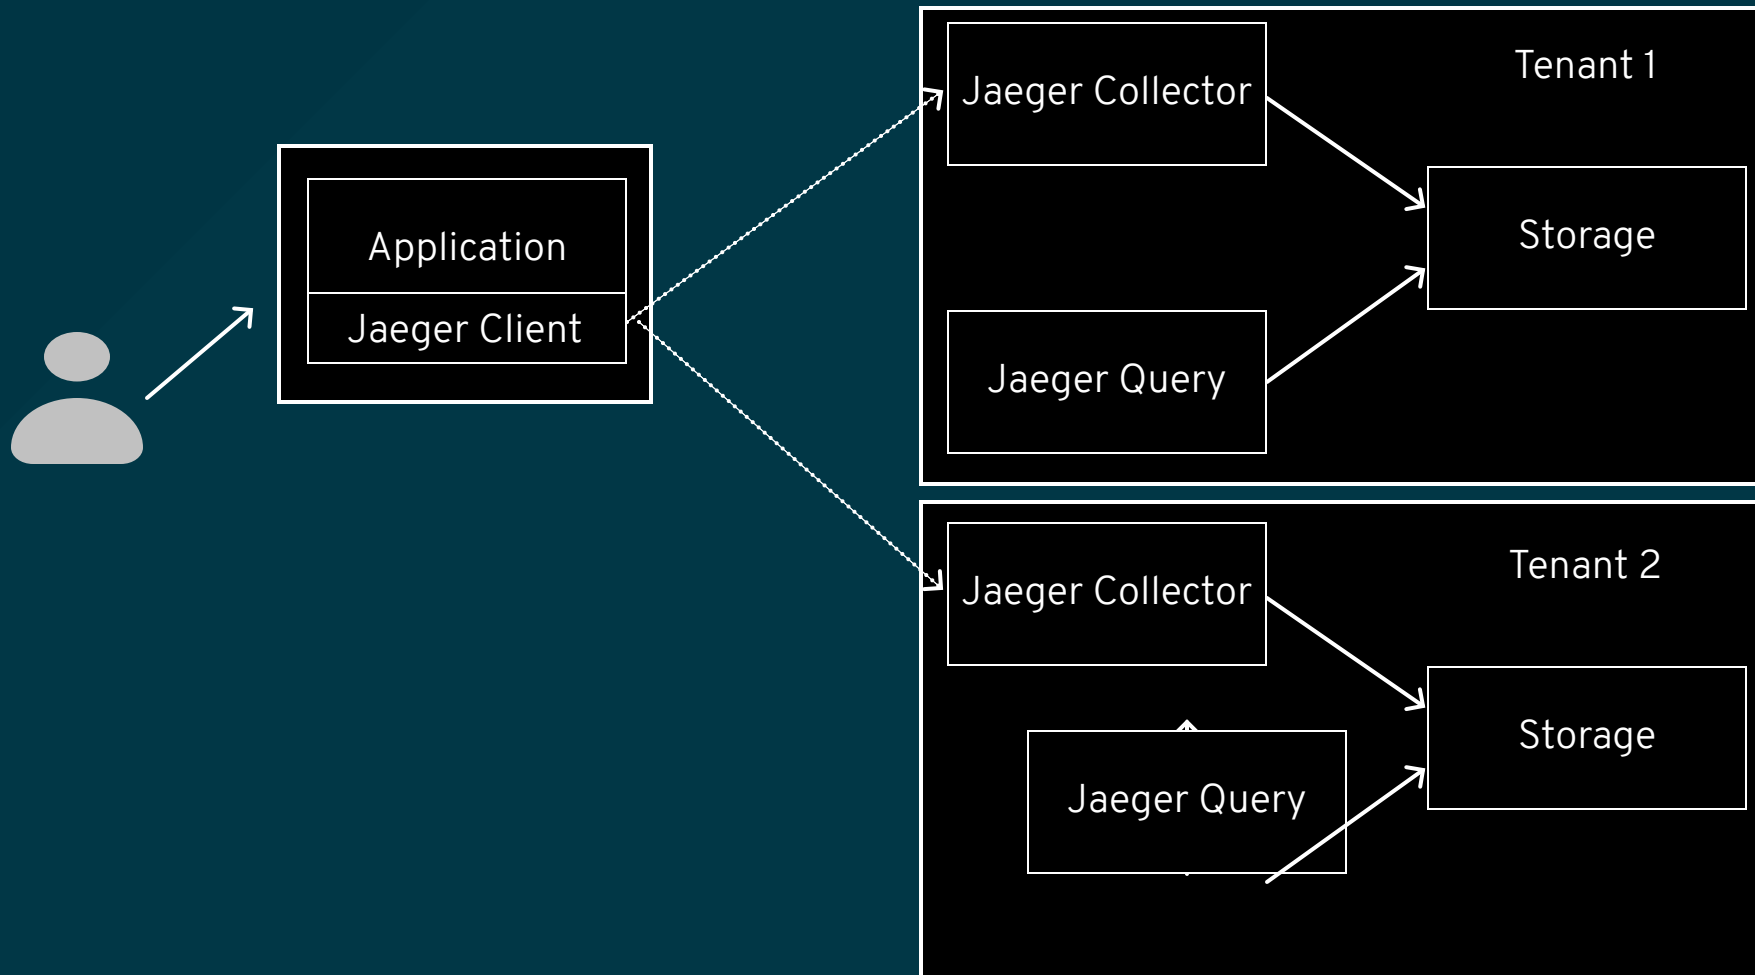
- * KUBERNETES NAMESPACES

- * CLIENT CAN IMPLEMENT DIFFERENT "REPORTERS"

DEPLOYMENT - MULTI TENANCY



DEPLOYMENT - MULTI TENANCY



DEPLOYMENT - SECURITY

* WHAT COULD BE SECURITY IN OUR CONTEXT?

** Should the agent require authentication, or just ingest spans?

** Should the collector require authentication? If multi tenancy is involved, should it enforce user vs. tenant?

** Should the UI be protected?

** Should the Query server be protected?

* EVERYONE SEEMS TO HAVE A DIFFERENT REQUIREMENT

* REMEMBER: COMPONENTS ARE LIGHTWEIGHT!

DEPLOYMENT - SECURITY

- * MICROSERVICES APPROACH: DELEGATE !
- * SECURITY PROXIES
- * CLIENT CAN IMPLEMENT DIFFERENT "REPORTERS"

ROADMAP

- * DATA PIPELINE

- * ADAPTIVE SAMPLING

- * STORAGE PLUGIN

JOIN OUR COMMUNITY

- * WE HAVE TONS OF QUESTIONS, HELP US FIND ANSWERS!
- * WE MEET ONCE EVERY TWO WEEKS
- * GITTER, MAILING LIST

Q&A