Trace Coordinator

Ahmad Faour

Polytechnique Montréal

DORSAL Laboratory

Agenda

- Introduction
- Trace Coordinator
- Connect logically the traces
- Distributed Analysis
- Future work
- Conclusion

Introduction - Limitation of Trace Compass Server

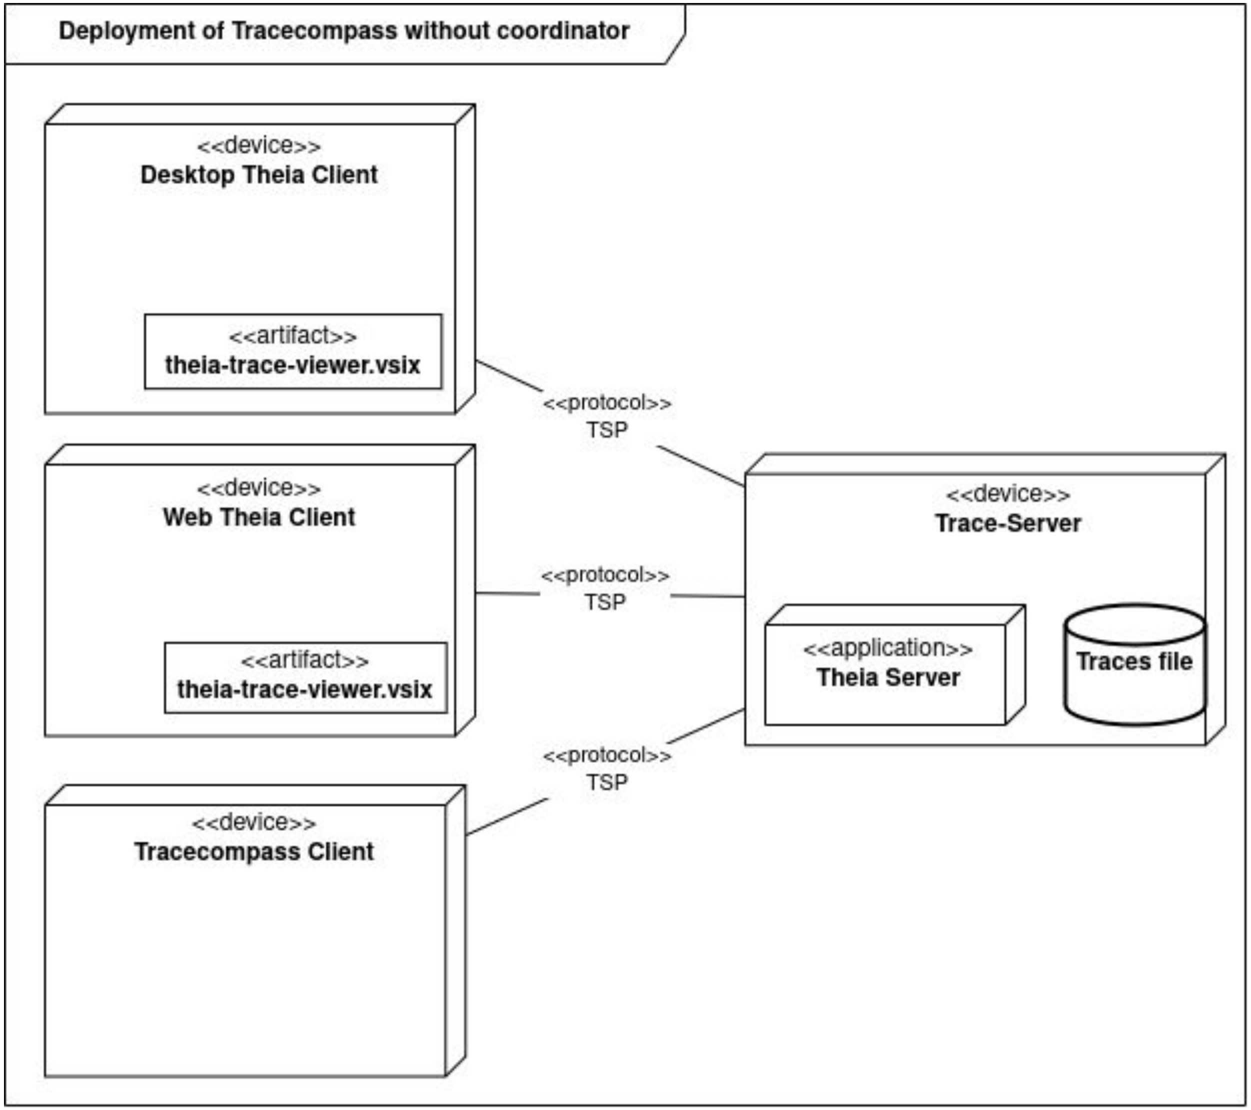

- Different Clients
- Trace Server
- Trace Server Protocol (TSP)
- Traces file

Limitation:

- Scale Horizontally
- All trace files must be uploaded on the same node

Introduction - Use cases

- Target use cases
  - High Performance Computing with MPI Cluster
  - Microservice with Kubernetes Cluster
- Other interesting use cases
  - LTTng Log Rotation
  - Similarity of queries
  - Client-Server

Introduction - Challenges about distributed tracing

- Number and size of traces (Solved)
- Connect logically the traces
  - What are the available traces (Solved)
  - Where are the traces (Solved)
  - How to connect the traces
- Distributed analysis
  - Aggregation over time (Horizontally) (Solved)
  - Aggregation of trace data (Vertically) (Solved)
  - Follow dependencies between processes (Critical Path)*
- Visualisation (Solved)

*Distributed computation of critical path (from Pierre-Frédérick Denys)

Reminder - Trace Coordinator

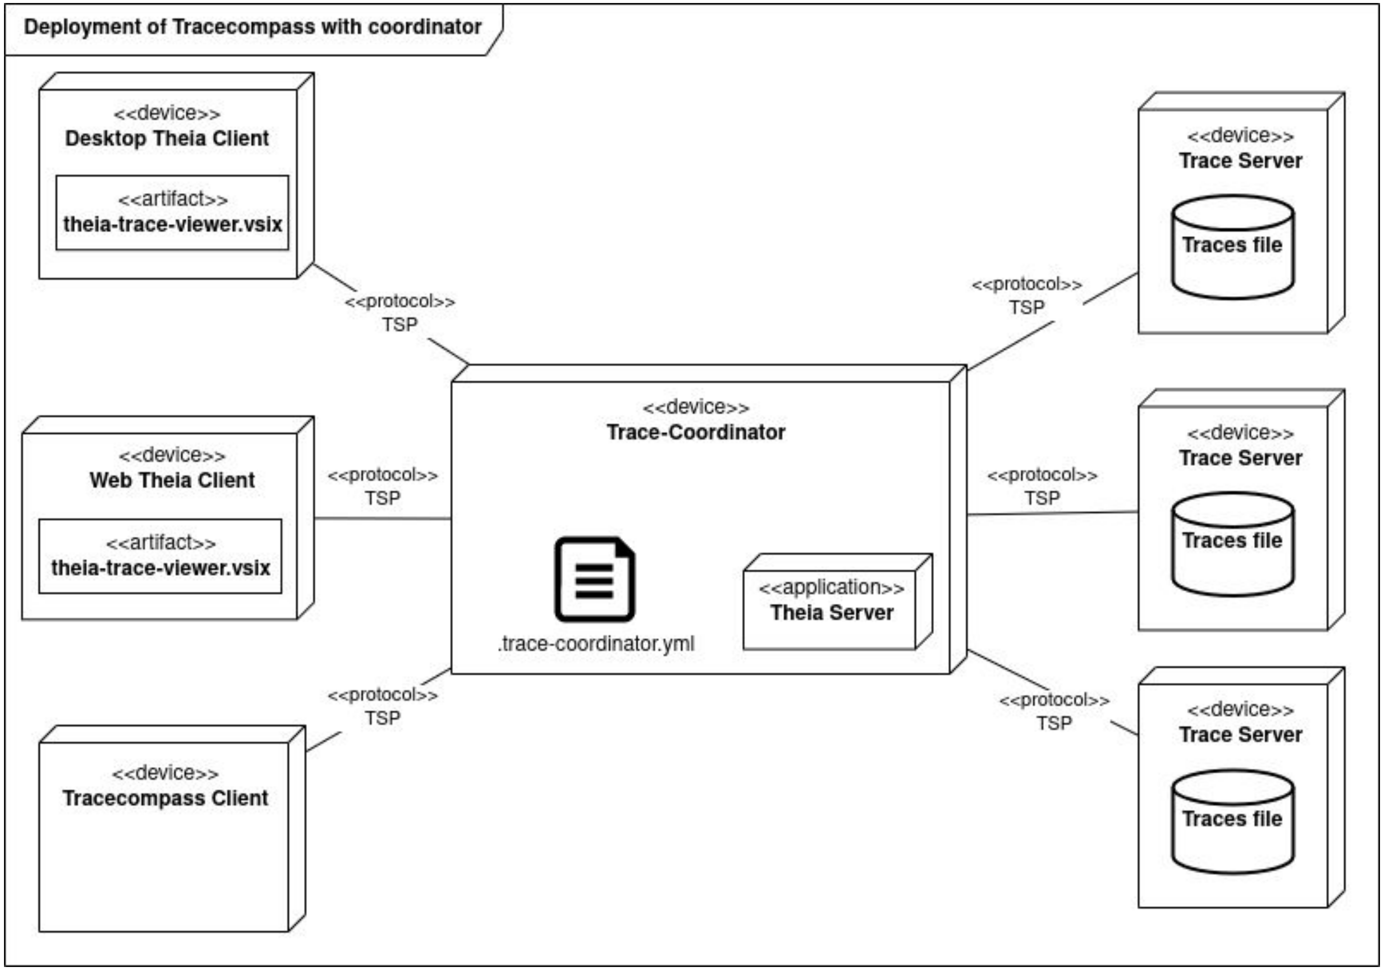

- Different Clients
- Trace Server Protocol (TSP)
- Many Trace Servers
- Trace Coordinator
- Distributed trace files

Connect logically the traces

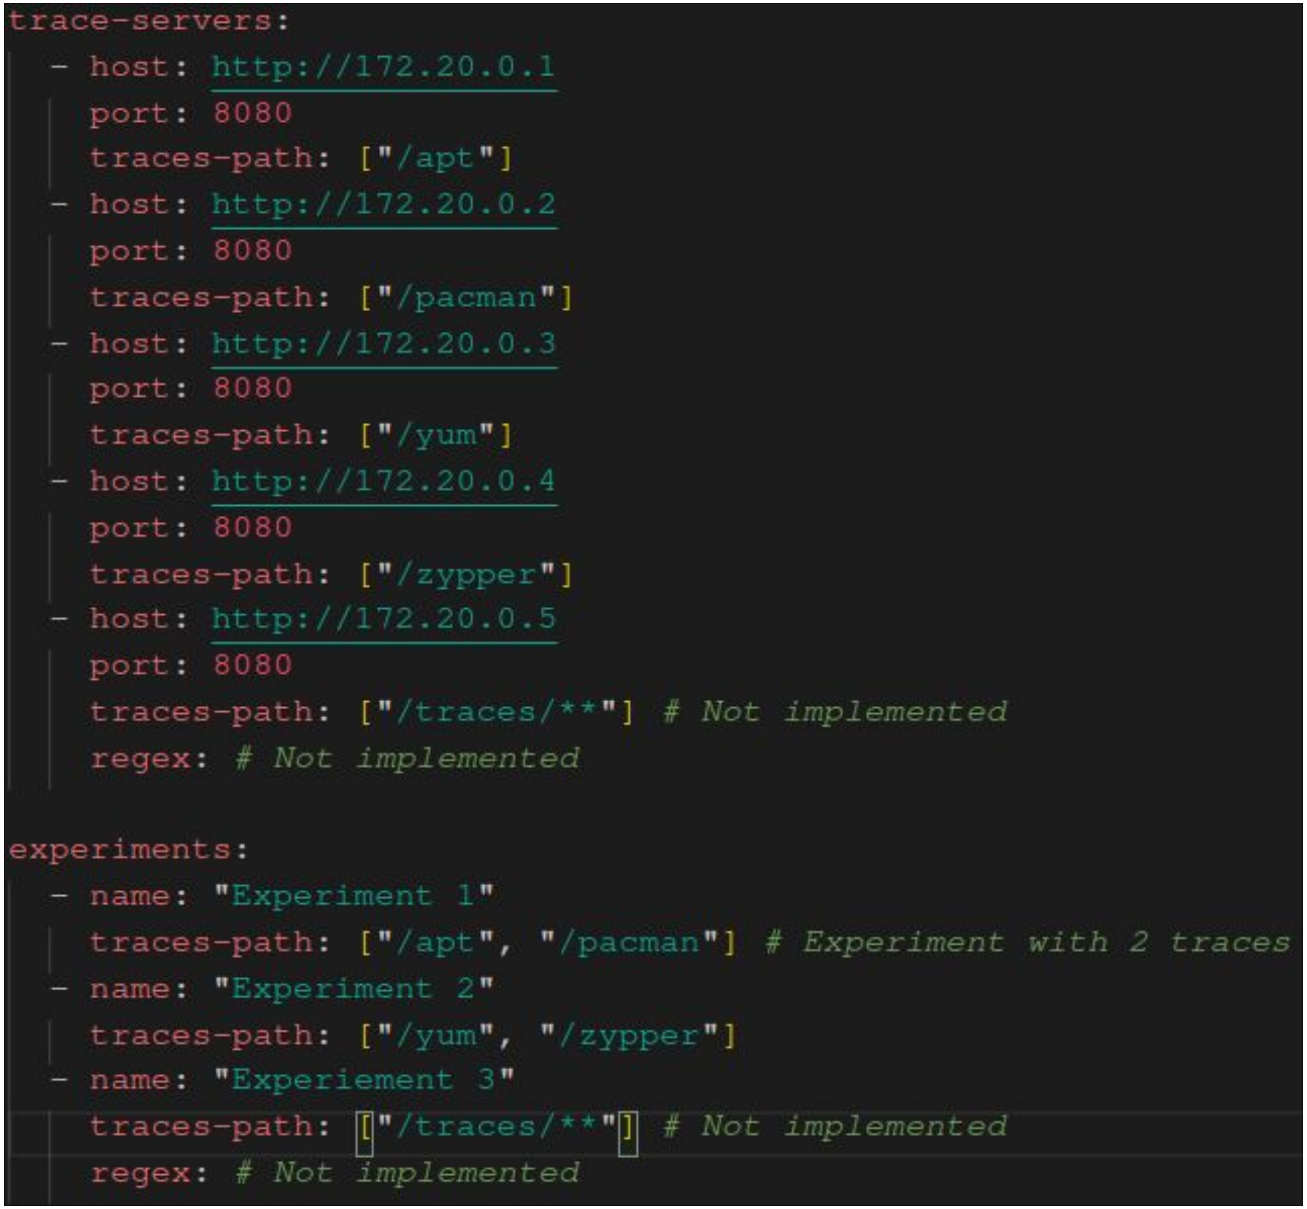

```
trace-servers:
  - host: http://172.20.0.1
    port: 8080
    traces-path: ["/apt"]
  - host: http://172.20.0.2
    port: 8080
    traces-path: ["/pacman"]
  - host: http://172.20.0.3
    port: 8080
    traces-path: ["/yum"]
  - host: http://172.20.0.4
    port: 8080
    traces-path: ["/zypper"]
  - host: http://172.20.0.5
    port: 8080
    traces-path: ["/traces/**"] # Not implemented
    regex: # Not implemented

experiments:
  - name: "Experiment 1"
    traces-path: ["/apt", "/pacman"] # Experiment with 2 traces
  - name: "Experiment 2"
    traces-path: ["/yum", "/zypper"]
  - name: "Experiement 3"
    traces-path: ["/traces/**"] # Not implemented
    regex: # Not implemented
```

- What are the available traces
- Where are the traces
- How to connect the traces

Distributed analysis - Aggregation of trace data (Vertically)

- Generate new Id
  - XY (Series, Tree)
  - TimeGraph (Arrows, States, Tree)
- Each trace server is assigned a range of id (those id are transparent for the Trace Server)

$$step = \frac{INT32\_MAX}{NUMBER\_OF\_TRACE\_SERVER}$$

$$range = [step * traceServerId, step * (traceServerId + 1)[$$

Visualisation

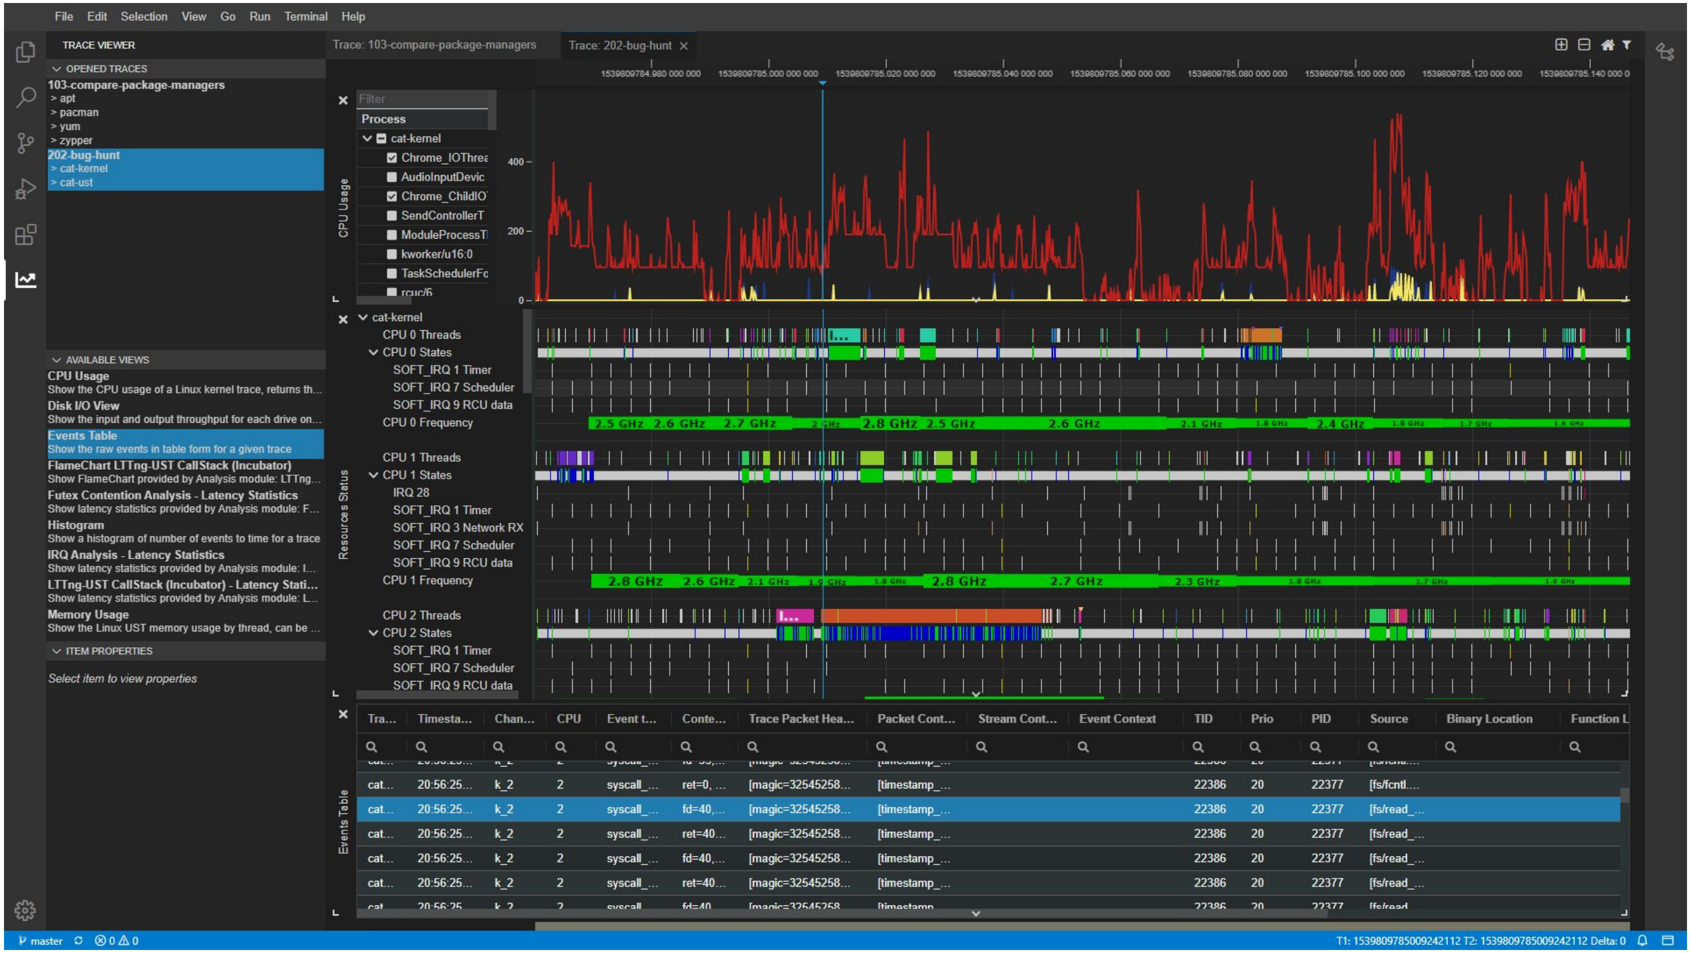

TRACE VIEWER

OPENED TRACES

103-compare-package-managers
> apt
> pacman
> yum
> zypper
202-bug-hunt
> cat-kernel
> cat-ust

AVAILABLE VIEWS

CPU Usage
Show the CPU usage of a Linux kernel trace, returns th...
Disk I/O View
Show the input and output throughput for each drive on...
Events Table
Show the raw events in table form for a given trace
FlameChart LTTng-UST CallStack (Incubator)
Show FlameChart provided by Analysis module: LTTng...
Futex Contention Analysis - Latency Statistics
Show latency statistics provided by Analysis module: F...
Histogram
Show a histogram of number of events to time for a trace
IRQ Analysis - Latency Statistics
Show latency statistics provided by Analysis module: I...
LTTng-UST CallStack (Incubator) - Latency Stati...
Show latency statistics provided by Analysis module: L...
Memory Usage
Show the Linux UST memory usage by thread, can be ...

ITEM PROPERTIES

Select item to view properties

Tra...	Timesta...	Chan...	CPU	Event t...	Conte...	Trace Packet Hea...	Packet Cont...	Stream Cont...	Event Context	TID	Prio	PID	Source	Binary Location	Function L
cat...	20:56:25...	k_2	2	syscall_...	ret=0, ...	[magic=32545258...	[timestamp_...			22386	20	22377	[fs/fcntl....		
cat...	20:56:25...	k_2	2	syscall_...	fd=40,...	[magic=32545258...	[timestamp_...			22386	20	22377	[fs/read_...		
cat...	20:56:25...	k_2	2	syscall_...	ret=40...	[magic=32545258...	[timestamp_...			22386	20	22377	[fs/read_...		
cat...	20:56:25...	k_2	2	syscall_...	fd=40,...	[magic=32545258...	[timestamp_...			22386	20	22377	[fs/read_...		
cat...	20:56:25...	k_2	2	syscall_...	ret=40...	[magic=32545258...	[timestamp_...			22386	20	22377	[fs/read_...		

T1: 1539809785009242112 T2: 1539809785009242112 Delta: 0

- Client that implement TSP (Theia - Trace Viewer Extension)
- Experiments must be created beforehand

Future work

- How to connect the traces
- Aggregation over time (Horizontally)
- Follow dependencies between processes (Critical Path)*

*Distributed computation of critical path (from Pierre-Frédérick Denys)

Conclusion

- Extends parallelisation to all kinds of analysis
- The Trace Coordinator supports the Trace Server Protocol (TSP)

Github Repository:

- tsp-java-client: Client side implementation, in Java, of the Trace Server Protocol
- Trace Coordinator: Trace Coordinator Project implement in Java

Q&A

Thank you for listening!