

LTTng & Trace Compass and Open Tracing API

Naser Ezzati, Polytechnique

May 2017

AGENDA

1 Introduction

2 A detailed web request analysis

3 Open Tracing API

4 Examples

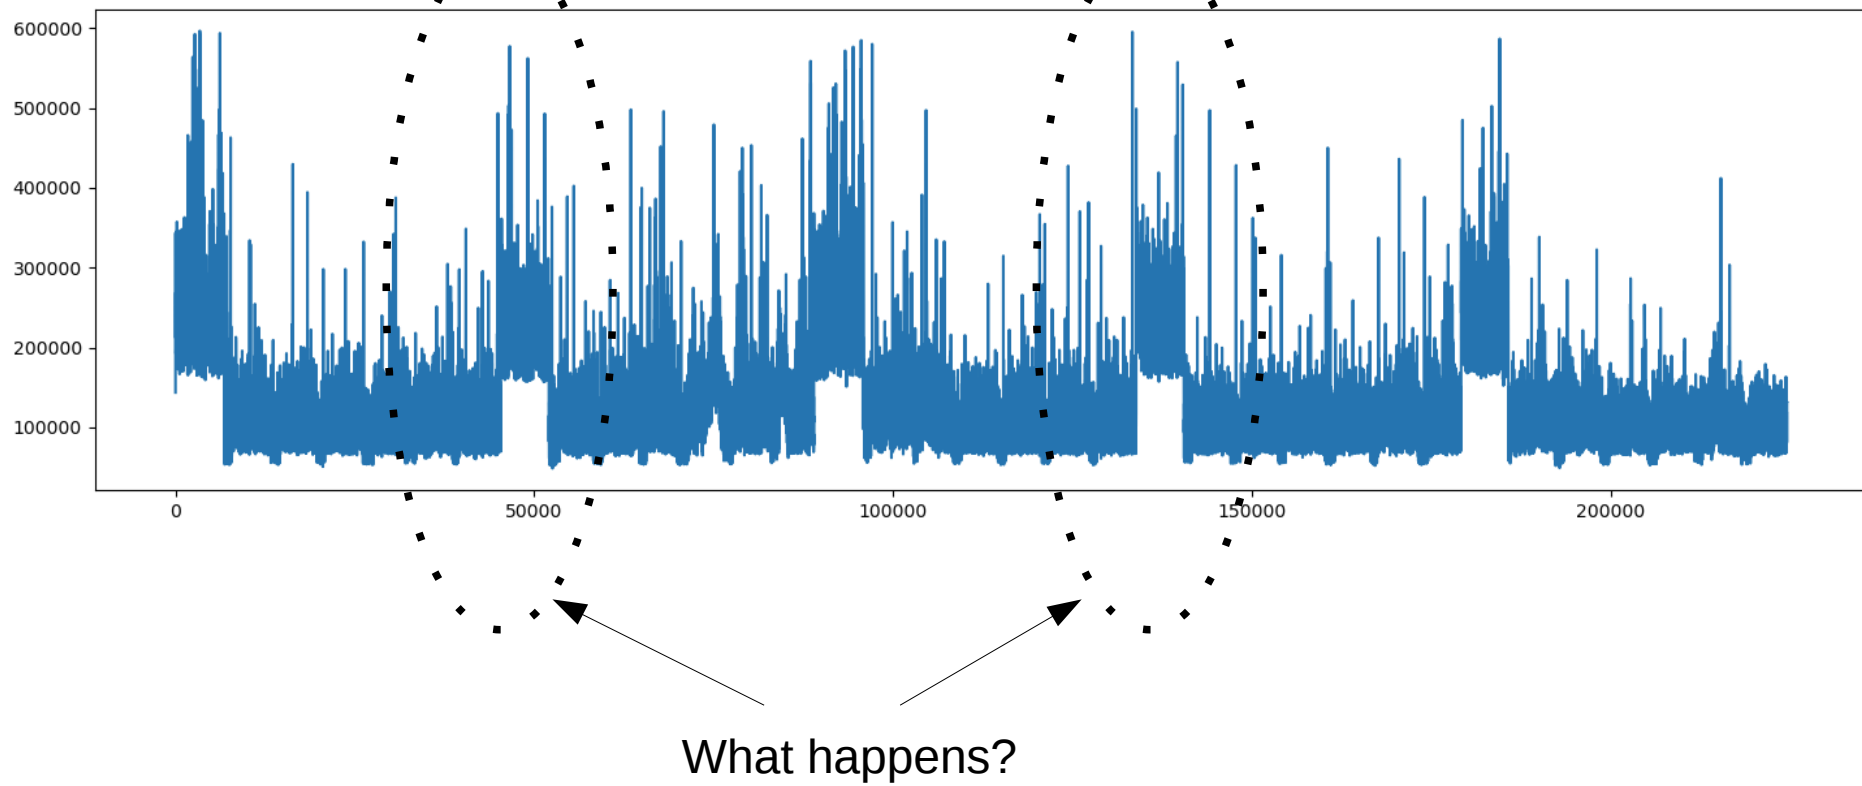
5 Conclusion

Recent Updates

- › Last time: LAMP Stack Instrumentation
- › APACHE,, PHP, MySQL
- › PHP (standard extension)
- › NGINX
- › PostgreSQL,
- › MongoDB (coming)
- ›
- › Linked Analysis
- › Usecase!

Example:

Response Time of a Web Application




There are slow-downs in the speed every few minutes. What are the reasons behind?!

PHP Request Anatomy

- PHP Is a scripting language
 - compiles any file you ask it to run, obtain OPCodes from compilation, run them, and trash them away immediately.
- Parse, compile, execute, forget
- Parse, compile, execute, forget
- Parse, compile, execute, forget
- ...


PHP "forgets" everything it's done in request N-1, when it comes to run request N.

- Even if it calls the same scripts several times.

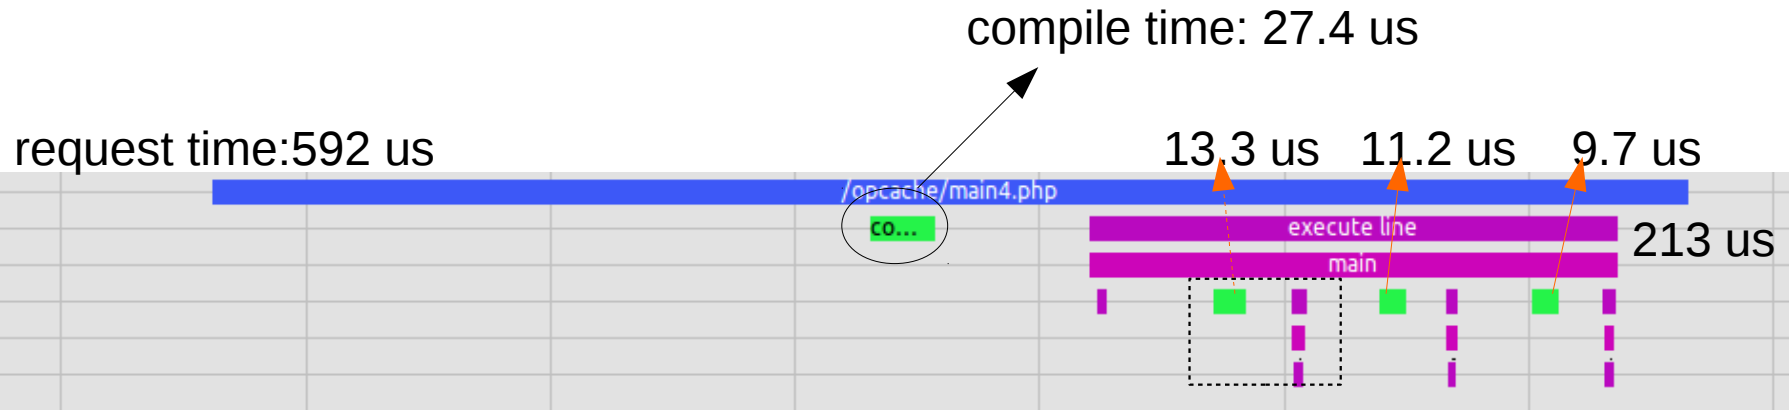


PHP Request Anatomy (2)

- Which one is the longest?
 - It depends!
 - Let's see what trace data gives us.

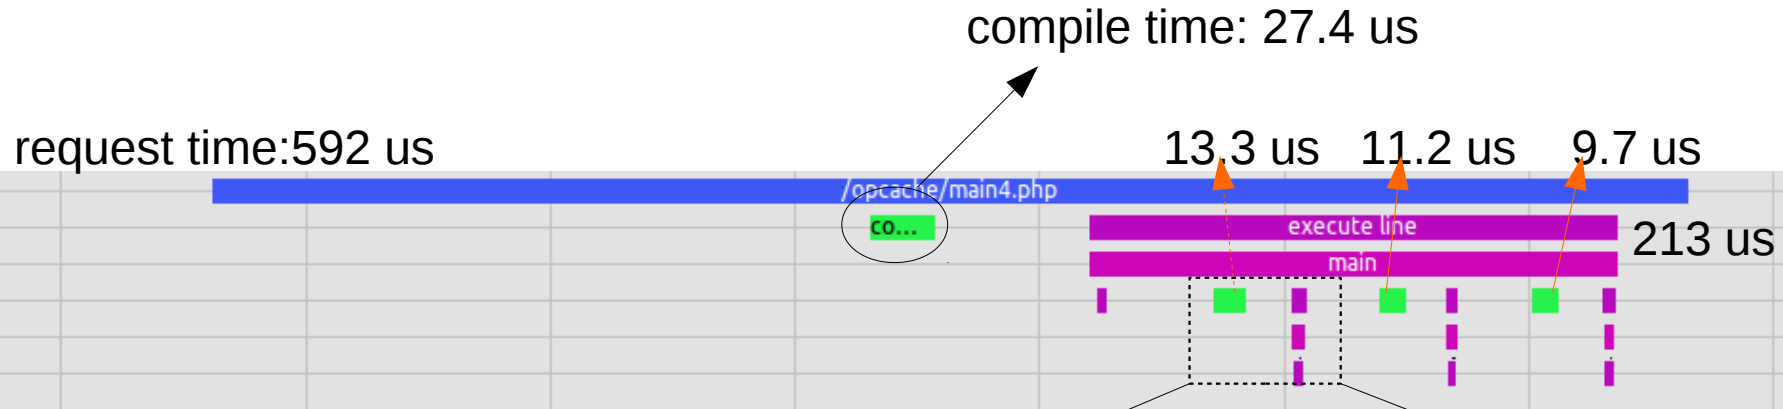


Compile Time: UST Events




```
1 <?php /*main4.php*/  
2  
3 $x = rand(0,1000);  
4  
5 echo $x.PHP_EOL;  
6 $xy = 123;  
7 include 'folder1/.$xy.'.php';  
8 include 'folder2/.$xy.'.php';  
9 include 'folder3/.$xy.'.php';  
10  
11 echo $x.PHP_EOL;  
12 ?>
```

Compile Time: UST Events



```
1 <?php /*main4.php*/  
2  
3 $x = rand(0,1000);  
4  
5 echo $x.PHP_EOL;  
6 $xy = 123;  
7 include 'folder1/.$xy.'.php';  
8 include 'folder2/.$xy.'.php';  
9 include 'folder3/.$xy.'.php';  
10  
11 echo $x.PHP_EOL;  
12 ?>
```



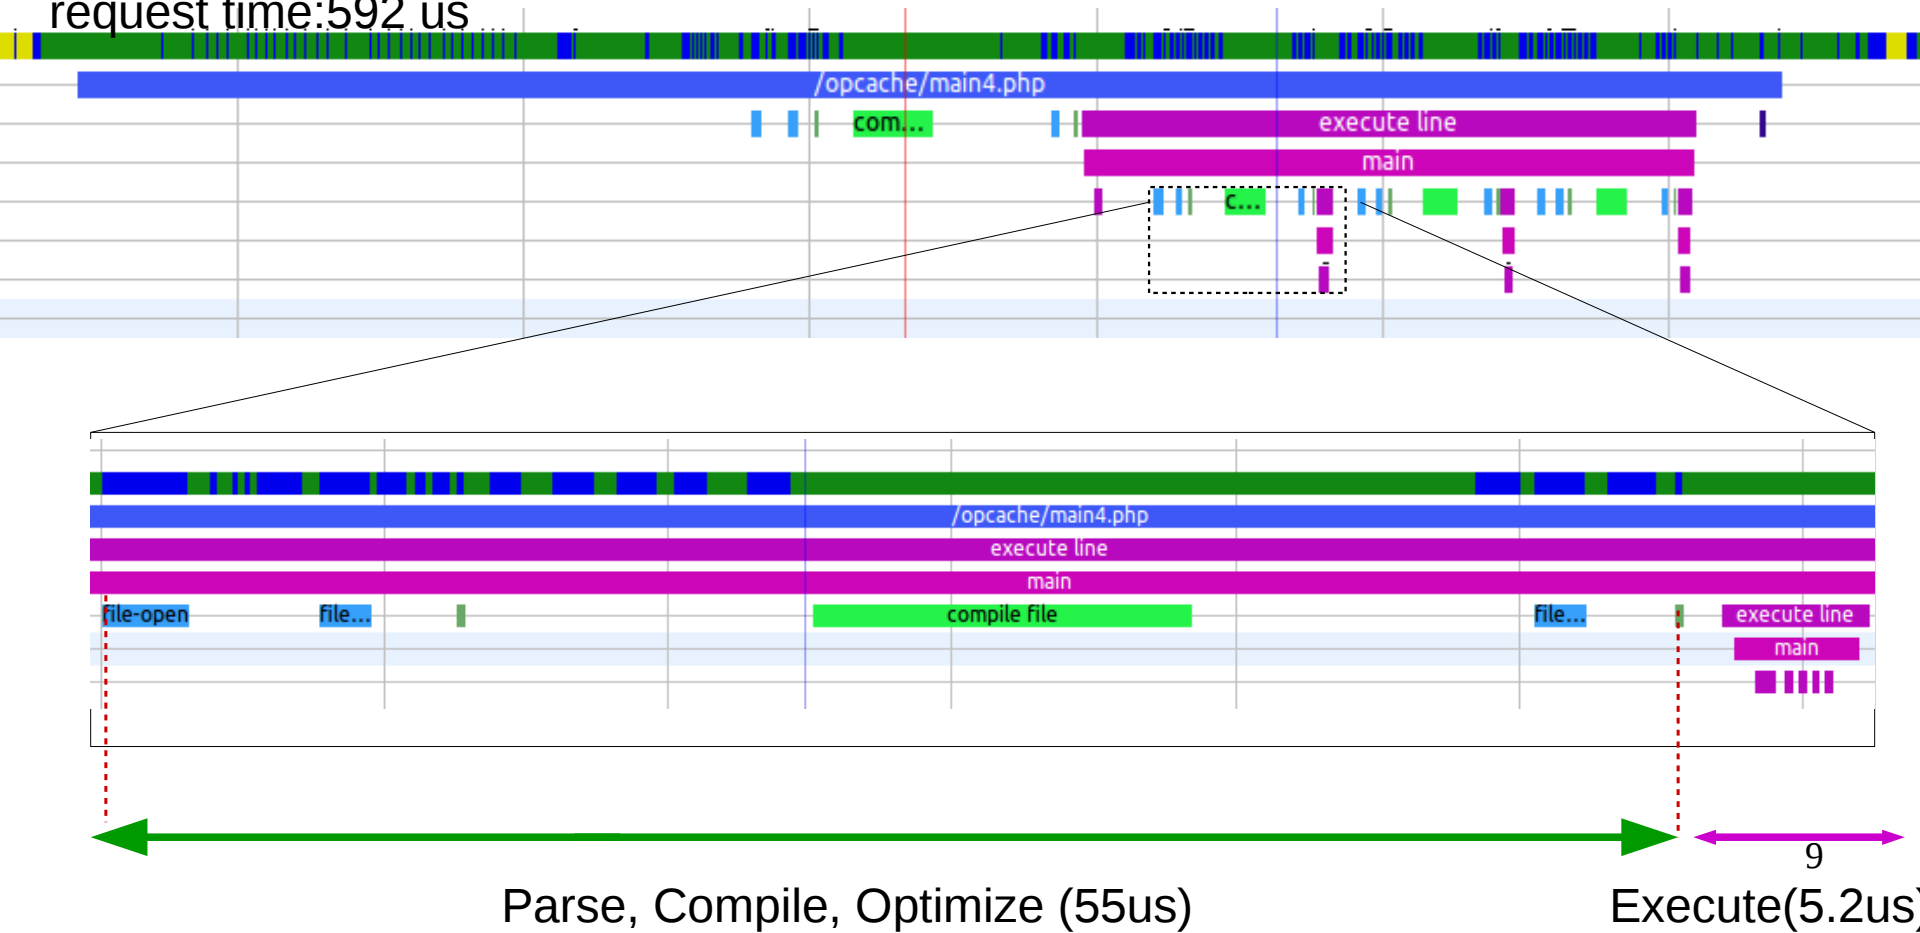
Compilation time: 60us (~ 10 % of the request time)

But, let's go deeper!


Kernel + UST Events

Now looks much more!


request time: 592 us



Solution: Opcode Cache (Opcache)

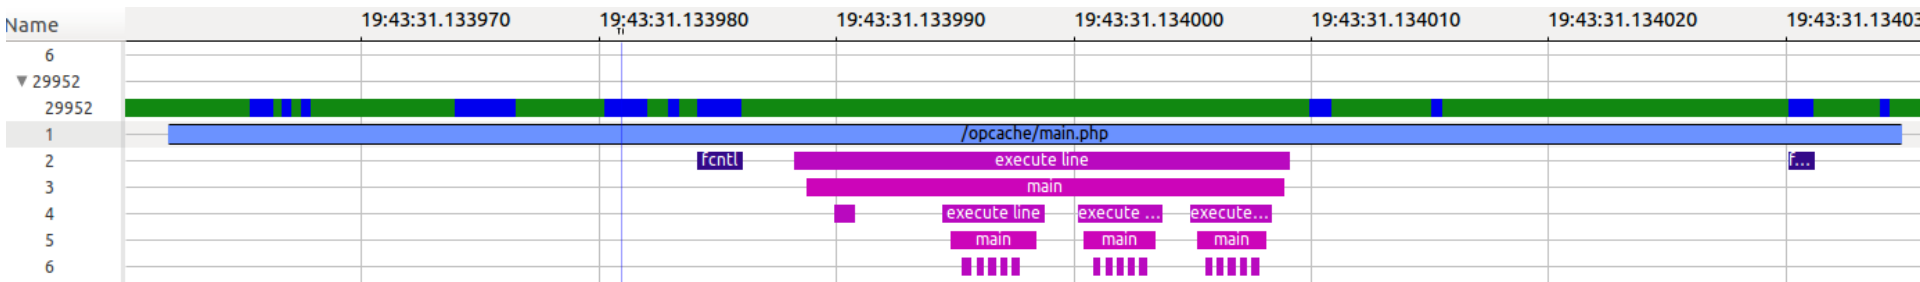


Cache at first run





Load from cache after

Solution: Opcode Cache (Opcache)





Loads from cache
Request time: 73 us
was: 592 us


Compile Time: Drupal



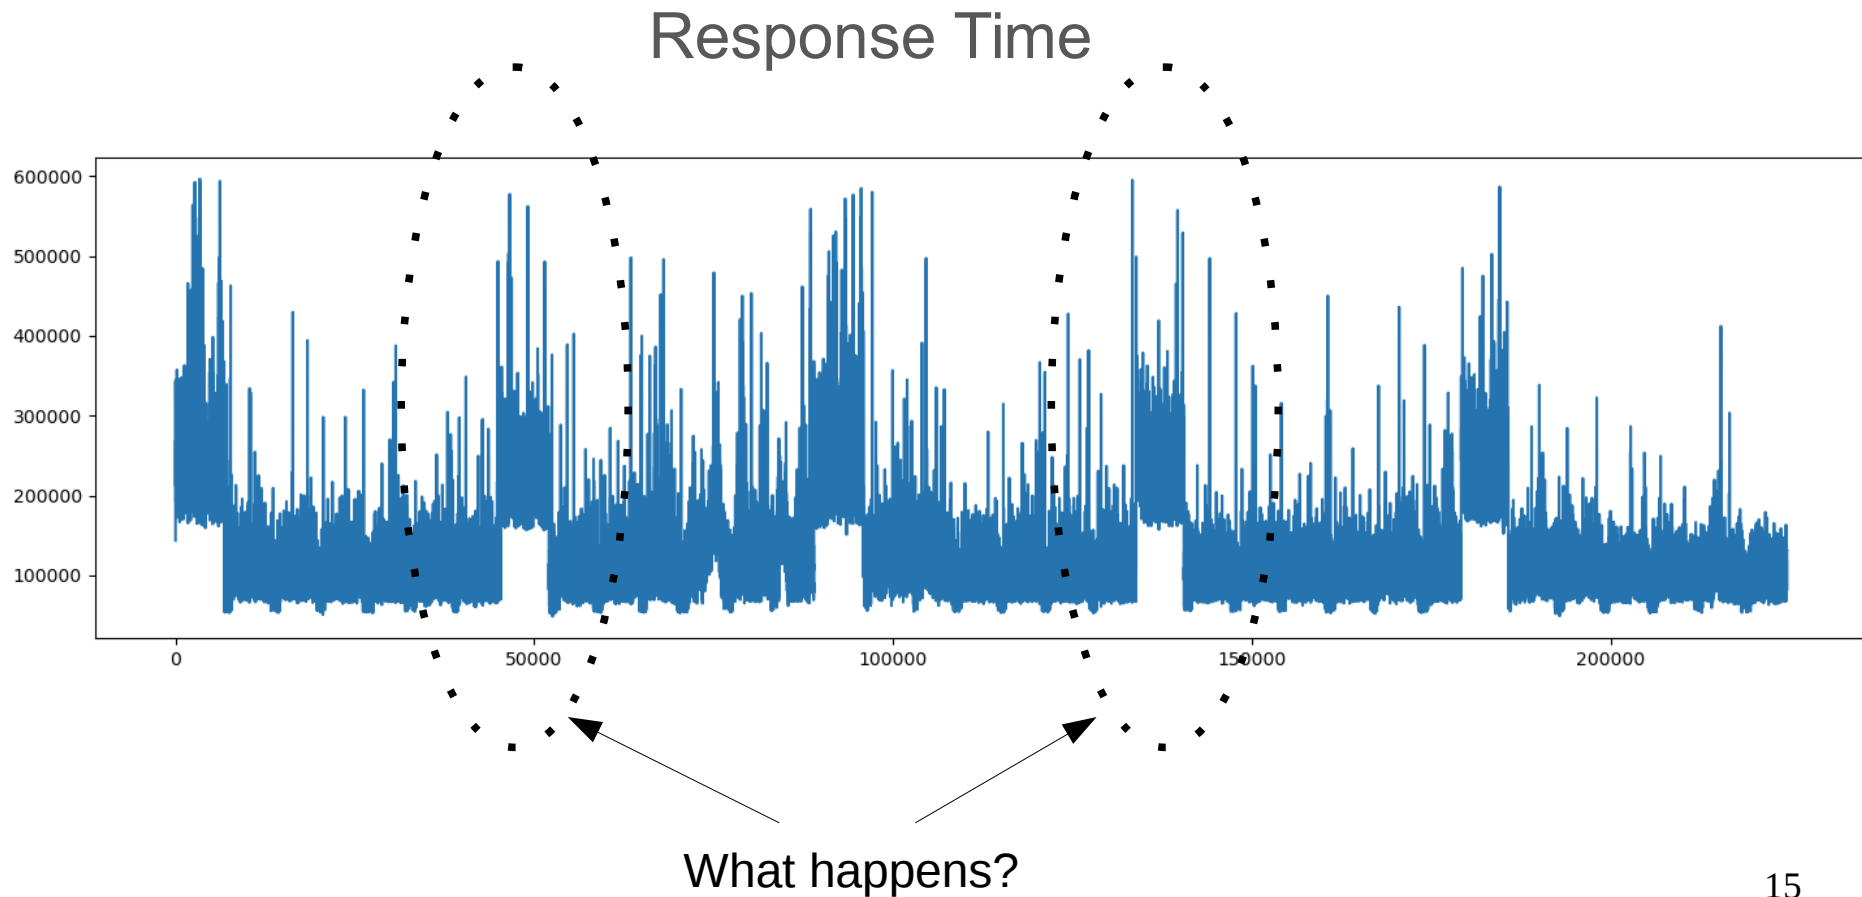
Compile Time: WordPress




Compile Time: MediaWiki


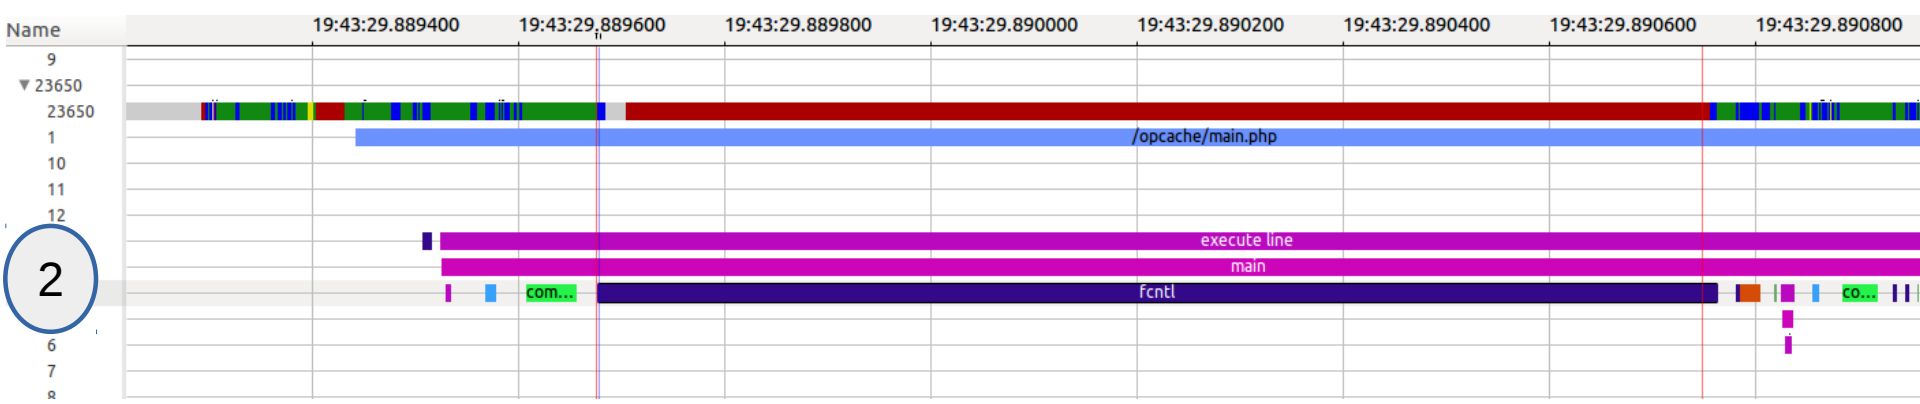
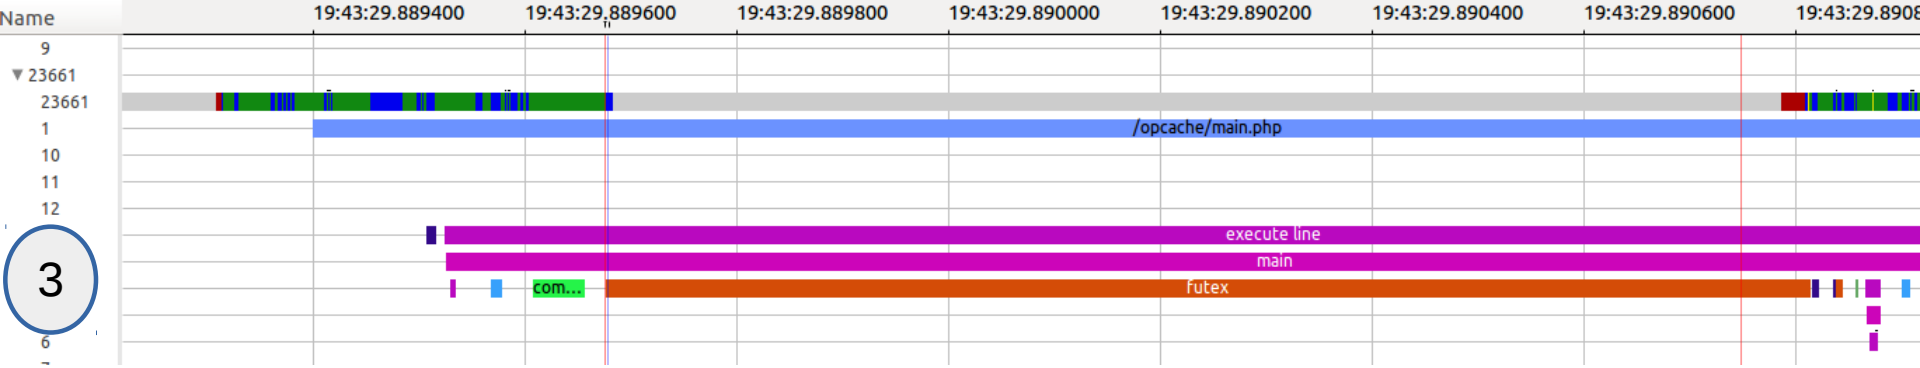


Let's back to our example




Response Time vs Compile Time






Critical Flow



Every PHP process that is willing to write into shared memory will lock every other process willing to write into shared memory as well.

Example2: Drupal Website



Not compilation issue anymore! What causes those latencies?

Php itself?

Database query?

The web server?


A background process (or another vm) in the system?

A problem in the network?


...

Maybe, we need more data?

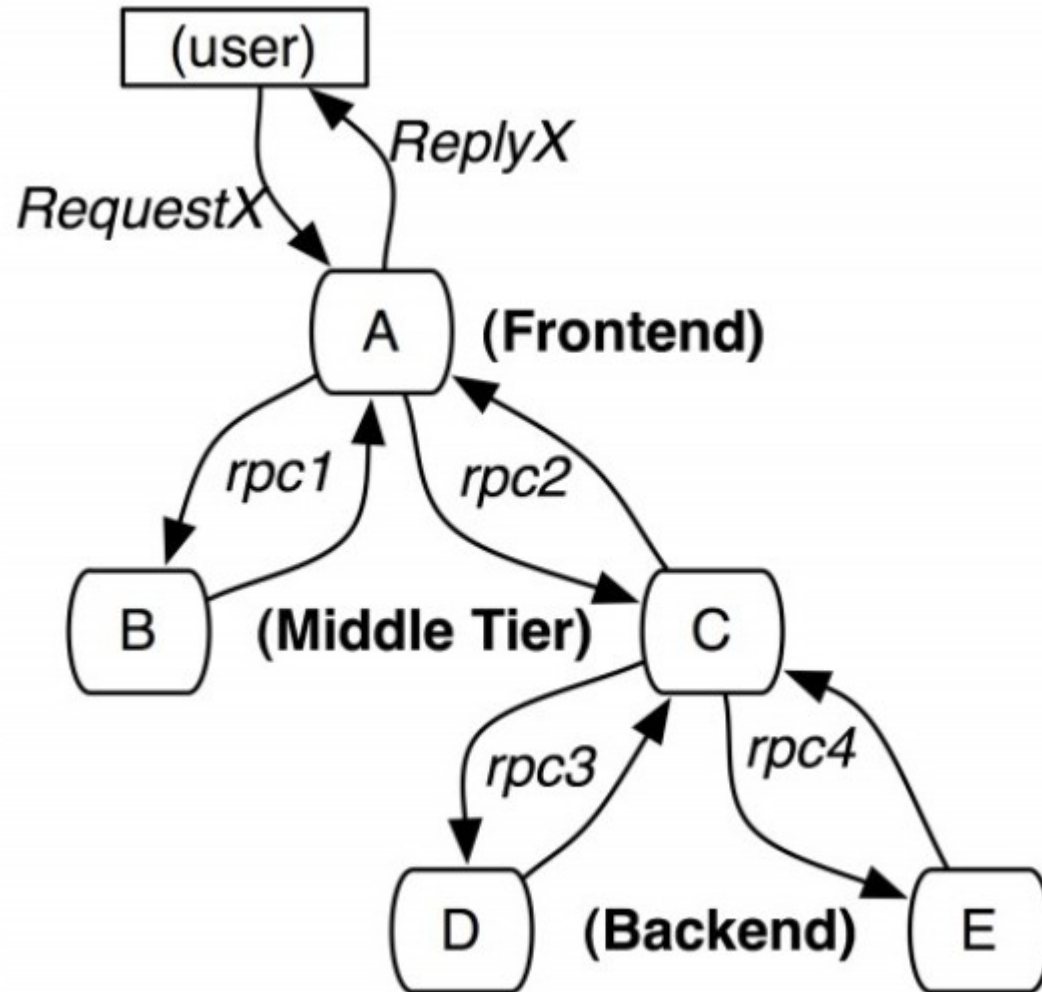
We trace all components and get more data !



How can I only follow a specific?



A General Case: Latency in a Distributed System



Causal Relationship (causal flow of control)

The image shows a terminal window with a black background and white text. The text consists of multiple lines of log entries, each starting with a timestamp and a process name. Several lines are highlighted with red boxes, and others with blue boxes. Blue arrows point from the red boxes to the blue boxes, indicating a causal relationship or flow of control between the events. The log entries include various system messages, such as 'http://internal.brn-travelplus.com', 'http://internal.brn-travelplus.com/...', and 'http://internal.brn-travelplus.com/...'. The text is dense and difficult to read due to the small font size and the large number of lines.

OpenTracing API


- The OpenTracing project provides a standard, portable API for distributed tracing instrumentation.
 - Ver 1.0 announced in Aug 2016.
- Consistent semantics across languages
 - Makes it easy for developers to add (or switch) tracing implementations with an $O(1)$ configuration change

Uses for OpenTracing

- Logging - Easy to output to any logging tool.
- Metrics/Alerting - Measure based on tags, span timing, log data.
- Critical Path Analysis - Drill down into request latency in very low granularity.
- Context Propagation - Use baggage to carry request and user ID's, etc.
- System Topology Analysis - Identify bottlenecks due to shared resources.

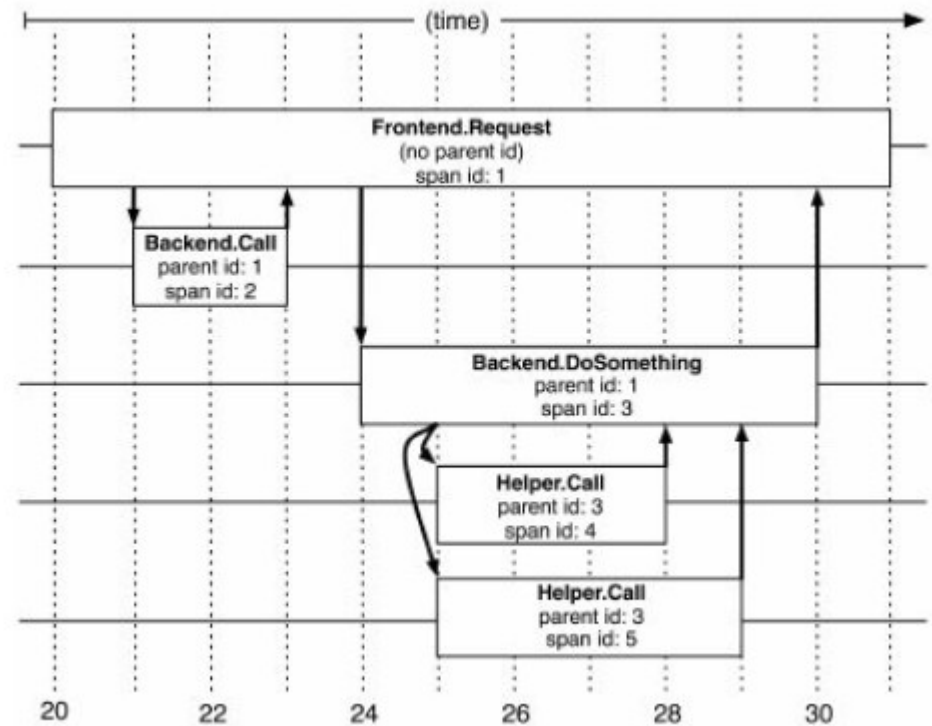
Trace Definition in OpenTracing

- At the highest level, a trace tells the story of a transaction or workflow as it propagates through a (potentially distributed) system.
- In OpenTracing, a trace is a directed acyclic graph (DAG) of "spans"
 - Named, timed operations representing a contiguous segment of work in that trace
 - Each component in a distributed trace will contribute its own span or spans




Data Model

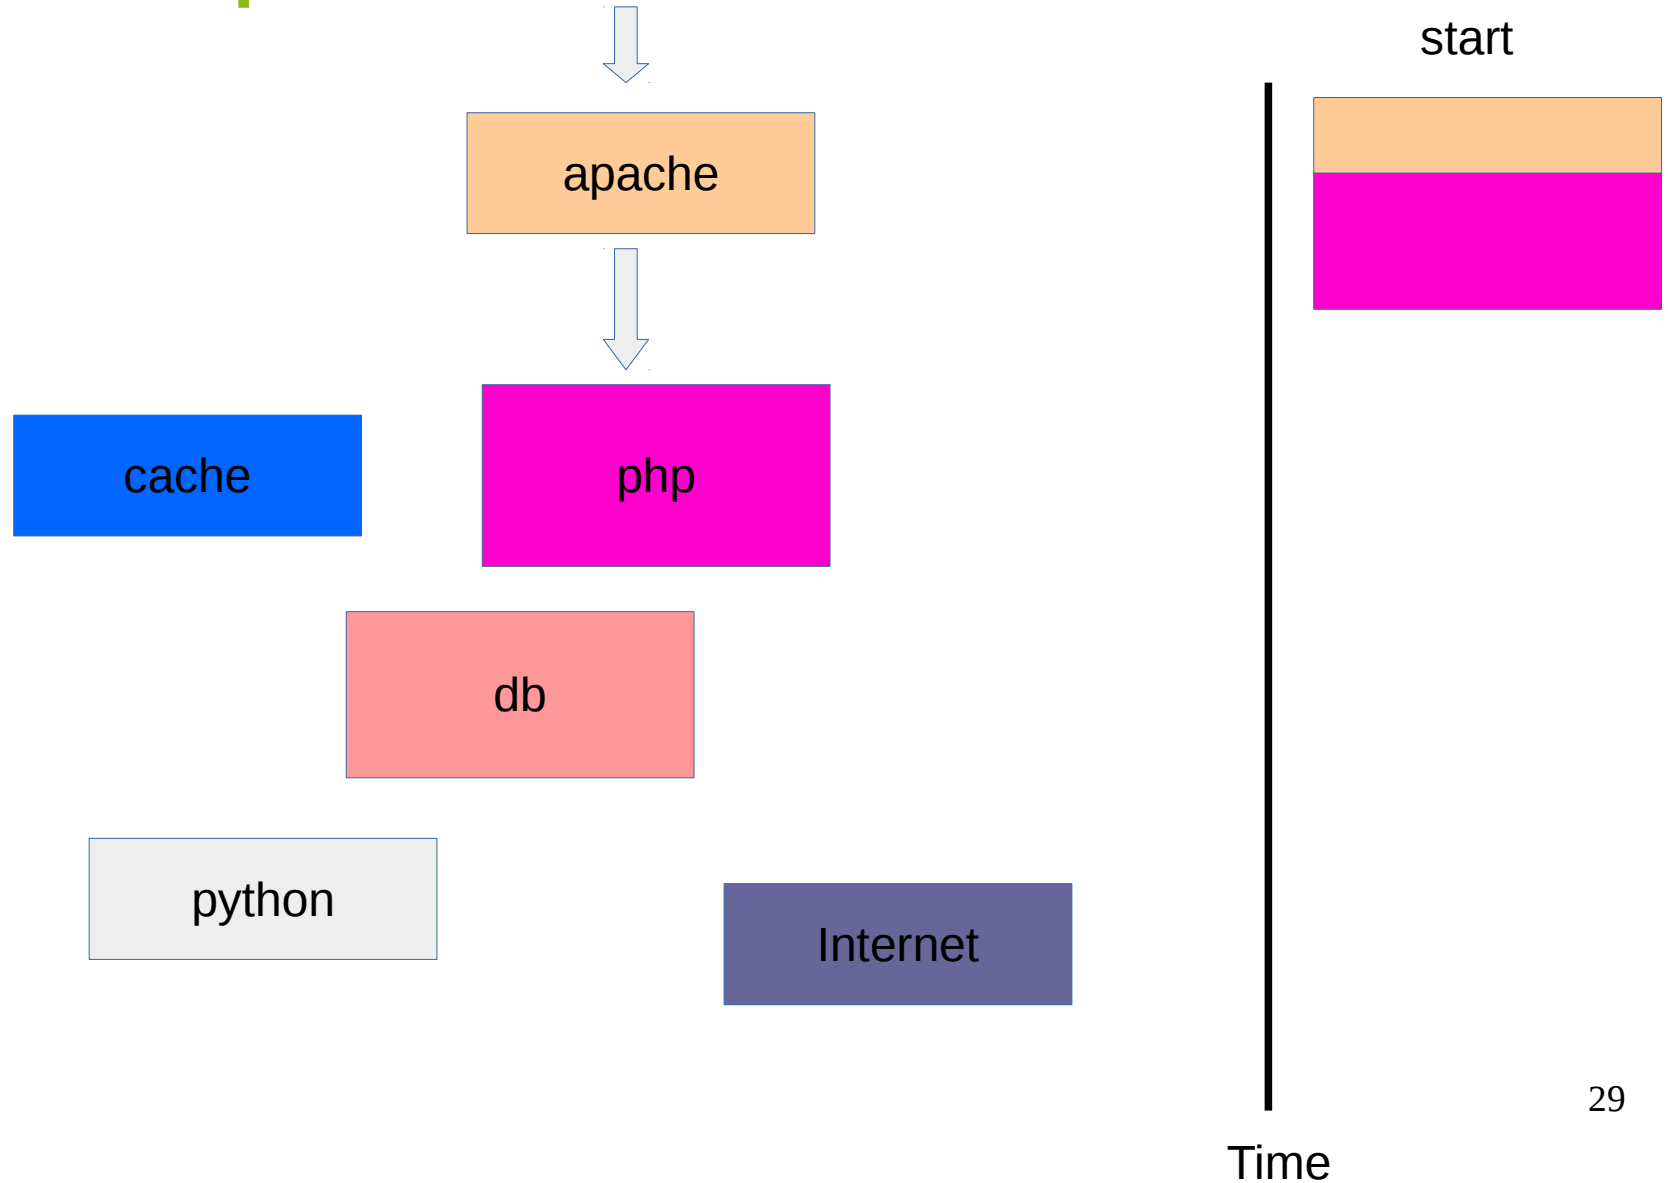
- Common libraries for several programming languages
 - Javascript
 - Python
 - Java
 - C
 - C#
 - Go
 - Ruby
 - PHP
- Trace as a tree of nested calls
 - Spans
 - Trace trees



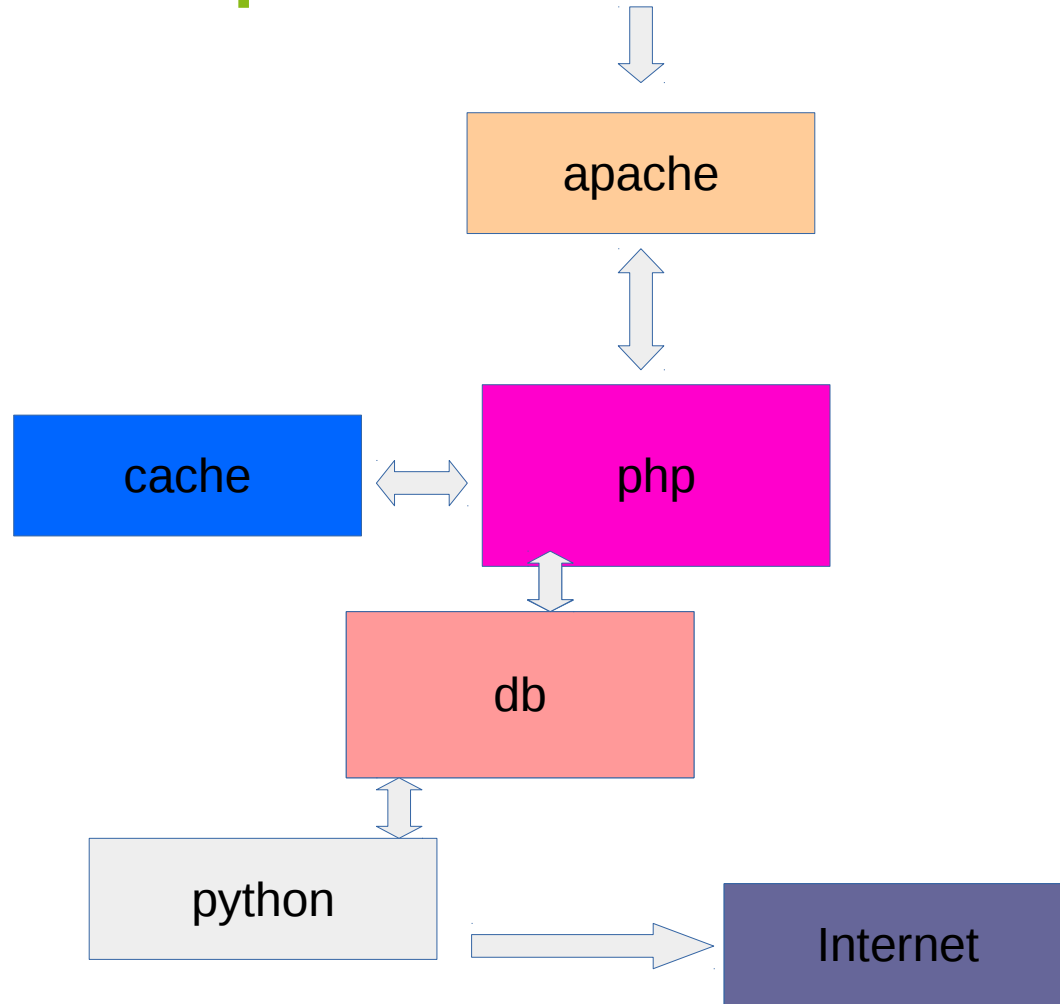
Causal Flow of Control



Example: Causal Flow of Control



Example: Causal Flow of Control



Timeline View

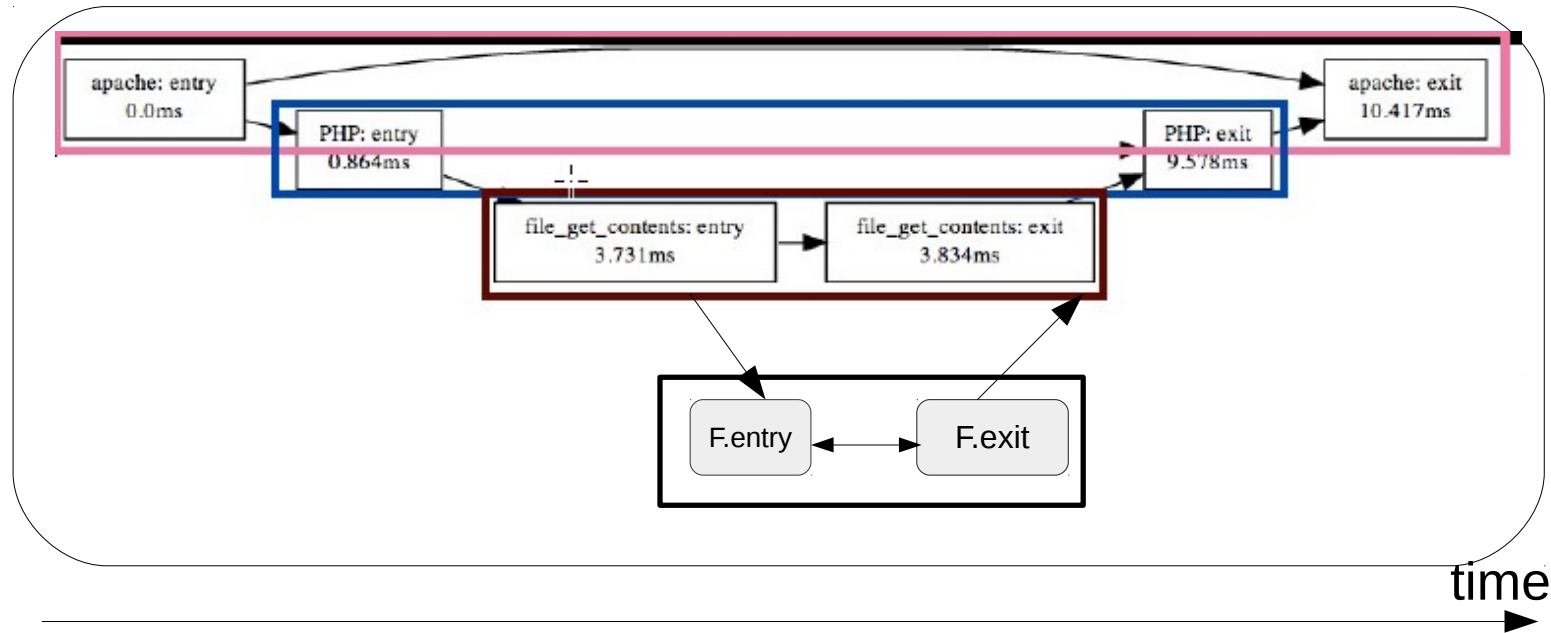
span1

span2

span3

...

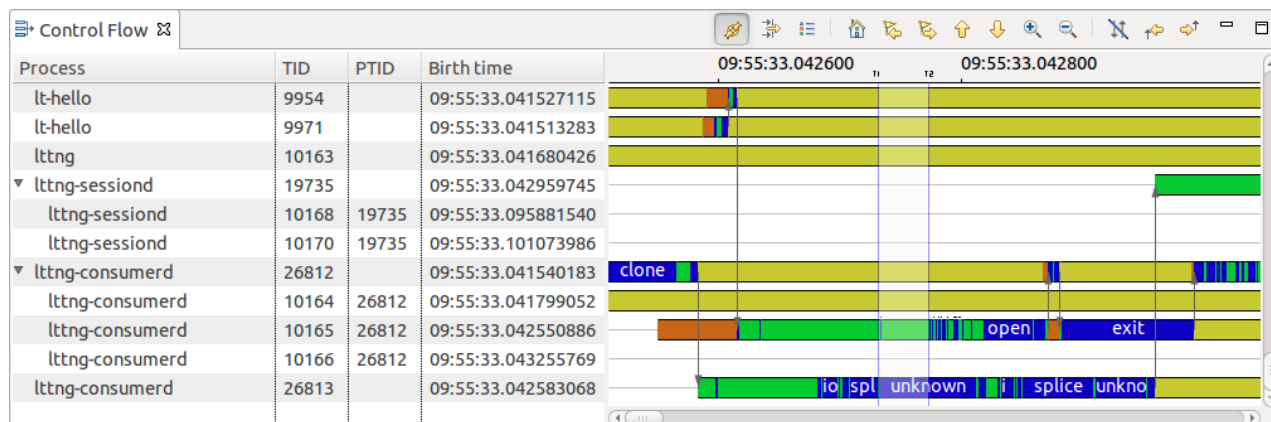
span n




- This type of visualization adds the **context of time**, the **hierarchy** of the services involved, and the serial or parallel nature of the process/task execution.
- This view **highlights the system's critical path**. By focusing on the critical path, attention can focus on the area of code where the most valuable improvements can be made.
 - For example, you might want to trace the resource allocation spans inside an API request down to the underlying blocking calls.

- Can we have this in LTTng and Trace Compass?
 - How to relate the traces/events/spans?
 - How to visualize them?



- We already have some notions of LINK in trace compass.
 - Arrows in control flow diagram
 - Arrows: scheduler switches from one process to another for a given CPU
 - Critical path analysis
 - Call stack
 - Function calls
 - **Why not link everything?**



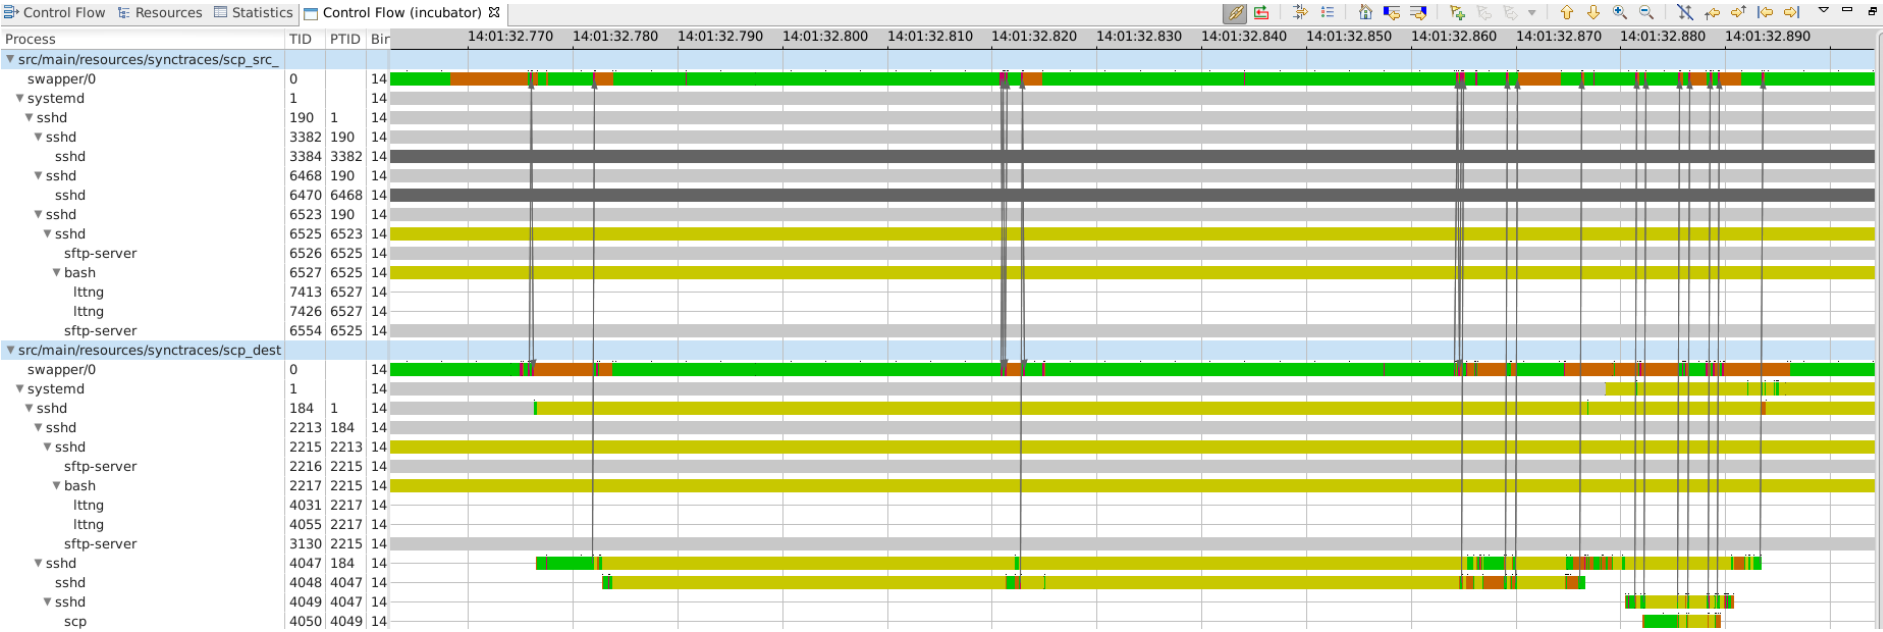
Example



Example (cntd)

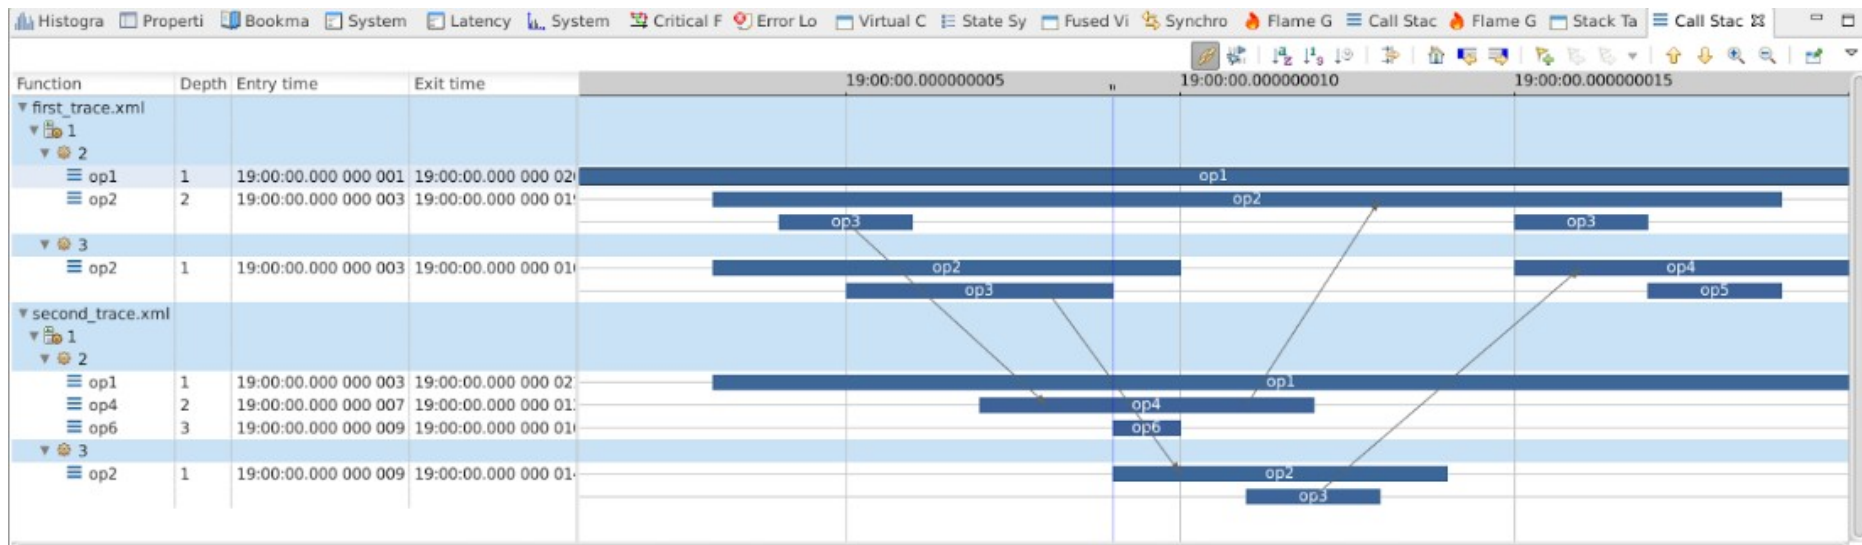


Linked Traces & Linked Events => Linked logics



Conclusion

- EventMatching Class in Trace Compass
 - Manual
 - Hardcoded links between event
 - Event a from trace 1 matches to event b from trace 2
- We need a standard way to define the semantic relationships
 - To define the events and fields that match to each other
 - Everyone can define and trace the links they need in the logics of their code/application/system
 - The standard can be Opentracing Api
- In LTTng:
 - To support the above standard
- In Trace Compass
 - A generic Event Matching class to analyse and visualize the links between events
 - View!



Q&A