

**KubeCon**



**CloudNativeCon**

**Europe 2019**



KubeCon



CloudNativeCon

Europe 2019

# What's the Performance Overhead?

Answering the Biggest Question in Tracing

*Gabriela Soria*

*Former Outreachy intern for CNCF*

# About me – Outreachy




KubeCon



CloudNativeCon

Europe 2019

[Outreachy](#) provides three-month paid internships to work in Free and Open Source Software (FOSS) for under-represented people.



OPENTRACING



@gabrielasoriag

# Agenda



KubeCon



CloudNativeCon

Europe 2019

- Introduction
- Benchmark tests
  - Scope and considerations
  - JMH Benchmark application example
  - Results
  - Conclusions
- Next steps
- Lessons learned

# Introduction – Microservices



KubeCon




CloudNativeCon

Europe 2019


*A monolithic application puts all its functionality into a single process...*




*... and scales by replicating the monolith on multiple servers*



*A microservices architecture puts each element of functionality into a separate service...*



*... and scales by distributing these services across servers, replicating as needed.*



# Distributed tracing – OpenTracing



KubeCon

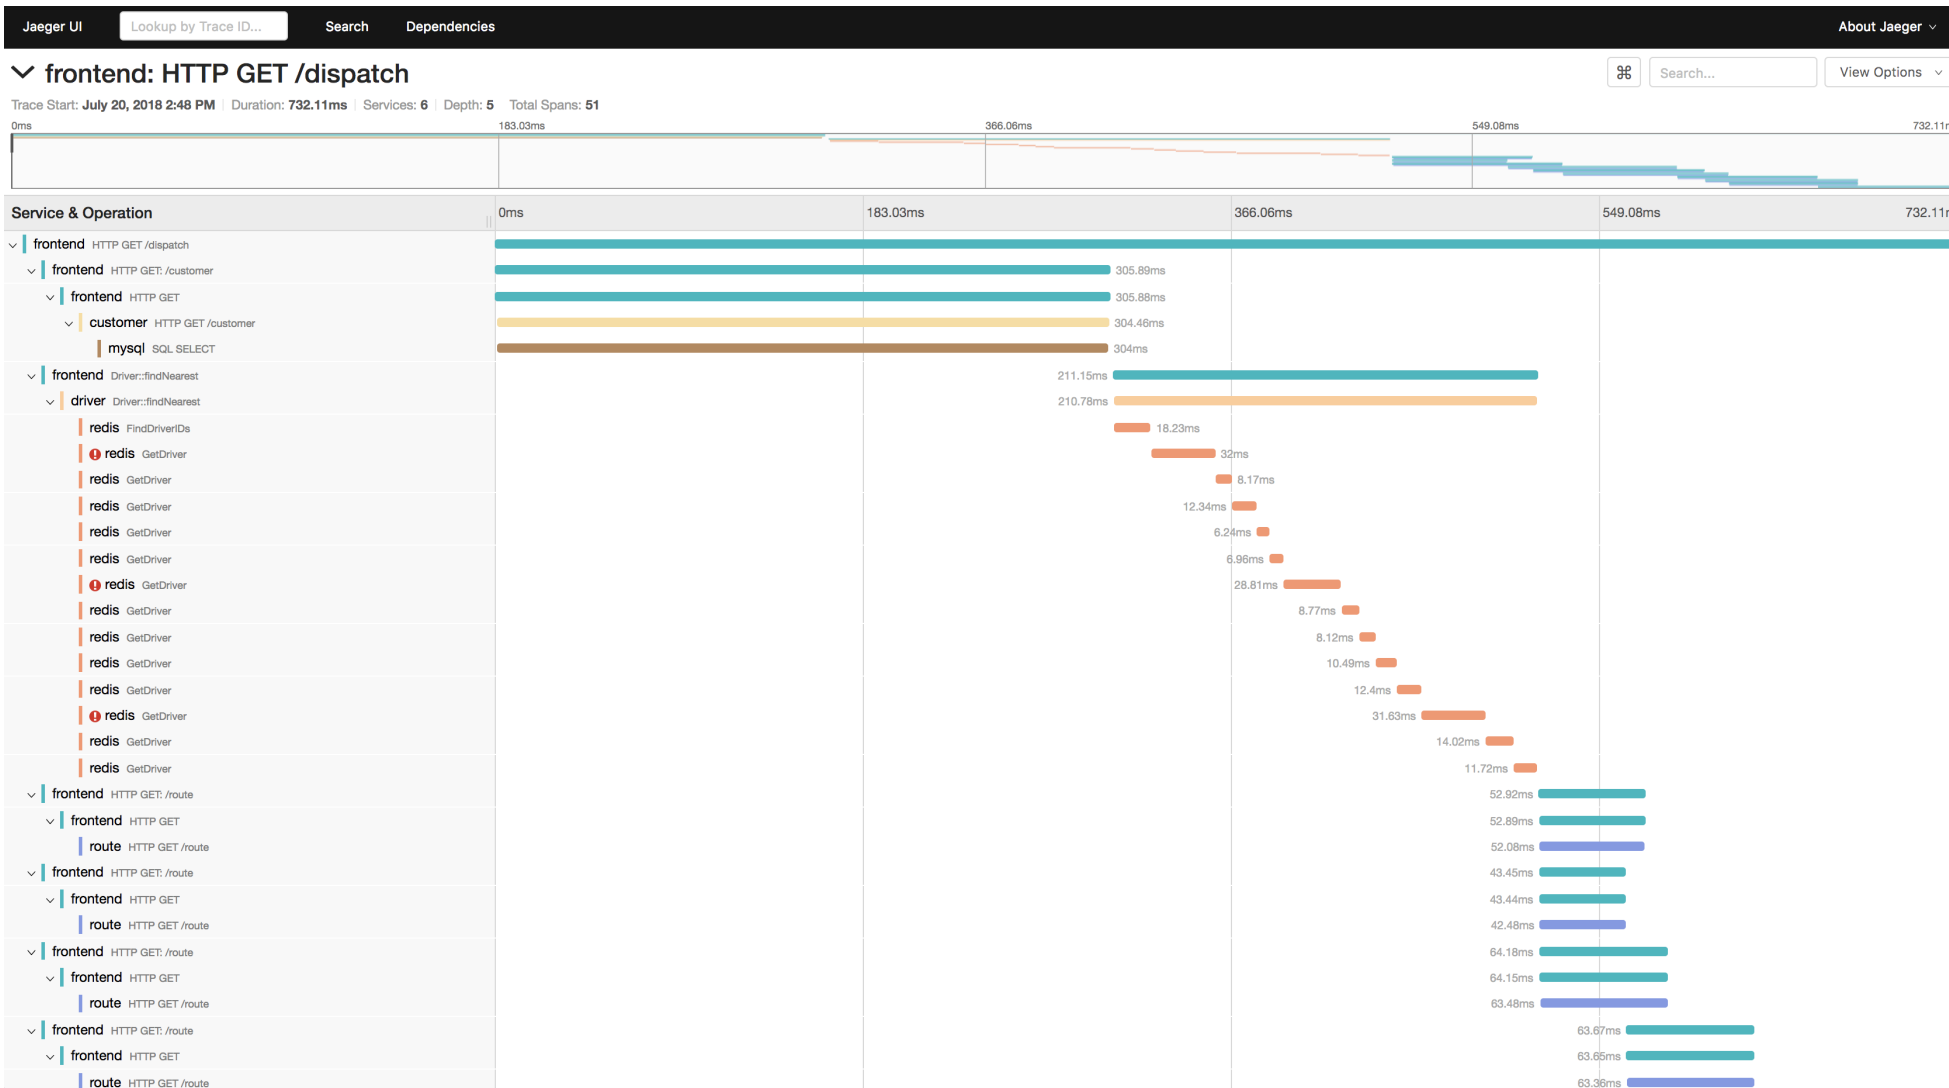


CloudNativeCon

Europe 2019

**Distributed tracing** profiles and monitor applications, especially those built using a microservices architecture.

**OpenTracing** is working towards creating more standardized APIs and instrumentation for distributed tracing.



# Scope and considerations



KubeCon



CloudNativeCon

Europe 2019

- We used Java Microbenchmark Harness (JMH)
- We used JMH Visualizer to present the results
- Tested Opentracing **Java** API only
- We tried to avoid network calls (in the scenarios that was possible)



# JMH Benchmark Application Example



KubeCon



CloudNativeCon

Europe 2019

- Create a JMH application:

```
mvn archetype:generate \  
-DinteractiveMode=false \  
-DarchetypeGroupId=org.openjdk.jmh \  
-DarchetypeArtifactId=jmh-java-benchmark-archetype \  
-DgroupId=org.sample \  
-DartifactId=jmh-examples \  
-Dversion=1.0
```

# JMH Benchmark Application Example



KubeCon



CloudNativeCon

Europe 2019

```
public class BenchmarkPetclinicSampleTime extends BenchmarkPetclinicBase {
```

```
@Benchmark
```

```
@BenchmarkMode(Mode.SampleTime)
```

```
public Owner noInstrumentation(StateVariablesNoInstrumentation state) {
```

```
    return findPetOwnerById(state);
```

```
}
```

```
@Benchmark
```

```
@BenchmarkMode(Mode.SampleTime)
```

```
public Owner noopTracer(StateVariablesNoopTracer state) {
```

```
    return findPetOwnerById(state);
```

```
}
```

```
@Benchmark
```

```
@BenchmarkMode(Mode.SampleTime)
```

```
public Owner jaegerTracer(StateVariablesJaeger state) {
```

```
    return findPetOwnerById(state);
```

```
}
```

```
....
```

@gabrielasoriag

# JMH Benchmark Application Example



KubeCon



CloudNativeCon

Europe 2019

```
public class BenchmarkPetclinicBase {

    public Owner findPetOwnerById(StateVariables state) {
        return state.petcontroller.findOwner(1);
    }
    @State(Scope.Benchmark)
    public static class StateVariables {
        public PetController petcontroller;
        public ConfigurableApplicationContext c;

        @TearDown(Level.Iteration)
        public void shutdownContext() {
            c.close();
        }

        public void initApplication() {
            c = SpringApplication.run(PetClinicApplication.class);
            petcontroller = c.getBean(PetController.class);
        }
    }
}
```

# JMH Benchmark Application Example



KubeCon



CloudNativeCon

Europe 2019

```
public static class StateVariablesNoInstrumentation extends StateVariables {  
    @Setup(Level.Iteration)  
    public void doSetup() {  
        System.setProperty("tracerresolver.disabled", Boolean.TRUE.toString());  
        initApplication();  
    }  
}
```

```
public static class StateVariablesJaeger extends StateVariables {  
    @Setup(Level.Iteration)  
    public void doSetup() {  
  
        System.setProperty(AbstractEnvironment.ACTIVE_PROFILES_PROPERTY_NAME,  
TracerImplementation.JAEGERTRACER);  
        initApplication();  
    }  
}
```

# JMH Benchmark Application Example



KubeCon



CloudNativeCon

Europe 2019

- Run the tests:

```
mvn clean install  
java -jar target/benchmarks.jar
```

|   |                                    |                      |
|---|------------------------------------|----------------------|
| ▼ | opentracing-benchmark-spring-cloud | Yesterday at 01:45   |
| ▶ | .idea                              | Today at 00:15       |
| ▶ | results-md                         | Yesterday at 02:55   |
| ↓ | README.md                          | Yesterday at 01:45   |
| ▶ | results-imgs                       | Yesterday at 01:39   |
| ▼ | results                            | Yesterday at 01:31   |
|   | jmh-2019-05-03-00-08-35.json       | May 3, 2019 at 09:25 |
|   | jmh-2019-05-03-01-05-59.json       | May 3, 2019 at 09:25 |
|   | jmh-2019-05-03-00-37-27.json       | May 3, 2019 at 09:25 |

# JMH Benchmark Application Example



KubeCon



CloudNativeCon

Europe 2019

- Visualize the results

<http://jmh.morethan.io/>

☰ JMH Visualizer

## Dropzone

Drop your JMH JSON report file(s) here!



*"JMH is a Java harness for building, running, and analysing nano/micro/milli/macro benchmarks written in Java and other languages targeting the JVM."*

@gabrielasoriag

# Results – Simple Java (Sample time)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkStringConcatenationSampleTime Sampling Time | 🔍 | ⚖️



@gabrielasoriag

# Results – Simple Java (Throughput)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkStringConcatenationThroughput Throughput | 🔍 | ⚖️



@gabrielasoriag



# Results – Spring boot (Sample time)



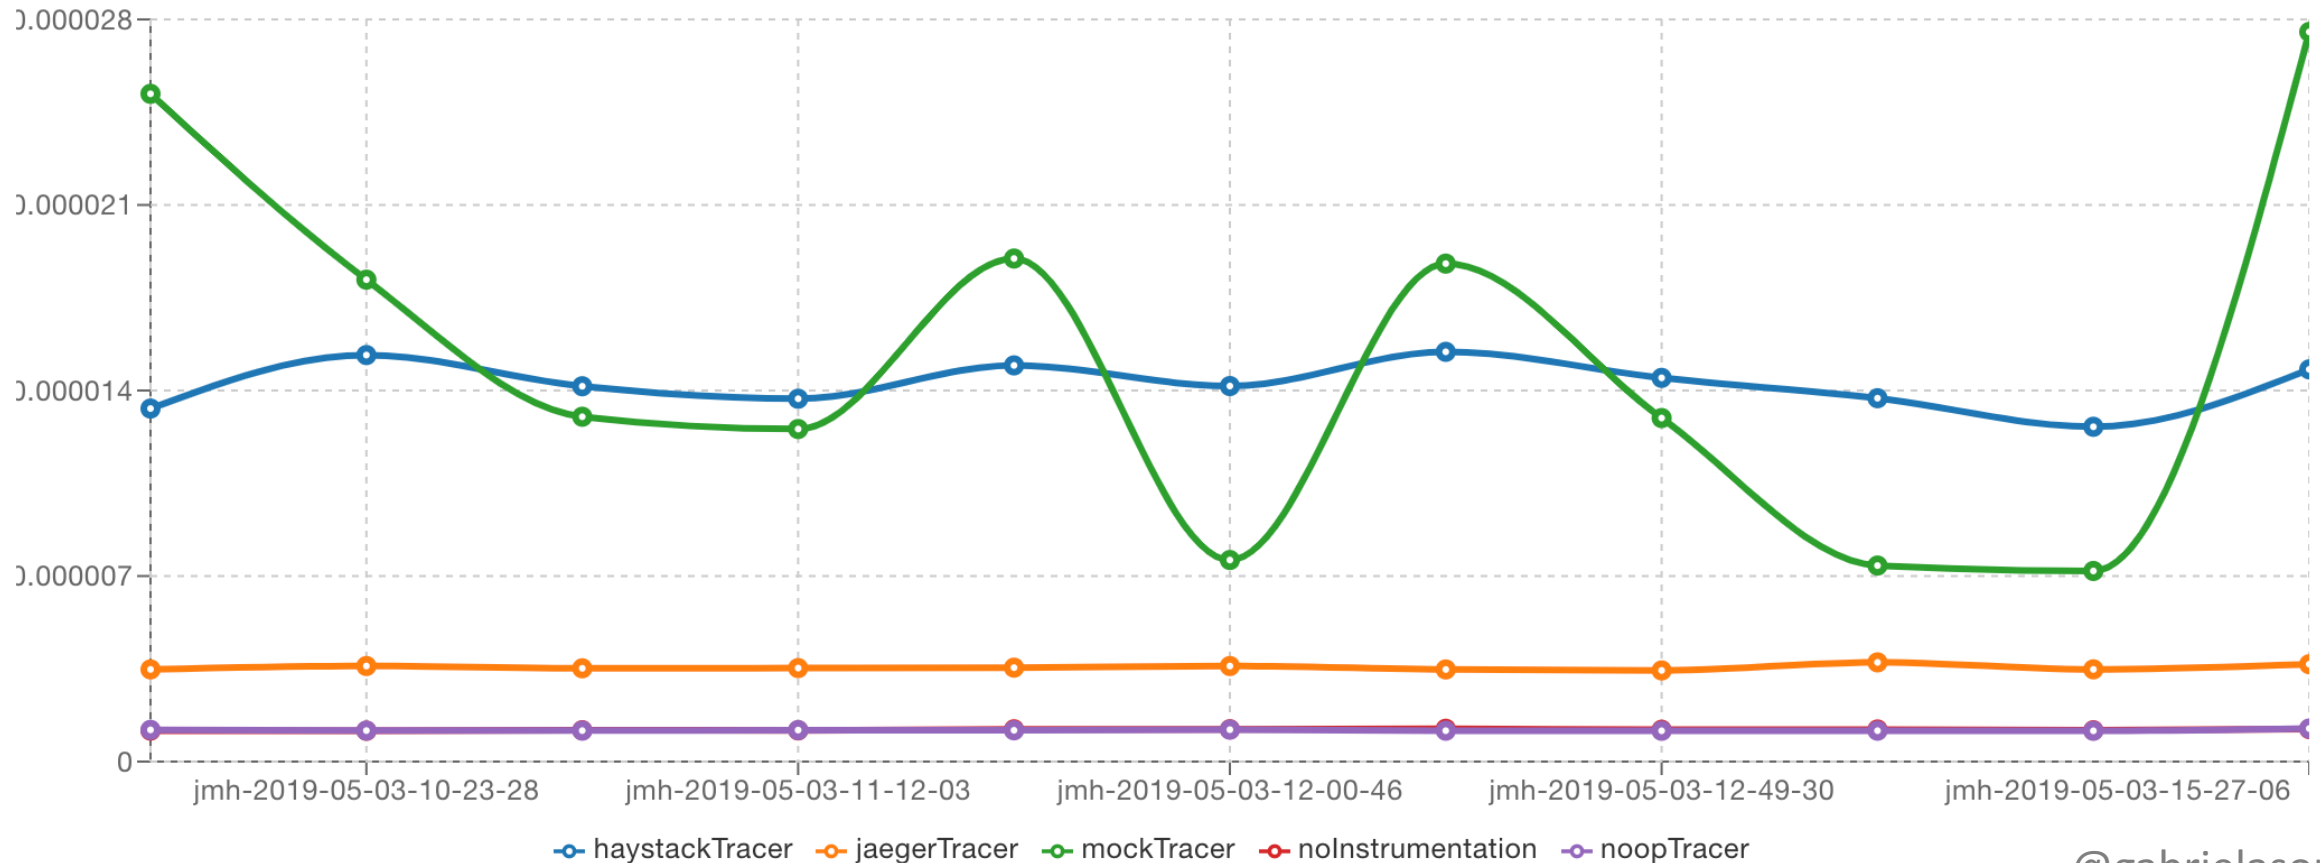
KubeCon



CloudNativeCon

Europe 2019

## BenchmarkBillingSampleTime Sampling Time | 🔍 | ⚖️



@gabrielasoriag

# Results – Spring boot (Throughput)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkBillingThroughput Throughput | 🔍 | ⚖️



@gabrielasoriag

# Results – Spring Cloud (Sample time)



KubeCon




CloudNativeCon

Europe 2019

## BenchmarkPetclinicSampleTime

Sampling Time



@gabrielasoriag

# Results – Spring Cloud (Throughput)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkPetclinicThroughput Throughput | 🔍 | ⚖️



@gabrielasoriag

# Results – JDBC (Sample time)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkCourseManagementSampleTime Sampling Time | 🔍 | ⚖️



@gabrielasoriag

# Results – JDBC (Throughput)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkCourseManagementThroughput Throughput | 🔍 | ⚖️



@gabrielasoriag

# Results – Servlet Filter (Sample time)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkSimpleServletSampleTime Sampling Time | 🔍 | ⚖️



@gabrielasoriag

# Results – Servlet Filter (Throughput)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkSimpleServletThroughput Throughput | 🔍 | ⚖️



@gabrielasoriag



# Results – JAX-RS (Sample time)



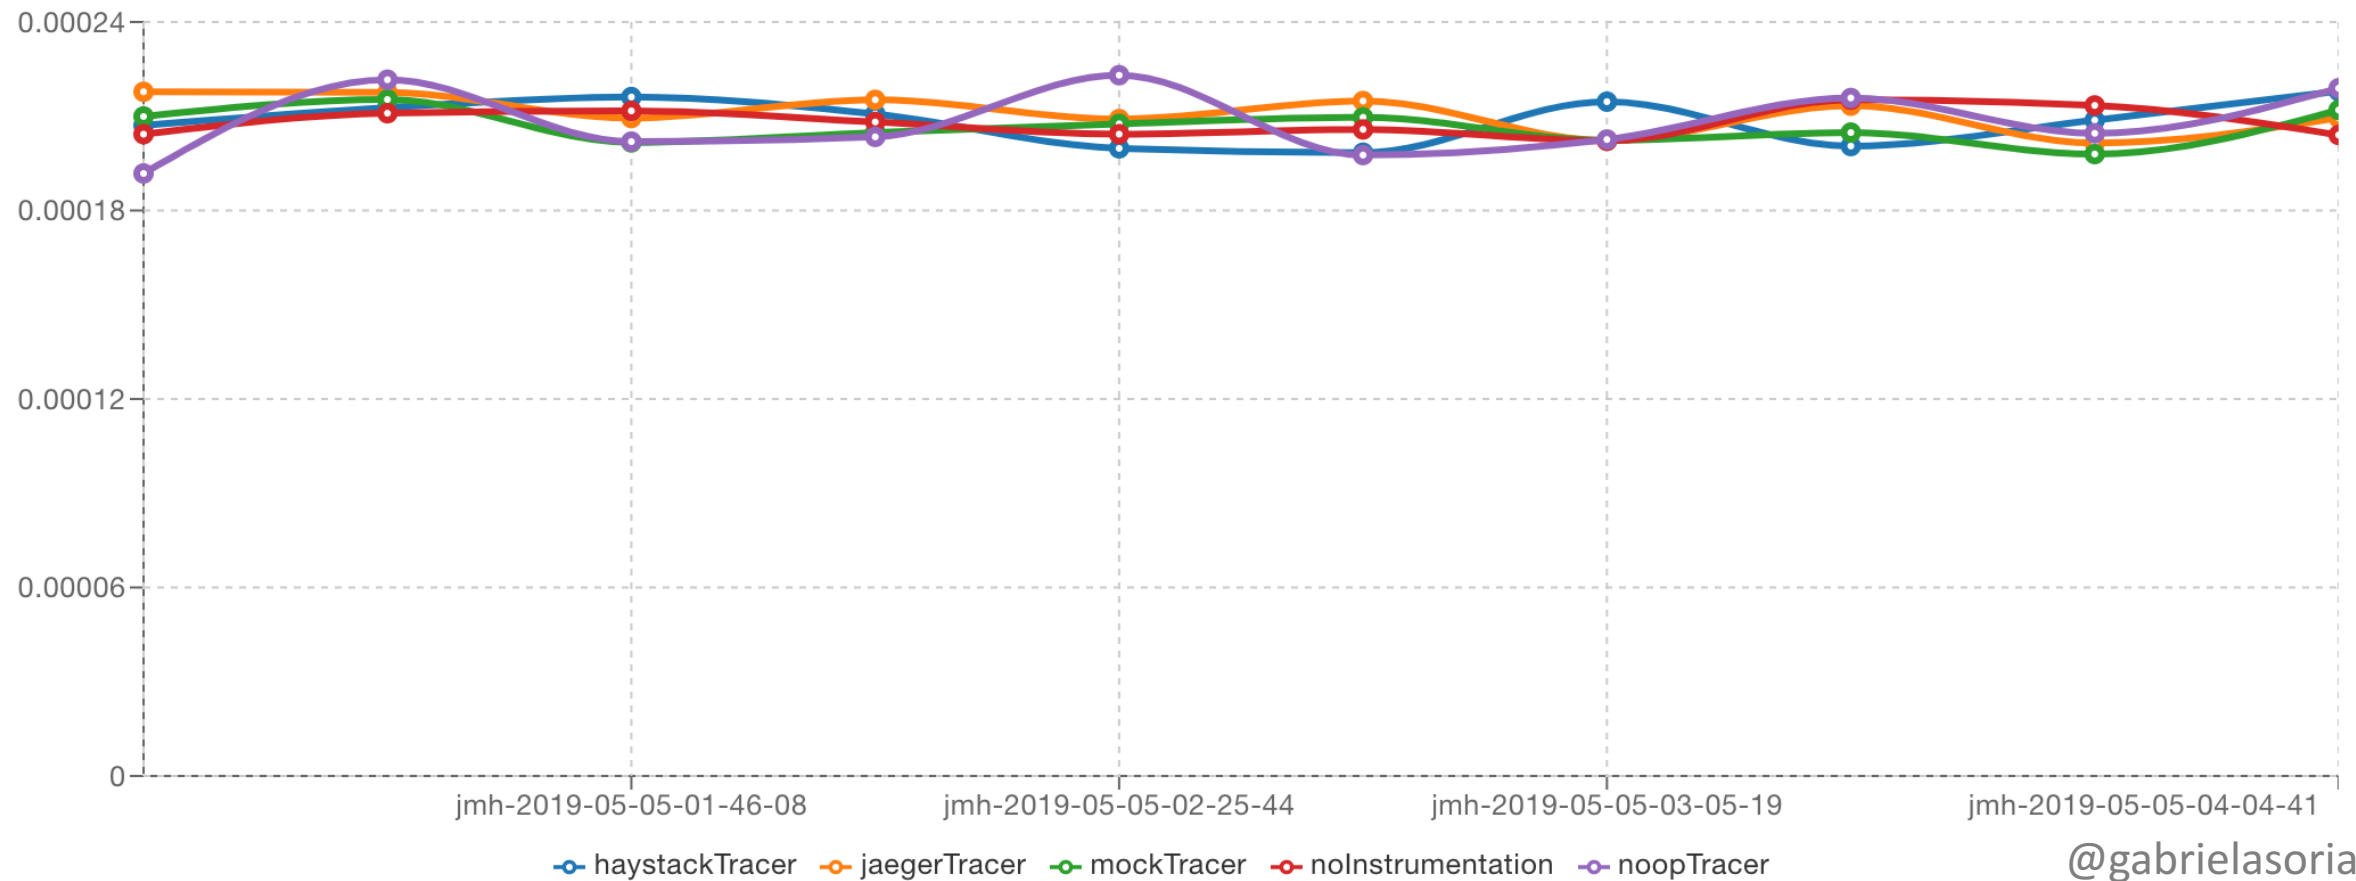
KubeCon



CloudNativeCon

Europe 2019

## BenchmarkCourseManagementSampleTime Sampling Time | 🔍 | ⚖️



@gabrielasoriag

# Results – JAX-RS (Throughput)




KubeCon



CloudNativeCon

Europe 2019

## BenchmarkCourseManagementThroughput Throughput | 🔍 | ⚖️



@gabrielasoriag

# Conclusions



KubeCon



CloudNativeCon

Europe 2019

- In simple java scenarios the throughput **decreases ~90%** and the sample time **increases ~440%**
- In scenarios that include calls through the framework (spring boot, spring cloud, JDBC), on average, the throughput **decreases 12%** and the sample time **increases 14%**.
- In the scenarios with **client calls through HTTP** (Servlet Filter, JAX-RS), the metrics show **no evidence of overhead**, as the deltas of throughput and sample time are not representative.

# Next steps



KubeCon



CloudNativeCon

Europe 2019

- Benchmark tests for Jaeger in different scenarios:
  - gRPC vs. Thrift
  - Agent (UDP) vs. Collector (gRPC)
  - Different backend config
- Re-run the tests after merging OpenTracing and OpenCensus

# Lessons learned



KubeCon



CloudNativeCon

Europe 2019

- Performance tests are tricky, the constant review of the code helps to improve the tests.
- You should ask yourself constantly if the benchmark test is measuring what you **really** want to measure.
- JMH
  - The method of the tests should return a value.
  - The garbage collector influences the results.
  - JMH performs dead code elimination (NoopTracer)

# Keep in touch :)




KubeCon



CloudNativeCon

Europe 2019



@gabrielasoriag



Opentracing Java Benchmarks

