



Combining OpenTracing and Kernel Tracing for Performance Analysis of Distributed Applications

Progress Report Meeting - May 6, 2019

DORSAL

Ecole polytechnique de Montréal

Loïc Gelle

Michel Dagenais

Context




OpenTracing: where does it help, where does it fail?

Key facts about OpenTracing

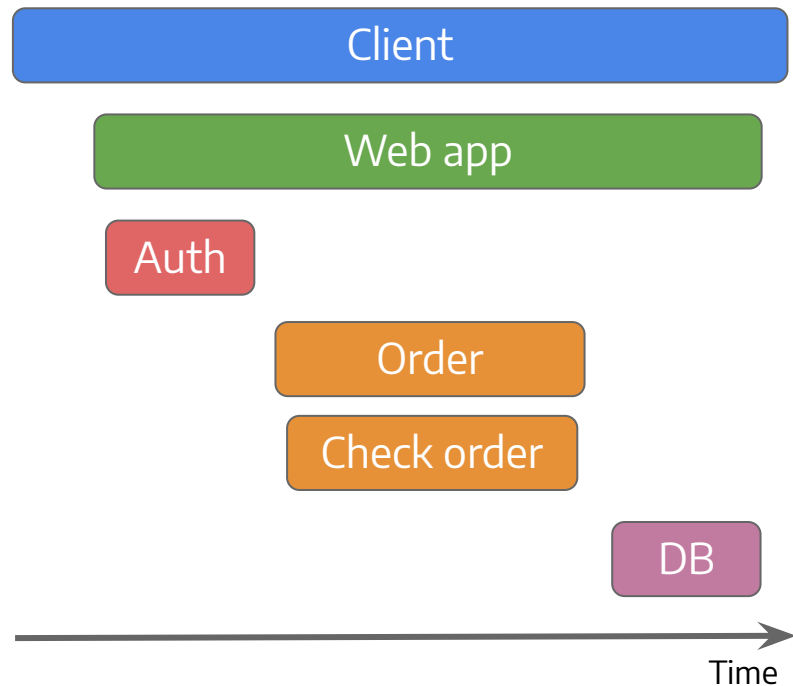
- An open-source **specification for distributed tracing**
- A **vendor-neutral API** for instrumenting libraries
 - API available for **popular languages** like Java, Go, C++, Python...
 - Lots of **libraries** like gRPC, NodeJS... are instrumented
- Many tracers (Jaeger, OpenZipkin, LightStep...) implement the OpenTracing specification
 - OpenTracing **leaves implementation details** to the tracers
 - Each tracer has **different purposes and analyses / UI**

Describing complex transactions

OpenTracing focuses on describing **tasks** instead of events.

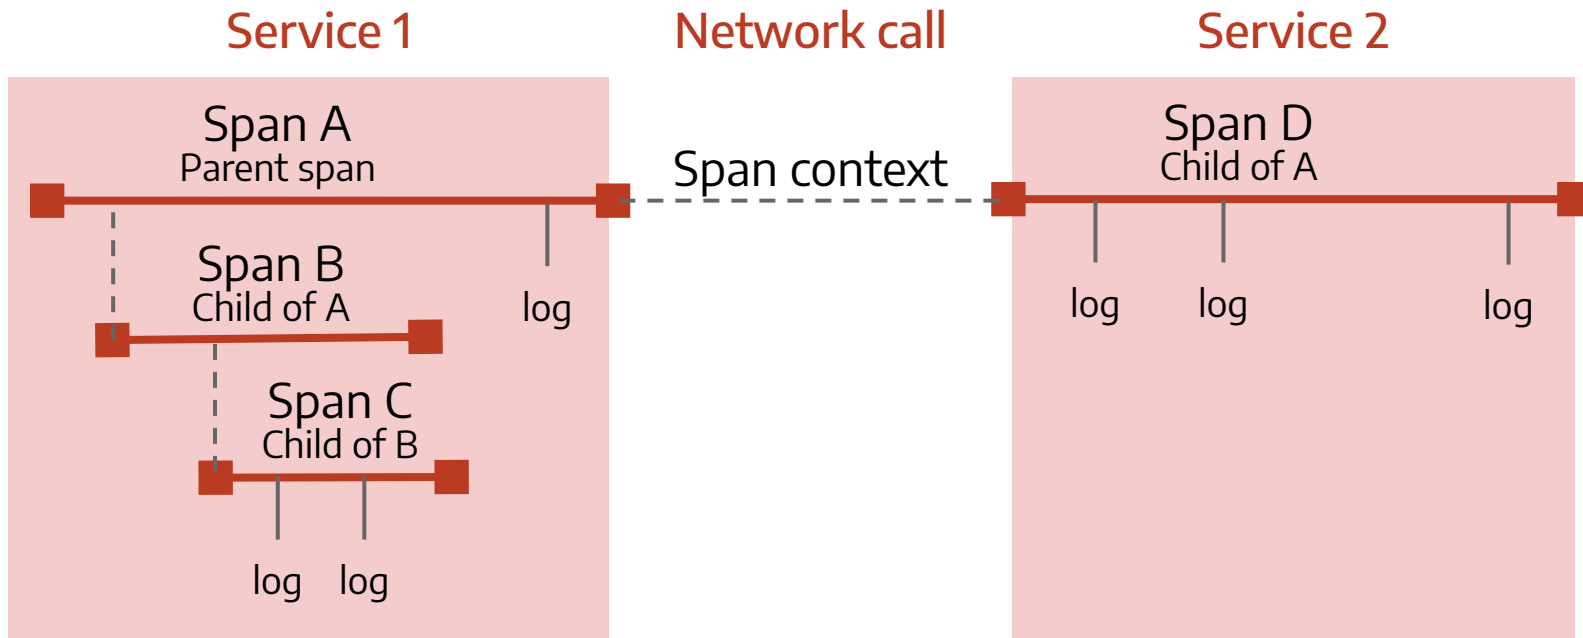


What the transaction looks like



What the trace looks like

Key concepts in OpenTracing

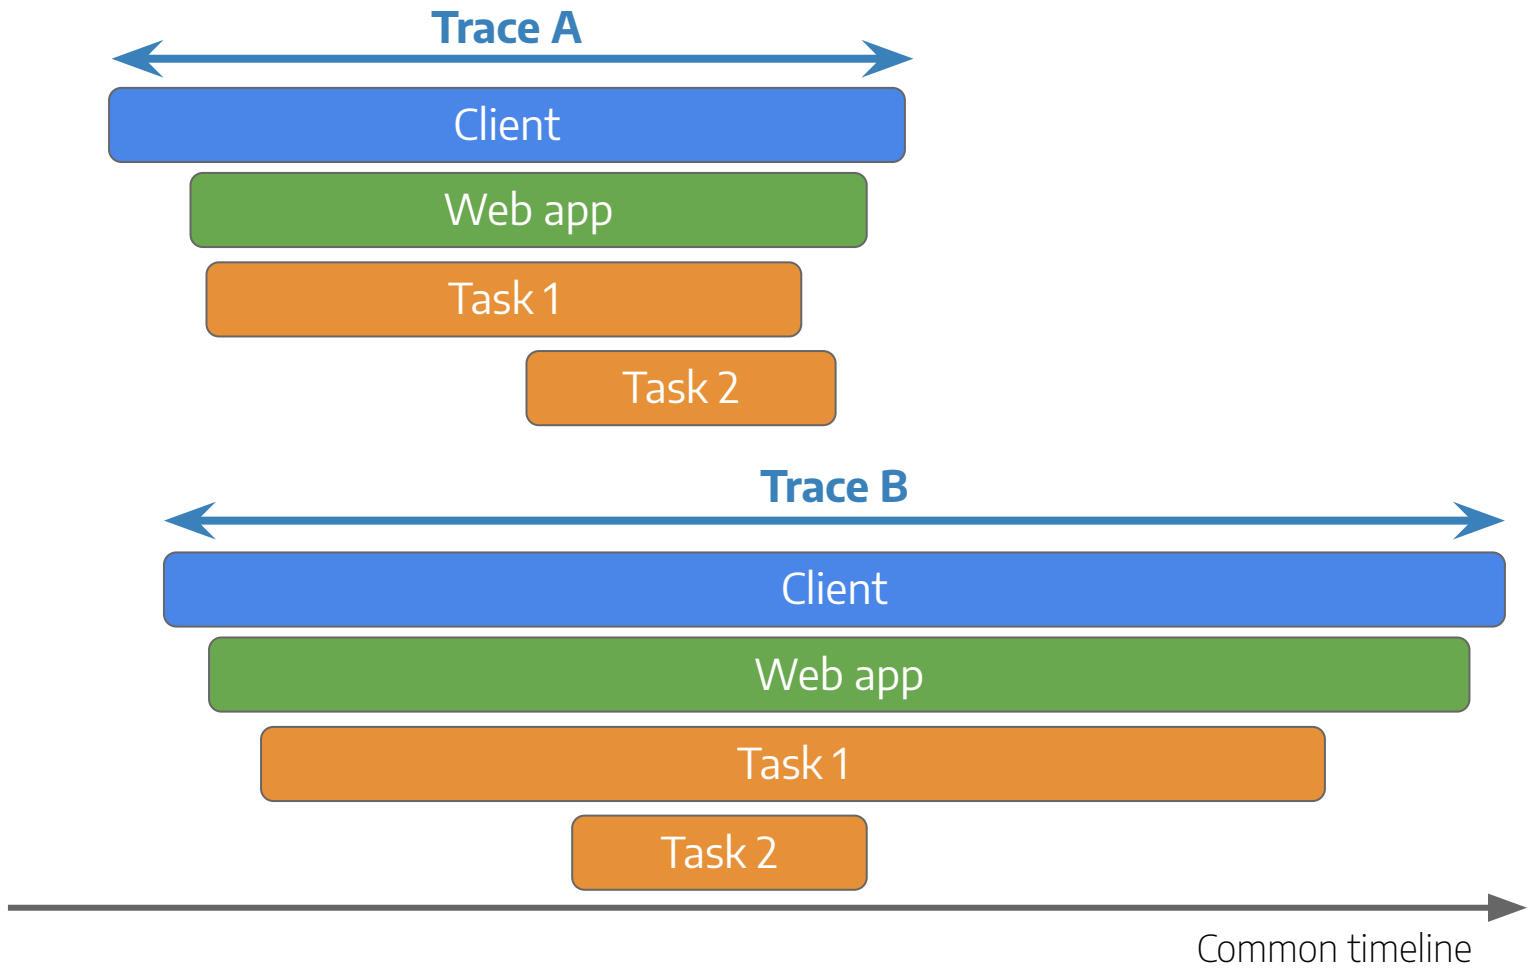


- A **span** has a name, a start, a duration, tags and attached logs.
- The **span context** identifies the trace; it is injected into requests.
- A **trace** is the recording of the whole transaction using the above!

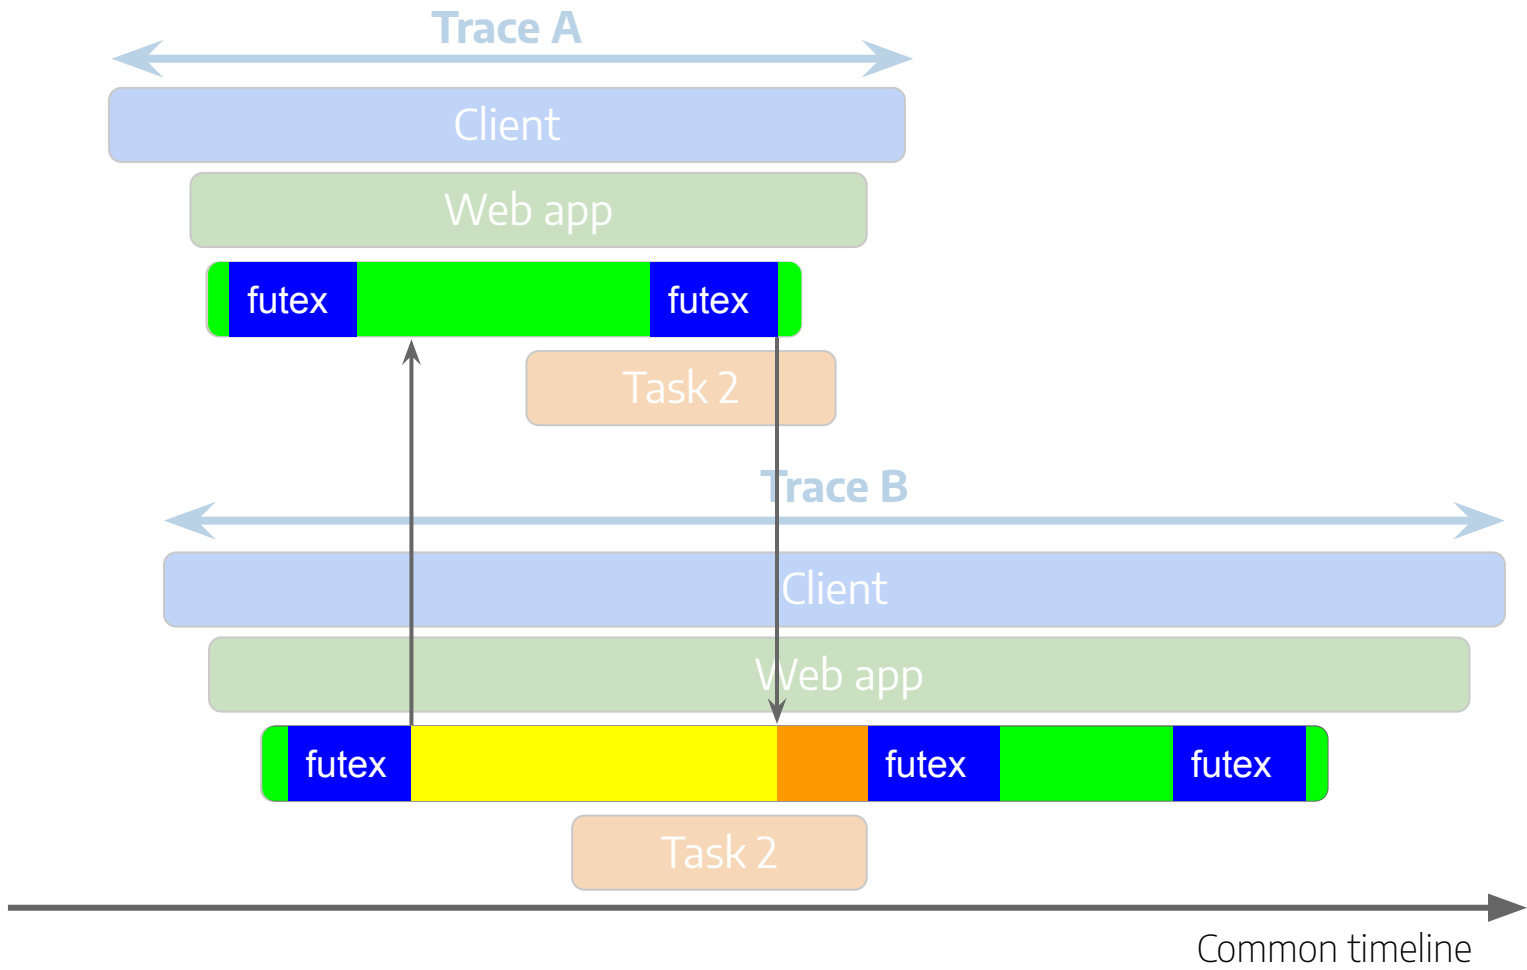
The benefits of OpenTracing

- The community is growing
- The traces provide useful **high-level context** for debugging applications
- The tracers provide the machinery to **collect the traces and display them**
- Use and deployment are fairly easy

Where does OpenTracing fail?



Same events, different perspective



The approach



Combining OpenTracing and kernel traces

Bridging the gap

- On the one side: threads, nanosecond-precise events

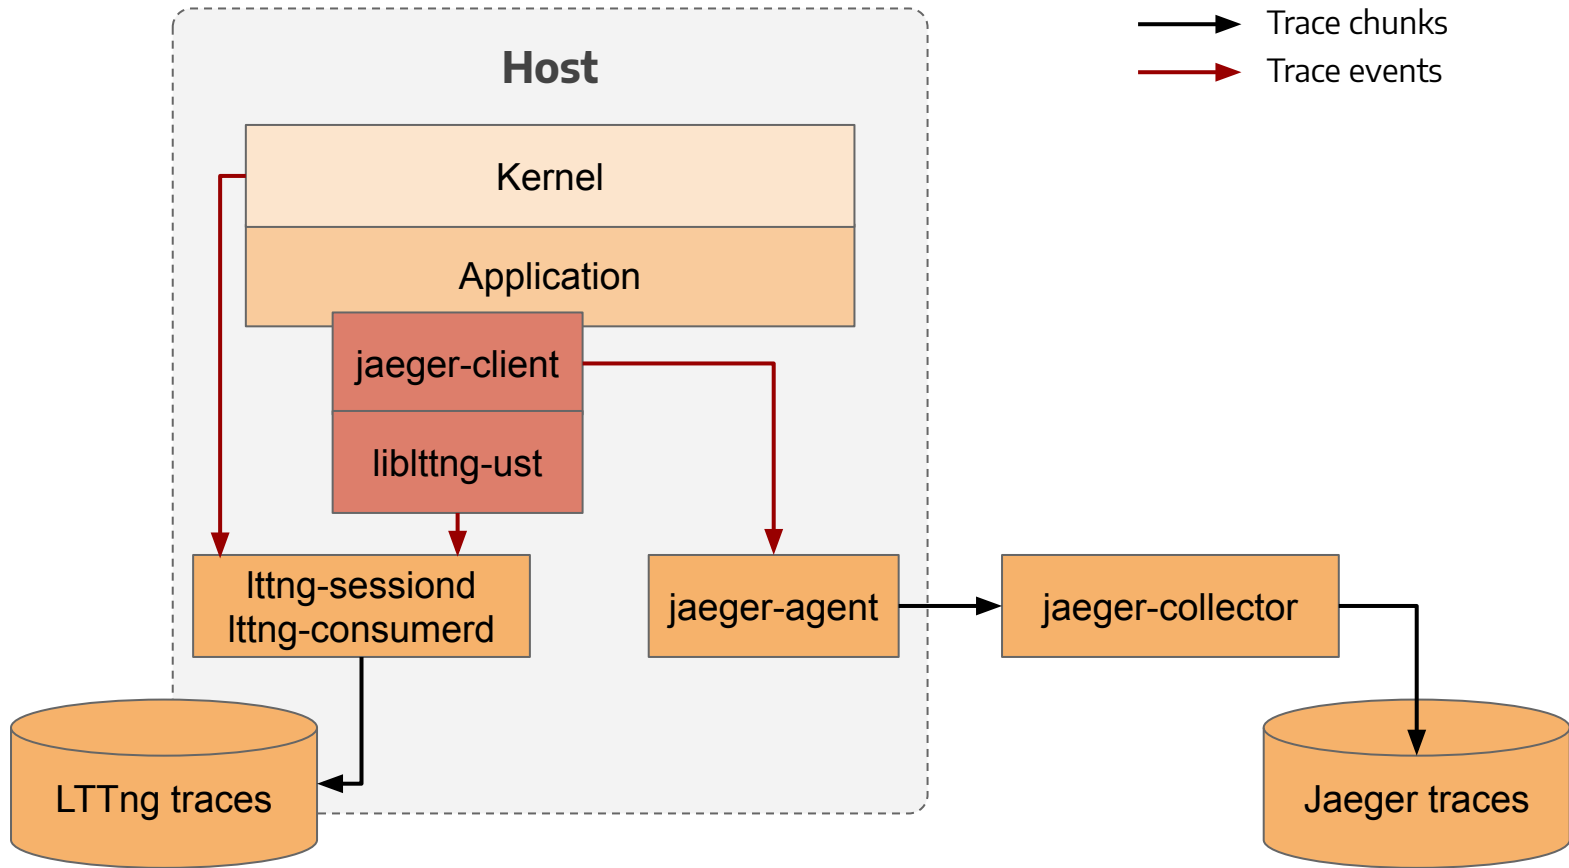


- On the other side: tasks, microsecond-precise events
 - We need to synchronize events
 - We need to relate tasks back to their thread(s)

Techniques for synchronization

- “Fake syscall” (*Google*)
- Kernel module + added LTTng kernel context (*Boston University*)
- Instrumentation of the OpenTracing tracer using LTTng-UST (*what we use*)

Collection of traces



Analyses



TraceCompass views

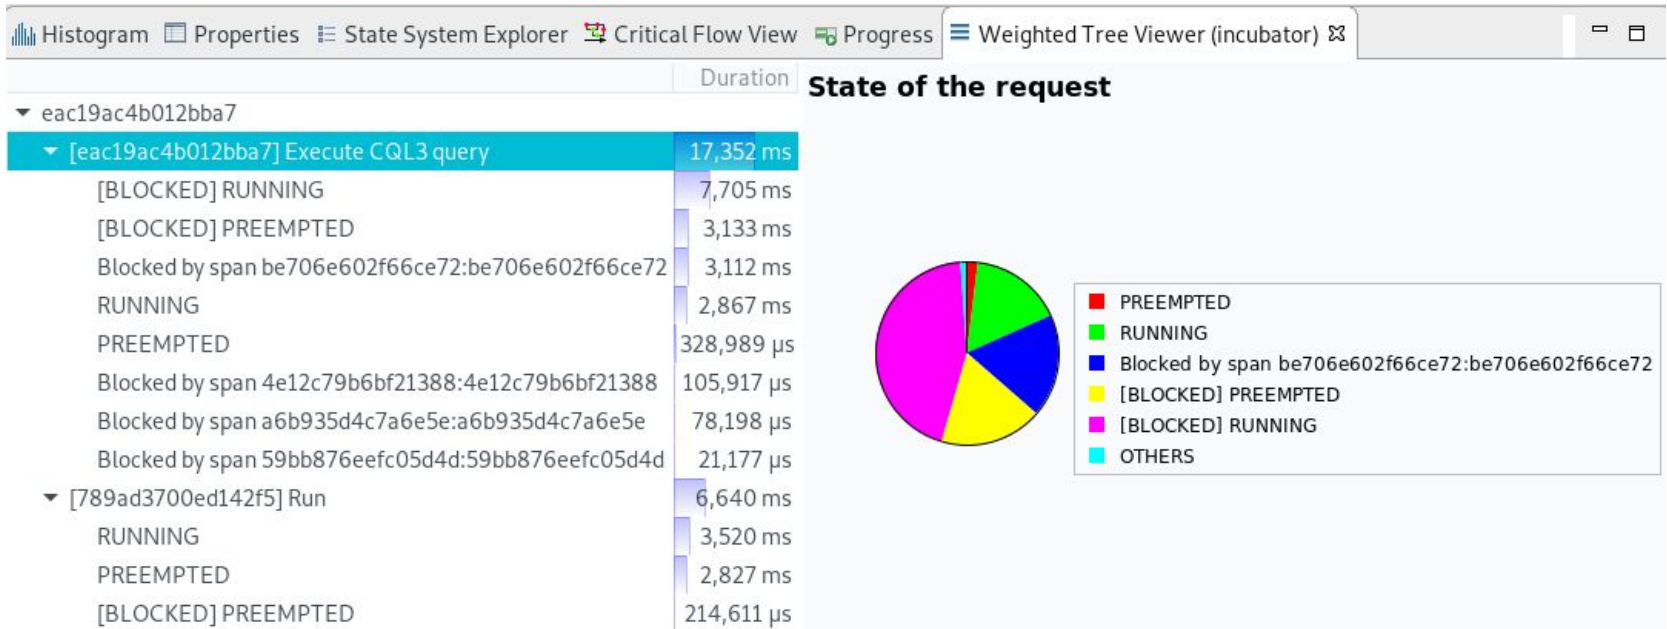
Proof of concept in TraceCompass

- Two views to validate the approach
 - Critical path of requests
 - Aggregated view per request of the critical path
- Based on prior work from Ericsson
- The instrumented application is Cassandra

Critical path of requests



Aggregated information



Conclusions and future work



Limitations and remarks

- We need developers to provide a good instrumentation of their application
- Analyses limited to a single machine
- The volume of the traces can be tough to handle and sampling is not straightforward
- Benchmarking has yet to be done

Future work

- Adapt the trace collection / analysis to applications hosted in containers
- Bring the analyses to UIs widely used by the OpenTracing community (Jaeger, Kibana)
- Work with the community to integrate the changes to the OpenTracing tracers



Thank you!
Questions, ideas, remarks?

 loic.gelle@polymtl.ca

 Github: [@loicgelle](https://github.com/loicgelle)